



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

# Self-Auditing Parameter-Efficient Fine-Tuning for Few-Shot 3D Medical Image Segmentation

Son Thai Ly<sup>1</sup> and Hien V. Nguyen<sup>1</sup>

Department of Electronics and Computer Engineering, University of Houston, USA  
stly@cougarnet.uh.edu      hvnguy35@central.uh.edu

**Abstract.** Adapting foundation models to new clinical sites remains challenging in practice. Domain shift and scarce annotations must be handled by experts, yet many clinical groups do not have ready access to skilled AI engineers to tune adapter designs and training recipes. As a result, adaptation cycles can stretch from weeks to months, particularly in few-shot settings. Existing PEFT methods either require manual adapter configuration or automated searches that are computationally infeasible in few-shot 3D settings. We propose **SEA-PEFT** (*self-auditing parameter-efficient fine-tuning*) to automate this process, moving towards realistic deployment in clinical environments. SEA-PEFT treats adapter configuration as an online allocation problem solved during fine-tuning rather than through manual, fixed-topology choices. SEA-PEFT uses a *search-audit-allocate* loop that trains active adapters, estimates each adapter’s Dice utility by momentarily toggling it off, and then reselects the active set under a parameter budget using a greedy knapsack allocator. Exponential Moving Average and Interquartile Range smoothing, together with a Finite-State Ranking controller, stabilize the loop and improve reliability in high-noise few-shot regimes. On TotalSegmentator and FLARE’22, SEA-PEFT improves mean Dice by 2.4–2.8 points over the strongest fixed-topology PEFT baselines across 1/5/10-shot settings while training <1% of parameters. For reproducibility purposes, we made our code publicly available at [https://github.com/tsly123/SEA\\_PEFT](https://github.com/tsly123/SEA_PEFT)

**Keywords:** Parameter-efficient fine-tuning · Medical image 3D segmentation · Few-shot learning · Foundation models · Adapter selection

## 1 Introduction

Recent foundation models (FMs) for 3D medical imaging, including Swin-UNETR [7], SuPreM [12], and FSEFT [18], have produced strong results in organ segmentation across CT and MRI. Adapting these models to a new clinical site is nevertheless non-trivial. Hospitals vary in scanner protocols, reconstruction kernels, and patient demographics, and this distribution shift requires site-specific fine-tuning. At the same time, volumetric mask annotation is labour-intensive, most sites lack dedicated AI staff, and 3D patch training imposes significant memory requirements. As a result, fine-tuning a medical FM in a new hospital typically requires an dedicated AI engineer, labelled data, and careful configuration choices that are unavailable in most clinical settings.


Parameter-Efficient Fine-Tuning (PEFT) [9,4,3,8,10] reduces the number of updated parameters and is therefore a natural candidate for low-resource clinical adaptation. The problem is that PEFT still requires three configuration decisions: which adapter type to use, where in the transformer to insert it, and what rank or bottleneck dimension to assign. These choices interact and depend on the target organ, dataset size, and acquisition protocol, so there is no universally good default. Offline search methods such as NOAH [25], AutoPEFT [26], and Fairtune [5] can find good configurations but require multiple end-to-end fine-tuning runs, which is impractical for few-shots with high-resolution 3D volumes.

We propose to solve the PEFT configuration problem online, during fine-tuning, using direct task-level feedback rather than weight-space proxies. The core idea is simple: to assess whether an adapter is useful, temporarily remove it and measure the change in performance on a held-out validation set. This on/off perturbation gives an unbiased, model-agnostic estimate of each adapter’s marginal performance contribution. Repeating this measurement across training, smoothing the estimates, and greedily selecting adapters under a parameter budget yields a configuration that improves as more audits accumulate, without requiring any offline sweep.

We proposed framework **SEA-PEFT** (SElf-Auditing PEFT). The search-audit-allocate loop trains active adapters for  $K$  steps (**Search**), estimates per-adapter utility via on/off toggling (**Audit**), and selects the active set under a global budget (**Allocate**). An Exponential Moving Average (EMA) + Interquartile Range (IQR) utility tracker stabilizes noisy audit estimates, and a Finite-State Machine (FSM) prevents premature adapter switching by requiring several consistent votes before committing a configuration change. After the loop converges, a guard-free knapsack re-solve selects the final configuration, which is then fine-tuned from scratch so that the reported model carries no exploratory weights from the search phase. **Contributions:** (i) We propose SEA-PEFT, an online PEFT configuration framework for few-shot 3D medical segmentation that selects adapter type, placement, and rank via task-level perturbation auditing during fine-tuning, with no offline search and no manual configuration; (ii) The novel couples with EMA+IQR utility tracker and FSM stabilizer as practical components for reliable adapter selection under few-shot noise; (iii) We provide both theoretical and empirical evidence on TotalSegmentator and FLARE datasets to demonstrate the effectiveness of SEA-PEFT.

