

ADAPTING FOUNDATION MODELS FOR FEW-SHOT MEDICAL IMAGE SEGMENTATION TASKS: ACTIVELY AND SEQUENTIALLY

Jingyun Yang, Guoqing Zhang, Jingge Wang, Yang Li

Shenzhen Key Laboratory of Ubiquitous Data Enabling
Tsinghua Shenzhen International Graduate School, Tsinghua University

ABSTRACT

Recent advances in foundation models have brought promising results in computer vision, including medical image segmentation. Fine-tuning foundation models on specific low-resource medical tasks has become a standard practice. However, ensuring reliable and robust model adaptation when the target task has a large domain gap and few annotated samples remains a challenge. Previous few-shot domain adaptation (FSDA) methods seek to bridge the distribution gap between source and target domains by utilizing auxiliary data. The selection and scheduling of auxiliaries are often based on heuristics, which can easily cause negative transfer. In this work, we propose an Active and Sequential domain AdaPtation (ASAP) framework for dynamic auxiliary dataset selection in FSDA. We formulate FSDA as a multi-armed bandit problem and derive an efficient reward function to prioritize training on auxiliary datasets that align closely with the target task, through a single-round fine-tuning. Empirical validation on diverse medical segmentation datasets demonstrates that our method achieves favorable segmentation performance, significantly outperforming the state-of-the-art FSDA methods. Code is available at ASAP.

Index Terms— Few-shot domain adaptation, auxiliary learning, active learning, medical image segmentation.

1. INTRODUCTION

Recent works like SwinUNet [1], MambaUNet [2] and MONAI [3] develop medical-tailored foundation models on large-scale medical image datasets. Intense interest has emerged in adapting these foundation models for specific medical image analysis tasks. However, the generalization capability of foundation models is limited by the large variability in training data, due to complex modalities, intricate anatomical structures, and wide-range object scales in medical images. Therefore, we seek to answer this critical question: how to effectively adapt these foundation models to our desired medical image processing tasks?

Unlike natural image analysis with large-scale labeled datasets, in medical image analysis, annotating disease-specific medical images is not only time-consuming but also

demanding specialty-oriented skills, leading to the problem of few-shot domain adaptation (FSDA). Most solutions to conventional domain adaptation problems either assume access to source data [4], which is not always feasible in real-world medical scenarios with various regulatory standards and ethical considerations, or they require a substantial amount of unlabeled target data to reduce the distribution gap across domains, as seen in unsupervised domain adaptation (UDA) methods [5]. FSDA, on the other hand, addresses the situation when only a limited number of target examples are available for training, whether labeled or unlabeled. Previous FSDA methods [6] propose to use intermediate/auxiliary domains to facilitate domain adaptation. However, this multi-step domain adaptation strategy requires fine-tuning the model twice or more. In this work, we propose to incorporate auxiliary datasets to solve the FSDA problem in a source-free manner through a single-round fine-tuning.

Training with auxiliary data introduces an inductive bias that helps models capture meaningful representations and reduces the risk of overfitting to spurious correlations [7]. Multi-task learning methods [8] cannot extend to a large number of tasks because the complexity of the search space will be exponentially explosive [9]. Other strategies in auxiliary learning and transfer learning hand-pick which auxiliary data to use based on heuristics [10] or metrics [11] prior to training, sometimes resulting in sub-optimal outcomes. Recent dynamic auxiliary learning works [7] propose to dynamically combine auxiliary objectives through task or data schedulers, but these methods involve complex and computationally demanding bi-level optimization steps.

To address the above issues, we propose an Active and Sequential domain AdaPtation (ASAP) framework for FSDA. Using a novel dynamic dataset selection strategy, the proposed framework prioritizes training on auxiliary datasets with similar solution spaces to the target task in a **single-round** computational complexity. Specifically, we formulate FSDA as a multi-armed bandit problem in active learning [12] and relate the set of auxiliary datasets to the arms. We introduce the classic trace upper confidence bound algorithm [13] to solve the multi-armed bandit problem. By balancing the trade-off between the exploration of unobserved arms and the exploitation of high-reward arms, we actively and sequen-

tially select the auxiliary dataset at each turn, maximizing their benefits. The reward functions we design add minimal memory and computational overhead.

