

DSFedMed: Dual-Scale Federated Medical Image Segmentation via Mutual Distillation Between Foundation and Lightweight Models

Hanwen Zhang, Qiaojin Shen, Yuxi Liu, Yuesheng Zhu, Guibo Luo*

Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology,
Shenzhen Graduate School, Peking University

hanwen@stu.pku.edu.cn, qjshen25@stu.pku.edu.cn, yuxiliu@stu.pku.edu.cn, zhuy@pku.edu.cn, luogb@pku.edu.cn

Abstract

Foundation Models (FMs) have demonstrated strong generalization across diverse vision tasks. However, their deployment in federated settings is hindered by high computational demands, substantial communication overhead, and significant inference costs. We propose DSFedMed, a dual-scale federated framework that enables mutual knowledge distillation between a centralized foundation model and lightweight client models for medical image segmentation. To support knowledge distillation, a set of high-quality medical images is generated to replace real public datasets, and a learnability-guided sample selection strategy is proposed to enhance efficiency and effectiveness in dual-scale distillation. This mutual distillation enables the foundation model to transfer general knowledge to lightweight clients, while also incorporating client-specific insights to refine the foundation model. Evaluations on five medical imaging segmentation datasets show that DSFedMed achieves an average 2 percent improvement in Dice score while reducing communication costs and inference time by nearly 90 percent compared to existing federated foundation model baselines. These results demonstrate significant efficiency gains and scalability for resource-limited federated deployments.

Code — <https://github.com/LMIAPC/DSFedMed>

Introduce

The rise of Foundation Models (FMs) (Bommasani et al. 2021) has brought a paradigm shift to Artificial Intelligence (AI) (Zhuang, Chen, and Lyu 2023). Models like Vision Transformers (ViTs) (Dosovitskiy et al. 2020; Steiner et al. 2021) and the Segment Anything Model (SAM) (Kirillov et al. 2023) demonstrate strong generalization across diverse vision tasks. Pre-trained on large-scale datasets (Guo and Yu 2022), these models offer transferable representations with promising applications in real-world domains such as medical imaging, where robustness and adaptability are critical.

While general-purpose FMs benefit from abundant web-scale data, domain-specific FMs, particularly in sensitive fields like healthcare, require access to high-quality private data that cannot be easily centralized due to strict privacy

*Corresponding author.

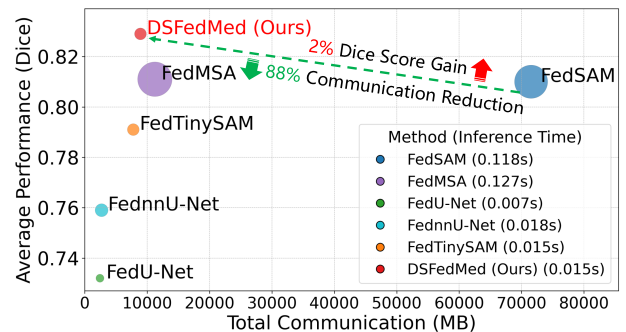


Figure 1: Comparison of federated segmentation methods in terms of communication cost and Dice performance. Bubble size represents inference time. Foundation model-based methods (e.g., FedSAM) achieve high accuracy but suffer from heavy communication and inference overhead. Lightweight models (e.g., FedU-Net) are efficient but sacrifice accuracy. Our proposed method, DSFedMed, achieves a better balance, offering both high accuracy and low computational and communication costs.

regulations (GDPR 2016; CCPA 2023) and institutional boundaries. This motivates the need for federated foundation models, where knowledge from distributed data silos can be aggregated without relying on direct access to local data.

Deploying foundation models directly in federated learning (FL) (McMahan et al. 2017; Konečný et al. 2016) presents significant challenges, especially in the medical imaging domain. Large foundation models require substantial computational resources and memory (Su, Li, and Xue 2024), which often exceed the capabilities of client devices commonly used in clinical environments (Abbas et al. 2024). Moreover, transmitting large model between clients and servers leads to heavy communication overhead, straining network bandwidth and increasing synchronization latency. Although parameter-efficient tuning methods like adapters can reduce training complexity to some extent (Malaviya, Shukla, and Lodha 2023; Cai et al. 2023; Chen et al. 2024), they do not solve the high inference cost and large memory requirement during deployment. These constraints are reflected in the practical trade-offs observed across federated

