

ProPL: Universal Semi-Supervised Ultrasound Image Segmentation via Prompt-Guided Pseudo-Labeling

Yaxiong Chen¹, Qicong Wang^{1*}, Chunlei Li², Jingliang Hu², Yilei Shi², Shengwu Xiong¹, Xiao Xiang Zhu³, Lichao Mou^{2†}

¹ Wuhan University of Technology

² MedAI Technology (Wuxi) Co. Ltd.

³ Technical University of Munich

{chenyaxiong, wqc2023302901, xiongsww}@whut.edu.cn, {chunlei.li, jingliang.hu, yilei.shi, lichao.mou}@medimagingai.com, xiaoxiang.zhu@tum.de

Abstract

Existing approaches for the problem of ultrasound image segmentation, whether supervised or semi-supervised, are typically specialized for specific anatomical structures or tasks, limiting their practical utility in clinical settings. In this paper, we pioneer the task of universal semi-supervised ultrasound image segmentation and propose ProPL, a framework that can handle multiple organs and segmentation tasks while leveraging both labeled and unlabeled data. At its core, ProPL employs a shared vision encoder coupled with prompt-guided dual decoders, enabling flexible task adaptation through a prompting-upon-decoding mechanism and reliable self-training via an uncertainty-driven pseudo-label calibration (UPLC) module. To facilitate research in this direction, we introduce a comprehensive ultrasound dataset spanning 5 organs and 8 segmentation tasks. Extensive experiments demonstrate that ProPL outperforms state-of-the-art methods across various metrics, establishing a new benchmark for universal ultrasound image segmentation.

Code — <https://github.com/WUTCM-Lab/ProPL>

Introduction

Medical ultrasound imaging has emerged as an indispensable diagnostic tool in clinical practice, offering real-time visualization, cost-effectiveness, and non-invasive examination capabilities. Within this context, ultrasound image segmentation plays a pivotal role in computer-aided diagnosis, surgical planning, and therapeutic assessment. However, training ultrasound image segmentation models necessitates large-scale annotated datasets, yet acquiring such datasets poses challenges. First, the annotation process is labor-intensive and time-consuming, requiring extensive expertise from medical professionals. Second, ultrasound images present inherent characteristics that make them more challenging to annotate compared to other medical imaging modalities such as CT and MRI, including speckle noise,

acoustic shadows, and tissue-dependent artifacts that often obscure anatomical boundaries.

Semi-supervised learning has emerged as a promising paradigm to address the scarcity of labeled data by effectively leveraging both labeled and unlabeled samples. In the context of semantic segmentation, semi-supervised approaches have demonstrated remarkable progress, primarily falling into two categories: pseudo-labeling methods and consistency-based methods. Pseudo-labeling approaches (Wang et al. 2022) generate pseudo targets for unlabeled data using model predictions with high confidence, subsequently utilizing these pseudo-labels for model training. Consistency-based methods (Yang et al. 2023), on the other hand, enforce consistent predictions across different perturbations or views of the same input, thereby encouraging the model to learn robust representations.

Despite these advances, a significant limitation persists in the field of medical image segmentation: most current approaches, whether fully supervised or semi-supervised, are designed for specific anatomical structures or diagnostic tasks. While recent works (Zhang et al. 2021; Ye et al. 2023; Liu et al. 2023) have explored universal segmentation frameworks capable of handling multiple anatomical structures simultaneously, these approaches primarily rely on fully supervised learning, thereby remaining constrained by the availability of labeled data. This motivates our investigation into universal semi-supervised medical image segmentation, aiming to develop a more flexible and data-efficient approach that can generalize across different anatomical structures while requiring minimal labeled data.

To this end, we introduce ProPL (prompt-guided pseudo-labeling), a framework for universal semi-supervised medical image segmentation. Our architecture comprises three key components: a shared vision encoder for generic visual feature extraction, a prompt encoder for task specification, and two independent decoders for segmentation prediction. The dual-decoder design facilitates mutual learning between the decoders through pseudo-label guidance. We propose a prompting-upon-decoding mechanism that integrates task-specific prompts with the decoders, enabling dynamic adaptation to different segmentation tasks. Further-

*Work done during an internship at MedAI Technology (Wuxi) Co. Ltd.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

more, we introduce an uncertainty-driven pseudo-label calibration (UPLC) module that refines the quality of pseudo-labels by considering prediction uncertainty, thereby improving the reliability of the self-training process.

To facilitate comprehensive evaluation and advance research in universal medical image segmentation, we compile and release a diverse ultrasound image dataset comprising 6,400 images spanning 5 different organs and encompassing 8 distinct segmentation tasks.

Our contributions are three-fold:

- We pioneer the investigation of universal semi-supervised ultrasound image segmentation, addressing a fundamental yet challenging problem in medical image analysis.
- We propose ProPL, featuring novel components including prompting-upon-decoding and UPLC, collectively enabling universal segmentation and learning from both labeled and unlabeled data.
- We present a comprehensive ultrasound image dataset that encompasses multiple organs and segmentation tasks. Extensive experiments demonstrate the efficacy of our approach over state-of-the-art methods.

Related Works

Single-Task Segmentation for Ultrasound Images

Ultrasound image segmentation has gained significant attention in recent years due to its critical role in medical diagnostics. U-Net (Ronneberger, Fischer, and Brox 2015) establishes a seminal architecture in medical image segmentation, inspiring numerous variants tailored for ultrasound imaging. Building on this foundation, CMU-Net (Tang et al. 2023) has been developed specifically for breast and thyroid ultrasound image segmentation tasks. Similarly, Liu et al. (2025) devise an asymmetric U-shaped network for breast lesion segmentation, while Tang and Ning (2025) introduce a low-level feature enhancement block integrated into U-Net to improve breast lesion segmentation performance. More recently, Transformers have been explored for ultrasound segmentation. CSwin-PNet (Yang and Yang 2023) employs a pyramid structure with Swin Transformer as the backbone, while He, Yang, and Xie (2023) design a hybrid CNN-Transformer network. LM-Net (Lu et al. 2024) presents a lightweight CNN-Transformer architecture specifically designed to enhance breast lesion segmentation in ultrasound imagery. However, training effective ultrasound image segmentation models typically requires large-scale annotated datasets, which remain challenging to obtain due to data acquisition difficulties and annotation costs.