## 2 Related Works

**PEFT for medical imaging.** Adapter-based methods [8,4], LoRA and variant [9,11,6], BitFit [3], and prompt tuning [13,14,10] all adapt pretrained models by updating a small parameter subset while keeping the backbone frozen. For 3D medical FMs, FSEFT [18,1], and FreqFit [15] apply fixed single-adapter designs to few-shot organ segmentation. In all of these works, adapter type, insertion site, and rank are chosen before training, either by hand or via grid search on a held-out set. SEA-PEFT eliminates this step by selecting and adjusting the configuration during training.



**Fig. 1.** SEA-PEFT iterates over a PEFT library with three steps: (1) Search selects a small adapter set under a parameter budget; (2) Audit runs lightweight on/off tests to estimate Dice-per-cost utility; (3) Allocate updates the active configuration via budget-aware rank adjustments. This towards rapid adaptation by clinical users without AI engineers.

**Automated PEFT configuration.** NOAH [25] unifies adapters, prompts, and LoRA in a supernet and runs a one-shot NAS to find a good prompt configuration. AutoPEFT [26] and Fairtune [5] use multi-objective Bayesian optimization to search over adapter type, placement, and rank combinations. Both approaches require instantiating and fully fine-tuning each candidate configuration, which is not feasible for high-resolution 3D segmentation at 1–10 shots. AdaLoRA [24] reallocates LoRA rank during training based on singular value magnitudes, avoiding full re-training, but it optimizes an internal weight-space criterion rather than directly measuring Dice utility per adapter. SEA-PEFT instead evaluates each adapter’s contribution to segmentation accuracy through on/off perturbation, which is both model-agnostic and directly aligned with the evaluation metric.

### 3 Self-Audit PEFT

#### 3.1 Problem Setup and Audit Space

Let  $f_\theta$  denote a frozen FM backbone. The *audit space*  $\mathcal{A} = \{a_i\}_{i=1}^N$  is a library of lightweight adapter candidates that can be activated, deactivated, or resized during training. Each unit  $a_i$  has trainable weights  $w_i$ , binary gate  $\alpha_i \in \{0, 1\}$ , and parameter cost  $c_i$ .

*Audit space.* Following standard practice for neural architecture search (NAS) studies [17, 22, 23], we adopt three adapter topologies: Serial (**SA**), Parallel (**PA**), and their composite **SAPA**. For layer output  $F(x)$  with projections  $W_{down} \in \mathbb{R}^{D \times d}$  and an up-projection  $W_{up} \in \mathbb{R}^{d \times D}$  ( $d \ll D$ ):

$$SA: y = \text{ReLU}(F(x) \cdot W_{down}) \cdot W_{up} \quad (1)$$

$$PA: y = \text{ReLU}(x \cdot W_{down}) \cdot W_{up} \quad (2)$$

$$SAPA: y = \text{ReLU}(F(x) \cdot W_{down}) \cdot W_{up} + \text{ReLU}(x \cdot W_{down}) \cdot W_{up} \quad (3)$$

**Objective.** Given sets  $\mathcal{D}_{train}$ ,  $\mathcal{D}_{val}$ , and budget  $P_{max}$ :

$$\max_{\alpha, w} \mathbb{E}_{(x, y) \sim \mathcal{D}_{val}} [\text{Dice}(f_{\theta, \alpha, w}(x), y)] \quad \text{s.t.} \quad \sum_i \alpha_i c_i \leq P_{max} \quad (4)$$

This mirrors a knapsack problem where each adapter  $a_i$  has cost  $c_i$  and value equal to its Dice improvement. The utilities are unknown and must be estimated from validation feedback, SEA-PEFT solves an online variant of Eq. (4).

### 3.2 Search–Audit–Allocate Cycle

**Search.** SEA-PEFT trains currently active adapters ( $\alpha_i = 1$ ) on  $\mathcal{D}_{train}$  while keeping the backbone frozen. After  $K$  steps, the framework transits to audit.