Extensive experiments on three public medical datasets validate the effectiveness of our proposed ASAP framework. We efficiently adapt pre-trained UNet [14], SwinUNet [1] and MambaUNet [2] from MONAI [3], a freely available open-source medical resource platform, for various target medical image segmentation tasks. Our method outperforms the FSDA auxiliary learning methods with lower computation costs. Our main contributions are as follows:

- **An active and sequential domain adaptation framework:** we propose a novel framework that incorporates auxiliary datasets to effectively adapt foundation models to the target medical segmentation task in a single-round fine-tuning, optimizing the use of public medical resources.
- **An exploration-exploitation balanced FSDA algorithm:** we design an efficient reward function and successfully apply the multi-armed bandit algorithm to dynamic auxiliary dataset selection through the ASAP framework.

2. METHODOLOGY

In this section, we will elaborate on the proposed active and sequential domain adaptation (ASAP) framework, shown in Fig. 1. First, we clarify the setting of few-shot domain adaptation with auxiliary datasets. Then we formulate it as a multi-armed bandit (MAB) problem and describe how we solve it.

2.1. Problem definition

For domain adaptation problems, the network is usually first trained on an adequate source domain dataset \mathcal{D}_S . We denote the pre-trained source model as Θ_s . Given a small quantity of data belonging to a target domain dataset $\mathcal{D}_T = \{(x_i^t, y_i^t)\}_{i=1}^m$, the goal is to adapt Θ_s to achieve high performance on \mathcal{D}_T with access to a set of available auxiliary datasets $\mathcal{D}_A = \{\mathcal{D}_{a_1}, \mathcal{D}_{a_2}, \dots, \mathcal{D}_{a_K}\}$. For all $a \in A$, \mathcal{D}_a is an individual auxiliary dataset.

In this work, we formulate the auxiliary data selection problem in FSDA as a Markov decision process by adopting the multi-armed bandit (MAB) setting [12]. MAB has been used in sequential experiment design in active learning, where the goal is to sequentially choose experiments to perform with the aim of maximizing some outcomes. The MAB learning paradigm involves an agent interacts with an environment over N turns by following a policy π . In our work, the environment consists of the target dataset \mathcal{D}_T , the set of auxiliary datasets \mathcal{D}_A , and the model f_θ . The agent learns a policy π that defines the selection strategy over all $\mathcal{D}_a \in \mathcal{D}_A$. At each turn t , the agent selects one of the environment’s K datasets $\mathcal{D}_a \in \mathcal{D}_A$ to jointly train with \mathcal{D}_T . The environment then updates the model f_θ . Accordingly, the agent receives a reward $R_{a,t}$ and uses it to update the policy π . Rewards for unplayed arms are not observed. The goal of the

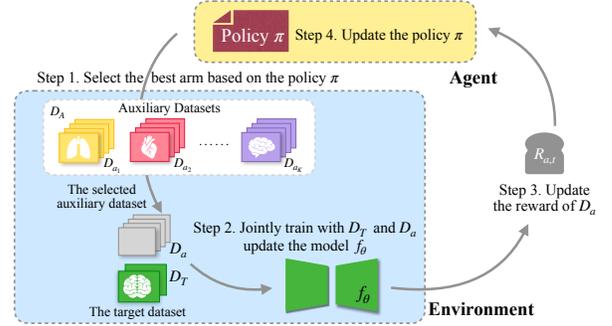


Fig. 1: Illustration of our active and sequential domain adaptation (ASAP) framework. The agent defines the policy π that determines which arm to pull. The environment includes the auxiliary data pool \mathcal{D}_A , the target dataset \mathcal{D}_T , and the model f_θ . At each turn t , ASAP executes the four shown steps.

agent is to adopt a policy π that selects actions that lead to the largest cumulative reward over N turns, $R = \sum_{t=1}^N R_t$.