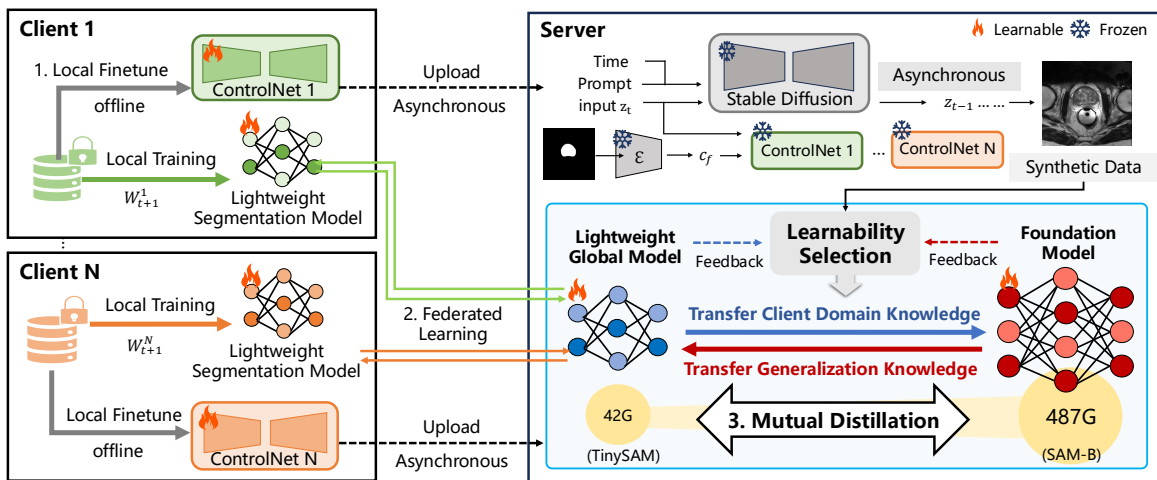


Figure 2: The Overview of DSFedMed.

ated segmentation methods. Figure 1 shows that existing approaches often prioritize either accuracy or efficiency, but rarely achieve both.

These practical limitations highlight the necessity for collaborative frameworks that integrate the strengths of large foundation models and lightweight models in a federated setting. Specifically, lightweight models deployed on client devices enable efficient, real-time inference grounded in domain-specific knowledge, while powerful foundation models hosted on the server offer comprehensive global insights and guidance. This hybrid federated structure is especially beneficial in medical imaging scenarios, where client devices are severely resource-constrained and privacy regulations restrict data sharing (Ali et al. 2022).

However, such dual-scale collaboration introduces several key challenges:

- How to effectively leverage the generalization capabilities of large-scale foundation models within highly resource-constrained, decentralized environments remains an open question.
- Due to privacy concerns and regulatory constraints, real medical data cannot be exchanged or centralized. This raises the need to enable knowledge transfer without access to real public datasets.
- Effective mutual knowledge distillation across heterogeneous model scales depends on selecting highly informative and reliable training samples, which makes learnability-guided sample selection crucial for efficient and accurate knowledge alignment.

In this paper, we propose **DSFedMed**, a dual-scale federated segmentation framework that addresses the challenges of deploying large foundation models in resource-constrained medical imaging environments. Specifically, our contributions:

- A dual-scale federated knowledge transfer framework enabling collaboration between server-side foundation

model and client-side lightweight model for medical image segmentation, with efficient client inference and communication for real-world deployment.

- A medical image generator that produces controllable and modality-adaptive samples aligned with the global data distribution based on ControlNet.
- A learnability-guided mutual knowledge distillation mechanism supported by generated samples to efficiently identify and transfer the most informative knowledge between heterogeneous models.

Our method provides a practical and scalable solution for integrating FMs into FL systems, offering new possibilities for privacy-compliant large-model deployment in real-world decentralized environments.

Method

Overview

We develop **DSFedMed**, a **Dual-Scale Federated Medical** image segmentation framework that enables collaborative knowledge transfer between a centralized foundation model and client-side lightweight models under federated setting. An overview of the framework is shown in Figure 2, which consists of three main components: (1) Controllable medical image generation based on client-specific generator. (2) Federated learning and alignment of dual-scale segmentation models. (3) Learnability-guided mutual distillation using generated samples.

In this framework, A lightweight segmentation model TinySAM (Shu et al. 2025), is deployed on each client for efficient local training and personalization. Meanwhile, a foundation model SAM-B (ViT-B/16), is hosted on the central server as a generalist model with strong generalization capacity. For simplicity, we refer to SAM-B as SAM throughout the paper.

To bridge the knowledge gap between these heterogeneous models while reducing exposure to sensitive data, we

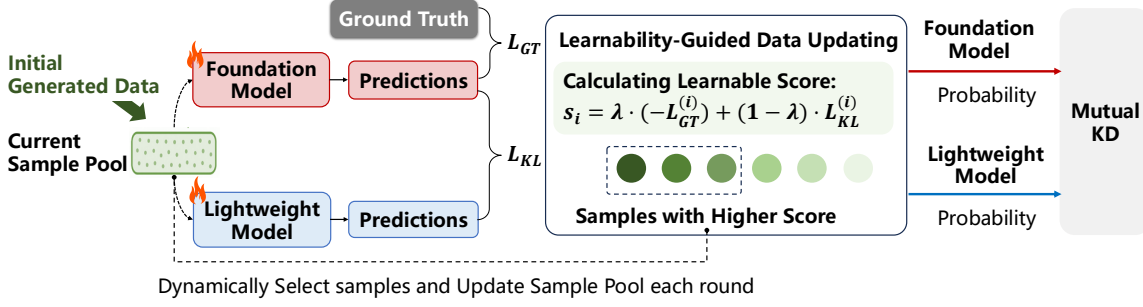


Figure 3: Illustration of proposed learnability-guided mutual distillation.

propose a mutual federated distillation framework supported by high-quality generated data. Specifically, each client finetunes a ControlNet (Zhang, Rao, and Agrawala 2023) module, initialized from released mask-conditioning weights, while keeping the backbone diffusion model frozen. This enables controllable and modality-adaptive medical image generation from structural masks. The ControlNet weights are then uploaded to the server and used to generate a global set of image-mask pairs, which are consistent in client-specific modality characteristics while preserving the structural information. These generated samples enable learnability-guided mutual distillation between server-side SAM and client-side TinySAM. By scoring each sample based on supervision reliability and model disagreement, we dynamically select the most informative subset to enhance SAM’s domain adaptation and TinySAM’s generalization, enabling efficient dual-scale collaboration.

