

# External Attention Assisted Multi-Phase Splenic Vascular Injury Segmentation with Limited Data

Yuyin Zhou, David Dreizin, Yan Wang, Fengze Liu, Wei Shen, and Alan L. Yuille, *Member, IEEE*

**Abstract**—The spleen is one of the most commonly injured solid organs in blunt abdominal trauma. The development of automatic segmentation systems from multi-phase CT for splenic vascular injury can augment severity grading for improving clinical decision support and outcome prediction. However, accurate segmentation of splenic vascular injury is challenging for the following reasons: 1) Splenic vascular injury can be highly variant in shape, texture, size, and overall appearance; and 2) Data acquisition is a complex and expensive procedure that requires intensive efforts from both data scientists and radiologists, which makes large-scale well-annotated datasets hard to acquire in general.

In light of these challenges, we hereby design a novel framework for multi-phase splenic vascular injury segmentation, especially with limited data. On the one hand, we propose to leverage external data to mine pseudo splenic masks as the spatial attention, dubbed *external attention*, for guiding the segmentation of splenic vascular injury. On the other hand, we develop a *synthetic phase augmentation* module, which builds upon generative adversarial networks, for populating the internal data by fully leveraging the relation between different phases. By jointly enforcing external attention and populating internal data representation during training, our proposed method outperforms other competing methods and substantially improves the popular DeepLab-v3+ baseline by more than 7% in terms of average DSC, which confirms its effectiveness.

**Index Terms**—splenic vascular injury segmentation, multi-phase CT, generative adversarial networks, attention

## I. INTRODUCTION

THE spleen is the most commonly injured abdominal organ after blunt trauma [1]–[3]. Associated splenic vascular injuries are risk factors for exsanguination, hemodynamic instability, and death. CT is the first-line imaging modality to

This work was supported by University of Maryland Accelerated Translational Incubator Pilot (ATIP) Grant and NIH K08 EB027141-01A1.

Y. Zhou is with the Department of Computer Science and Engineering at University of California, Santa Cruz, CA 95064 (e-mail: zhouyuyiner@gmail.com).

D. Dreizin is with the Dept of Radiology R Adams Cowley Shock Trauma Center and University of Maryland, Baltimore, MD 21201, USA (e-mail: daviddreizin@gmail.com).

Y. Wang is with Shanghai Key Laboratory of Multidimensional Information Processing, East China Normal University, Shanghai 200241, China (e-mail: wyanny9@gmail.com).

F. Liu and A.L. Yuille are with the Department of Computer Science, the Johns Hopkins University, Baltimore, MD 21218, USA (e-mail: {liufz13, alan.l.yuille}@gmail.com).

W. Shen is with MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, Shanghai, China, (e-mail: shenwei1231@gmail.com).




Fig. 1. Visual comparison of arterial and venous images (after alignment) as well as the manual segmentation of the splenic vascular injury. The injury volumes are different across phases.

screen for splenic injuries in this setting [4]. Vascular lesions include pseudoaneurysm and active bleeding. The degree of injury burden correlates with the need for interventions, namely angioembolization and splenectomy [3], [5]. It is therefore important to be able to localize and quantify the volume of vascular injury on CT [6]. Contrast kinetics of contained pseudoaneurysms (PSAs) follow the blood pool, and PSAs are most conspicuous on arterial phase images. These are typically either faint or imperceptible within the enhancing spleen on the portal phase. On the other hand, active arterial bleeding manifests as flame-shaped foci of hemorrhage that increase in size and decreases in density across phases [7]. Automated quantification of splenic vascular lesions would result in objective clinical imaging data for prognostication and personalized decision support to predict mortality, transfusion requirement, and need for intervention in this setting. Therefore, in this paper, our goal is to automatically segment the splenic vascular injury from multi-phase (*i.e.*, arterial and venous phase) images (examples illustrated in Fig. 1).

In recent years, deep learning has largely advanced the field of computer-aided diagnosis, especially in the field of biomedical image segmentation [8] [9] [10] [11]. However, we have identified several limitations in the current literature. Firstly, for challenging structures that show large inter-patient variation in terms of attributes such as shape, texture or size, even state-of-the-art architectures may yield less satisfactory performance in spite of their good representational power [12]. Secondly, deep learning based approaches are data-driven and therefore require a large and representative training set which consists of medical images as well as their corresponding voxel-wise label maps. However, in the cross-sectional medical imaging domain (particularly CT and MRI) where both images and annotations are expensive to acquire, rarely do we have such a perfectly-sized dataset to train a segmentation

model [13]. Data and labeling scarcity can be especially severe for applications such as splenic vascular injury segmentation, since collecting high-quality annotations from specially designated high-volume academic trauma centers can be even more challenging. More critically, these two challenges can potentially add to the difficulty of each other.

To tackle these challenges, in this paper, we present a novel framework for multi-phase splenic vascular injury segmentation. Specifically, our method consists of: 1) an external attention assisted segmentation model, where we fully exploit external data for mining the spatial attention from the affiliated target (*i.e.*, the spleen), referred to as **external attention**, for guiding the segmentation of splenic vascular injury; and 2) an internal **synthetic phase augmentation** module, where we populate the internal data representation by leveraging the relation between arterial and venous phases. By cooperatively enforcing external attention and augmenting the internal data representation, our method effectively addresses challenges arising from the large inter-patient variation of the splenic injury and data scarcity simultaneously.

In more concrete terms, the external attention assisted segmentation model aims to handle the variation of the discriminative features of splenic vascular injury, by allowing the network to focus on target structures of interest and neglect feature responses from more remote and irrelevant background regions. To achieve this goal, existing approaches apply self-attention mechanisms (*e.g.*, non-local operators) which exclusively rely on internal feature representations of the given target [14], [15]. We hereby provide a novel perspective with the introduction of external attention, wherein the attention is extracted from external data with available supervision from the normal spleen segmentation to guide the relatively harder splenic vascular injury segmentation. Mining external data has been known as a popular solution to address data scarcity [13]. This strategy has started to attract more and more research attention recently since it demonstrates great potential for boosting model performance and generalization [16]. Here our goal is to fully exploit external data not only for expanding the size of the training data but rather to offer attentive knowledge for guiding the detection of splenic vascular lesions.

Inspired by the high relevance between the location of a spleen and the associated injured regions (a splenic vascular injury resides either within or near the spleen), we formulate the unknown spleen<sup>1</sup> location as a latent variable for deriving the attention to facilitate the following training process. We note the integration of our internal splenic vascular injury dataset and the external dataset yields a partially-supervised setting, in that each set only has either the spleen or the injury labeled, but not both. An intuitive approach is to impute the missing category with pseudo-labels [17] and integrate both sets in a joint training paradigm. Nonetheless, this approach only addresses the label discrepancy between the internal and the external set without considering the correlation between different classes, *e.g.*, a vascular injury (one class) often resides close to and emanate from the associated spleen (another

<sup>1</sup>unknown spleen throughout this paper refers to the whole spleen on the internal splenic injury dataset, which is unknown due to that the spleen class is not labeled.

class). To fully exploit the spatial relationship between the spleen and its associated injury, we re-formulate the predicted splenic mask not only as pseudo-supervision but also to act as the spatial attention mask to reweight the importance of different voxels during training. Specifically, during the learning process, the pseudo-labels on the unknown spleen locations and the segmentation network parameters are alternately updated, making the derived attention mask iteratively refined to facilitate the subsequent training iterations. This allows the network to gradually focus exclusively on relevant foreground regions, so as to largely benefit the segmentation for the splenic vascular injury [12].

To further address the limited training data, especially in the context of multi-phase data, we develop a synthetic phase augmentation module by exploiting the underlying relationship between different imaging phases. We use CycleGAN [18], to build a phase translation model, and use the learned transformations to generate synthetic phases. The real and the generated multi-phase images are then jointly trained using our proposed external attention assisted segmentation model. Consequently, the augmented synthetic examples not only greatly enrich the training set but also benefit the knowledge integration from both phases.

In this paper, we have curated a patient cohort of 55 consecutive multi-phase CT studies with splenic vascular injury annotations, which, to the best of our knowledge, is the largest available set to date for this injury type. By evaluating this dataset comprehensively, our approach can consistently outperform prior arts by a large margin. The main contributions of this paper are summarized as follows:

- We present a novel approach for multi-phase splenic vascular injury segmentation, which is a common and clinically important entity, yet has been rarely studied.
- We establish a new attention mechanism, where the attention is explicitly extracted from external data with available supervision for associated affiliated targets, referred to as *external attention*.
- We introduce *synthetic phase augmentation*, where synthetic phases are created to be jointly trained with real phases to further benefit the knowledge integration from limited training data.
- Extensive experiments conducted under various settings suggest that our approach not only significantly outperforms other competing methods for splenic vascular injury segmentation, but also well generalizes to liver tumor segmentation and pancreatic tumor segmentation.

## II. RELATED WORKS

### A. Training with External Data

To enlarge the size of the training data and populate the training distribution, one popular strategy is to employ external unlabeled data or heterogeneous labeled datasets. This requires no additional manual efforts but can generally lead to imperfect training data [13].