2.2. Deriving an efficient reward function

To ensure that the decision-making process adds minimal memory and computational overhead, we derive rewards from the model’s intrinsic information and the optimized losses, rather than relying on an external model or metric that requires extra training. To achieve positive transfer during sequential adaptation, we design the reward with two considerations in mind: positive for model convergence and positive for joint training with the target task.

Formally, at turn t for the auxiliary dataset \mathcal{D}_a the reward of positive for model convergence is defined as:

$$\mathcal{R}_{a,t}^{PM} = -\mathcal{L}_{a,t} = -\mathcal{L}(f_{\theta_t}, \mathcal{D}_a). \quad (1)$$

Let $\nabla_a = \nabla_{\theta} \mathcal{L}(f_{\theta_t}, \mathcal{D}_a)$ be the auxiliary dataset gradient and $\nabla_T = \nabla_{\theta} \mathcal{L}(f_{\theta_t}, \mathcal{D}_T)$ be the target dataset gradient, we denote the reward of positive for joint training with \mathcal{D}_T as:

$$\mathcal{R}_{a,t}^{PT} = \frac{\nabla_a \cdot \nabla_T}{\|\nabla_a\|_2 \|\nabla_T\|_2}. \quad (2)$$

Overall, at turn t the reward of the auxiliary dataset \mathcal{D}_a is defined as:

$$\mathcal{R}_{a,t} = \alpha \mathcal{R}_{a,t}^{PM} + (1 - \alpha) \mathcal{R}_{a,t}^{PT}, \quad (3)$$

where α is a time-variant weight that decreases with each selection turn. Considering that $\mathcal{R}_{a,t}$ relies on the loss and gradients, which are intrinsic to the model, $\mathcal{R}_{a,t}$ is naturally training-efficient.

2.3. The decision-making policy

To optimize a decision-making policy, we propose to adopt a trace upper confidence bound (UCB) algorithm [13] where the agent greedily selects arms according to their upper confidence bound. Formally, after pulling the arm a at turn t , the agent receives a observed reward $R_{a,t}$ and then calculate the estimated mean reward as:

$$\hat{R}_a = (1 - \beta) \hat{R}_a + \beta R_{a,t}, \quad (4)$$

where β is the smoothing factor [15]. Accordingly, we define the upper confidence bound based on Hoeffding’s inequality [13] for arm a at turn t being played n_a times:

$$UCB_{a,t} = \begin{cases} \infty, & \text{if } n_a = 0 \\ \hat{R}_a + \sqrt{\frac{2 \ln t}{n_a}}, & \text{otherwise.} \end{cases} \quad (5)$$

This allows us to balance the exploitation of arms with a high predicted reward and the exploration of areas with high uncertainty. The proposed algorithm is shown in Algorithm. 1.