Efficient Data Generation

To enable distillation without accessing real datasets, we construct an offline synthetic data generation pipeline based on ControlNet and Stable Diffusion (Rombach et al. 2022). Building on the general insight from prior work (Jin et al. 2025) that synthetic data can support distillation, our framework allows clients to finetune only lightweight ControlNet modules using private image-mask pairs while freezing the diffusion backbone. This reduces training cost and memory usage. A low VRAM mode is also available for 8GB GPUs. The framework enables client-specific generation without altering the pre-trained diffusion model structure.

Given a structural mask m as condition c , ControlNet integrates this control into the diffusion process by extending the frozen diffusion network $F(x; \Phi)$ with a trainable branch Φ_C , connected via zero-initialized 1×1 convolution layers $Z(\cdot; \Phi_{z1}), Z(\cdot; \Phi_{z2})$:

$$y_c = F(x; \Phi) + Z(F(x + Z(c_f; \Phi_{z1}); \Phi_C); \Phi_{z2}) \quad (1)$$

Initially, $y_c = F(x; \Phi)$, preserving stable feature flow; during training, Φ_C learns to conditionally modulate generation via encoded condition $c_f = E(c)$, achieving fine structural alignment with high image fidelity. After offline training, each client uploads its personalized ControlNet parameters $\Phi_C^{(k)}$ to the server. The server then produces image-mask

pairs from all clients’ ControlNets by sampling local masks $m \sim \mathcal{M}_k$ for each client k :

$$\tilde{D} = \bigcup_{k=1}^K \left\{ (\tilde{x}_k, m) \mid \tilde{x}_k = g^{\text{ctrl}}(m; \Phi_C^{(k)}), m \sim \mathcal{M}_k \right\} \quad (2)$$

This generated dataset enables knowledge distillation without relying on real data.

Collaborative Training of Dual-Scale Model

Each client $k \in \{1, \dots, K\}$ participates in the federated training of the lightweight TinySAM model f_k^{Tiny} on private data $\mathcal{D}_k^{\text{real}}$ by minimizing segmentation loss \mathcal{L}_{sup} :

$$\min_{\theta_k} \mathcal{L}_{\text{sup}}(f_k^{\text{Tiny}}(x_k), y_k), (x_k, y_k) \sim \mathcal{D}_k^{\text{real}} \quad (3)$$

After each local training iteration, clients upload their updated model parameters θ_k to the server. The server performs aggregation via federated averaging, weighting each client’s parameters by its dataset size:

$$\theta^{\text{Tiny}} = \sum_{k=1}^K \frac{n_k}{n} \theta_k, n_k = |\mathcal{D}_k^{\text{real}}|, n = \sum_{k=1}^K n_k \quad (4)$$

Meanwhile, the server uses the uploaded ControlNet to produce a global generated dataset \tilde{D} as described above. The global SAM model f^{SAM} is finetuned on \tilde{D} :

$$\min_{\theta^{\text{SAM}}} \mathcal{L}_{\text{sup}}(f^{\text{SAM}}(\tilde{x}), m), (\tilde{x}, m) \in \tilde{D} \quad (5)$$

Although TinySAM and SAM are trained on real and generated data respectively, the client-specific ControlNet modules ensure consistent semantic styles across generated samples, thereby fostering a shared representation space that facilitates knowledge alignment between two model scales.

Learnability-Guided Mutual Distillation

To enable effective knowledge exchange between heterogeneous models, we propose a learnability-guided mutual distillation mechanism that iteratively selects and updates a dynamic sample pool from initially generated data, as illustrated in Figure 3. The process begins with an initial generated dataset $\tilde{D} = \{(x_i, m_i)\}_{i=1}^N$, where each image x_i

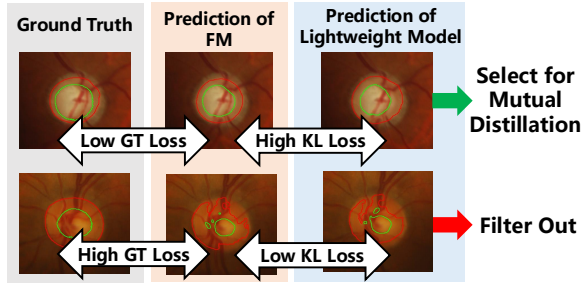


Figure 4: Examples of Learnability-Guided Selection.

is generated from a structural mask m_i as described above. Rather than using the entire dataset uniformly, we maintain a current sample pool that is continuously updated each round through a dynamic selection strategy guided by sample-wise informativeness. To evaluate each sample’s value for mutual distillation, we compute a learnability score x_i by combining two complementary metrics: the ground truth loss (GT-loss) with respect to SAM’s prediction, and the mutual KL divergence measuring the discrepancy between SAM and TinySAM outputs:

$$s_i = \lambda \cdot (-L_{GT}^{(i)}) + (1 - \lambda) \cdot L_{KL}^{(i)}, \lambda \in [0, 1] \quad (6)$$

where:

$$\begin{aligned} \mathcal{L}_{GT}^{(i)} &= \ell(f_{SAM}(x_i), m_i), \\ \mathcal{L}_{KL}^{(i)} &= D_{KL}(p_{SAM}(x_i) \parallel p_{TinySAM}(x_i)) \\ &\quad + D_{KL}(p_{TinySAM}(x_i) \parallel p_{SAM}(x_i)) \end{aligned} \quad (7)$$

Here, L_{GT} ensures that the sample provides reliable SAM supervision, while L_{KL} prioritizes samples where predictions diverge, indicating a greater potential for improvement between dual-scale models.