Single-Task Semi-Supervised Image Segmentation

Current semi-supervised learning approaches are primarily categorized into consistency-based and pseudo-labeling methods. Among consistency-based approaches, FixMatch (Sohn et al. 2020) generates pseudo-labels on weakly-augmented unlabeled images and trains models to predict these pseudo-labels on strongly-augmented versions of the same images. Zhao et al. (2023) propose random intensity-based augmentation combined with adaptive

label-aided CutMix-based augmentation. UniMatch (Yang et al. 2023) introduces a dual-stream augmentation strategy within the FixMatch framework to explore broader augmentation spaces. For pseudo-labeling methodologies, self-training frameworks (Yang et al. 2022) utilize teacher models trained on labeled datasets to generate pseudo-labels on unlabeled images, subsequently using high-confidence pseudo-labels to supervise student models on augmented unlabeled images. U²PL (Wang et al. 2022) extends this approach by incorporating unreliable pseudo-labels into training, thereby maximizing the utilization of unlabeled data. However, existing approaches for medical images remain limited to specific organs or tasks, lacking generalizability across diverse medical imaging scenarios.

Universal Medical Image Segmentation

Recent advances in universal medical image segmentation focus on unifying diverse segmentation tasks, imaging modalities, and clinical objectives within single frameworks. The segment anything model (SAM) has garnered significant attention in natural image segmentation due to its outstanding performance. Several efforts have adapted SAM for medical imaging: SAMUS (Lin et al. 2023) augments SAM with a two-branch CNN via cross-branch attention for ultrasound image segmentation, while DB-SAM (Qin et al. 2024) introduces a dual-branch framework that adapts SAM for universal segmentation across 2D/3D medical images. However, SAM-based methods rely on additional cues such as point, scribble, or mask prompts, which complicates practical deployment. Alternative approaches have incorporated task priors for universal segmentation. DoDNet (Zhang et al. 2021) presents a dynamic on-demand network using a single encoder-decoder architecture with a dynamic head, encoding tasks through one-hot vectors. UniSeg (Ye et al. 2023) substitutes the one-hot encoding in DoDNet with learnable prompts, enabling task-specific decoders. More recently, Liu et al. (2023) integrate language supervision by fusing CLIP-derived text embeddings rather than learnable prompts into the segmentation pipeline. Despite these advances in universal segmentation frameworks, these approaches remain fully supervised, thereby being constrained by labeled data availability.

Universal Semi-Supervised Image Segmentation

Universal semi-supervised medical image segmentation addresses challenges posed by limited labeled data and anatomical structure specificity. In computer vision for natural images, Kalluri et al. (2019) introduce a framework for universal semi-supervised semantic segmentation that minimizes supervised losses alongside within- and cross-domain unsupervised losses, while introducing a novel feature alignment objective based on pixel-aware entropy regularization. VerSemi (Zeng et al. 2025) proposes an approach to integrate various tasks into a unified model with a broad label space, enabling utilization of large volumes of unlabeled data. To date, limited work exists on this topic, with no research focusing specifically on universal semi-supervised segmentation for ultrasound images, which represents the focus of our investigation in this paper.

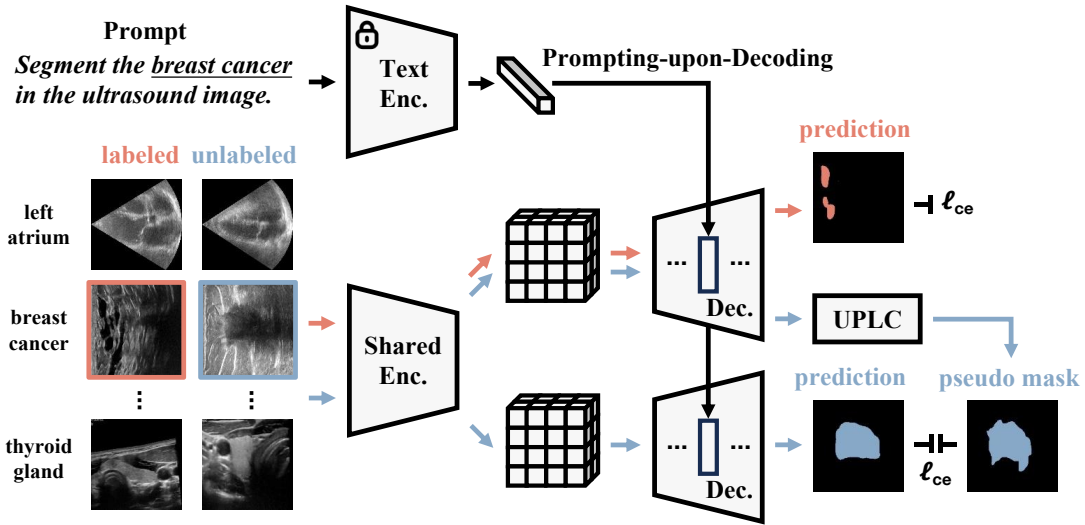


Figure 1: Overview of our proposed universal semi-supervised model ProPL.

Methodology

Network Architecture

Let f_θ represent our proposed network parameterized by θ . For the i -th segmentation task, we define the labeled images and their corresponding ground truth masks as $\mathcal{D}_i^s = \{\mathbf{x}_{i,j}^s, \mathbf{y}_{i,j}^s\}$, where $\mathbf{x}_{i,j}^s$ denotes the j -th labeled sample. Similarly, we define unlabeled data as $\mathcal{D}_i^u = \{\mathbf{x}_{i,j}^u\}$, where $\mathbf{x}_{i,j}^u$ represents the j -th unlabeled example for the i -th task. Our ProPL framework simultaneously leverages both labeled and unlabeled data during training. As illustrated in Figure 1, the framework consists of four key components: a shared vision encoder \mathcal{F}_{ve} , a prompt encoder \mathcal{F}_{pe} , and two decoders \mathcal{G}_{sd} and \mathcal{G}_{pd} .

The shared vision encoder \mathcal{F}_{ve} in our model utilizes ConvNeXt-Tiny (Liu et al. 2022), which comprises four stages constructed using enhanced ResNet blocks. Given an input ultrasound image $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$, the encoder extracts hierarchical visual representations from the four stages. These multi-scale features are defined as $\mathbf{v}_k \in \mathbb{R}^{\frac{H}{d_k} \times \frac{W}{d_k} \times C_k}$, $k = 1, 2, 3, 4$, where C_k denotes the number of channels at stage k , and H and W represent the height and width of the input image, respectively.

Each of our dual decoders consists of four stages, with skip connections between the encoder and decoders implemented following UNETR (Hatamizadeh et al. 2022b). In addition, sub-pixel convolutions are employed for up-sampling. Similar to the encoder features \mathbf{v}_k , we denote the corresponding features at each decoding stage as $\mathbf{z}_k \in \mathbb{R}^{\frac{H}{d_k} \times \frac{W}{d_k} \times C_k}$.

For task-specific prompting, we leverage BERT (Devlin et al. 2018) to extract textual features $\mathbf{t} \in \mathbb{R}^{L \times D}$ from a given task prompt, where D represents the feature dimension and L denotes the length of the prompt. This encoder transforms task-specific prompts into continuous representations that the model can interpret, enabling the model to comprehend ongoing segmentation tasks via our proposed

prompting-upon-decoding module.