Algorithm 1 The MAB decision-making policy

Require: $\mathcal{D}_A, \mathcal{D}_T$: Auxiliary and target datasets

Require: f_θ : Parameterized model

Require: α, β : Decaying and smoothing factors

- 1: **Initialize** $f_{\theta_0} = \Theta_s$
 - 2: **Initialize** the information of each arm $a \in A$
 $\forall a \in A: n_a = 1,$
 - 3: $\nabla_a = \nabla_\theta \mathcal{L}(f_{\theta_0}, \mathcal{D}_a), \nabla_T = \nabla_\theta \mathcal{L}(f_{\theta_0}, \mathcal{D}_T),$
 $\mathcal{R}_{a,0}^{PM} = -\mathcal{L}(f_{\theta_0}, \mathcal{D}_a), \mathcal{R}_{a,0}^{PT} = \cos(\nabla_a, \nabla_T),$
 $\hat{\mathcal{R}}_a = 0.5\mathcal{R}_{a,0}^{PM} + 0.5\mathcal{R}_{a,0}^{PT}$
 - 4: **for** $t = 1, 2, \dots, N$ **do**
 - 5: Calculate the upper confidence bound for each arm
 - 6: $a^* = \arg \max_{a \in A} \left(\hat{\mathcal{R}}_a + \sqrt{\frac{2 \ln t}{n_a}} \right)$
 - 7: Select the auxiliary dataset \mathcal{D}_{a^*}
 - 8: $n_{a^*} \leftarrow n_{a^*} + 1$
 - 9: $\nabla_T \leftarrow \nabla_\theta \mathcal{L}(f_{\theta_{t-1}}, \mathcal{D}_T)$
 - 10: $\nabla_{a^*} \leftarrow \nabla_\theta \mathcal{L}(f_{\theta_{t-1}}, \mathcal{D}_{a^*})$
 - 11: Update model parameters w.r.t. $\nabla_T + \nabla_{a^*}$
 - 12: Update the reward of the pulled arm
 - 13: $R_{a^*,t} = \alpha R_{a^*,t}^{PM} + (1 - \alpha) R_{a^*,t}^{PT}$
 - 14: $\hat{\mathcal{R}}_{a^*} \leftarrow (1 - \beta) \hat{\mathcal{R}}_{a^*} + \beta R_{a^*,t}$
 - 15: Release memory $\nabla_{a^*} \leftarrow 0$
 - 16: **end for**
 - 17: **return** f_θ
-

3. EXPERIMENTS AND RESULTS

To showcase the flexibility of our ASAP framework, we conduct extensive experiments on MRI and CT datasets covering various modalities and anatomical regions.

3.1. Datasets and Implementation Details

For MRI experiments, we construct the auxiliary datasets pool based on FeTS 2022 [16] (brain tumor segmentation) and iSeg2019 [17] (brain tissue segmentation). The auxiliary task pool consists of 30 datasets, each with a sample size exceeding 30. For the target datasets, we use two brain 3D MRI segmentation datasets: the periventricular leukomalacia (PVL) dataset [10], characterized by tissue reduction in periventricular and manually delineated on each slice of the patient’s T2 MRI images, and the White Matter Hyperintensity (WMH) dataset [18] which segments white matter hyperintensities on FLAIR MRI images. For CT experiments, we construct 30 datasets from TotalSegmentator (TOS) [19]

as the auxiliary datasets, based on label diversity and density. TOS is a whole-body-segmented 3D CT dataset that contains 117 main default tasks. For the target datasets, we experiment with vessel and liver segmentation tasks from a benchmark 3D CT dataset, the Medical Segmentation Decathlon (MSD) dataset [20]. For each auxiliary dataset, we use at most 30 training examples. For each target task, all the experiments are conducted under the 1-way 3-shot scenario using 5-fold cross-validation. We experimented with 5, 3, and 2 target samples, and found that using 3 samples yields satisfactory results in few-shot settings while also reducing the size requirements of the target dataset. We implement all methods on pre-trained UNet [14], SwinUNet [1], and MambaUNet [2], from MONAI [3]. For each kind of foundation model, we use consistent hyperparameters for all compared methods. The models are trained using a single A800 80GB GPU for a maximum of 400 epochs with a learning rate of 0.0003 and weight decay of 1e-7.

3.2. Performance evaluation

We compare our framework with state-of-the-art few-shot domain adaptation methods: 1) direct fine-tuning (FT) the source model on the target dataset, 2) GMS [11] identifies one best auxiliary dataset to aid the target based on gradient magnitude similarity, 3) a mixed-batch multi-task learning (MTL) framework [8] utilizes all auxiliary data simultaneously, 4) a dynamic auxiliary learning (DAL) method [9] adaptively samples the auxiliary data to jointly train with the target dataset based on gradient alignment. We evaluate the target segmentation performance using the Dice score and the mean IoU. A quantitative analysis of model adaptation performance on MRI and CT datasets is detailed in Table 1. The proposed ASAP framework outperforms all the baselines on all datasets, across modalities and anatomical regions. We also present the WMH and liver segmentation results on MambaUNet of different methods, clearly demonstrating the enhancements our method brings to the target few-shot medical image segmentation tasks, as shown in Fig 2.