In terms of unlabeled data, Bai *et al.* [19] present a self-training-based method for cardiac MR image segmentation, where the network parameters and the pseudo-labels were

alternatively updated. Zhou *et al.* [20], [21] further propose to co-train multiple networks where the pseudo-labels are refined by exploiting the consensus of network predictions in the ensemble. To make the learned models more robust, consistency-based methods [22], [23] and uncertainty-driven approaches [24], [25] are proposed for different medical image classification and segmentation tasks.

For incorporating external labeled datasets, Lenga *et al.* [26] demonstrate that joint training on multiple heterogeneous datasets leads to superior performance for chest X-ray classification. To properly handle the variation of image characteristics and labeling protocols among different datasets, strategies for bridging domain gaps, imputing missing annotations, regularizing and stabilizing predictions are proposed for obtaining a universal feature representation [16], [27]–[33].

However, in this paper, we provide a different perspective, that is, we propose to utilize external data to incorporate attention for learning the affiliated target.

### B. Multi-phase Learning

Medical images acquired in different phases usually contain complementary information, hence radiologists usually rely on analyzing multi-phase data for better image interpretations. Recently, researchers have explored different strategies for combining multi-phase images, such as joint training [34]–[37], feature-level fusion [38], [39], and generative models [40]. In this paper, similar to [40], we also use generative models to create synthetic multi-phase images which can be then used for training the segmentation model.

### C. Synthetic Data Augmentation

To alleviate overfitting and improve model performance, researchers propose to use Generative Adversarial Networks (GANs) for synthetic data augmentation. Many studies such as [41]–[43] have suggested that generative adversarial networks generate convincing appearances of CT. For instance, Frid-Adar *et al.* [41] show that the synthetic lesions by GANs are meaningful in appearance, and using the generated examples as augmentation achieves a significant improvement of 7% for liver lesion classification. Sandfort *et al.* [42] demonstrate that the synthetic non-contrast images (using CycleGAN, a variant of GANs) appear convincing - even when significant abnormalities are present in the contrast CT scans. Different from these methods, here we use CycleGAN to exploit the multi-phase information and generate synthetic imaging phases to benefit the segmentation of multi-phase splenic injury. The generated images are augmented in the training set to benefit the knowledge integration from multi-phase images, especially from a limited set.

### D. Attention Mechanisms

The attention mechanism has been widely applied to many vision problems. Wang *et al.* [14] propose to model long-range relationships and design a non-local operator accordingly. In the field of medical image analysis, Zhou *et al.* [12], [44] propose a multi-stage framework where the first stage explicitly extracts the saliency region, which is then used for facilitating the segmentation in the second stage for better detection of small organs such as the pancreas. To further

reduce the computational resources and model parameters, Schlemper *et al.* [15] propose additive attention gate modules which are integrated in the skip connections. Additionally, attention modules are proposed to be used at multiple resolutions which are then fused for prostate segmentation [45].

Different from the works above, we propose to incorporate attention for the splenic vascular injury segmentation by leveraging the spatial relevance between the spleen and the injury. Therefore, the attention in our study is derived by mining the unknown spleen class with the help of additional external data.

## III. METHODOLOGY

### A. Overview

In the context of multi-phase splenic vascular injury segmentation, the data and the associated annotations can be categorized into two sets based on the imaging phases:  $\mathcal{A} = \{\mathbf{X}_A \in \mathbb{R}^{D_A}, \mathbf{Y}_A \in \mathbb{I}_{int}^{D_A}\}$  and  $\mathcal{V} = \{\mathbf{X}_V \in \mathbb{R}^{D_V}, \mathbf{Y}_V \in \mathbb{I}_{int}^{D_V}\}$ , where  $\mathbf{X}_A$ ,  $\mathbf{X}_V$  and  $\mathbf{Y}_A$ ,  $\mathbf{Y}_V$  denote the image and the annotation of dimension  $D_A$ ,  $D_V$  respectively; the subscripts A and V, denote *arterial* and *venous* phases in CT imaging protocols. Note that unlike previous studies [38], [46], [47], here the dataset is not paired, *i.e.*, the arterial phase and the venous phase of the same patient do not have identical annotations ( $\mathbf{Y}_A \neq \mathbf{Y}_V$ ).  $\mathbb{I}_{int} = \{\text{background, splenic vascular injury}\}$  denotes the labeling space of the internal multi-phase splenic vascular injury dataset. Additionally, we also introduce an external dataset  $\mathcal{S}_{ext} = \{\mathbf{X}_{ext} \in \mathbb{R}^{D_E}, \mathbf{Y}_{ext} \in \mathbb{I}_{ext}^{D_E}\}$  where the spleen supervision is available, *i.e.*,  $\mathbb{I}_{ext} = \{\text{background, spleen}\}$ . Image  $\mathbf{X}_{ext}$  and annotation  $\mathbf{Y}_{ext}$  are of the same dimension  $D_E$ . Note that the external dataset only contains normal abdominal CT scans, *i.e.*, there is no splenic vascular injury available, whereas in the internal dataset, only voxel-wise splenic vascular injury annotations are given. Therefore, our goal of using both datasets to boost splenic vascular injury segmentation also belongs to the partially-supervised setting.

Fig. 2 illustrates the overall pipeline, which can be summarized into the following two stages. In the first stage, we train a spleen segmentation model  $\mathbf{f}_s[\cdot; \theta_s]$  exclusively on the external set  $\mathcal{S}_{ext}$ , in order to obtain a good initialization of the external attention (Initial Splenic Attention Extraction). For a given input  $\mathbf{X}_\phi$ , the corresponding pseudo splenic mask  $\hat{\mathbf{Y}}_{\phi,s}$  and the attention map  $\mathbf{W}_\phi$  (of the same size as  $\mathbf{X}_\phi$ ) are first initialized by inferring on the trained model and later updated in the following stage. Here we also train a domain translation model to generate synthetic external data as augmentation. In the second stage, a phase translation model is trained to exploit the relationship between different phases, and the learned mapping functions are used to create synthetic arterial and venous phases (Synthetic Phase Augmentation). The real and synthetic phases are then jointly trained with the attention map applied on top of the loss function as spatial weights, where the network  $\mathbf{f}[\cdot; \Theta]$ , the pseudo splenic mask  $\hat{\mathbf{Y}}_{\phi,s}$ , and the attention map  $\mathbf{W}_\phi$  are alternately updated (Iterative Splenic Attention Refinement). Below, we will elaborate on each component along with technical details.

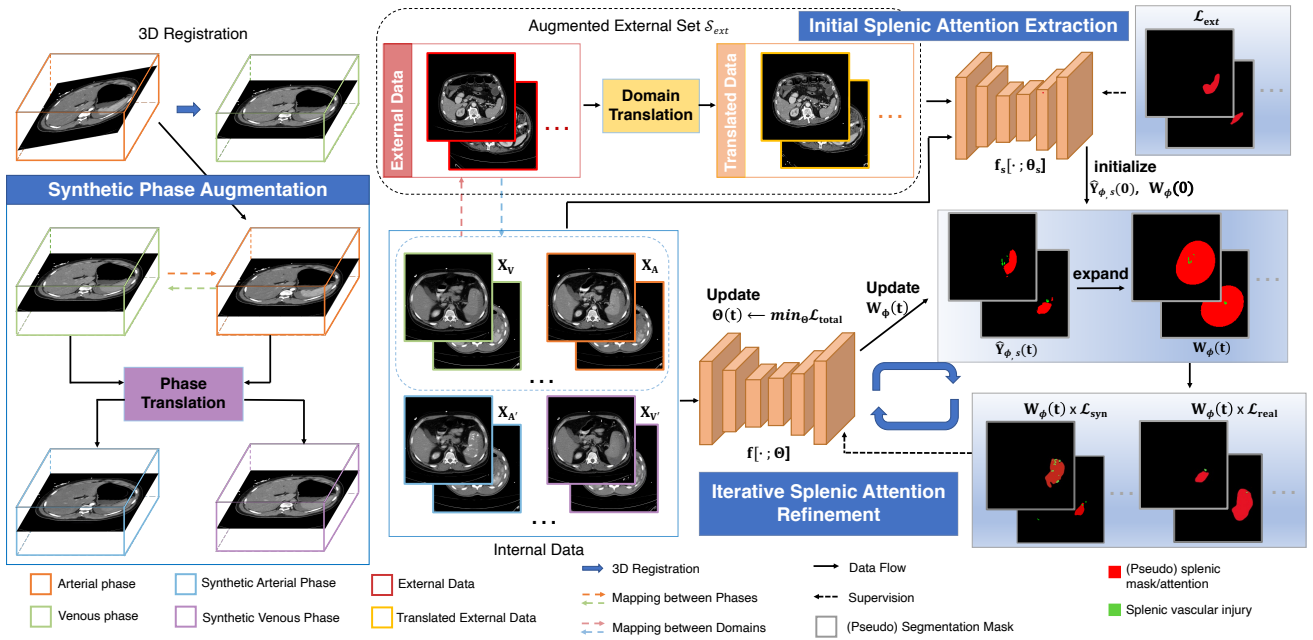


Fig. 2. The overall framework. We first train a spleen segmentation model  $f_s[\cdot; \theta_s]$  on the external set  $\mathcal{S}_{ext}$ . For a given input  $X_{\phi}$ , the corresponding pseudo splenic mask  $\hat{Y}_{\phi,s}$  and the attention map  $W_{\phi}$  are first initialized by inferring on  $f_s[\cdot; \theta_s]$ . A domain translation model is also trained for generating synthetic external data as augmentation. Secondly, a phase translation model is trained for creating synthetic arterial and venous phases. The real and synthetic phases are then jointly trained with the attention map applied on top of the losses as spatial weights, during which the attention  $W_{\phi}$  can be iteratively refined. We can see that the attention mask can well enclose the injured regions by expanding.