Prompting-upon-Decoding

To incorporate task-specific information, we integrate \mathbf{t} with the decoders. First, we perform self-attention on \mathbf{z}_k :

$$\mathbf{z}'_k = \mathbf{z}_k + \text{LN}(\text{MHSA}(\mathbf{z}_k)), \quad (1)$$

where $\text{MHSA}(\cdot)$ represents multi-head self-attention, and $\text{LN}(\cdot)$ denotes layer normalization.

Subsequently, \mathbf{t} is processed through a 1D convolution, followed by a linear projection to align the dimension of textual features with that of the visual representation \mathbf{z}_k :

$$\boldsymbol{\tau} = \text{Conv}(\mathbf{t})\mathbf{W}, \quad (2)$$

where \mathbf{W} is a learnable mapping matrix. $\boldsymbol{\tau} \in \mathbb{R}^{M \times C_i}$ is the resulting text embedding.

Finally, we employ multi-head cross-attention (MHCA) to inject the prompt information into our model:

$$\mathbf{h}_k = \mathbf{z}'_k + \alpha \text{LN}(\text{MHCA}(Q = \mathbf{z}'_k, K = \boldsymbol{\tau}, V = \boldsymbol{\tau})), \quad (3)$$

where α is a learnable parameter that modulates the influence of the prompt features.

Training and Pseudo-Label Calibration

We present the training of our pseudo-labeling-based framework for universal semi-supervised ultrasound image segmentation, followed by the proposed UPLC mechanism that enhances pseudo-label quality.

Pseudo-Labeling For each labeled image $\mathbf{x}_{i,j}^s \in \mathcal{D}_i^s$, we process both the image through the vision encoder \mathcal{F}_{ve} and its corresponding task prompt through the prompt encoder \mathcal{F}_{pe} . The decoder \mathcal{G}_{sd} , incorporating the prompting-upon-decoding module, then generates a segmentation mask. We optimize this process using a combined loss function that uses binary cross-entropy and Dice losses to minimize the discrepancy between the predicted and ground truth masks.

Method		BC		FH		LA		LV		MD		OT		TG		TN		mDice	mIoU
		Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU		
♠	U-Net	64.15	52.25	94.76	90.44	80.07	69.78	87.26	78.00	76.80	63.08	67.96	56.48	78.07	67.89	52.27	40.18	75.17	64.76
	PSPNet	64.56	52.45	93.13	87.96	76.27	65.04	86.23	76.57	73.80	59.43	60.64	48.44	75.89	65.22	51.84	39.97	72.80	61.89
	DeepLabv3+	65.90	54.21	94.58	90.07	74.43	62.47	84.98	74.78	71.25	56.37	65.48	53.56	76.20	65.10	56.00	44.00	73.60	62.57
	UNet++	50.49	38.75	87.18	79.12	71.45	58.17	82.53	71.43	68.04	52.71	57.17	43.46	72.06	60.52	43.48	30.67	66.55	54.35
	SegFormer	70.85	59.30	95.02	90.89	74.59	62.39	86.41	76.66	72.11	57.07	69.90	58.14	76.73	66.34	54.76	42.01	75.05	64.10
	Swin-Unet	65.63	53.39	95.35	91.33	77.06	65.03	85.20	74.83	67.69	51.82	67.97	56.22	75.28	63.80	51.19	38.53	73.17	61.87
	Swin UNETR	53.81	41.05	86.54	78.34	69.37	56.63	84.19	73.50	68.74	53.34	60.78	48.18	74.12	62.59	40.43	27.53	67.25	55.15
	HiFormer	55.93	43.17	92.97	87.52	74.26	61.78	86.33	76.56	72.58	57.75	64.70	51.66	77.95	67.10	52.19	38.43	72.11	60.50
	H2Former	41.20	27.89	85.03	75.60	55.67	40.43	82.34	70.84	55.73	39.29	60.77	47.70	68.24	56.28	26.91	16.58	59.49	46.83
	Mamba-UNet	65.38	53.73	96.54	93.47	78.47	68.47	86.92	77.51	77.66	64.16	62.74	52.76	78.85	69.42	61.03	49.18	75.95	66.09
♥	FixMatch	33.92	26.23	87.92	79.51	26.10	16.74	78.35	65.60	63.64	47.65	33.85	22.17	73.29	63.06	29.86	22.13	53.37	42.89
	U ² PL	64.48	52.46	93.83	89.05	71.58	57.75	83.52	72.63	70.99	56.21	67.87	55.18	78.68	68.04	49.27	37.97	72.53	61.16
	ST++	39.85	29.38	83.12	72.52	60.85	47.39	71.54	58.16	56.61	42.37	40.56	31.27	66.82	54.35	31.33	21.92	56.34	44.67
	UniMatch	43.36	33.41	87.64	78.99	34.54	23.05	79.25	66.68	58.30	42.35	67.88	55.17	74.68	64.49	51.71	39.64	62.17	50.47
	AugSeg	15.66	10.15	82.27	71.92	18.88	11.73	50.94	35.20	36.17	23.04	27.06	18.15	41.72	28.16	21.11	13.73	36.73	26.51
	DDFP	71.86	58.72	95.85	94.82	85.05	75.37	88.74	80.68	77.48	62.03	73.04	68.43	82.68	79.17	66.56	61.99	80.16	72.65
	CSC-PA	49.04	36.97	87.16	79.84	76.06	63.72	83.66	72.70	74.16	59.92	60.96	48.90	78.60	68.26	46.79	34.70	69.55	58.13
♣	DoDNet	73.06	57.55	94.94	90.37	29.79	17.50	61.97	44.90	19.44	10.77	79.42	65.87	85.59	74.82	65.93	49.17	63.77	51.37
	CLIP-UM	71.20	55.28	95.21	90.87	43.20	27.55	36.54	22.35	23.96	13.61	80.20	66.95	85.13	74.11	69.93	53.76	63.17	50.56
♦	Univ-full	43.68	32.59	86.07	77.49	21.89	12.76	43.88	30.23	30.55	18.62	52.37	39.26	57.12	45.05	43.58	31.80	47.39	35.98
	ProPL	73.71	63.76	95.65	92.17	86.63	77.58	91.11	83.94	81.31	69.05	73.49	63.83	81.21	71.14	65.96	54.22	81.13	71.96

Table 1: Performance comparison of single-task (supervised ♠, semi-supervised ♥) and universal (supervised ♣, semi-supervised ♦) models across eight ultrasound image segmentation tasks under the 1/16 data partition. Metrics reported as Dice (%) per task, mean Dice (mDice) and mean IoU (mIoU). Best results in **bold**. Tasks: breast cancer (BC), fetal head (FH), left atrium (LA), left ventricle (LV), myocardium (MD), ovarian tumor (OT), thyroid gland (TG), and thyroid nodule (TN).