3.2.1. Effectiveness of exploring and exploiting

The a-priori dataset selection method GMS is inferior to ours because it relies solely on exploiting the relations determined prior to training, but never exploring, e.g., as observed in the

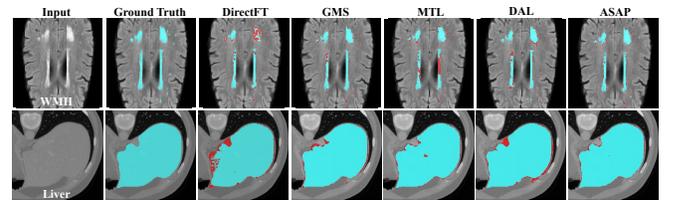


Fig. 2: Visualization of different domain adaptation methods performance of two specific target tasks: WMH segmentation on MRI images and liver segmentation on CT images, both using MambaUNet. The pixels highlighted in red represent incorrect predictions.

Table 1: Results of different domain adaptation strategies on MRI (left) and CT (right) datasets for three models. The segmentation evaluation metrics are the Dice score and the mean IoU score. Bold number: best score.

Target	Method	UNet		SwinUNet		MambaUNet	
		Dice (%)	mIoU (%)	Dice (%)	mIoU (%)	Dice (%)	mIoU (%)
PVL	Direct FT	23.48	13.82	18.35	10.48	27.20	16.18
	GMS (NIPS' 20)	23.50	13.56	20.41	11.86	27.45	16.72
	MTL (MIA' 23)	20.79	11.70	14.62	8.41	19.02	11.09
	DAL (NIPS' 24)	26.76	15.99	19.65	11.26	26.92	16.23
	ASAP (Ours)	27.95	16.80	22.16	13.02	30.06	18.48
WMH	Direct FT	57.03	40.78	64.24	48.79	62.35	45.51
	GMS (NIPS' 20)	52.86	39.40	64.54	50.91	65.78	50.36
	MTL (MIA' 23)	58.13	41.29	53.69	38.83	53.81	38.76
	DAL (NIPS' 24)	62.74	46.75	67.16	51.47	63.62	48.68
	ASAP (Ours)	66.94	53.10	67.62	52.04	71.06	56.04

Target	Method	UNet		SwinUNet		MambaUNet	
		Dice (%)	mIoU (%)	Dice (%)	mIoU (%)	Dice (%)	mIoU (%)
Vessel	Direct FT	48.70	33.13	41.31	26.64	49.08	33.39
	GMS (NIPS' 20)	48.74	33.85	36.16	22.51	46.62	40.94
	MTL (MIA' 23)	22.64	13.12	32.23	19.58	35.05	21.86
	DAL (NIPS' 24)	45.78	30.33	41.24	26.81	44.54	30.91
	ASAP (Ours)	49.22	33.55	48.05	32.81	49.27	33.34
Liver	Direct FT	79.79	66.97	84.86	74.63	84.40	73.67
	GMS (NIPS' 20)	90.88	83.50	86.99	77.53	90.04	82.20
	MTL (MIA' 23)	81.15	71.52	85.56	75.80	91.26	84.11
	DAL (NIPS' 24)	91.66	84.22	84.06	73.05	87.71	78.50
	ASAP (Ours)	92.10	85.61	87.27	78.67	92.91	86.95

vessel experiment on the right side of Table 1. In contrast, the multi-task learning (MTL) framework continuously explores all auxiliary data but never exploits knowledge of their relation to the target, leading to unsatisfactory results, e.g., as observed in the WMH experiment on the left side of Table 1, with significantly increasing training time—up to 34 times longer than direct fine-tuning. By balancing the trade-off between exploration and exploitation, our ASAP achieves a 24.39% gain on WMH compared to MTL, and a 13.18% gain on vessel segmentation compared to GMS in Dice scores.