## B. Initial Splenic Attention Extraction

The spleen is a small organ compared to the whole abdominal region. In comparison, the splenic vascular injury is even smaller. In our internal dataset, the fraction of the splenic vascular injury, relative to the entire volume, is often much smaller than 0.005%. This largely increases the difficulty of direct segmentation or even localization from the entire CT volume, which coincides with previous findings for pancreatic cyst segmentation [48]. To deal with this problem, Zhou *et al.* [48] suggest that starting from the pancreas mask can largely increase the chance of accurately segmenting the cyst and propose to perform cyst segmentation based on the pancreatic region, which is relatively easy to detect. However, directly applying this method would require voxel-wise annotations for the spleen, which are not available in our internal dataset. Instead, we have sought to  $\mathcal{S}_{ext}$  for effective splenic attention extraction to guide the splenic vascular injury segmentation in the following stage (Sec. III-D). To bridge the gap between internal and external domains, we also train a **domain translation model** and use the translated data as augmentation. In other words, the external set  $\mathcal{S}_{ext}$  used for training the initial attention extraction model here is the union of both the external data and the translated data (see Fig. 2). The translated data are generated from the external data by training a domain translation model which learns the mapping between the internal and the external domains. The technical details for training the domain translation model are essentially the same as used in the phase translation model (see Sec. III-C).

To obtain a good feature extractor for the spleen, we first optimize the cross-entropy loss exclusively on the external set

as follows:

$$-\sum_{l \in \mathbb{L}_{ext}} \sum_j \mathbb{1}(y_{ext,l}^j) \log p_{ext,l}^j, \quad (1)$$

where  $p_{ext,l}^j$  denotes the probability of class  $l$  on the  $j$ -th voxel.

To facilitate the detection of the splenic vascular injury in the later stage (Sec. III-D), for a given input  $X_{\phi}$  ( $\phi$  is the imaging phase), we first compute the prediction for the unknown spleen class as the initial pseudo splenic mask  $\hat{Y}_{\phi,s}^{(0)}$ . Considering that the splenic vascular injury can sit outside the spleen, we further relax the predicted boundary to compute the initial splenic attention map  $W_{\phi}^{(0)}$  via a transformation function  $r(\cdot)$ :

$$W_{\phi}^{(0)} = r(\hat{Y}_{\phi,s}^{(0)}, \sigma) = \mathbb{I}[(\hat{Y}_{\phi,s}^{(0)} \star N(\sigma)) \geq \rho], \quad (2)$$

where  $N(\sigma)$  is the Gaussian kernel with the standard deviation of  $\sigma$ .  $\star$  denotes the convolution operator and  $\rho$  is the threshold for generating the attention mask.  $\mathbb{I}(\cdot)$  is the indicator function. As can be seen from Fig. 2, the transformed attention mask can successfully enclose exterior injuries through expanding.

## C. Synthetic Phase Augmentation via Generative Adversarial Networks

Due to the different imaging protocols,  $\mathcal{A}$  and  $\mathcal{V}$  demonstrate two distributions of different appearances. Meanwhile, they are also highly correlative since different imaging phases are still corresponding to the same patient. Thereby, how to incorporate multi-phase information to boost model performance has become a promising direction to explore [34]–[38], [40]. The goal is thus to train a segmentation model for both  $\mathcal{A}$  and  $\mathcal{V}$  simultaneously, so that the information in these two phases are considered in a collaborative way.




Fig. 3. Real images vs. synthetic images. Synthesized injury spots are inpainted onto the generated images.

Following [40], we use the CycleGAN model [18] which aims to learn mapping functions between domains, to build relation between the two phases, referred to as the phase translation model. Our goal here is to learn mapping functions  $g : \mathbf{X}_A \rightarrow \mathbf{X}_V$  and  $h : \mathbf{X}_V \rightarrow \mathbf{X}_A$ . Before training the phase translation model, we first perform **3D registration** where all arterial phase CT scans of the same patient are pre-registered to the venous phase using DEEDS [49]. The goal of the mapping function  $g(\cdot)$  is to generate images  $g(\mathbf{X}_A)$  indistinguishable from  $\mathbf{X}_V$  and minimise the loss function in Eqn. (3), where the discriminator  $D_V(\cdot)$  is trying to distinguish the generator's generated images and maximize the overall loss:

$$\min_g \max_{D_V} \mathcal{L}_{A \rightarrow V} = \min_g \max_{D_V} (\mathbb{E}_{\mathbf{X}_V \sim \mathcal{V}} [\log D_V(\mathbf{X}_V)] + \mathbb{E}_{\mathbf{X}_A \sim \mathcal{A}} [\log(1 - D_V(g(\mathbf{X}_A)))]), \quad (3)$$

To learn the mapping function  $h(\cdot)$ , a similar adversarial loss *w.r.t.*  $h(\cdot)$  and its discriminator  $D_A(\cdot)$  is optimized in the same manner, *i.e.*,  $\min_h \max_{D_A} \mathcal{L}_{V \rightarrow A}$ . Additionally, a cycle-consistency loss is imposed:

$$\mathcal{L}_{cyc} = \mathbb{E}_{\mathbf{X}_A \sim \mathcal{A}} [\|\mathbf{h}(g(\mathbf{X}_A)) - \mathbf{X}_A\|_1] + \mathbb{E}_{\mathbf{X}_V \sim \mathcal{V}} [\|\mathbf{g}(h(\mathbf{X}_V)) - \mathbf{X}_V\|_1]. \quad (4)$$

Therefore, the final objective function for training the phase translation model  $g(\cdot)$  and  $h(\cdot)$  is:

$$\mathcal{L}_{A \rightarrow V} + \mathcal{L}_{V \rightarrow A} + \lambda \mathcal{L}_{cyc}, \quad (5)$$

where  $\lambda$  is the hyper-parameter for balancing the loss terms.

Fig. 3 illustrates a few pairs of examples of the original image and the generated synthetic image. We can see that our phase translation model have successfully brought contrast and texture changes to the original images, which in turn make them appear more similar to images from the other phase. This means that the new texture of the splenic vascular injury can be inpainted onto the generated images. Therefore, unlike previous studies [42], [43], here the real and the synthetic data will be jointly trained using the subsequent segmentation network in a self-learning manner. Self-learning, or self-training, usually consists the following steps: 1) train a teacher model on labeled images and use it to propagate labels from the labeled to the unlabeled data; 2) then use the larger set combined of both labeled and pseudo labeled data for training a student model. Then this process can be iterated by putting back the student as the teacher. This approach has been successfully applied to multiple natural and medical imaging

tasks [17], [19], [20], [50], [51]. In our method, similarly, we train the external attention assisted model in a similar manner where it propagates labels from the internal labeled images to the synthetic images, and then using the larger, newly labeled set for training. The detailed training algorithm is illustrated next.

#### D. Iterative Attention Refinement via Unknown Class Mining

Motivated by the high relevance between a spleen and its associated injury, we propose to formulate the splenic mask as spatial attention to guide the learning process of the splenic vascular injury. Unlike previous attention-based methods [15], here we aim to mine additional information from the unknown spleen category by incorporating external data with available splenic supervision. An intuitive approach is to leverage the power of pseudo-labels [17]. However, here we provide a new perspective—to exploit the predicted splenic mask not only as pseudo-supervision in the training process but also to provide attention to let the network focus exclusively on relevant foreground regions. With a proper transformation function to account for exterior injuries (Sec. III-B), an attention mask can be computed from the pseudo splenic mask to reweight the importance of different voxels. Since the attention is originated from external data, we refer to our model as external attention assisted segmentation model, denoted by  $\mathbf{f}[\cdot; \Theta]$ . As the pseudo-labels of the unknown voxels and the network parameters  $\Theta$  are alternately updated, the derived external attention can be iteratively refined in the learning process. Specifically at the  $t$ -th update, by applying the attention mask  $\mathbf{W}_\phi^{(t)}$  as spatial weights to the internal loss terms computed on the splenic vascular injury dataset, the loss on real arterial and venous phases (A/V) can be written as:

$$\mathcal{L}_{real} = - \sum_{\phi \in \{A, V\}} \sum_j w_\phi^j \cdot [\mathbb{1}(\hat{y}_{\phi, s}^j) \log p_{\phi, s}^j + \mathbb{1}(\hat{y}_{\phi, i}^j) \log p_{\phi, i}^j + \mathbb{1}(\hat{y}_{\phi, b}^j) \log p_{\phi, b}^j], \quad (6)$$

where the subscripts  $s, i, b$  stand for the spleen, injury and background class, respectively.  $p_\phi^j$  is the  $j$ -th output of the computed probability map  $\mathbf{f}[\mathbf{X}_\phi; \Theta]$ .  $y_{\phi, i}^j$  denotes the  $j$ -th voxel of the splenic injury manual segmentation mask  $\mathbf{Y}_{\phi, i}$ . For real data where the injury annotation is given, we need to re-estimate the pseudo-label of the spleen class ( $\hat{\mathbf{Y}}_{\phi, s}$ ) and the background class ( $\hat{\mathbf{Y}}_{\phi, b}$ ). Therefore for non-injury regions,  $\hat{y}_{\phi, s}^j = 1, \hat{y}_{\phi, b}^j = 0$  if  $p_{\phi, s}^j > p_{\phi, b}^j$ , otherwise  $\hat{y}_{\phi, s}^j = 0, \hat{y}_{\phi, b}^j = 1$ .  $w_\phi^j$  is the  $j$ -th voxel of the attention mask  $\mathbf{W}_\phi^{(t)}$ .