Method		BC		FH		LA		LV		MD		OT		TG		TN		mDice	mIoU
		Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU		
♠	U-Net	71.45	60.47	96.00	92.54	84.24	74.56	89.34	81.09	79.45	66.68	76.10	65.34	81.36	71.84	61.74	49.64	79.96	70.27
	PSPNet	70.14	58.54	94.89	90.80	82.22	71.92	87.96	79.12	78.27	65.01	70.46	58.42	79.82	70.36	57.37	45.64	77.64	67.48
	DeepLabv3+	72.04	61.13	95.46	91.62	81.93	71.42	87.23	78.02	77.08	63.44	73.87	62.62	78.79	69.05	58.77	46.54	78.15	67.98
	UNet++	53.68	41.60	89.16	82.08	74.20	61.98	86.35	76.61	72.09	57.29	62.25	49.27	76.89	65.87	45.61	32.42	70.03	58.39
	SegFormer	76.08	65.34	95.89	92.35	80.31	69.21	87.58	78.44	76.31	62.37	76.37	65.44	78.56	68.75	63.02	50.85	79.27	69.09
	Swin-Unet	74.29	63.17	96.11	92.60	82.04	71.03	87.85	78.71	74.80	60.37	75.14	64.05	78.68	68.36	63.25	50.84	79.02	68.64
	Swin UNETR	59.21	46.99	89.72	83.10	77.48	66.01	86.43	76.80	73.58	59.13	64.92	51.97	78.70	68.34	46.32	33.50	72.05	60.73
	HiFormer	62.80	50.92	95.32	91.35	80.41	69.40	88.15	79.30	76.89	63.03	71.14	59.32	79.94	69.87	58.84	46.03	76.69	66.15
	H2Former	59.23	46.96	93.23	88.07	75.87	63.78	87.13	77.77	62.49	46.11	64.33	51.21	72.89	61.29	48.14	35.29	70.41	58.81
	Mamba-UNet	73.78	63.86	96.71	93.78	83.30	73.16	88.13	79.41	79.99	67.38	78.04	68.22	83.20	75.25	66.06	54.35	81.15	71.93
♥	FixMatch	55.54	43.89	88.47	80.42	38.56	26.68	80.52	68.12	63.10	47.15	36.46	25.13	75.22	65.55	42.41	32.22	60.04	48.65
	U ² PL	62.19	50.78	95.10	91.10	77.27	64.73	86.33	76.60	76.71	63.06	70.91	59.08	80.93	71.22	54.43	42.48	75.48	64.88
	ST++	51.08	39.76	88.24	80.83	64.02	51.48	79.21	66.82	65.72	50.75	58.03	46.85	72.83	61.03	43.44	31.73	65.32	53.66
	UniMatch	58.65	46.63	88.04	79.48	39.34	27.29	78.04	64.97	60.68	44.52	64.83	51.24	75.06	65.26	56.04	43.36	65.09	52.84
	AugSeg	14.74	9.64	85.45	75.95	43.24	29.57	58.94	42.51	29.55	18.13	43.63	31.96	46.03	32.30	22.36	14.90	42.99	31.87
	DDFP	72.54	62.50	97.07	95.13	85.72	76.57	90.23	83.32	80.51	66.86	76.54	71.78	84.07	80.32	68.81	63.52	81.94	75.00
	CSC-PA	49.04	36.97	87.16	79.84	76.06	63.72	83.66	72.70	74.16	59.92	60.96	48.90	78.60	68.26	46.79	34.70	69.55	58.13
♣	DoDNet	78.15	64.14	96.73	93.67	30.50	17.99	43.80	28.04	25.34	14.51	82.81	70.67	90.18	82.11	74.68	59.59	65.27	53.84
	CLIP-UM	79.17	65.52	96.49	93.22	17.84	9.79	64.22	47.30	43.86	28.09	84.51	73.18	89.58	81.13	76.18	61.53	68.98	57.47
♦	Univ-full	48.87	37.31	85.13	75.87	19.76	11.22	58.51	43.11	19.97	11.27	56.66	45.41	56.99	46.04	44.62	33.22	48.81	37.93
	ProPL	75.03	65.29	96.25	93.20	87.16	78.85	92.29	85.90	84.13	73.08	77.05	66.91	83.61	75.59	71.30	59.83	83.35	74.83

Table 2: Performance comparison of single-task (supervised ♠, semi-supervised ♥) and universal (supervised ♣, semi-supervised ♦) models across eight ultrasound image segmentation tasks under the 1/8 data partition.

For an unlabeled image $x_{i,j}^u \in \mathcal{D}_i^u$, we first extract feature representations using \mathcal{F}_{ve} and generate a pseudo segmentation mask through \mathcal{G}_{sd} . These features are subsequently processed by \mathcal{G}_{pd} to predict another segmentation mask. We apply the same combined loss function to enforce consistency between these two masks.

UPLC While pseudo-labels are valuable for semi-supervised learning, they often contain inherent noise that can impact model performance. To address this limitation,

we introduce UPLC for pseudo-label refinement.

Given an unlabeled image x^u (omitting subscripts i and j for notational simplicity), we apply N stochastic perturbations to its encoded visual representation v_4 before processing through \mathcal{G}_{sd} . This generates an ensemble of pseudo segmentation masks $\{\hat{y}_i\}$. We then compute the empirical