**Audit.** A mini-batch  $\mathcal{M}^{(t)} \subset \mathcal{A}$  of  $M \ll N$  adapters is sampled using a *active/inactive*: empirically 30% from currently active adapters (exploitation) and 70% from inactive ones (exploration), plus an  $\varepsilon$ -exploration term up-weighting rarely probed adapters. For each  $i \in \mathcal{M}^{(t)}$ :

$$u_i = [\text{Dice}(f_{\theta, \alpha}, \mathcal{D}_{val}) - \text{Dice}(f_{\theta, \alpha \setminus i}, \mathcal{D}_{val})] / c_i \quad (5)$$

$$r_i^{(t)} = \text{median}_w(\tilde{u}_i^t) - \lambda_S \cdot \text{IQR}_w(\tilde{u}_i^t) \quad \text{with} \quad \tilde{u}_i^t = (1 - \beta)u_i^t + \beta\tilde{u}_i^{(t-1)} \quad (6)$$

*Utility smoothing.* Raw utilities are noisy because only  $M$  adapters are probed per cycle. SEA-PEFT applies a two-stage filter as in Eq. (6).

Each adapter maintains a history deque of 3–5 audits. High  $r_i^{(t)}$  signals stable Dice improvement; large IQR suppresses unstable adapters.

**Proposition 1 (Convergence of the EMA+IQR utility estimator).** *Let  $u_i^{(t)} = \mu_i^{(t)} + \xi_i^{(t)}$  where  $\mathbb{E}[\xi_i^{(t)} | \mathcal{F}_{t-1}] = 0$ ,  $\text{Var}(\xi_i^{(t)}) \leq \sigma^2$ , and  $\mu_i^{(t)}$  drifts by at most  $\delta$  between consecutive audits. After  $t$  audits with Exponential Moving Average (EMA) decay  $\beta \in (0, 1)$ :*

$$\left| \mathbb{E}[\tilde{u}_i^{(t)}] - \mu_i^{(t)} \right| \leq \frac{\delta\beta}{1-\beta} + \beta^t \left| \mathbb{E}[\tilde{u}_i^{(0)}] - \mu_i^{(0)} \right|, \quad \text{Var}(\tilde{u}_i^{(t)}) \leq \frac{(1-\beta)\sigma^2}{1+\beta} \quad (7)$$

Furthermore, over a sliding history window of audits, the sample Interquartile Range (IQR)  $\text{IQR}_H(\tilde{u}_i^{(t)})$  concentrates around the population IQR of the noise distribution with fluctuations of order  $O(\sigma_{\text{eff}}/\sqrt{H})$ , where  $\sigma_{\text{eff}} = \sigma\sqrt{(1-\beta)/(1+\beta)}$  is the effective noise scale of the smoothed estimates. The robust score  $r_i^{(t)} = \text{median}_H(\tilde{u}_i^{(t)}) - \lambda_S \cdot \text{IQR}_H(\tilde{u}_i^{(t)})$ , with IQR shrinkage weight  $\lambda_S > 0$ , suppresses adapters whose utility is highly variable across audits.

*Proof (Proof sketch).* The bias bound follows by expanding the EMA recursion, applying  $|\mu_i^{(t-k)} - \mu_i^{(t)}| \leq k\delta$ , and bounding  $(1-\beta) \sum_{k=1}^{\infty} k\beta^k = \delta\beta/(1-\beta)$ . The variance bound uses independence of  $\{\xi_i^{(t)}\}$  and  $\sum_{k \geq 0} \beta^{2k} = (1-\beta^2)^{-1}$ . IQR concentration follows standard order-statistics theory under bounded variance.

**Lemma 1 (Audit coverage).** *Under any sampling strategy with exploration rate  $\varepsilon > 0$ , audit set  $\mathcal{M}^{(t)}$  (size  $M$ ) at cycle  $t$ , and  $|\mathcal{A}| = N$ , every adapter  $a_i \in \mathcal{A}$*

$$\mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[a_i \in \mathcal{M}^{(t)}] \right] \geq \rho_c T, \quad \rho_c = \frac{\varepsilon M}{N} > 0 \quad (8)$$

satisfying the positive coverage condition required for sublinear regret in [19]. That is, every adapter is audited  $\Omega(T)$  times regardless of its initial activity state.

*Proof (Proof sketch).* At each audit cycle the  $\varepsilon$ -exploration term assigns a minimum selection probability of  $\varepsilon M/N$  to every adapter independently of its activity gate  $\alpha_i$ . Summing over  $T$  cycles and taking expectations gives the bound directly. Without  $\varepsilon$ -exploration,  $\rho_c = 0$  and inactive adapters at initialization are never audited, making the final knapsack re-solve arbitrarily suboptimal.

**Allocate.** Adapters are greedily activated in descending density order  $d_i = r_i/c_i$  until  $\sum_i \alpha_i c_i \leq P_{\max}$ . Replacements are accepted iff improvement exceeds  $\mu_{\text{eff}}$ .