Similarly, the loss on synthetic arterial and venous phases ( $A'/V'$ ) can be written as:

$$\mathcal{L}_{syn} = - \sum_{\phi \in \{A', V'\}} \sum_j w_\phi^j \cdot [\mathbb{1}(\hat{y}_{\phi, s}^j) \log p_{\phi, s}^j + \mathbb{1}(\hat{y}_{\phi, i}^j) \log p_{\phi, i}^j + \mathbb{1}(\hat{y}_{\phi, b}^j) \log p_{\phi, b}^j], \quad (7)$$

where  $\hat{y}_{\phi, s}^j, \hat{y}_{\phi, i}^j$  denotes the  $j$ -th label of the pseudo splenic mask  $\hat{\mathbf{Y}}_{\phi, s}$  and pseudo splenic injury mask  $\hat{\mathbf{Y}}_{\phi, i}$  respectively. Consequently, the  $j$ -th label of the background class can be computed as  $\hat{y}_{\phi, b}^j = 1 - \hat{y}_{\phi, s}^j - \hat{y}_{\phi, i}^j$ . Unlike typical synthetic

---

**Algorithm 1:** External Attention Assisted Training

---

**Input :** Internal image  $\mathbf{X}_\phi$  ( $\phi = \{A, V, A', V'\}$ ) and labelmap  $\mathbf{Y}_{\phi,i}$ ;  
 External image  $\mathbf{X}_{ext}$  and labelmap  $\mathbf{Y}_{ext,s}$ ;  
 Max number of iterations  $T$ ;  
 Hyper-parameters  $\alpha, \beta$ ;

**Output:** Network  $\mathbf{f}[\cdot; \Theta]$ ;

- 1 Training a spleen segmentation model  $\mathbf{f}_s[\cdot; \theta_s]$  on the external set;
  - 2 Extract initial splenic attention as  $\mathbf{W}_\phi^{(0)} = \mathbf{r}(\hat{\mathbf{Y}}_{\phi,s}^{(0)}, \sigma)$ , where  $\hat{\mathbf{Y}}_{\phi}^{(0)} = \arg \max_{l \in \mathbb{L}_{ext}} \mathbf{f}_s[\mathbf{X}_\phi; \theta_s]$ ;
  - 3  $t \leftarrow 0$ ;
  - 4 **repeat**
  - 5      $t \leftarrow t + 1$ ;
  - 6     Fix  $\mathbf{W}_\phi = \mathbf{W}_\phi^{(t-1)}$  and  $\hat{\mathbf{Y}}_{\phi,s} = \hat{\mathbf{Y}}_{\phi,s}^{(t-1)}$ ;
  - 7     Compute  $\mathcal{L}_{real}$  and  $\mathcal{L}_{syn}$  on the internal set by Eqn. (6) & Eqn. (7);
  - 8     Compute  $\mathcal{L}_{ext}$  on the external set by Eqn. (8);
  - 9     Compute  $\mathcal{L}_{total}$  by Eqn. (9);
  - 10    Update the segmentation model  $\mathbf{f}[\mathbf{X}_\phi; \Theta^{(t)}]$  by  $\Theta^{(t)} \leftarrow \min_{\Theta} \mathcal{L}_{total}$ ;
  - 11     $\hat{\mathbf{Y}}_{\phi}^{(t)} \leftarrow \arg \max_{l \in \mathbb{L}_{int} \cup \mathbb{L}_{ext}} \mathbf{f}[\mathbf{X}_\phi; \Theta^{(t)}]$ ;
  - 12     $\mathbf{W}_\phi^{(t)} = \mathbf{r}(\hat{\mathbf{Y}}_{\phi,s}^{(t)}, \sigma)$ ;
  - 13 **until**  $t = T$ ;
- Return:**  $\Theta = \Theta^{(T)}$ .
- 

data augmentation where the augmented training samples are often assumed to have the same labelmaps as the original data, here we use the pseudo splenic injury mask  $\hat{\mathbf{Y}}_{\phi,i}$  for computing the loss on the generated phases, following [19], [20]. This is due to that synthetic injuries can be inpainted onto generated images, making the original labelmaps no longer applicable for injury segmentation on these generated images (Sec. III-C, Fig. 3). In addition, the loss on the external set can be computed as follows:

$$\mathcal{L}_{ext} = - \sum_j [\mathbb{1}(y_{ext,s}^j) \log p_{ext,s}^j + \mathbb{1}(y_{ext,b}^j) \log p_{ext,b}^j], \quad (8)$$

where  $p_{ext}^j$  is the  $j$ -th output of the computed probability map  $\mathbf{f}[\mathbf{X}_{ext}; \Theta]$ . Therefore, the overall loss function becomes a weighted sum of both internal and external losses:

$$\mathcal{L}_{total} = \alpha \mathcal{L}_{real} + (1 - \alpha) \mathcal{L}_{syn} + \beta \mathcal{L}_{ext}, \quad (9)$$

where we introduce a coefficient  $\alpha \in [0.0, 1.0]$  to adjust the weight of the original and synthetic data. Specially,  $\alpha = 1.0$  indicates that synthetic phases are not used in the training.  $\beta$  is used for balancing the weight between the internal and external data.

During the learning process, the attention mask  $\mathbf{W}_\phi$  is initialized by Eqn. (2), and then iteratively refined during the training process. After the  $t$ -th update of the pseudo-labels, the attention mask  $\mathbf{W}_\phi^{(t)}$  is then expanded (Sec. III-B) for reweighting the loss function to further facilitate the following training process to learn  $\Theta^{(t+1)}$  (Eqn. (6) & Eqn. (7)).  $\hat{\mathbf{Y}}_{\phi}^{(t)}$

and  $\mathbf{W}_\phi^{(t)}$  can be written as:

$$\hat{\mathbf{Y}}_{\phi}^{(t)} = \arg \max_{l \in \mathbb{L}_{int} \cup \mathbb{L}_{ext}} \mathbf{f}[\mathbf{X}_\phi; \Theta^{(t)}], \quad (10)$$

$$\mathbf{W}_\phi^{(t)} = \mathbf{r}(\hat{\mathbf{Y}}_{\phi,s}^{(t)}, \sigma), \quad (11)$$

where  $t = 0, \dots, T - 1$  is the index of iteration. As indicated by Eqn. (10) & (11), in our framework, the attention map can be iteratively refined by exploiting the unknown spleen class. The overall training procedure is illustrated in Algorithm 1. It consists of two training stages: 1) **Initial Splenic Attention Extraction**, where the initial attention is extracted by training a spleen segmentation model  $\theta_s$  on the external data; 2) **Iterative Attention Refinement**, where the attention mask is initialized by the initial splenic attention mask. Then we train another segmentation network  $\Theta$  on the internal dataset with the attention map applied on top of the loss function as spatial weights. The network parameter  $\Theta$  and the applied attention mask are alternately updated to refine the attention during the training. Each stage is trained end-to-end but the two stages are trained separately since the only purpose of stage 1 is to get the initialization of the attention map for training stage 2.

## IV. EXPERIMENTS

### A. Dataset

**Splenic Vascular Injury CT Dataset.** We collect 55 consecutive multi-phase CT studies with splenic vascular injury annotations, where we randomly partitioned the studies into the training and evaluation split and made sure that no patients overlap between each split, following the experimental settings in [20]–[22], [24], [52]. Specially, 35 cases are used for training and the remaining 20 cases are used for evaluation. All patients were scanned with 40, 64, or dual source 128 CT scanners and archived at 1.5-3 mm section thickness. Images were acquired in the arterial and portal venous phases from the dome of the diaphragm through the greater trochanters.

**Pancreatic Tumor Segmentation CT Dataset.** We use the 282 abdominal CT scans from the Medical Segmentation Decathlon (MSD) challenge [53], referred to as the pancreas tumor dataset. All cases are pathological cases with tumor annotation, where were collected in the portal venous phase.