During training, we dynamically select samples based on updated scores using a fixed selection rate, refreshing the sample pool each round. This sample selection process can be viewed as an implicit mutual distillation, where the most informative examples guide the knowledge transfer between models:

$$\mathbb{E}_{x \sim a(x)} [D_{KL}(p_{SAM}(x) \parallel p_{TinySAM}(x)) + D_{KL}(p_{TinySAM}(x) \parallel p_{SAM}(x))] \quad (8)$$

where the sampling distribution $a(x) \propto \exp(s_i)$ biases towards samples that maximize mutual knowledge transfer, both SAM and TinySAM are fully trainable during this process. This learnability-guided scheme accelerates convergence by focusing on informative and reliable supervision.

Overall, our mutual distillation framework exploits a curated subset to transfer general representations from SAM to TinySAM, enhancing the lightweight model’s generalization, while injecting domain-specific knowledge from TinySAM back to SAM via disagreement samples. The example of learnability-guided selection is shown in Figure 4.

Type	Method
Centralized	SAM
Federated Foundation Model	FedSAM FedMSA
Federated Lightweight Segmentation Model	FedU-Net FednnU-Net FedTinySAM
Dual-Scale FL (Server: FMs; Client: Lightweight Models)	DSFedMed

Table 1: Categorization of compared baselines.

Experiments

Datasets

We constructed five Non-IID federated datasets from public sources, covering diverse modalities and tasks. Following the preprocessing pipeline of MSA (Wu et al. 2025), all input images were resized to 1024×1024 to retain detail, and the corresponding segmentation masks are standardized to 256×256 resolution for efficiency.

- **Fundus:** Collected from independent medical centers and captured by different scanning devices (Sivaswamy et al. 2015; Fumero et al. 2011; Orlando et al. 2020), targeting segmentation of the Optic Disc (OD) and Optic Cup (OC).
- **Prostate Cancer:** Derived from publicly available MRI datasets collected across multiple medical centers (Litjens et al. 2014; Lemaître et al. 2015).
- **Nuclei:** Sourced from several nuclei segmentation datasets (Gamper et al. 2019, 2020; Kumar et al. 2017; Naylor et al. 2018; Verma et al. 2021), with samples from different tissue types in the PanNuke dataset distributed across clients to introduce heterogeneity.
- **ISIC:** Derived from the ISIC 2017 dataset (Codella et al. 2018) for skin lesion segmentation. To introduce heterogeneity, the training set was partitioned into three subsets based on skin tone (Bevan and Atapour-Abarghouei 2022), each assigned to a distinct client.
- **CHAOS:** Sourced from a classic benchmark providing paired CT and MRI data for liver segmentation (Kavur et al. 2021). Two modalities were allocated to separate clients to simulate cross-modality federated learning.

For the real-world multi-center Fundus, Prostate, and Nuclei datasets, we use leave-one-domain-out cross-validation, holding out each client as the test domain in turn while others train collaboratively. Since ISIC and CHAOS are not naturally multi-center, we manually split them and adopt a standard federated learning setup with all clients training concurrently and evaluation on an independent public test set.

Compared Baselines

As shown in Table 1, we compare DSFedMed with the following baselines. Unless otherwise specified, both SAM and MSA are implemented with ViT-B/16 (Dosovitskiy et al. 2020; Steiner et al. 2021). When using the larger variant ViT-L/16, we denote the models as SAM-L and MSA-L.

Datasets	Fundus-OD		Fundus-OC		Prostate		Nuclei		ISIC		CHAOS	
Method	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice	IoU
SAM (ceiling)	0.914	0.850	<u>0.782</u>	<u>0.673</u>	0.801	0.762	0.617	0.467	0.840	0.755	0.938	0.895
FedSAM	<u>0.926</u>	<u>0.869</u>	0.771	0.661	0.814	0.781	<u>0.623</u>	<u>0.480</u>	0.822	0.735	0.902	0.848
FedMSA	<u>0.926</u>	0.868	0.764	0.647	0.799	0.764	0.617	0.472	<u>0.840</u>	<u>0.755</u>	0.917	0.864
FedU-Net	0.896	0.820	0.733	0.600	0.517	0.490	0.590	0.435	0.797	0.708	0.860	0.788
FednnU-Net	0.907	0.837	0.743	0.614	0.590	0.516	0.603	0.450	0.819	0.726	0.894	0.843
FedTinySAM	0.917	0.852	0.738	0.620	0.756	0.723	0.620	0.473	0.832	0.747	0.883	0.831
DSFedMed	0.949	0.906	0.797	0.686	<u>0.813</u>	<u>0.777</u>	0.644	0.498	0.842	0.758	<u>0.927</u>	<u>0.876</u>

Table 2: Segmentation results of different methods on five medical datasets. The best results are highlighted in bold and the suboptimal ones are underlined.