**Finite-State Machine Stabilizer.** Naïve greedy allocation allows adapters to flip on/off at nearly every cycle, causing *configuration chatter* that slows convergence and destabilizes EMA utility estimates. The FSM requires  $\tau_{\text{act}}$  consecutive consistent “votes” before committing an activity change (and  $\tau_{\text{rank}}$  for rank):

$$\alpha_i^{(t+1)} \leftarrow \hat{\alpha}_i^{(t+1)} \quad \text{only if } c_i^{\text{act}} \geq \tau_{\text{act}} \quad (9)$$

with  $\tau_{\text{act}}, \tau_{\text{rank}} \approx 2-3$ . FSM reduces reconfigurations to  $T_c \approx 20-30$  per run, suppressing noise-driven instability while preserving online adaptability

**Proposition 2 (FSM reduces structural error).** *Let  $T_c$  denote the number of audit cycles in which the active set changes, and let  $\hat{\alpha}_i^{(t+1)} \in \{0, 1\}$  be the activity proposed by the greedy allocator at cycle  $t$ . The FSM maintains a per-adapter vote counter  $c_i^{\text{act}}$ , incremented when  $\hat{\alpha}_i^{(t+1)} \neq \alpha_i^{(t)}$  and reset otherwise, and commits a change only when  $c_i^{\text{act}} \geq \tau_{\text{act}}$  consecutive consistent proposals are observed. This gives:*

$$T_c \leq \left\lfloor \frac{T}{\tau_{\text{act}}} \right\rfloor \quad (10)$$

*In the regret decomposition adapting [19], combined with the structural error term  $O(\varepsilon P_{\max} T_c)$  in the regret decomposition, the FSM reduces structural error by a factor of  $\tau_{\text{act}}$ , tightening the knapsack approximation guarantee of [20].*

*Proof (Proof sketch).* A configuration change at cycle  $t$  requires  $\tau_{\text{act}}$  consecutive cycles with a consistent new proposal from the allocator. Hence at least  $\tau_{\text{act}}$  cycles must elapse between any two changes, giving  $T_c \leq \lfloor T/\tau_{\text{act}} \rfloor$  by direct counting. Hence, reducing the structural error term  $O(\varepsilon P_{\max} T_c)$  by a factor of  $\tau_{\text{act}}$  and tightening the final approximation guarantee.

**Final Re-solve and Re-fine-tuning.** After  $T$  steps, SEA-PEFT performs one *guard-free* knapsack re-solve on stable scores  $\{r_i^{(T)}\}$  using density-greedy plus a single-swap local improvement pass, yielding  $\mathcal{S}_{\text{final}}^*$ . Adapter weights and optimizer states are then re-initialized, and a standard fine-tuning run with only the selected adapters active is performed. This two-phase design (online search + focused re-finetune) ensures the final model is not burdened by exploratory weights from the audit phase.

## 4 Experiments

**Audit space.** Following the empirical experience in prior studies [9, 4, 2, 15, 18, 1], we fine-tune the following: For self-attention layer: we attached LoRA [9] with rank

$\in \{2, 4, 8, 16\}$  with SA, PA, and SAPA. Feed-forward layer: we attached LoRA [9] with rank  $\in \{2, 4, 8, 16\}$  and AdaptFormer [4] with bottleneck  $\in \{4, 8, 16, 32\}$ , both with SA, PA, and SAPA. Normalization layer: we attached Affine-LN [2].

**Datasets and backbones.** We evaluate on nine abdominal organs (spleen, left kidney, gallbladder, esophagus, liver, pancreas, stomach, duodenum, aorta) under 1-, 5-, and 10-shot regimes. Binary segmentation is performed on TotalSegmentator [21] and multi-class on FLARE [16], using pre-trained FSEFT [18] and SuPreM [12] on frozen Swin-UNETR [7] backbones.

Note that, NOAH [25] and AutoPEFT [26] require full tuning per candidate configuration. Over the audit space at 1–10 shots, this would require hundreds of 3D training runs, SEA-PEFT estimates adapter utility via on/off perturbation within a single run, making direct comparison infeasible and arguably unfair.

**Implementation.** After the search–audit–allocate loop converges, following [18, 1, 15], the selected configuration is fine-tuned image size of  $96 \times 96 \times 96$  patches with batch size 1, AdamW,  $\text{lr} = 1 \times e^{-3}$ , 200 epochs with early stopping. Augmentation includes intensity shifts and 90 degree rotations. All experiments run on a single V100 32 GB GPU. Results are averaged across 3 random seeds on hold-out test set. The metric presented is Dice.