**Liver Tumor Segmentation CT Dataset.** We use the 131 contrast-enhanced 3D abdominal CT scans from the 2017 Liver Tumor Segmentation Challenge<sup>2</sup>, where the voxel-wise annotations from both liver and tumor classes are provided. The dataset was acquired by different scanners and protocols at six clinical sites, with a largely varying in-plane resolution of 0.55 ~ 1.0 mm and slice spacing of 0.45 ~ 6.0 mm.

**External Abdominal CT Dataset.** We use the 30 abdominal CT scans (3779 axial contrast-enhanced abdominal clinical CT images in total) collected under the venous phase imaging released by MICCAI 2015 Multi-Atlas Abdomen Labeling Challenge<sup>3</sup> as the external set  $\mathcal{S}_{ext}$ . For each case, 13 anatomical structures are annotated, including spleen, liver, pancreas.

<sup>2</sup>[https://competitions.codalab.org/competitions/17094#participate-get\\_data](https://competitions.codalab.org/competitions/17094#participate-get_data)

<sup>3</sup><https://www.synapse.org/#!Synapse:syn3193805/wiki/217789>

## B. splenic vascular injury Segmentation

1) *Technical details*: Our experiments were performed on the whole CT scan and the implementations are based on TensorFlow. All experiments were run on Titan Xp GPU.

For data pre-processing, similar to [10], [20], [54], we simply truncated the raw intensity values to be within the range of the soft tissue CT window, *i.e.*,  $[-125, 275]$  HU, and then normalized each raw CT case to  $[0.0, 255.0]$ . Random rotation of  $[0, 15^\circ]$  is used as an online data augmentation. We follow prior works [10], [55] and partition all CT scans into 2D slices and our implementations are based on the widely popular 2D DeepLab-v3+ backbone<sup>4</sup> [56]. We follow [12] and only train with images which contain spleen/splenic injury. In the testing stage, since the groundtruth is not available, all images will be tested. This applies to all datasets which were used in this study. A *poly* learning policy is applied with an initial learning rate of 0.08 with a decay power of 0.9. We follow [10], [12], [44] to use ImageNet pretrained model for initialization. Following [57], we also apply data augmentation by randomly scaling the input images (from 1.0 to 2.5) during training. During the testing stage, multi-scale fusion is applied (scale factors are  $\{1.0, 1.5, 2.0, 2.5\}$ ) by taking at each position the average response across the different scales.

We first train a segmentation network  $\theta_s$  for initial splenic attention extraction on the external set  $\mathcal{S}_{ext}$ . We set the number of training iteration as 50,000 with a batch size of 16. Then for synthetic phase augmentation, we use the official PyTorch implementation of CycleGAN<sup>5</sup> and set the coefficient  $\lambda = 10$  in Eqn. (5). We use the Adam solver with a batch size of 2. The networks were trained from scratch with an initial learning rate of 0.0004 and a *linear* decay policy. The whole network is trained for 50 epochs. Even though our collected splenic vascular injury CT dataset only contains 55 multi-phase CT, we want to emphasize that the actual number of training samples are 5,500 axial slices in each imaging phase. We find this number is sufficient to train the CycleGAN model.

The model trained on the external data  $\theta_s$  is then used to generate attention maps on both real and synthetic phases. We first compute the prediction for the unknown spleen class, and further expand the predicted boundary by a Gaussian filter with the standard deviation of  $\sigma = 32$ . To compute the initial splenic attention map, we binarize the result by using a cut-off threshold of  $\rho = 0.005$  in Eqn. (2). The splenic attention map  $\mathbf{W}_\phi$  and the pseudo splenic mask  $\hat{\mathbf{Y}}_{\phi,s}$  are updated for  $T = 2$  times. And during each update, the real and synthetic phases, as well as the external set, are trained for 40,000 iterations with a batch size of 16. Unless otherwise specified, hyper-parameters  $\alpha$  and  $\beta$  are set as 0.5 and 0.2, respectively.

2) *Evaluation Metric*: The accuracy of segmentation is evaluated by the Dice-Sørensen coefficient (DSC). Since the injury volumes are quite different across phases, we evaluate and report the average DSC score together with the standard deviation over all testing cases are reported on each individual phase, as in [40].

<sup>4</sup><https://github.com/tensorflow/models/tree/master/research/deeplab>

<sup>5</sup><https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix>

TABLE I

PERFORMANCE COMPARISON (AVERAGE DSC  $\pm$  STANDARD DEVIATION, %) ON THE SPLENIC VASCULAR INJURY DATASET.

Method	Venous	Arterial
U-Net [58]	46.23 $\pm$ 19.20	43.12 $\pm$ 20.05
V-Net [59]	43.56 $\pm$ 18.77	47.22 $\pm$ 19.34
C2F [12]	47.96 $\pm$ 16.89	51.43 $\pm$ 17.85
RSTN [44]	50.37 $\pm$ 15.32	53.24 $\pm$ 16.58
Attention U-Net [15]	48.72 $\pm$ 16.55	52.88 $\pm$ 16.24
Ours	<b>54.77 <math>\pm</math> 13.27</b>	<b>58.54 <math>\pm</math> 14.90</b>

3) *Results*: Performance comparison with different methods is summarized in Table I and Fig. 5(1). We compare our method with 2 popular medical image segmentation approaches (U-Net [9] and V-Net [59]) and 3 internal attention based methods (C2F [12], RSTN [44], and attention U-Net [15]). C2F and RSTN use a multi-stage framework where the first stage explicitly extracts visual attention cues, which are then fed to the second stage, whereas attention U-Net uses attention gates in the network based on non-local operators [14]. Here “internal attention” refers to attention mechanisms which are established without any assistance from external data. It is observed that, without assistance from external datasets, C2F, RSTN, and attention U-Net achieve better results on both arterial and venous phases than U-Net and V-Net, showing that attention mechanisms from internal data can already be beneficial for splenic vascular injury segmentation. Out of these three methods, the best result is achieved by RSTN, which yields an average DSC of 50.37% and 53.24% on arterial and venous phases, respectively. Meanwhile, our learning framework which leverages external splenic attention can further boost the performance to 54.77% and 58.54%, which outperform internal attention based methods by a large margin of more than 4%. Our method also achieves the lowest standard deviation on both phases, demonstrating that the improvement is general and consistent.

## C. Ablation Study

We conduct ablations to analyze the influence of different designs, components and hyper-parameters in our approach.

1) *On the effectiveness of synthetic phase augmentation*: We perform experiments by varying the input imaging phase  $\phi$  from  $\{A, V, A', V', ext\}$ . Here we define *Joint* as directly training on the union of both imaging phases. And the trained model is then used for evaluating results on both arterial and venous phases. *Separate* refers to the setting where only single-phase images are available, therefore the training and evaluation is performed separately on either the arterial or the venous phase. From Table II, we can observe that *Joint* outperforms *Separate* significantly, due to the involvement of both phases in the learning process.

Built upon *Joint*, the model performance can be further enhanced when training with additional synthetic phase images as augmentation (indicated by *SynPhaseAug* in Table II). Here, we note our segmentation network is jointly trained on both the real and the synthetic data in a self-learning manner, *i.e.*, we use pseudo-labels of the injury class on synthetic phases  $A'/V'$  computed by the *Joint* model for training rather than the annotated labels from the original real phases  $A/V$ . As

TABLE II

ABLATION STUDY ON VARIOUS COMPONENTS OF OUR FRAMEWORK FOR MULTI-PHASE SPLENIC VASCULAR INJURY SEGMENTATION.

Method	$\phi = A$	$\phi = V$	$\phi = A'$	$\phi = V'$	$\phi = ext$	Venous	Arterial
Separate (single-phase)	✓					-	42.38 ± 22.18
		✓				40.02 ± 21.88	-
Joint (multi-phase)	✓	✓				47.06 ± 17.55	50.05 ± 17.29
+ SynPhaseAug (w/o self-learning)	✓	✓	✓	✓		50.45 ± 16.06	52.19 ± 16.67
+ SynPhaseAug	✓	✓	✓	✓		51.36 ± 15.64	54.55 ± 15.10
+ External Attention (w/o SynPhaseAug)	✓	✓	✓	✓	✓	52.89 ± 14.88	56.01 ± 15.14
+ External Attention	✓	✓	✓	✓	✓	54.77 ± 13.27	58.54 ± 14.90