Method	BC		FH		LA		LV		MD		OT		TG		TN		mDice	mIoU	
	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU			
♠	U-Net	78.55	68.75	96.66	93.66	86.32	77.48	90.30	82.78	82.43	70.73	79.77	70.29	86.51	78.42	69.64	58.62	83.77	75.09
	PSPNet	76.21	65.58	96.39	93.18	85.18	76.40	90.11	82.34	81.76	69.71	76.65	66.10	84.78	78.89	66.23	54.26	82.16	73.31
	DeepLabv3+	76.78	66.82	96.46	93.32	83.84	74.54	89.64	81.58	80.59	68.09	79.20	69.11	85.48	76.91	67.04	55.48	82.38	73.23
	UNet++	61.56	49.22	92.11	86.69	82.36	72.39	88.59	79.97	77.69	64.32	65.42	53.00	83.25	73.05	50.04	37.07	75.13	64.46
	SegFormer	78.35	68.36	96.37	93.19	83.80	74.33	89.54	81.46	79.86	66.98	78.98	68.61	84.44	75.35	71.15	58.95	82.81	73.40
	Swin-Unet	77.66	67.19	96.41	93.14	84.99	75.22	87.19	77.74	78.74	65.55	78.69	68.00	82.43	72.74	71.07	58.90	82.15	72.31
	Swin UNETR	65.53	53.83	93.15	88.21	83.21	73.10	88.29	79.55	78.39	65.23	67.68	55.35	84.26	75.35	50.53	37.84	76.38	66.06
	HiFormer	73.82	63.40	96.09	92.66	83.90	74.42	89.33	81.20	80.15	67.37	76.33	65.19	83.53	74.57	65.45	53.16	81.08	71.50
	H2Former	70.78	59.47	95.15	91.19	83.96	74.56	89.00	80.84	77.78	64.18	74.82	63.03	81.41	72.14	58.64	46.14	78.94	68.94
	Mamba-UNet	78.59	69.59	97.06	94.38	85.22	76.38	89.87	82.00	82.23	70.41	81.87	72.47	86.99	79.26	71.57	60.60	84.18	75.64
♥	FixMatch	56.00	44.37	84.99	75.70	43.88	30.57	81.99	70.36	62.79	47.19	64.92	51.57	76.06	66.29	55.27	43.46	65.74	53.69
	U ² PL	69.16	57.80	95.73	92.10	80.61	69.06	88.14	79.36	78.36	65.19	72.14	60.40	82.83	73.68	59.24	46.47	78.28	68.01
	ST++	58.47	46.33	92.58	87.30	66.34	52.22	81.58	70.16	70.99	55.84	64.66	53.14	76.04	64.26	47.63	35.65	69.79	58.11
	UniMatch	60.95	48.82	88.63	80.60	42.65	29.91	79.96	67.60	63.92	48.17	67.88	55.17	76.32	66.54	59.50	46.97	67.48	55.47
	AugSeg	35.07	25.60	88.28	80.23	48.25	32.79	63.57	47.54	20.01	11.72	47.25	33.54	43.54	30.21	36.56	26.51	47.82	36.02
	DDFP	71.86	58.72	96.85	94.82	85.95	75.37	88.74	80.68	77.48	62.03	73.04	68.43	82.68	79.17	66.56	61.99	80.40	72.65
	CSC-PA	49.91	37.38	89.05	82.16	75.62	63.21	84.69	74.12	74.41	60.08	60.95	48.44	78.37	67.90	45.51	33.21	69.81	58.31
♣	DoDNet	79.72	66.28	97.08	94.33	34.87	21.11	58.29	41.13	75.69	60.88	85.86	75.22	92.47	86.00	77.67	63.47	75.21	63.55
	CLIP-UM	79.21	65.57	96.58	93.39	84.71	73.48	75.43	60.55	71.25	55.34	85.49	74.66	88.95	80.11	77.69	63.52	82.41	70.83
♦	Univ-full	55.71	43.52	85.81	76.71	38.66	25.05	53.10	37.33	39.89	26.65	62.60	51.05	63.77	52.12	46.84	35.75	55.80	43.52
	ProPL	77.18	67.91	96.56	93.64	88.72	80.83	93.00	87.07	85.24	74.70	78.31	68.86	89.52	82.14	75.44	64.50	85.50	77.46

Table 3: Performance comparison of single-task (supervised ♠, semi-supervised ♥) and universal (supervised ♣, semi-supervised ♦) models across eight ultrasound image segmentation tasks under the 1/4 data partition.




Figure 2: Comparison of our prompting-upon-decoding method against established prompting approaches implemented in CLIP-UM (Liu et al. 2023) and DoDNet (Zhang et al. 2021).

mean and variance of these masks:

$$\mu = \frac{1}{N} \sum_{i=1}^N \hat{y}_i, \quad \gamma = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i - \mu)^2. \quad (4)$$

The variance γ provides an uncertainty measure for the pseudo-labels. We leverage this uncertainty to modulate supervision signals for \mathcal{G}_{pd} by defining rectified pseudo-labels as:

$$\hat{y} = \exp(-\gamma) \odot \mu, \quad (5)$$

where \odot denotes element-wise multiplication.

Experiments

Datasets and Evaluation Metrics

We evaluate our network on eight ultrasound image segmentation tasks spanning 6,400 samples, with 800 samples per task. Our dataset comprises breast cancer images from Dataset B (Yap et al. 2018) and BUSI (Walid et al. 2020), fetal head images from HC18 (van den Heuvel et al. 2018), left atrium and left ventricle data from CAMUS (Leclerc et al. 2019), myocardium data from CAMUS and HMC-QU (Degerli et al. 2021), ovarian tumor images from MMOTU (Zhao et al. 2022), thyroid gland images from TN3K (Gong et al. 2023), and thyroid nodule data from TN3K and DDTI (Pedraza et al. 2015). We maintain a 3:1 train-test split ratio

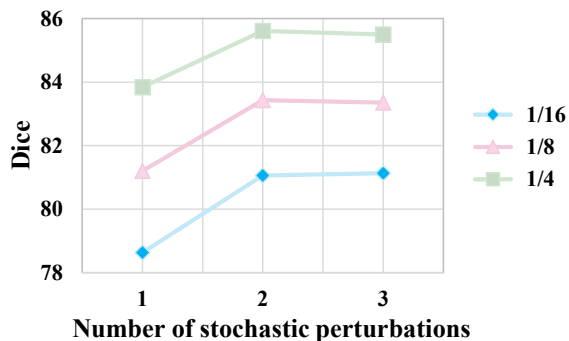


Figure 3: Impact of varying the number of stochastic perturbations in UPLC.




Figure 4: Performance-efficiency trade-off analysis across various models.