#### 4.1 Results and Discussion

**Main results.** Tables 1 reports few-shot adaptation results on TotalSegmentator [21] and FLARE’22 [16] using the final configurations selected by our *Search–Audit–Allocate* loop. Across all shot regimes and both datasets, SEA-PEFT consistently achieves the highest mean Dice, outperforming fixed-topology PEFT baselines without any manual tuning of adapter type, placement, or rank. On TotalSegmentator, it attains 78.33/79.90/80.29 Dice in the 1-, 5-, and 10-shot settings, a clear improvement over the strongest baseline. On FLARE’22, SEA-PEFT achieves 75.77 (1-shot) and 76.33 (5-shot), continuing the same upward trend as supervision increases. These gains demonstrate that online self-auditing reliably generalizes across distinct clinical datasets.

Results on pre-trained SuPreM [12] in Table 2 reinforce the same pattern observed on TotalSegmentator and FLARE’22. SEA-PEFT achieves the highest average Dice in both the 5-shot (77.89) and 10-shot (79.93) settings, with consistent gains on gallbladder, esophagus, and pancreas over Spatial Adapter [18]. Notably, the margin over the best fixed baseline widens from 0.88 points at 5-shot to 3.58 points at 10-shot, mirroring the scaling trend across all benchmarks: as more training examples are available, audit signals become less noisy, the allocator converges to better configurations. The advantage of dynamic over static capacity allocation grows.

**Runtime and Practicality.** Tab. 3 reports the end-to-end cost of SEA-PEFT, covering the full search–audit–allocate process and a final refinement run. Across 1/5/10-shot settings, adaptation completes within 2.5–6.5 hours while consistently using only 0.2% trainable parameters. This efficiency stems from SEA-PEFT’s online design: each audit cycle evaluates only a small subset of adapters, so compute scales smoothly with data rather than exploding with

**Table 1.** Few-shot efficient adaptation on TotalSegmentator [21] and FLARE [16] with pre-trained FSEFT [18]. **Bold** and underscore indicates the best and second best result.