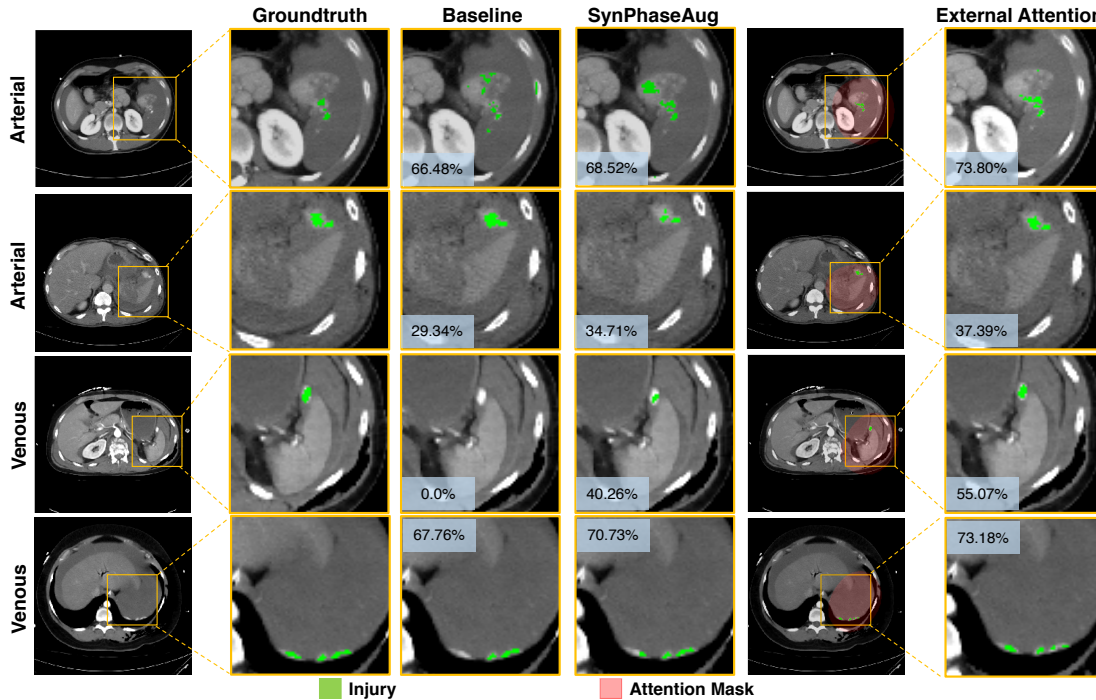


Fig. 4. Qualitative comparison. Synthetic phase augmentation already yields a performance improvement, while external attention can further enhance the segmentation accuracy. The improvement can be witnessed for both arterial and venous phases. The generated attention mask well aligns with the region-of-interest.

mentioned in Sec. III-C, our intuition lies in that a number of fake injury regions can be inpainted on the generated images (see examples in Fig. 3), which makes the usage of pseudo-labels more reasonable than the original labels.

Our experimental results also validate this intuition, *i.e.*, synthetic phase augmentation without self-learning (indicated by *SynPhaseAug (w/o self-learning)* in Table II) results in performance degradation as large as 2.36% on the arterial phase in segmentation accuracy. With self-learning, the performance improvement compared with *Joint* can be more than 4% in terms of average DSC, boosting the performance on the venous and the arterial phase from 47.06% and 50.05% to 51.36% and 54.55% respectively. We have also provided the performance comparison in boxplot (see Fig. 5(2)). This large performance improvement suggests that the joint training with synthetic phase augmentation in a self-learning manner not only enables an augmented training set, but also allows a more thorough knowledge integration from both phases. The benefits of synthetic phase augmentation are also demonstrated in Fig. 4.

2) *On the effectiveness of external attention*: For multi-phase splenic vascular injury segmentation, our external attention mechanism is built upon synthetic phase augmentation by

default. All real and generated images form our internal set whereas an external set of abdominal CT scans is used for extracting splenic attention. We also use CycleGAN to train a domain translation model between the internal and the external set, for the purpose of augmentation. All hyperparameter settings are the same as those used for synthetic phase augmentation.

After initial attention map generation, the attention map, pseudo-labels as well as network weights are iteratively updated. Throughout the iterations, we find that the quality of the pseudo splenic mask can be improved. To quantify this, we calculate the recall rate as the ratio between the number of retrieved injury voxels and the number of all injury voxels. The recall rates are improved from 94.55% and 96.37% to 98.96% and 99.79% on the venous and the arterial phase respectively. We also show a set of qualitative examples of refined pseudo spleen masks as well as the attention masks in Fig. 6, which demonstrates the high quality of the refined masks.

As aforementioned, compared to *Joint*, *i.e.*, directly training with multi-phase images, training with additional synthetic phases (*SynPhaseAug*) already suggests large benefits. Nevertheless, from Table II, we observe exploiting external attention can further lead to a performance gain of 3.41% and 3.99%



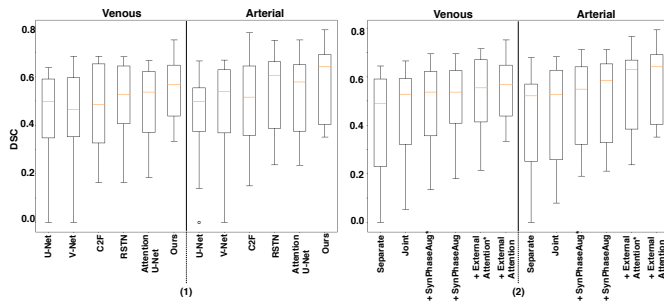


Fig. 5. Performance comparison (DSC) in box plots of multi-phase splenic injury segmentation: (1) Comparison with different segmentation methods. Our method outperforms all other competitors in both Venous and Arterial phases. (2) Ablation study on various components of our framework. Both synthetic phase augmentation and external attention are crucial for improving the model performance. For simplicity, SynPhaseAug\* refers to SynPhaseAug (w/o self-learning) and External Attention\* refers to External Attention (w/o SynPhaseAug).




Fig. 6. Qualitative examples of refined pseudo spleen masks as well as the expanded attention masks.

on each phase compared with *SynPhaseAug*. The performance comparison in boxplot (see Fig. 5(2)) further suggests the significance of the improvement. We also demonstrate the effectiveness of the proposed external attention in a set of qualitative examples, *e.g.*, as shown in Fig. 4, the computed attention masks are well aligned with the region-of-interest, which effectively reduces the irrelevant responses from the background region during the learning process. Last but not the least, it is also worth mentioning that even without synthetic phase augmentation, pure external attention (indicated by *External Attention (w/o SynPhaseAug)* in Table II) also achieves superior results than *Joint* by a large margin of 5.83% and 5.96% in the arterial and venous phase, respectively. The generalization of pure attention is also well justified in Sec. IV-D and Sec. IV-E.

3) *Diagnosis on hyper-parameters  $\alpha$  and  $\beta$* :  $\alpha$  is used for balancing the weight between the real and synthetic data, whereas  $\beta$  is used for adjusting the weight between internal and external data during the joint training. Here we provide an ablation analysis to study the importance of  $\alpha$  and  $\beta$  in Eqn. (9). By varying the value of  $\alpha$  and  $\beta$  from 0.0 - 1.0 and 0.0 - 0.5 respectively, the results are summarized in Fig. 7. We first fix  $\beta = 0.2$  and show how different choices of  $\alpha$  can affect the overall performance. From the results, we can see that  $\alpha = 0.5$  achieves the best result, which suggests that features from real and synthetic data should be equally learned to yield the best model. Meanwhile, we also note that we should avoid making the portion of synthetic imaging data over 0.7 ( $\alpha > 0.7$ ) during training. Drawing  $\alpha$  near 0.5 yields




Fig. 7. The importance of hyper-parameters  $\alpha$  and  $\beta$ .




Fig. 8. Ablation study on the number of training cases.

the most robust segmentation results on both testing phases. On the contrary,  $\alpha = 0.0$  (*i.e.*, only real images are used for training) or  $\alpha = 1.0$  (*i.e.*, only synthetic images are used for training) achieves the worst results.

We then fix  $\alpha = 0.5$  and show how different choices of  $\beta$  will affect the overall performance. We observe that  $\beta = 0.2$  achieves the best performance. Meanwhile, we note that our approach is not sensitive to this hyper-parameter when  $\beta$  is drawn between 0.1-0.3.

4) *The number of training cases*: We also reduce the number of training cases from 35 to 15 and 25 to see the effects of synthetic phase augmentation and external attention with even less supervision. The results are summarized in Fig. 8. Both *SynPhaseAug* and *External Attention* leads to larger performance gain with fewer training cases.

This suggests that both synthetic phase augmentation and external attention are beneficial for multi-phase splenic vascular injury segmentation, especially with limited data.

5) *The number of iteration T*: Next, we show how the segmentation results vary with different iteration number  $T$  in Table III.  $T = 0$  means that no external attention is applied. We can see that at first the performance increases when  $T$  goes up, and then begins to saturates after  $T$  reaches 2.

#### D. Generalization to Liver Tumor Segmentation

Following [52], we also report the average DSC on the 131 abdominal CT scans of the liver tumor segmentation CT dataset, with a random split of 100 training cases and 31

TABLE III  
ABLATION STUDY ON THE NUMBER OF ITERATION  $T$ .

$T$	Venous	Arterial
0	51.36 $\pm$ 15.64	54.55 $\pm$ 15.10
1	54.15 $\pm$ 14.13	57.08 $\pm$ 15.01
2	54.77 $\pm$ 13.27	58.54 $\pm$ 14.90
3	55.05 $\pm$ 13.29	57.96 $\pm$ 14.77
4	54.29 $\pm$ 13.53	58.70 $\pm$ 14.85

cases for validation. We also test on smaller training sets, *i.e.*, training with only 10 cases and 20 cases (indicated by *10% supervised* and *20% supervised*), to show the generalization of our approach on liver tumor segmentation, especially with limited training data. Similar to previous settings, only the annotation of the liver tumor is utilized in our training. Specifically, for our approach (indicated by “ours”), the 30 cases of the external abdominal CT dataset  $S_{ext}$  along with the liver supervision are used for extracting liver attention. Here, the external attention is instantiated without applying synthetic phase augmentation (*i.e.*,  $\alpha$  is set as 1.0 in Eqn. (9)) since multi-phase information is not accessible for the liver tumor dataset (*i.e.*, only venous phase images are provided). All other implementation details are the same as described in Sec. IV-B.1.