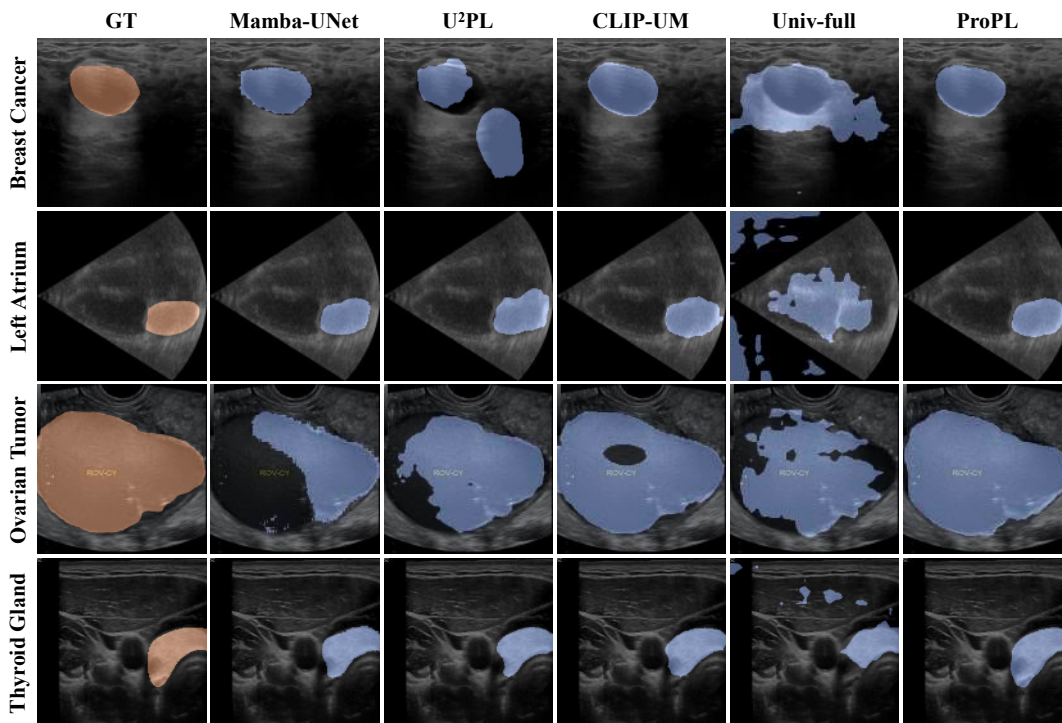


Figure 5: Qualitative segmentation results comparing the proposed method against competitive approaches across four segmentation tasks: breast cancer, left atrium, ovarian tumor, and thyroid gland.

across all tasks.

Following standard semi-supervised learning protocols, we conduct experiments using 1/16, 1/8, and 1/4 of the training data as labeled samples, with the remainder serving as unlabeled data. We evaluate performance using the Dice similarity coefficient and mean Intersection over Union (mIoU).

Implementation Details

We implement our framework in PyTorch, conducting all experiments on a single NVIDIA GeForce RTX 4090 GPU. All the images are resized to 224×224 pixels and augmented through random rotation and scaling. We train with a batch

size of 16 using SGD (momentum = 0.9, weight decay = $1e - 5$). The learning rate follows a polynomial schedule: $lr = init_{lr} \times (1 - \frac{iter}{max_{iter}})^{power}$ where $init_{lr} = 0.001$ and $power = 0.9$. Models are trained for 200 epochs. We train the model for 200 epochs. Within UPLC, we introduce perturbations via dropout (rate 0.3) and draw two stochastic perturbation instances.

For text prompts, we use minimal instructions such as “Segment the breast cancer in the ultrasound image.” The term “breast cancer” can be replaced with other targets, including the left atrium, ovarian tumor, or thyroid gland.

PuD	UPLC	BC		FH		LA		LV		MD		OT		TG		TN		mDice	mIoU
		Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU		
-	-	67.76	56.36	95.57	91.87	26.78	20.06	38.92	29.58	36.92	27.54	72.34	72.34	79.88	69.96	63.19	50.42	60.17	52.27
-	✓	72.47	62.22	95.35	91.70	17.75	14.56	41.41	33.47	41.26	31.81	73.54	63.44	80.80	79.56	69.72	54.41	61.54	53.90
✓	-	67.67	56.90	95.18	91.26	83.52	73.45	88.56	80.08	77.85	64.31	73.53	62.32	80.80	70.80	61.95	49.39	78.63	68.56
✓	✓	73.71	63.76	95.65	92.17	86.63	77.58	91.11	83.94	81.31	69.05	73.49	63.83	81.21	71.14	65.96	54.22	81.13	71.96

Table 4: Ablation study across eight ultrasound image segmentation tasks under the 1/16 data partition. PuD denotes prompting-upon-decoding.

PuD	UPLC	BC		FH		LA		LV		MD		OT		TG		TN		mDice	mIoU
		Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU		
-	-	71.86	61.65	95.55	91.96	21.13	15.33	63.78	53.18	26.71	18.95	74.25	63.60	84.04	74.44	65.83	53.91	62.89	54.13
-	✓	70.30	59.79	94.10	89.61	16.35	10.51	68.97	54.24	40.30	26.78	74.93	64.75	76.56	65.98	64.01	52.17	63.19	52.98
✓	-	68.97	58.23	96.35	93.32	87.68	78.98	91.78	85.10	83.35	71.94	74.38	64.13	80.54	72.19	66.57	53.97	81.20	72.23
✓	✓	75.03	65.29	96.25	93.20	87.16	78.85	92.29	85.90	84.13	73.08	77.05	66.91	83.61	75.59	71.30	59.83	83.35	74.83

Table 5: Ablation study across eight ultrasound image segmentation tasks under the 1/8 data partition.

PuD	UPLC	BC		FH		LA		LV		MD		OT		TG		TN		mDice	mIoU
		Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU		
-	-	70.09	58.36	94.66	90.25	20.02	12.07	70.18	54.81	39.95	25.98	74.35	62.87	78.92	67.88	64.18	51.57	64.04	52.97
-	✓	70.71	59.18	94.93	90.87	31.71	19.59	66.32	50.35	36.49	23.37	74.75	63.96	75.57	63.50	66.91	54.73	64.67	53.19
✓	-	75.15	65.33	96.37	93.36	87.73	79.41	92.65	86.46	84.28	73.29	78.31	68.46	87.19	79.39	69.02	57.38	83.84	75.39
✓	✓	77.18	67.91	96.56	93.64	88.72	80.83	93.00	87.07	85.24	74.70	78.31	68.86	89.52	82.14	75.44	64.50	85.50	77.46

Table 6: Ablation study across eight ultrasound image segmentation tasks under the 1/4 data partition.

Comparison with State-of-the-Art Methods

We compare ProPL against four categories of methods:

- Single-task supervised: U-Net (Ronneberger, Fischer, and Brox 2015), PSPNet (Zhao et al. 2017), Swin-Unet (Cao et al. 2023), Mamba-UNet (Wang et al. 2024b), DeepLabv3+ (Chen et al. 2018), SegFormer (Xie et al. 2021), Swin UNETR (Hatamizadeh et al. 2022a), HiFormer (Heidari et al. 2023), H2Former (He et al. 2023), and UNet++ (Zhou et al. 2018)
- Single-task semi-supervised: FixMatch (Sohn et al. 2020), UniMatch (Yang et al. 2023), U²PL (Wang et al. 2022), AugSeg (Zhao et al. 2023), DDFP (Wang et al. 2024a), CSC-PA (Ding et al. 2025), and ST++ (Yang et al. 2022)
- Universal supervised: DoDNet (Zhang et al. 2021) and CLIP-UM (Liu et al. 2023)
- Universal semi-supervised: Univ-full (Kalluri et al. 2019)

Note that Zeng et al. (2025) present pioneering work for universal semi-supervised medical image segmentation on CT/MRI images and achieve impressive results. However, in our experiments, it performs mediocly on ultrasound images, so we did not include it. Table 1 presents results across all eight segmentation tasks under the 1/16 partition protocol. ProPL demonstrates substantial improvements over all competitors, surpassing the best single-task supervised model by 5.18% in mean Dice and 5.87% in mIoU, the top single-task semi-supervised model by 0.97%, the leading universal supervised model by 17.36% and 20.59%, and Univ-full by 33.74% and 35.98%. We conduct additional

comparison experiments under 1/8 and 1/4 data partitions (Tables 2 and 3). Under 1/8 data partitioning, our model outperforms the second-best model by 2.2% in mean Dice, while under 1/4 data partitioning, our model achieves improvements of 1.32% in mean Dice and 1.82% in mean IoU compared to the second-best model.