3.2.2. Effectiveness of the efficient reward function

Compared to static dataset selection methods GMS and MTL, DAL offers a relatively better solution by dynamically selecting the auxiliary data based on gradient alignment. However, our reward function, with the consideration of \mathcal{R}^{PM} term, the reward of positive for model convergence, serves as a more effective guide to enable the model to deliver superior performance. Meanwhile, the proposed reward functions only rely on the losses and gradients, which are intrinsic to the model, making it naturally training-efficient: it took 7.93 hours to adapt the MambaUNet for the target task vessel segmentation, compared to 33.49 hours in MTL, 8.10 hours in DAL, and 7.82 hours in GMS, the a-priori dataset selection method, with the same batch size of 4 and input size of $32 \times 240 \times 240$. As per our policy, we update only the selected arm’s reward during training, which keeps the additional complexity stable, irrespective of the size of the auxiliary data pool.

3.2.3. Investigating the active and sequential training dynamics.

A closer look at the selected auxiliary tasks illustrates the active and sequential adaptation training mechanism, visualized in Fig 3. We show the selected auxiliary datasets at different turns for target tasks WMH segmentation on MRI images and liver segmentation on CT images. Interestingly, the policy does not initially sample the task experientially similar to the target. Instead, it sequentially selects the auxiliary dataset that progressively aligns with the target. Despite lacking access to the source domain, we can still effectively narrow the domain discrepancy by following this step-by-step knowledge acquisition, demonstrating the strength of the active and sequential domain adaptation framework.

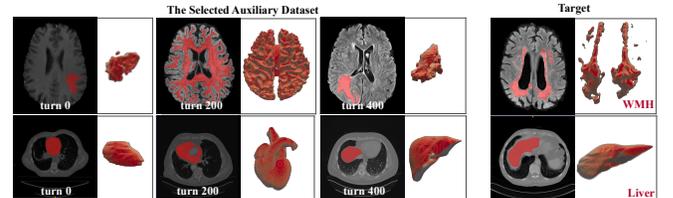


Fig. 3: A 3D visualization of the active and sequential training process. The figure shows the selected auxiliary datasets at turn 0, turn 200, and turn 400, for two specific target tasks. The images are presented alongside their ground truth.

4. CONCLUSION

We propose a novel active and sequential domain adaptation (ASAP) framework to adapt foundation models for the few-shot medical image segmentation tasks. With our desiderata in mind, the proposed ASAP achieve: 1) no requirement for access to the source domain or a substantial amount of target data, 2) incorporation of auxiliary data with dynamic scheduling of prioritized learning, adding minimal extra memory and computational overhead, 3) effective and efficient adaptation of foundational models, leading to strong performance on the target task. The results of experiments on three benchmark medical datasets show that the proposed method achieves state-of-the-art performance in foundation model adaptation problems within the realm of medical image processing. We believe our proposed approach will better leverage public medical resources, including foundation models and available auxiliary datasets, to tailor a model for the desired few-shot target task in a fast and scalable way.

5. ACKNOWLEDGMENTS

This work is supported in part by the Natural Science Foundation of China (Grant 62371270) and Shenzhen Key Laboratory of Ubiquitous Data Enabling (No.ZDSYS20220527171406015)

6. REFERENCES

- [1] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang, “Swin-UNET: UNet-like pure transformer for medical image segmentation,” in *European conference on computer vision*. Springer, 2022, pp. 205–218.