The performance comparison is summarized in Table IV, where we can see that our approach achieves competitive results compared with existing methods, including V-Net [59], 3D ResNet-18 [52], U-Net [58], nn-UNet [60], DeepLab V3+ [56]. It is also noteworthy to mention that by fully exploiting external attention, our approach demonstrates greater superiority with less supervision. For instance, our method outperforms 3D ResNet-18 and U-Net by a large margin of 3.63% and 9.68% under 20% supervision. Under 10% supervision, our performance surpasses those of ResNet-18 and U-Net by even further of 4.05% and 11.99%. Compared with the state-of-the-art nn-UNet, our method obtains similar results under 100% supervision, and outperforms it by 2.06% and 2.88% under 10% and 20% supervision. This promising result suggests the generalization of the proposed external attention.

### E. Generalization to Pancreatic Tumor Segmentation

We also report the segmentation results on the pancreas tumor segmentation dataset. Following the evaluation protocol in [21], we partition the whole dataset into 200 cases for training and 81 cases for validation. In the training stage, we use the tumor annotation as supervision; in the testing stage, our goal is to segment the tumor region from the whole CT. Other settings are similar to the liver tumor segmentation experimental settings, *i.e.*, we use the 30 cases of the external abdominal CT dataset  $S_{ext}$  along with the pancreas supervision for extracting pancreas attention for our approach (denoted as “ours”). Again, synthetic data augmentation is not applicable here since only single-phase images are provided in the pancreas tumor dataset. We compare our method with ResDSN [54], the state-of-the-art architecture on pancreatic ductal adenocarcinoma detection, and V-Net [59], nn-UNet [60], DeepLab V3+ [56]. We also evaluate our approach with 20 and 40 training cases (*i.e.*, *10% supervised* and *20% supervised*), to demonstrate the effectiveness of external attention with limited supervision.

The performance comparison is summarized in Table V. Our method achieves superior performance under all settings. Notably, when fewer training cases are provided, the proposed external attention demonstrates greater significance compared with prior arts. For instance, our segmentation accuracy is




Fig. 9. Qualitative examples of generated synthetic phase images.

TABLE IV  
COMPARISON (DSC, %) ON THE LIVER TUMOR DATASET.

Method	100% supervised	20% supervised	10% supervised
V-Net [59]	63.02	48.40	41.92
3D ResNet-18 (MV) [52]	65.65	53.15	48.90
3D ResNet-18 (Single) [52]	64.00	50.39	43.98
U-Net [58]	62.24	47.10	40.96
nn-UNet [60]	65.78	54.72	50.07
DeepLab V3+ [56]	64.86	51.58	45.62
Ours	65.98	56.78	52.95

TABLE V  
COMPARISON (DSC, %) ON THE PANCREAS TUMOR DATASET.

Method	100% supervised	20% supervised	10% supervised
V-Net [59]	47.53	38.97	35.68
ResDSN [54]	49.22	41.72	38.25
nn-UNet [60]	51.03	44.12	42.31
DeepLab V3+ [56]	49.96	42.38	39.52
Ours	51.63	46.33	44.85

4.61% and 6.6% higher than ResDSN with 20% and 10% supervision respectively.

## V. DISCUSSION

We have presented a novel approach for splenic vascular injury segmentation from multi-phase trauma CT scans. Extensive experiments are conducted on our curated dataset and show the effectiveness of our method. To the best of our knowledge, this is the first study that addresses this important clinical use case with deep learning to-date. We hope our study can inspire more researchers to develop machine learning techniques for trauma radiology, which has not drawn much research attention to-date.

In our work, we propose to exploit external data for iterative splenic attention refinement, which can be served as important

guidance for splenic vascular injury detection, via mining the unknown class. This principle can be widely applied to different tasks. For instance, we have also demonstrated the effectiveness of using additional abdominal CT with pancreas/liver supervision to facilitate the relatively harder tumor segmentation. We have also shown more results generated from CycleGAN in Fig. 9 to further demonstrate the effectiveness of our proposed synthetic phase augmentation module. We hope our method can offer insights for future research regarding detecting affiliated lesions of an organ.

There are several limitations for our proposed work. For one, the current study does not consider the segmentation of multi-phase splenic injury from multi-site CT yet. Future work should also increase the scale of the current study as well as conducting further characterization of the injury type.

## VI. CONCLUSION

This paper addresses the problem of multi-phase splenic vascular injury segmentation, especially with limited data. Specifically, we propose to mine external attention to enhance the associated injury segmentation. We additionally introduce synthetic phase augmentation for populating the training set, which further enhances the learning process from multi-phase images. Extensive experiments have shown the effectiveness of both external attention and synthetic phase augmentation, especially with fewer training cases. The generalization of our approach is further validated on both liver tumor segmentation and pancreas tumor segmentation.