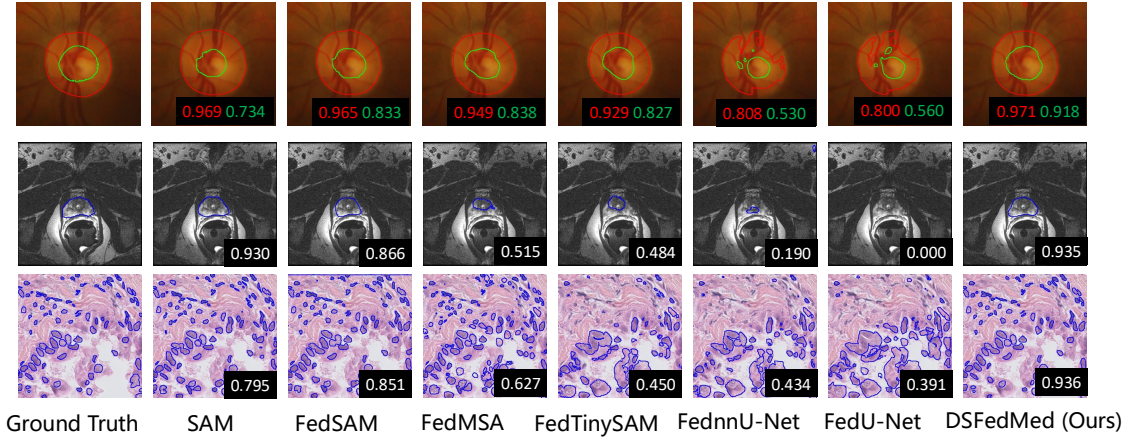


Figure 5: The example of datasets and the comparison in the results of SAM, FedSAM, FedMSA, FedU-Net, FednnU-Net, FedTinySAM and DSFedMed. The average Dice of the method on the dataset is marked in the bottom right corner.




Figure 6: Asynchronous training design of DSFedMed. (Left: The distribution of asynchronous training time. Right: Extra training time of DSFedMed and FedSAM for adding a new client.)

- **Centralized**, in which the server’s foundation model SAM is finetuned using the datasets combining private datasets of involved clients and the public dataset.
- **Federated Foundation Model**, on each client we adopt FedSAM to finetune the full SAM and FedMSA (Liu, Luo, and Zhu 2024) to efficiently finetune the MSA adapter and decoder.

- **Lightweight Federated Learning Baselines**, we adopt FedU-Net, FednnU-Net and FedTinySAM as our federated segmentation baselines. Both of them are implemented under the FedAvg framework (McMahan et al. 2017). U-Net (Ronneberger, Fischer, and Brox 2015) is a widely used architecture for biomedical image segmentation, while nnU-Net (Isensee et al. 2018) is a more robust extension of U-Net. TinySAM is a lightweight variant of SAM designed for edge efficiency.

Efficiency Analysis

Segmentation Result

Across all tasks and clients, DSFedMed consistently outperforms lightweight FL methods. DSFedMed achieves an average Dice improvement of 3.8% across five datasets compared to FedTinySAM. Moreover, DSFedMed outperforms not only lightweight models but also strong foundation model baselines: it achieves a 1.4% gain over centralized SAM, and surpasses FedSAM and the FedMSA by 1.9% each. By leveraging high-quality generated data to facilitate mutual distillation, DSFedMed bridges the gap with foundation models and even exceeds their results in federated settings. We provide the results of the experiment in Table 2 and show the visualization results in Figure 5.

Method	Inference Model	Communication Overhead	Inference Time (per image)	Synchronous Training Time	Asynchronous Training Time	Average Dice Score
FedSAM	SAM-B	71,538 MB	0.118 s	1,411.67 min	/	0.810
FedMSA	MSA-B	11,230 MB	0.127 s	1,160.00 min	/	0.811
FedU-Net	U-Net	2,430 MB	0.007 s	59.50 min	/	0.732
FednnU-Net	nnU-Net	2,690 MB	0.018 s	345.68 min	/	0.759
FedTinySAM	TinySAM	7,770 MB	0.015 s	198.02 min	/	0.791
DSFedMed	TinySAM	8,920 MB	0.015 s	198.02 min	455.5 min	0.829

Table 3: Efficiency analysis of FL baselines and DSFedMed.