| Setting             | Method  | Spl              | lKid         | Gall         | Eso          | Liv          | Pan          | Sto          | Duo          | Aor          | Avg.         |              |
|---------------------|---------|------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| TotalSeg (binary)   | 1-shot  | Full Fine-tuning | 91.31        | 88.00        | 76.58        | 49.11        | <b>98.80</b> | 77.31        | 68.54        | <b>58.24</b> | <b>81.17</b> | 75.90        |
|                     |         | BitFit           | 88.42        | 87.89        | 75.49        | 48.23        | 92.00        | 74.92        | 70.27        | 49.95        | 72.03        | 73.24        |
|                     |         | LoRA             | 90.38        | 88.68        | 77.63        | 44.54        | 92.48        | 78.13        | 72.19        | 54.75        | 70.28        | 74.34        |
|                     |         | AdaptFormer      | 89.85        | 88.75        | 81.22        | 47.64        | 89.95        | 68.97        | 70.63        | 57.45        | 76.47        | 74.55        |
|                     |         | Affine-LN        | 91.62        | 88.87        | 79.34        | 45.57        | 90.17        | 75.05        | 67.43        | 49.76        | 71.78        | 73.29        |
|                     |         | SEA-PEFT         | <b>93.32</b> | <b>89.94</b> | <b>85.04</b> | <b>50.13</b> | 93.42        | <b>79.68</b> | <b>75.75</b> | 56.95        | 80.71        | <b>78.33</b> |
|                     | 5-shot  | Full Fine-tuning | 90.47        | 89.02        | 71.88        | <b>54.33</b> | <b>93.75</b> | 78.20        | 59.26        | 66.27        | <b>91.20</b> | 77.15        |
|                     |         | BitFit           | 91.46        | 88.26        | 71.37        | 50.50        | 92.96        | 57.94        | 71.22        | 56.91        | 73.65        | 72.70        |
|                     |         | LoRA             | 92.18        | 89.08        | 73.69        | 48.98        | 93.18        | 78.27        | 72.73        | 60.18        | 78.43        | 76.30        |
|                     |         | AdaptFormer      | 90.32        | 89.49        | 74.93        | 45.54        | 93.16        | 71.36        | 74.01        | 64.62        | 83.22        | 76.29        |
|                     |         | Affine-LN        | 91.75        | 90.19        | 76.86        | 47.51        | 93.30        | 77.25        | 74.15        | 63.56        | 74.66        | 76.58        |
|                     |         | SEA-PEFT         | <b>93.72</b> | <b>90.68</b> | <b>81.21</b> | 52.25        | 93.42        | <b>80.23</b> | <b>76.56</b> | <b>66.30</b> | 84.76        | <b>79.90</b> |
|                     | 10-shot | Full Fine-tuning | 92.97        | 80.85        | 72.96        | <b>56.56</b> | 92.36        | 78.67        | 61.27        | 66.56        | <b>90.15</b> | 76.93        |
|                     |         | BitFit           | 92.51        | 89.41        | 76.48        | 50.57        | 92.54        | 79.40        | 67.09        | 65.32        | 76.89        | 76.69        |
|                     |         | LoRA             | 91.91        | 90.10        | 82.66        | 49.36        | 93.45        | 80.86        | 67.97        | 61.36        | 82.47        | 77.79        |
|                     |         | AdaptFormer      | 90.36        | 89.91        | 73.67        | 49.19        | 89.50        | 71.13        | 66.85        | 64.95        | 80.48        | 75.12        |
|                     |         | Affine-LN        | 92.20        | 86.02        | 79.58        | 50.27        | 89.98        | 77.64        | 69.15        | 67.64        | 83.53        | 77.33        |
|                     |         | SEA-PEFT         | <b>93.35</b> | <b>90.57</b> | <b>83.96</b> | 51.70        | <b>94.23</b> | <b>81.74</b> | <b>72.55</b> | <b>67.83</b> | 86.72        | <b>80.29</b> |
| FLARE (multi-organ) | 1-shot  | LoRA             | 87.81        | 76.22        | 55.09        | <b>74.33</b> | <b>95.05</b> | 82.60        | 74.72        | 36.03        | 90.31        | 74.68        |
|                     |         | Adaptformer      | 83.69        | <b>76.77</b> | 54.49        | 72.13        | 94.24        | 80.01        | 73.85        | 43.86        | 90.79        | 74.43        |
|                     |         | Affine-LN        | 85.46        | 76.55        | <b>55.91</b> | 72.23        | 94.85        | 82.20        | 73.14        | <b>44.66</b> | <b>91.57</b> | 75.17        |
|                     |         | SEA-PEFT         | <b>88.33</b> | 74.69        | 55.57        | 74.15        | 94.61        | <b>83.34</b> | <b>79.87</b> | 40.70        | 90.70        | <b>75.77</b> |
|                     | 5-shot  | LoRA             | 86.56        | 76.15        | 54.24        | 74.46        | 94.93        | 83.01        | 75.17        | 43.66        | <b>92.01</b> | 75.58        |
|                     |         | Adaptformer      | 81.50        | 75.36        | <b>55.86</b> | 74.53        | 94.66        | 82.87        | <b>76.15</b> | 46.89        | 91.43        | 75.47        |
|                     |         | Affine-LN        | 85.73        | 76.00        | 53.71        | 75.20        | 94.62        | 82.04        | 74.69        | <b>47.24</b> | 91.38        | 75.62        |
|                     |         | SEA-PEFT         | <b>87.15</b> | <b>77.97</b> | 54.96        | <b>76.02</b> | <b>95.05</b> | <b>83.04</b> | 75.70        | 45.20        | 91.93        | <b>76.33</b> |
|                     | 10-shot | LoRA             | 84.00        | <b>77.05</b> | 54.72        | 74.02        | 94.64        | 82.66        | 76.31        | 46.11        | 91.59        | 75.68        |
|                     |         | Adaptformer      | 87.08        | 76.96        | 55.69        | 74.78        | 93.69        | 82.45        | 77.07        | 46.29        | 91.38        | 76.15        |
|                     |         | Affine-LN        | <b>87.68</b> | 76.99        | <b>55.91</b> | <b>75.61</b> | 95.04        | <b>83.57</b> | <b>79.27</b> | <b>51.88</b> | 91.69        | <b>77.52</b> |
|                     |         | SEA-PEFT         | 87.22        | 76.68        | 55.71        | 75.55        | <b>95.17</b> | 83.27        | 79.23        | 48.82        | <b>91.69</b> | 77.03        |

configuration space size. In practice, this yields high-quality configurations in hours, moving SEA-PEFT towards realistic deployment in clinical environments where engineers are most likely unavailable and model adaptation must fit within tight timelines.