## REFERENCES

- [1] R. A. Kozar, M. Crandall, K. Shanmuganathan, B. L. Zarzaur, M. Coburn, C. Cribari, K. Kaup, K. Schuster, G. T. Tominaga, A. P. A. Committee *et al.*, "Organ injury scaling 2018 update: spleen, liver, and kidney," *Journal of Trauma and Acute Care Surgery*, 2018.
- [2] H. Marmery, K. Shanmuganathan, M. T. Alexander, and S. E. Mirvis, "Optimization of selection for nonoperative management of blunt splenic injury: comparison of mdct grading systems," *American Journal of Roentgenology*, 2007.
- [3] A. Bhangu, D. Nepogodiev, N. Lal, and D. M. Bowley, "Meta-analysis of predictive factors and outcomes for failure of non-operative management of blunt splenic trauma," *Injury*, 2012.
- [4] D. Dreizin and F. Munera, "Blunt polytrauma: evaluation with 64-section whole-body ct angiography," *Radiographics*, 2012.
- [5] S. R. Smith, L. Morris, S. Spreadborough, W. Al-Obaydi, M. D'Auria, H. White, and A. J. Brooks, "Management of blunt splenic injury in a uk major trauma centre and predicting the failure of non-operative management: a retrospective, cross-sectional study," *European Journal of Trauma and Emergency Surgery*, 2018.
- [6] M. Gavant, M. Schurr, P. Flick, M. Croce, T. Fabian, and R. Gold, "The potential role of multi-feature ct volumetry for personalized outcome prediction and decision support after blunt splenic trauma," *RSNA*, 2020.
- [7] A. R. Boscak, K. Shanmuganathan, S. E. Mirvis, T. R. Fleiter, L. A. Miller, C. W. Sliker, S. D. Steenburg, and M. Alexander, "Optimizing trauma multidetector ct protocol for blunt splenic injury: need for arterial and portal venous phase scans," *Radiology*, 2013.
- [8] Q. Dou, H. Chen, L. Yu, L. Zhao, J. Qin, D. Wang, V. C. Mok, L. Shi, and P.-A. Heng, "Automatic detection of cerebral microbleeds from mr images via 3d convolutional neural networks," *TMI*, 2016.
- [9] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *MICCAI*, 2015.
- [10] H. Roth, L. Lu, A. Farag, A. Sohn, and R. Summers, "Spatial Aggregation of Holistically-Nested Networks for Automated Pancreas Segmentation," in *MICCAI*, 2016.
- [11] W. Zhu, Y. Huang, L. Zeng, X. Chen, Y. Liu, Z. Qian, N. Du, W. Fan, and X. Xie, "Anatomynet: Deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy," *Medical physics*, 2019.
- [12] Y. Zhou, L. Xie, W. Shen, Y. Wang, E. K. Fishman, and A. L. Yuille, "A fixed-point model for pancreas segmentation in abdominal ct scans," in *MICCAI*, 2017.
- [13] N. Tajbakhsh, L. Jeyaseelan, Q. Li, J. N. Chiang, Z. Wu, and X. Ding, "Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation," *Medical Image Analysis*, p. 101693, 2020.
- [14] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *CVPR*, 2018.
- [15] J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, B. Glocker, and D. Rueckert, "Attention gated networks: Learning to leverage salient regions in medical images," *Medical image analysis*, 2019.
- [16] L. Luo, L. Yu, H. Chen, Q. Liu, X. Wang, J. Xu, and P.-A. Heng, "Deep mining external imperfect data for chest x-ray disease screening," *TMI*, 2020.
- [17] D.-H. Lee, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," in *ICML Workshop on challenges in representation learning*, 2013.
- [18] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *ICCV*, 2017.
- [19] W. Bai, O. Oktay, M. Sinclair, H. Suzuki, M. Rajchl, G. Tarroni, B. Glocker, A. King, P. M. Matthews, and D. Rueckert, "Semi-supervised learning for network-based cardiac mr image segmentation," in *MICCAI*, 2017.
- [20] Y. Zhou, Y. Wang, P. Tang, S. Bai, W. Shen, E. Fishman, and A. Yuille, "Semi-supervised 3d abdominal multi-organ segmentation via deep multi-planar co-training," in *WACV*, 2019.
- [21] Y. Xia, D. Yang, Z. Yu, F. Liu, J. Cai, L. Yu, Z. Zhu, D. Xu, A. Yuille, and H. Roth, "Uncertainty-aware multi-view co-training for semi-supervised medical image segmentation and domain adaptation," *Medical Image Analysis*, 2020.
- [22] X. Li, L. Yu, H. Chen, C.-W. Fu, L. Xing, and P.-A. Heng, "Transformation-consistent self-ensembling model for semisupervised medical image segmentation," *TNNLS*, 2020.
- [23] Q. Liu, L. Yu, L. Luo, Q. Dou, and P. A. Heng, "Semi-supervised medical image classification with relation-driven self-ensembling model," *TMI*, 2020.
- [24] L. Yu, S. Wang, X. Li, C.-W. Fu, and P.-A. Heng, "Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation," in *MICCAI*, 2019.
- [25] X. Wang, F. Tang, H. Chen, L. Luo, Z. Tang, A.-R. Ran, C. Y. Cheung, and P. A. Heng, "Ud-mil: Uncertainty-driven deep multiple instance learning for oct image classification," *JBHI*, 2020.
- [26] M. Lenga, H. Schulz, and A. Saalbach, "Continual learning for domain adaptation in chest x-ray classification," in *MIDL*, 2020.
- [27] C. Huang, H. Han, Q. Yao, S. Zhu, and S. K. Zhou, "3d u<sup>2</sup>-net: A 3d universal u-net for multi-domain medical image segmentation," in *MICCAI*, 2019.
- [28] K. Dmitriev and A. E. Kaufman, "Learning multi-class segmentations from single-class datasets," in *CVPR*, 2019.
- [29] S. Guendel, S. Grbic, B. Georgescu, S. Liu, A. Maier, and D. Comaniciu, "Learning to recognize abnormalities in chest x-rays with location-aware dense networks," in *Iberoamerican Congress on Pattern Recognition*, 2018.
- [30] X. Fang and P. Yan, "Multi-organ segmentation over partially labeled datasets with multi-scale feature abstraction," *TMI*, 2020.
- [31] R. Huang, Y. Zheng, Z. Hu, S. Zhang, and H. Li, "Multi-organ segmentation via co-training weight-averaged models from few-organ datasets," in *MICCAI*, 2020.
- [32] G. Shi, L. Xiao, Y. Chen, and S. K. Zhou, "Marginal loss and exclusion loss for partially supervised multi-organ segmentation," *arXiv*, 2020.
- [33] Y. Tang, R. Gao, Y. Chen, D. Gao, M. R. Savona, R. G. Abramson, S. Bao, Y. Huo, and B. A. Landman, "Learning from dispersed manual annotations with an optimized data weighting policy," *Journal of Medical Imaging*, 2020.
- [34] K. Yan, J. Cai, Y. Zheng, A. P. Harrison, D. Jin, Y.-B. Tang, Y.-X. Tang, L. Huang, J. Xiao, and L. Lu, "Learning from multiple datasets with heterogeneous and partial labels for universal lesion detection in ct," *arXiv*, 2020.
- [35] A. Raju, C.-T. Cheng, Y. Huo, J. Cai, J. Huang, J. Xiao, L. Lu, C. Liao, and A. P. Harrison, "Co-heterogeneous and adaptive segmentation from multi-source and multi-phase ct imaging data: A study on pathological liver and lesion segmentation," *arXiv*, 2020.
- [36] Y. Huo, J. Cai, C.-T. Cheng, A. Raju, K. Yan, B. A. Landman, J. Xiao, L. Lu, C.-H. Liao, and A. Harrison, "Harvesting, detecting, and characterizing liver lesions from large-scale multi-phase ct data via deep dynamic texture learning," *arXiv*, 2020.

- [37] L. Zhang, Y. Shi, J. Yao, Y. Bian, K. Cao, D. Jin, J. Xiao, and L. Lu, "Robust pancreatic ductal adenocarcinoma segmentation with multi-institutional multi-phase partially-annotated ct scans," *arXiv*, 2020.
- [38] Y. Zhou, Y. Li, Z. Zhang, Y. Wang, A. Wang, E. K. Fishman, A. L. Yuille, and S. Park, "Hyper-pairing network for multi-phase pancreatic ductal adenocarcinoma segmentation," in *MICCAI*, 2019.
- [39] Y. Xia, Q. Yu, W. Shen, Y. Zhou, and A. L. Yuille, "Detecting pancreatic adenocarcinoma in multi-phase ct scans via alignment ensemble," in *MICCAI*, 2020.
- [40] H. Zheng, L. Xie, T. Ni, Y. Zhang, Y.-F. Wang, Q. Tian, E. Fishman, and A. Yuille, "Phase collaborative network for two-phase medical image segmentation," *arXiv*, 2018.
- [41] M. Frid-Adar, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "Synthetic data augmentation using gan for improved liver lesion classification," in *ISBI*, 2018.
- [42] V. Sandfort, K. Yan, P. J. Pickhardt, and R. M. Summers, "Data augmentation using generative adversarial networks (cycleGAN) to improve generalizability in ct segmentation tasks," *Scientific reports*, 2019.
- [43] C. Bowles, L. Chen, R. Guerrero, P. Bentley, R. Gunn, A. Hammers, D. A. Dickie, M. V. Hernández, J. Wardlaw, and D. Rueckert, "Gan augmentation: Augmenting training data using generative adversarial networks," *arXiv preprint arXiv:1810.10863*, 2018.
- [44] Q. Yu, L. Xie, Y. Wang, Y. Zhou, E. K. Fishman, and A. L. Yuille, "Recurrent saliency transformation network: Incorporating multi-stage visual cues for small organ segmentation," in *CVPR*, 2018.
- [45] Y. Wang, H. Dou, X. Hu, L. Zhu, X. Yang, M. Xu, J. Qin, P.-A. Heng, T. Wang, and D. Ni, "Deep attentive features for prostate segmentation in 3d transrectal ultrasound," *TMI*, 2019.
- [46] M. Havaei, A. Davy, D. Warde-Farley, A. Biard, A. Courville, Y. Bengio, C. Pal, P. Jodoin, and H. Larochelle, "Brain Tumor Segmentation with Deep Neural Networks," *Medical Image Analysis*, 2017.
- [47] K. Kamnitsas, C. Ledig, V. Newcombe, J. Simpson, A. Kane, D. Menon, D. Rueckert, and B. Glocker, "Efficient Multi-Scale 3D CNN with Fully Connected CRF for Accurate Brain Lesion Segmentation," *arXiv*.
- [48] Y. Zhou, L. Xie, E. Fishman, and A. Yuille, "Deep Supervision for Pancreatic Cyst Segmentation in Abdominal CT Scans," in *MICCAI*, 2017.
- [49] M. P. Heinrich, M. Jenkinson, M. Brady, and J. A. Schnabel, "Mrf-based deformable registration and ventilation estimation of lung ct," *TMI*, 2013.
- [50] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, "Self-training with noisy student improves imagenet classification," in *CVPR*, 2020.
- [51] L. Zhang, V. Gopalakrishnan, L. Lu, R. M. Summers, J. Moss, and J. Yao, "Self-learning to detect and segment cysts in lung ct images without manual annotation," in *ISBI*, 2018.
- [52] Y. Xia, F. Liu, D. Yang, J. Cai, L. Yu, Z. Zhu, D. Xu, A. Yuille, and H. Roth, "3d semi-supervised learning with uncertainty-aware multi-view co-training," in *WACV*, 2020.
- [53] A. L. Simpson, M. Antonelli, S. Bakas, M. Bilello, K. Farahani, B. Van Ginneken, A. Kopp-Schneider, B. A. Landman, G. Litjens, B. Menze *et al.*, "A large annotated medical image dataset for the development and evaluation of segmentation algorithms," *arXiv*, 2019.
- [54] Z. Zhu, Y. Xia, L. Xie, E. K. Fishman, and A. L. Yuille, "Multi-scale coarse-to-fine segmentation for screening pancreatic ductal adenocarcinoma," in *MICCAI*, 2019.
- [55] Y. Zhou, L. Xie, E. K. Fishman, and A. L. Yuille, "Deep supervision for pancreatic cyst segmentation in abdominal ct scans," in *MICCAI*, 2017.
- [56] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," 2018.
- [57] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *PAMI*, 2018.
- [58] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3d u-net: learning dense volumetric segmentation from sparse annotation," in *MICCAI*, 2016.
- [59] F. Milletari, N. Navab, and S. Ahmadi, "V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation," in *3DV*, 2016.
- [60] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnu-net: a self-configuring method for deep learning-based biomedical image segmentation," *Nature methods*, vol. 18, no. 2, pp. 203–211, 2021.