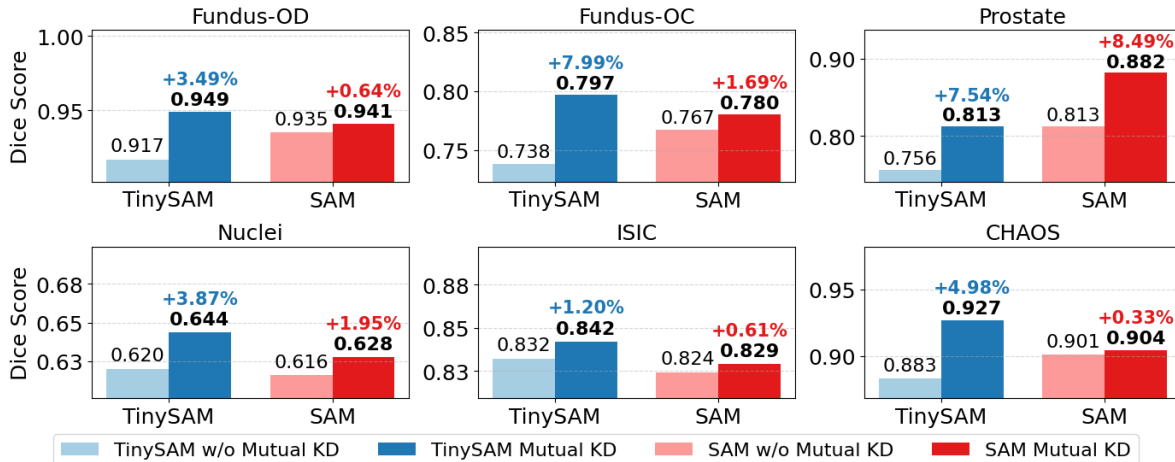


Figure 7: Visualization of mutual enhancement across dual-scale models.

As shown in Table 3, we compare federated learning methods in terms of total communication overhead, inference latency, training time, and segmentation performance. Communication overhead is computed cumulatively over 100 communication rounds, and training time is reported on the Prostate dataset.

FedSAM achieves a Dice score of 0.810 but is constrained by high communication cost, long training time, and slow inference. FedMSA improves efficiency via adapter-based tuning but incurs additional inference cost. Lightweight methods (e.g., FedU-Net, FednnU-Net) offer fast and efficient deployment but lack sufficient accuracy.

In contrast, DSFedMed strikes an optimal balance between segmentation accuracy and deployment efficiency, outperforming competing methods comprehensively. Compared to FedSAM, it reduces communication cost by a substantial 88% while boosting the Dice score by 2%, achieving highest accuracy of 0.829. Additionally, DSFedMed maintains a low inference time of 0.015 seconds, demonstrating both high accuracy and practical viability.

The asynchronous training phase includes client-side generator training and server-side data generation and mutual knowledge distillation, as shown in Figure 6. Most asynchronous computations, including data generation and knowledge distillation, occur on the server without clients needing to stay online. This design enables new clients to join without full retraining. The server-side processing is ef-

ficient and can be performed on a single GPU, such as an NVIDIA A800 or A6000.

Mutual Enhancement of Dual-Scale Model

To assess the effectiveness of our dual-scale distillation framework, we compare the performance of both the lightweight model (TinySAM) and the foundation model (SAM) with and without mutual distillation. Experiments across five datasets demonstrate that mutual distillation consistently enhances both model scales as shown in Figure 7.

TinySAM achieves notable improvements, especially under domain shift, with Dice scores increasing on datasets like Prostate (+5.7%) and Fundus-OC (+5.9%). Performance gains are also observed on ISIC and CHAOS, confirming its enhanced generalization. Meanwhile, SAM improves on client-specific data, such as Prostate and Fundus-OD, indicating better domain adaptation.

These results demonstrate effective mutual knowledge exchange: SAM provides semantic priors to guide TinySAM, while TinySAM offers localized feedback to refine SAM.

Ablation Studies

We conducted an ablation study to evaluate the contributions of Mutual Knowledge Distillation (Mutual KD) and Learnability-Guided sample selection (LG selection) as shown in Table 4. Replacing unidirectional distillation, Mutual KD consistently improves segmentation across all five

Method		Fundus-OD		Fundus-OC		Prostate		Nuclei		ISIC		Chaos		time(s)
Mutual KD	LG Selection	Dice	IOU	Dice	IOU	Dice	IOU	Dice	IOU	Dice	IOU	Dice	IOU	
		0.944	0.896	0.778	0.662	0.780	0.741	0.627	0.485	0.823	0.732	0.893	0.828	399.010
✓		0.946	0.900	0.781	0.656	0.791	0.754	0.634	0.488	0.828	0.738	0.916	0.869	399.349
	✓	0.945	0.900	0.788	0.674	0.794	0.764	0.636	0.491	0.839	0.752	0.911	0.852	209.827
✓	✓	0.949	0.906	0.797	0.686	0.813	0.777	0.644	0.498	0.842	0.758	0.927	0.876	211.113

Table 4: Ablation study of key components in the proposed method. The best results are highlighted in bold.

Dual-Scale FL (Server / Client)	Fundus-OD		Fundus-OC	
	Dice	IoU	Dice	IoU
SAM-B / U-Net	0.940	0.891	0.739	0.613
SAM-B / TinySAM	0.949	0.906	0.797	0.686
MSA-L / U-Net	0.934	0.882	0.781	0.661
MSA-L / TinySAM	0.944	0.897	0.795	0.681

Table 5: Performance of proposed Dual-Scale FL framework with different server/client model scales. SAM-B and MSA-L denote Base and Large variants, respectively.

Datasets	Fundus		Prostate		Nuclei	
	FID	FLD	FID	FLD	FID	FLD
Pix2Pix	77.19	8.27	143.96	6.87	172.12	9.82
DDPM	34.86	9.33	86.57	4.96	156.65	8.22
ControlNet	7.83	2.98	62.77	1.57	148.73	4.64

Table 6: Quantitative evaluation of the generated data. The best results are highlighted in bold.

datasets, highlighting the advantage of bidirectional knowledge exchange. LG selection alone boosts accuracy and cuts training time by nearly 50%. Their combination achieves the highest Dice and IoU scores with training efficiency matching LG selection alone, confirming their complementary roles in our framework.