Furthermore, our model also delivers strong performance on other metrics. For example, in terms of HD95 (lower is better), our method attains 20.59, outperforming the second-best result by 1.06 under the 1/8 setting.

Figure 5 shows qualitative comparisons between our model, Mamba-UNet, U²PL, CLIP-UM, and Univ-full.

Discussion

Task-Specific Prompting Removing task prompts leads to performance drops of 19.59% and 18.06% in mean Dice and mIoU under 1/16 data partition, respectively (see Table 4). We can also observe similar degradation under 1/8 and 1/4 data partitions (see Tables 5 and 6).

Prompting-upon-Decoding Our approach outperforms existing prompting methods implemented in CLIP-UM (Liu et al. 2023) and DoDNet (Zhang et al. 2021) across all eight tasks (Figure 2), demonstrating gains of 3.05% and 1.16% in mean Dice, respectively.

UPLC Impact The absence of pseudo-label calibration reduces performance by 2.5% and 3.4% in average Dice and IoU under 1/16 data partition (cf. Table 4). We can also see consistent performance degradation under 1/8 and 1/4 data partitions (cf. Tables 5 and 6).

Stochastic Perturbations in UPLC We investigate the impact of varying stochastic perturbation counts in UPLC.

As shown in Figure 3, optimal performance is achieved with two perturbations under 1/4 and 1/8 partition protocols (85.61% and 83.43% respectively) and three perturbations under the 1/16 protocol (81.13%). GPU memory constraints prevented testing beyond four perturbations. Based on these findings, we adopt two stochastic perturbations in our implementation.

Increasing N by 1 incurs an additional 2.43 GB of memory and 27.6 s per epoch, which remains practically acceptable.

We also conduct comparative experiments on random perturbation strategies, evaluating Dropout against Gaussian noise. The results show that Gaussian noise performs worse, with a 2.58 mIoU drop.

Memory Footprint We compare ProPL’s memory efficiency against representative methods including U-Net, Mamba-UNet, U²PL, DoDNet, CLIP-UM, and Univ-full. As shown in Figure 4, ProPL achieves the highest Dice while maintaining a modest memory footprint of 712 MB, demonstrating optimal performance-efficiency trade-off.

Prompting Trade-Off We use minimal prompts at negligible generation cost. Removing them causes collapse (Tables 4, 5, 6), making extra-computation metrics meaningless. Compared to one-hot prompting, ours adds only 18 s/epoch but improves results—an acceptable trade-off.

Generalization Public ultrasound datasets are scarce. On an in-house external set (700 images), the model achieves 73.33 mIoU—comparable to internal results—showing good generalization.

Conclusion

In this paper, we address the challenging problem of universal semi-supervised ultrasound image segmentation by introducing ProPL, a framework that leverages prompt-guided pseudo-labeling to effectively segment multiple anatomical structures while utilizing both labeled and unlabeled data. Extensive experiments on our newly compiled ultrasound dataset demonstrate that ProPL consistently outperforms state-of-the-art methods.

Acknowledgments

This work was supported in part by the National Key Research and Development Program of China (Grant 2022ZD0160604), the Hainan Province “Nanhai New Star” Technology Innovation Talent Platform Project (Grant NHXXRCXM202361), the Sanya Science and Technology Special Fund (Grant 2024KJFX034), the NSFC (Grant 62571381), the General Program of the Hubei Natural Science Foundation (Grant 2025AFB615), and the Key Research and Development and Achievement Transformation Program of the Inner Mongolia Autonomous Region (Grant 2025YFHH0078).

References

Cao, H.; Wang, Y.; Chen, J.; Jiang, D.; Zhang, X.; Tian, Q.; and Wang, M. 2023. Swin-Unet: Unet-Like Pure Transformer for Medical Image Segmentation. In *Proceedings of the European Conference on Computer Vision*, 205–218.

Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In *Proceedings of the European Conference on Computer Vision*, 801–818.

Degerli, A.; Zabihi, M.; Kiranyaz, S.; Hamid, T.; Mazhar, R.; Hamila, R.; and Gabbouj, M. 2021. Early Detection of Myocardial Infarction in Low-Quality Echocardiography. *IEEE Access*, 9: 34442–34453.

Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4171–4186.

Ding, Z.; Chen, G.; Zhang, Q.; Wu, H.; and Qin, J. 2025. CSC-PA: Cross-Image Semantic Correlation via Prototype Attentions for Single-Network Semi-Supervised Breast Tumor Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15632–15641.

Gong, H.; Chen, J.; Chen, G.; Li, H.; Li, G.; and Chen, F. 2023. Thyroid Region Prior Guided Attention for Ultrasound Segmentation of Thyroid Nodules. *Computers in Biology and Medicine*, 155: 106389.

Hatamizadeh, A.; Nath, V.; Tang, Y.; Yang, D.; Roth, H. R.; and Xu, D. 2022a. Swin UNETR: Swin Transformers for Semantic Segmentation of Brain Tumors in MRI Images. In *International MICCAI Brainlesion Workshop*, 272–284.

Hatamizadeh, A.; Tang, Y.; Nath, V.; Yang, D.; Myronenko, A.; Landman, B.; Roth, H.; and Xu, D. 2022b. UNETR: Transformers for 3D Medical Image Segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 574–584.

He, A.; Wang, K.; Li, T.; Du, C.; Xia, S.; and Fu, H. 2023. H2Former: An Efficient Hierarchical Hybrid Transformer for Medical Image Segmentation. *IEEE Transactions on Medical Imaging*, 42(9): 2763–2775.

He, Q.; Yang, Q.; and Xie, M. 2023. HCTNet: A Hybrid CNN-Transformer Network for Breast Ultrasound Image Segmentation. *Computers in Biology and Medicine*, 155: 106629.

Heidari, M.; Kazerouni, A.; Soltany, M.; Azad, R.; Aghdam, E. K.; Cohen-Adad, J.; and Merhof, D. 2023. HiFormer: Hierarchical Multi-Scale Representations Using Transformers for Medical Image Segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 6202–6212.

Kalluri, T.; Varma, G.; Chandraker, M.; and Jawahar, C. 2019. Universal Semi-Supervised Semantic Segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 5259–5270.