- [2] Ziyang Wang, Jian-Qing Zheng, Yichi Zhang, Ge Cui, and Lei Li, “Mamba-unet: Unet-like pure visual mamba for medical image segmentation,” *arXiv preprint arXiv:2402.05079*, 2024.
- [3] M Jorge Cardoso, Wenqi Li, Richard Brown, Nic Ma, Eric Kerfoot, Yiheng Wang, Benjamin Murrey, Andriy Myronenko, Can Zhao, Dong Yang, et al., “Monai: An open-source framework for deep learning in healthcare,” *arXiv preprint arXiv:2211.02701*, 2022.
- [4] Róger Bermúdez-Chacón, Pablo Márquez-Neila, Mathieu Salzmann, and Pascal Fua, “A domain-adaptive unet for electron microscopy image segmentation,” in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. IEEE, 2018, pp. 400–404.
- [5] Fuping Wu and Xiahai Zhuang, “Unsupervised domain adaptation with variational approximation for cardiac segmentation,” *IEEE Transactions on Medical Imaging*, vol. 40, no. 12, pp. 3555–3567, 2021.
- [6] Yanyang Gu, Zongyuan Ge, C Paul Bonnington, and Jun Zhou, “Progressive transfer learning and adversarial domain adaptation for cross-domain skin disease classification,” *IEEE journal of biomedical and health informatics*, vol. 24, no. 5, pp. 1379–1393, 2019.
- [7] Aviv Navon, Idan Achituve, Haggai Maron, Gal Chechik, and Ethan Fetaya, “Auxiliary learning by implicit differentiation,” *arXiv preprint arXiv:2007.02693*, 2021.
- [8] Simon Graham, Quoc Dang Vu, Mostafa Jahanifar, Shan E Ahmed Raza, Fayyaz Minhas, David Snead, and Nasir Rajpoot, “One model is all you need: multi-task learning enables simultaneous histology image segmentation and classification,” *Medical Image Analysis*, vol. 83, pp. 102685, 2023.
- [9] Alon Albalak, Colin A Raffel, and William Yang Wang, “Improving few-shot generalization by exploring and exploiting auxiliary data,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [10] Jingyun Yang, Jie Hu, Yicong Li, Heng Liu, and Yang Li, “Joint pvl detection and manual ability classification using semi-supervised multi-task learning,” in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VII 24*. Springer, 2021, pp. 453–463.
- [11] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn, “Gradient surgery for multi-task learning,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 5824–5836, 2020.
- [12] William G Macready and David H Wolpert, “Bandit problems and the exploration/exploitation tradeoff,” *IEEE Transactions on evolutionary computation*, vol. 2, no. 1, pp. 2–22, 1998.
- [13] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer, “Finite-time analysis of the multiarmed bandit problem,” *Machine learning*, vol. 47, pp. 235–256, 2002.
- [14] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*. Springer, 2015, pp. 234–241.
- [15] Lai Wei and Vaibhav Srivastava, “Nonstationary stochastic multiarmed bandits: Ucb policies and minimax regret,” *arXiv preprint arXiv:2101.08980*, 2021.
- [16] Sarthak Pati, Ujjwal Baid, Maximilian Zenk, Brandon Edwards, Micah Sheller, G Anthony Reina, Patrick Foley, Alexey Gruzdev, Jason Martin, Shadi Albarqouni, et al., “The federated tumor segmentation (fets) challenge,” *arXiv preprint arXiv:2105.05874*, 2021.
- [17] Yue Sun, Kun Gao, Zhengwang Wu, Guannan Li, Xiaopeng Zong, Zhihao Lei, Ying Wei, Jun Ma, Xiaoping Yang, Xue Feng, et al., “Multi-site infant brain segmentation algorithms: The iseg-2019 challenge,” *IEEE Transactions on Medical Imaging*, vol. 40, no. 5, pp. 1363–1376, 2021.
- [18] Hugo J Kuijf, J Matthijs Biesbroek, Jeroen De Bresser, Rutger Heinen, Simon Andermatt, Mariana Bento, Matt Berseth, Mikhail Belyaev, M Jorge Cardoso, Adria Casamitjana, et al., “Standardized assessment of automatic segmentation of white matter hyperintensities and results of the wmh segmentation challenge,” *IEEE transactions on medical imaging*, vol. 38, no. 11, pp. 2556–2568, 2019.
- [19] Jakob Wasserthal, Hanns-Christian Breit, Manfred T Meyer, Maurice Pradella, Daniel Hinck, Alexander W Sauter, Tobias Heye, Daniel T Boll, Joshy Cyriac, Shan Yang, et al., “Totalsegmentator: robust segmentation of 104 anatomic structures in ct images,” *Radiology: Artificial Intelligence*, vol. 5, no. 5, 2023.
- [20] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, et al., “The medical segmentation decathlon,” *Nature communications*, vol. 13, no. 1, pp. 4128, 2022.