Performance Across Different Model Scales

The result in Table 5 shows that the proposed Dual-Scale Federated Learning framework consistently improves segmentation performance across combinations of server and client models on Fundus dataset. We evaluate with SAM-B (ViT-B/16) and Medical SAM Adapter-L (MSA-L, ViT-L/16) as server-side models, and U-Net and TinySAM as lightweight client models, demonstrating the framework’s general applicability across model scales.

Sensitivity to Hyperparameters

Experimental results show that sample selection consistently outperforms the baseline using all samples (selection rate = 1.0), achieving stable improved accuracy across selection rates from 0.2 to 0.8 while reducing training cost. Varying the weighting parameter λ between 0.1 and 0.9 induces only modest fluctuations in segmentation performance, with optimal outcomes observed near $\lambda = 0.5$. These findings confirm the effectiveness and robustness of our framework. Figure 8 summarizes the results.




Figure 8: Sensitivity to hyperparameters.

Quality of Generated Data

In terms of generated data, we measured and evaluated three methods: pix2pix(Isola et al. 2017), DDPM(Ho, Jain, and Abbeel 2020), and ControlNet(Zhang, Rao, and Agrawala 2023), using Fréchet Inception Distance (FID) (Heusel et al. 2017) and the Feature Likelihood Difference (FLD) (Jiralerpong et al. 2023) to assess the quality of synthetic data. FID measures feature distribution similarity, while FLD assesses novelty, realism, and diversity via density estimation. ControlNet achieves the best balance of precision and controllability, outperforming the others. Results are in Table 6.

Conclusion

In this work, we proposed DSFedMed, a Dual-Scale Federated learning framework that enables collaboration between a centralized foundation model and distributed lightweight client models for Medical image segmentation. By integrating an efficient data generator and a learnability-guided sample selection strategy, our approach facilitates efficient mutual knowledge distillation without requiring access to real medical data or direct communication of foundation model.

Our results demonstrate the practicality and scalability of DSFedMed in resource-constrained federated settings, offering a viable path toward deploying foundation models in privacy-sensitive domains such as healthcare. The method balances global semantic generalization and local domain adaptability, addressing both performance and deployment limitations inherent in medical federated learning.

Nevertheless, several open challenges remain. These include accelerating data generation stage and validating the framework in real clinical settings. Future work may also extend DSFedMed to multimodal scenarios.

Acknowledgments

This work is supported by National Key Research and Development Program of China (2024YFE0203100), Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology (Grant No. 2024B1212010006), Shenzhen Science and Technology Program (JCYJ20230807120800001), and the project supplementary funding of National Innovation 2030 Major S&T Project of China (2020AAA0104203).

References

- Abbas, S. R.; Abbas, Z.; Zahir, A.; and Lee, S. W. 2024. Federated learning in smart healthcare: a comprehensive review on privacy, security, and predictive analytics with IoT integration. In *Healthcare*, volume 12, 2587. MDPI.
- Ali, M.; Naem, F.; Tariq, M.; and Kaddoum, G. 2022. Federated learning for privacy preservation in smart healthcare systems: A comprehensive survey. *IEEE journal of biomedical and health informatics*, 27(2): 778–789.
- Bevan, P. J.; and Atapour-Abarghouei, A. 2022. Detecting melanoma fairly: Skin tone detection and debiasing for skin lesion classification. In *MICCAI Workshop on Domain Adaptation and Representation Transfer*, 1–11. Springer.
- Bommasani, R.; Hudson, D. A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M. S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Cai, D.; Wu, Y.; Wang, S.; Lin, F. X.; and Xu, M. 2023. Efficient federated learning for modern nlp. In *Proceedings of the 29th annual international conference on mobile computing and networking*, 1–16.
- CCPA. 2023. California Consumer Privacy Act of 2018. Assembly Bill No. 375. As amended by the California Privacy Rights Act (CPRA), effective January 1, 2023. CCPA.
- Chen, H.; Zhang, Y.; Krompass, D.; Gu, J.; and Tresp, V. 2024. Feddat: An approach for foundation model finetuning in multi-modal heterogeneous federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 11285–11293.
- Codella, N. C.; Gutman, D.; Celebi, M. E.; Helba, B.; Marchetti, M. A.; Dusza, S. W.; Kallou, A.; Liopyris, K.; Mishra, N.; Kittler, H.; et al. 2018. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*, 168–172. IEEE.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Fumero, F.; Alayón, S.; Sanchez, J. L.; Sigut, J.; and Gonzalez-Hernandez, M. 2011. RIM-ONE: An open retinal image database for optic nerve evaluation. In *2011 24th international symposium on computer-based medical systems (CBMS)*, 1–6. IEEE.
- Gamper, J.; Alemi Koohbanani, N.; Benet, K.; Khuram, A.; and Rajpoot, N. 2019. Pannuke: an open pan-cancer histology dataset for nuclei instance segmentation and classification. In *European congress on digital pathology*, 11–19. Springer.
- Gamper, J.; Koohbanani, N. A.; Benes, K.; Graham, S.; Jahanifar, M.; Khuram, S. A.; Azam, A.; Hewitt, K.; and Rajpoot, N. 2020. Pannuke dataset extension, insights and baselines. *arXiv preprint arXiv:2003.10778*.
- GDPR. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council. Official Journal of the European Union. General Data Protection Regulation (GDPR).
- Guo, X.; and Yu, H. 2022. On the domain adaptation and generalization of pretrained language models: A survey. *arXiv preprint arXiv:2211.03154*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Isensee, F.; Petersen, J.; Klein, A.; Zimmerer, D.; Jaeger, P. F.; Kohl, S.; Wasserthal, J.; Koehler, G.; Norajitra, T.; Wirkert, S.; et al. 2018. nnu-net: Self-adapting framework for u-net-based medical image segmentation. *arXiv preprint arXiv:1809.10486*.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.
- Jin, H.; Liu, S.; Cong, C.; Feng, Q.; Liu, Y.; Huang, L.; and Hu, Y. 2025. FedWSIDD: Federated Whole Slide Image Classification via Dataset Distillation. *arXiv:2506.15365*.
- Jiralerspong, M.; Bose, J.; Gemp, I.; Qin, C.; Bachrach, Y.; and Gidel, G. 2023. Feature likelihood divergence: evaluating the generalization of generative models using samples. *Advances in Neural Information Processing Systems*, 36: 33095–33119.
- Kavur, A. E.; Gezer, N. S.; Barış, M.; Aslan, S.; Conze, P.-H.; Groza, V.; Pham, D. D.; Chatterjee, S.; Ernst, P.; Özkan, S.; et al. 2021. CHAOS challenge-combined (CT-MR) healthy abdominal organ segmentation. *Medical image analysis*, 69: 101950.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4015–4026.
- Konečný, J.; McMahan, H. B.; Yu, F. X.; Richtárik, P.; Suresh, A. T.; and Bacon, D. 2016. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*.
- Kumar, N.; Verma, R.; Sharma, S.; Bhargava, S.; Vahadane, A.; and Sethi, A. 2017. A dataset and a technique for generalized nuclear segmentation for computational pathology. *IEEE transactions on medical imaging*, 36(7): 1550–1560.