**SEA-PEFT vs. Random Configurations.** To examine whether SEA-PEFT’s gains arise from principled selection rather than larger parameter counts, we run a controlled diagnostic on three representative organs, as shown in Fig. 2. SEA-PEFT consistently selects configurations near the upper-left frontier—achieving higher Dice with comparable or fewer parameters. The gap widens as the number of shots increases: with more stable gradients, the audit signal becomes cleaner and SEA-PEFT towards stronger, non-redundant adapter subsets, while random configurations remain broadly scattered. The experiment confirms that SEA-PEFT’s improvements come from targeted, task-level utility estimation rather than sheer parameter volume.



Fig. 3 shows qualitative results across all eight organs. SEA-PEFT produces masks that more closely match the ground truth on small, low-contrast structures: duodenum and gallbladder are largely missed by fixed-topology PEFT but partially recovered by SEA-PEFT, while esophagus and left kidney (yellow arrows) show cleaner boundary delineation. On large, high-contrast organs such as liver

**Table 2.** Few-shot efficient adaptation on TotalSegmentator with pre-trained SuPreM [12].

| Setting | Method               | Gall         | Eso          | Pan          | Duo          | Aor          | Avg.         |
|---------|----------------------|--------------|--------------|--------------|--------------|--------------|--------------|
| 5-shot  | LoRA [9]             | 52.50        | 46.43        | 66.86        | 54.15        | 73.33        | 58.65        |
|         | Adaptformer [4]      | 60.17        | 51.79        | 76.73        | 74.49        | <b>93.12</b> | 71.26        |
|         | Affine-LN [2]        | 76.10        | 50.04        | 75.46        | 71.91        | 90.91        | 72.88        |
|         | Spatial Adapter [18] | 85.08        | 55.56        | 78.84        | <b>78.17</b> | 87.40        | 77.01        |
|         | <b>SEA-PEFT</b>      | <b>86.67</b> | <b>57.32</b> | <b>78.87</b> | 75.86        | 90.74        | <b>77.89</b> |
| 10-shot | LoRA [9]             | 59.25        | 55.33        | 77.72        | 73.89        | 80.59        | 69.35        |
|         | Adaptformer [4]      | 77.78        | 54.10        | 76.05        | 77.58        | 93.25        | 75.75        |
|         | Affine-LN [2]        | 80.84        | 55.80        | 76.98        | 75.66        | 92.50        | 76.35        |
|         | Spatial Adapter [18] | 81.61        | 56.24        | 77.69        | <b>79.54</b> | 84.66        | 75.95        |
|         | <b>SEA-PEFT</b>      | <b>88.51</b> | <b>57.88</b> | <b>80.98</b> | 78.95        | <b>93.35</b> | <b>79.93</b> |

**Table 3.** Average times of end-to-end search-audit-allocate and re-fine-tune with optimal configurations for all experiments.

| Setting | Total steps | Avg. Time | Avg. #Params (%) |
|---------|-------------|-----------|------------------|
| 1-shot  | 6000        | 2.5 hours | 0.2              |
| 5-shot  | 8000        | 4.5 hours | 0.2              |
| 10-shot | 12000       | 6.5 hours | 0.2              |

**Fig. 2.** SA-PEFT vs. random PEFT configurations under fixed parameter budgets. Across 1/5/10-shot settings, SEA-PEFT’s gains (Orange) come from principled selection rather than simply using more parameters.**Fig. 3.** Qualitative segmentation comparison between fixed-topology PEFT and SEA-PEFT across eight organs.

and spleen, both methods perform comparably, consistent with the quantitative results in Tab. 1.

## 5 Conclusion

We presented SEA-PEFT, a framework that automates PEFT training and configuration search for few-shot 3D medical segmentation by treating adapter selection as an online allocation problem. Each adapter’s utility is estimated by on/off Dice perturbation on a validation set, smoothed via EMA and IQR filtering, and used to greedily allocate parameter budget across training. An FSM stabilizer prevents configuration chatter in noisy few-shot settings. On TotalSegmentator and FLARE’22, SEA-PEFT consistently outperforms fixed-topology PEFT baselines across 1/5/10-shot regimes using <1% trainable parameters.