Leclerc, S.; Smistad, E.; Pedrosa, J.; Østvik, A.; Cervenansky, F.; Espinosa, F.; Espeland, T.; Berg, E. A. R.; Jodoin, P.-M.; Grenier, T.; Lartizien, C.; D’hooge, J.; Lovstakken, L.;

- and Bernard, O. 2019. Deep Learning for Segmentation Using An Open Large-Scale Dataset in 2D Echocardiography. *IEEE Transactions on Medical Imaging*, 38(9): 2198–2210.
- Lin, X.; Xiang, Y.; Zhang, L.; Yang, X.; Yan, Z.; and Yu, L. 2023. Samus: Adapting Segment Anything Model for Clinically-Friendly and Generalizable Ultrasound Image Segmentation. *arXiv preprint arXiv:2309.06824*, 4(11).
- Liu, J.; Shao, J.; Xu, S.; Tang, Z.; Liu, W.; Li, Z.; Wang, T.; and Bian, X. 2025. Asym-UNet: An Asymmetric U-Shape Network for Breast Lesions Ultrasound Images Segmentation. *Biomedical Signal Processing and Control*, 99: 106822.
- Liu, J.; Zhang, Y.; Chen, J.; Xiao, J.; Lu, Y.; A Landman, B.; Yuan, Y.; Yuille, A.; Tang, Y.; and Zhou, Z. 2023. CLIP-Driven Universal Model for Organ Segmentation and Tumor Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 21152–21164.
- Liu, Z.; Mao, H.; Wu, C.; Christoph, F.; Trevor, D.; and Sainig, X. 2022. A ConvNet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11976–11986.
- Lu, Z.; She, C.; Wang, W.; and Huang, Q. 2024. LM-Net: A Light-Weight and Multi-Scale Network for Medical Image Segmentation. *Computers in Biology and Medicine*, 168: 107717.
- Pedraza, L.; Vargas, C.; Narváez, F.; Durán, O.; Muñoz, E.; and Romero, E. 2015. An Open Access Thyroid Ultrasound Image Database. In *International Symposium on Medical Information Processing and Analysis*, 188–193.
- Qin, C.; Cao, J.; Fu, H.; Khan, F. S.; and Anwer, R. M. 2024. DB-SAM: Delving into High Quality Universal Medical Image Segmentation. In *Medical Image Computing and Computer Assisted Intervention*, 498–508.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer Assisted Intervention*, 234–241.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020. FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. In *Advances in Neural Information Processing Systems*, 596–608.
- Tang, F.; Wang, L.; Ning, C.; Xian, M.; and Ding, J. 2023. CMU-Net: A Strong ConvMixer-Based Medical Ultrasound Image Segmentation Network. In *IEEE 20th International Symposium on Biomedical Imaging*, 1–5.
- Tang, R.; and Ning, C. 2025. MLFEU-NET: A Multi-Scale Low-Level Feature Enhancement Unet for Breast Lesions Segmentation in Ultrasound Images. *Biomedical Signal Processing and Control*, 100: 106931.
- van den Heuvel, T. L. A.; de Bruijn, D.; de Korte, C. L.; and Ginneken, B. v. 2018. Automated Measurement of Fetal Head Circumference Using 2D Ultrasound Images. *PLOS ONE*, 13(8): 1–20.
- Walid, A.-D.; Mohammed, G.; Hussien, K.; and Aly, F. 2020. Dataset of Breast Ultrasound Images. *Data in Brief*, 28: 104863.
- Wang, X.; Bai, H.; Yu, L.; Zhao, Y.; and Xiao, J. 2024a. Towards the Uncharted: Density-Descending Feature Perturbation for Semi-Supervised Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3303–3312.
- Wang, Y.; Wang, H.; Shen, Y.; Fei, J.; Li, W.; Jin, G.; Wu, L.; Zhao, R.; and Le, X. 2022. Semi-Supervised Semantic Segmentation Using Unreliable Pseudo-Labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4248–4257.
- Wang, Z.; Zheng, J.; Zhang, Y.; Cui, G.; and Li, L. 2024b. Mamba-UNet: UNet-Like Pure Visual Mamba for Medical Image Segmentation. *arXiv preprint arXiv:2402.05079*.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. In *Advances in Neural Information Processing Systems*, 12077–12090.
- Yang, H.; and Yang, D. 2023. CSwin-PNet: A CNN-Swin Transformer Combined Pyramid Network for Breast Lesion Segmentation in Ultrasound Images. *Expert Systems with Applications*, 213: 119024.
- Yang, L.; Qi, L.; Feng, L.; Zhang, W.; and Shi, Y. 2023. Revisiting Weak-to-Strong Consistency in Semi-Supervised Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7236–7246.
- Yang, L.; Zhuo, W.; Qi, L.; Shi, Y.; and Gao, Y. 2022. ST++: Make Self-Training Work Better for Semi-Supervised Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4268–4277.
- Yap, M. H.; Pons, G.; Martí, J.; Ganau, S.; Sentís, M.; Zwigelaar, R.; Davison, A. K.; and Martí, R. 2018. Automated Breast Ultrasound Lesions Detection Using Convolutional Neural Networks. *IEEE Journal of Biomedical and Health Informatics*, 22(4): 1218–1226.
- Ye, Y.; Xie, Y.; Zhang, J.; Chen, Z.; and Xia, Y. 2023. UniSeg: A Prompt-Driven Universal Segmentation Model as Well as A Strong Representation Learner. In *Medical Image Computing and Computer Assisted Intervention*, 508–518.
- Zeng, Q.; Xie, Y.; Lu, Z.; Lu, M.; Wu, Y.; and Xia, Y. 2025. Segment Together: A Versatile Paradigm for Semi-Supervised Medical Image Segmentation. *IEEE Transactions on Medical Imaging*, 44(7): 2948–2959.
- Zhang, J.; Xie, Y.; Xia, Y.; and Shen, C. 2021. DoD-Net: Learning to Segment Multi-Organ and Tumors from Multiple Partially Labeled Datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1195–1204.
- Zhao, H.; Shi, J.; Qi, X.; Wang, X.; and Jia, J. 2017. Pyramid Scene Parsing Network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2881–2890.
- Zhao, Q.; Lyu, S.; Bai, W.; Cai, L.; Liu, B.; Wu, M.; Sang, X.; Yang, M.; and Chen, L. 2022. A Multi-Modality

Ovarian Tumor Ultrasound Image Dataset for Unsupervised Cross-Domain Semantic Segmentation. *arXiv preprint arXiv:2207.06799*.

Zhao, Z.; Yang, L.; Long, S.; Pi, J.; Zhou, L.; and Wang, J. 2023. Augmentation Matters: A Simple-Yet-Effective Approach to Semi-Supervised Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11350–11359.

Zhou, Z.; Rahman Siddiquee, M. M.; Tajbakhsh, N.; and Liang, J. 2018. UNet++: A Nested U-Net Architecture for Medical Image Segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, 3–11.