- Lemaître, G.; Martí, R.; Freixenet, J.; Vilanova, J. C.; Walker, P. M.; and Meriaudeau, F. 2015. Computer-aided detection and diagnosis for prostate cancer based on mono and multi-parametric MRI: a review. *Computers in biology and medicine*, 60: 8–31.
- Litjens, G.; Toth, R.; Van De Ven, W.; Hoeks, C.; Kerkstra, S.; Van Ginneken, B.; Vincent, G.; Guillard, G.; Birbeck, N.; Zhang, J.; et al. 2014. Evaluation of prostate segmentation algorithms for MRI: the PROMISE12 challenge. *Medical image analysis*, 18(2): 359–373.
- Liu, Y.; Luo, G.; and Zhu, Y. 2024. Fedfms: Exploring federated foundation models for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 283–293. Springer.
- Malaviya, S.; Shukla, M.; and Lodha, S. 2023. Reducing communication overhead in federated learning for pre-trained language models using parameter-efficient finetuning. In *Conference on Lifelong Learning Agents*, 456–469. PMLR.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Naylor, P.; Laé, M.; Reyat, F.; and Walter, T. 2018. Segmentation of nuclei in histopathology images by deep regression of the distance map. *IEEE transactions on medical imaging*, 38(2): 448–459.
- Orlando, J. I.; Fu, H.; Breda, J. B.; Van Keer, K.; Bathula, D. R.; Diaz-Pinto, A.; Fang, R.; Heng, P.-A.; Kim, J.; Lee, J.; et al. 2020. Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Medical image analysis*, 59: 101570.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Shu, H.; Li, W.; Tang, Y.; Zhang, Y.; Chen, Y.; Li, H.; Wang, Y.; and Chen, X. 2025. Tinsam: Pushing the envelope for efficient segment anything model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 20470–20478.
- Sivaswamy, J.; Krishnadas, S.; Chakravarty, A.; Joshi, G.; Tabish, A. S.; et al. 2015. A comprehensive retinal image dataset for the assessment of glaucoma from the optic nerve head analysis. *JSM Biomedical Imaging Data Papers*, 2(1): 1004.
- Steiner, A.; Kolesnikov, A.; Zhai, X.; Wightman, R.; Uszkoreit, J.; and Beyer, L. 2021. How to train your vit? data, augmentation, and regularization in vision transformers. *arXiv preprint arXiv:2106.10270*.
- Su, S.; Li, B.; and Xue, X. 2024. Fedra: A random allocation strategy for federated tuning to unleash the power of heterogeneous clients. In *European Conference on Computer Vision*, 342–358. Springer.
- Verma, R.; Kumar, N.; Patil, A.; Kurian, N. C.; Rane, S.; Graham, S.; Vu, Q. D.; Zwager, M.; Raza, S. E. A.; Rajpoot, N.; et al. 2021. MoNuSAC2020: A multi-organ nuclei segmentation and classification challenge. *IEEE Transactions on Medical Imaging*, 40(12): 3413–3423.
- Wu, J.; Wang, Z.; Hong, M.; Ji, W.; Fu, H.; Xu, Y.; Xu, M.; and Jin, Y. 2025. Medical sam adapter: Adapting segment anything model for medical image segmentation. *Medical image analysis*, 102: 103547.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.
- Zhuang, W.; Chen, C.; and Lyu, L. 2023. When foundation model meets federated learning: Motivations, challenges, and future directions. *arXiv preprint arXiv:2306.15546*.