**Acknowledgments.** This work was supported in part by the National Institutes of Health (NIH) under Grant 5R01DK134055

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Baklouti, G., Silva-Rodríguez, J., Dolz, J., Bahig, H., Ben Ayed, I.: Regularized low-rank adaptation for few-shot organ segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 517–527. Springer (2025) 2, 5, 6
2. Basu, S., Hu, S., Massiceti, D., Feizi, S.: Strong baselines for parameter-efficient few-shot fine-tuning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 11024–11031 (2024) 5, 6, 8
3. Cai, H., Gan, C., Zhu, L., Han, S.: Tinytl: Reduce memory, not parameters for efficient on-device learning. NIPS **33**, 11285–11297 (2020) 2
4. Chen, S., Ge, C., Tong, Z., Wang, J., Song, Y., Wang, J., Luo, P.: Adaptformer: Adapting vision transformers for scalable visual recognition. arXiv:2205.13535 (2022) 2, 5, 6, 8
5. Dutt, R., Bohdal, O., Tsaftaris, S.A., Hospedales, T.: Fairtune: Optimizing parameter efficient fine tuning for fairness in medical image analysis. arXiv preprint arXiv:2310.05055 (2023) 2, 3
6. Gao, Z., Wang, Q., Chen, A., Liu, Z., Wu, B., Chen, L., Li, J.: Parameter-efficient fine-tuning with discrete fourier transform. arXiv preprint arXiv:2405.03003 (2024) 2
7. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In: International MICCAI brainlesion workshop. pp. 272–284. Springer (2021) 1, 6
8. Houshy, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S.: Parameter-efficient transfer learning for nlp. In: ICML. pp. 2790–2799. PMLR (2019) 2
9. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. arXiv:2106.09685 (2021) 2, 5, 6, 8
10. Jia, M., Tang, L., Chen, B.C., Cardie, C., Belongie, S., Hariharan, B., Lim, S.N.: Visual prompt tuning. In: ECCV. pp. 709–727. Springer (2022) 2
11. Kopiczko, D.J., Blankevoort, T., Asano, Y.M.: Vera: Vector-based random matrix adaptation. arXiv preprint arXiv:2310.11454 (2023) 2
12. Li, W., Yuille, A., Zhou, Z.: How well do supervised models transfer to 3d image segmentation. In: The Twelfth International Conference on Learning Representations. vol. 1 (2024) 1, 6, 8
13. Li, X.L., Liang, P.: Prefix-tuning: Optimizing continuous prompts for generation. arXiv preprint arXiv:2101.00190 (2021) 2
14. Liu, X., Ji, K., Fu, Y., Tam, W., Du, Z., Yang, Z., Tang, J.: P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). pp. 61–68 (2022) 2

15. Ly, S.T., Nguyen, H.V.: Frequency strikes back: Boosting parameter-efficient foundation model adaptation for medical imaging. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 258–267. Springer (2025) [2](#), [5](#), [6](#)
16. Ma, J., Zhang, Y., Gu, S., Ge, C., Mae, S., Young, A., Zhu, C., Yang, X., Meng, K., Huang, Z., et al.: Unleashing the strengths of unlabelled data in deep learning-assisted pan-cancer abdominal organ quantification: the flare22 challenge. *The Lancet Digital Health* **6**(11), e815–e826 (2024) [6](#), [7](#)
17. Ru, R., Esperanca, P., Carlucci, F.M.: Neural architecture generator optimization. *Advances in Neural Information Processing Systems* **33**, 12057–12069 (2020) [3](#)
18. Silva-Rodríguez, J., Dolz, J., Ayed, I.B.: Towards foundation models and few-shot parameter-efficient fine-tuning for volumetric organ segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 213–224. Springer (2023) [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
19. Streeter, M., Golovin, D.: An online algorithm for maximizing submodular functions. *Advances in Neural Information Processing Systems* **21** (2008) [4](#), [5](#)
20. Sviridenko, M.: A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters* **32**(1), 41–43 (2004) [5](#)
21. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., et al.: Totalsegmentator: robust segmentation of 104 anatomic structures in ct images. *Radiology: Artificial Intelligence* **5**(5), e230024 (2023) [6](#), [7](#)
22. Xie, S., Kirillov, A., Girshick, R., He, K.: Exploring randomly wired neural networks for image recognition. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1284–1293 (2019) [3](#)
23. Yang, A., Esperança, P.M., Carlucci, F.M.: Nas evaluation is frustratingly hard. arXiv preprint arXiv:1912.12522 (2019) [3](#)
24. Zhang, Q., Chen, M., Bukharin, A., Karampatziakis, N., He, P., Cheng, Y., Chen, W., Zhao, T.: Adalora: Adaptive budget allocation for parameter-efficient fine-tuning. arXiv preprint arXiv:2303.10512 (2023) [3](#)
25. Zhang, Y., Zhou, K., Liu, Z.: Neural prompt search. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024) [2](#), [3](#), [6](#)
26. Zhou, H., Wan, X., Vulić, I., Korhonen, A.: Autopeft: Automatic configuration search for parameter-efficient fine-tuning. *Transactions of the Association for Computational Linguistics* **12**, 525–542 (2024) [2](#), [3](#), [6](#)