

000 DUAL-KERNEL ADAPTER: EXPANDING SPATIAL
 001 HORIZONS FOR DATA-CONSTRAINED MEDICAL IMAGE
 002 ANALYSIS
 003

006 **Anonymous authors**

007 Paper under double-blind review

010 ABSTRACT

013 Adapters have become a widely adopted strategy for efficient fine-tuning of large
 014 pretrained models, particularly in resource-constrained settings. However, their
 015 performance under extreme data scarcity—common in medical imaging due to high
 016 annotation costs, privacy regulations, and fragmented datasets—remains underex-
 017 plored. In this work, we present the first comprehensive study of adapter-based
 018 fine-tuning for large pretrained models in low-data medical imaging scenarios. We
 019 find that, contrary to their promise, conventional Adapters can degrade performance
 020 under severe data constraints, performing even worse than simple linear probing
 021 when trained on less than 1% of the corresponding training data. Through sys-
 022 tematic analysis, we identify a sharp reduction in Effective Receptive Field (ERF)
 023 as a key factor behind this degradation. Motivated by these findings, we propose
 024 the Dual-Kernel Adapter (DKA), a lightweight module that expands spatial con-
 025 text via large-kernel convolutions while preserving local detail with small-kernel
 026 counterparts. Extensive experiments across diverse classification and segmentation
 027 benchmarks show that DKA significantly outperforms existing Adapter methods,
 028 establishing new leading results in both data-constrained and data-rich regimes.
 029 Code is submitted.

030 1 INTRODUCTION

033 The rapid proliferation of large pretrained models has significantly advanced various fields such
 034 as natural language processing (Chowdhary & Chowdhary, 2020) and computer vision (Szeliski,
 035 2022), yet it has also amplified challenges related to computational overhead (Thompson et al., 2020),
 036 memory consumption (Mahendran, 2021), and the complexity of downstream adaptation (Jiang et al.,
 037 2024), especially when deploying these models in specialized domains like medical imaging (Strubell
 038 et al., 2020; Ji et al., 2021; Wang et al., 2025). To address these issues, adapter-based fine-tuning
 039 (Wang et al., 2020b; Wu et al., 2025; Gong et al., 2023; Hu et al., 2024; Chen et al., 2024) has
 040 emerged as a popular strategy, enabling efficient adaptation by adjusting only a small subset of model
 041 parameters rather than performing full fine-tuning.

042 In medical imaging, assembling large, well-annotated datasets is notoriously costly: expert radiolo-
 043 gists must painstakingly delineate structures in high-resolution 2-D and 3-D scans, and inter-observer
 044 variability further inflates the annotation burden (Suzuki, 2017; Ritter et al., 2011). Strict privacy
 045 regulations such as HIPAA (U.S. Department of Health & Human Services, 2003) and the GDPR (Eu-
 046 ropean Union, 2016), coupled with heterogeneous institutional policies, severely limit data sharing,
 047 fragmenting what little data exist (Schäfer et al., 2023; Wang et al., 2021). As a result, many clinically
 048 important downstream tasks still operate in a pronounced low-data regime. This scenario naturally
 049 leads to the critical question:

050 *Can standard Adapter perform effectively in medical imaging tasks under constrained data?*

051
 052 In this paper, we first provide a comprehensive evaluation and analysis of applying conventional
 053 Adapter (Houlsby et al., 2019) under constrained-data scenarios across various datasets and backbone
 architectures. Our study reveals several key insights:

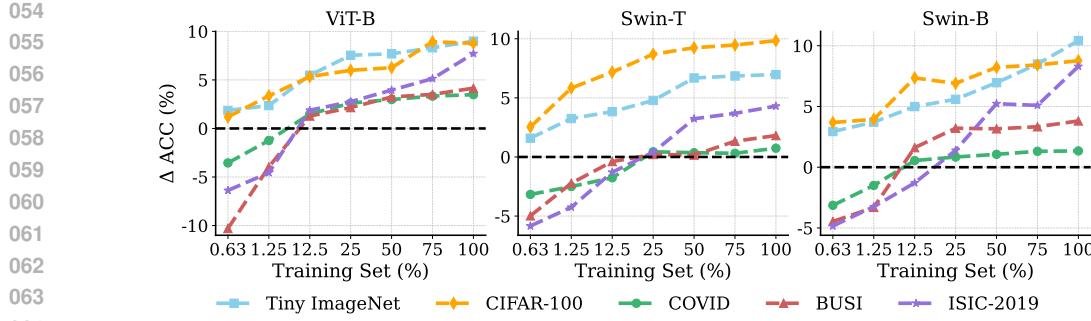


Figure 1: **Performance of Adapter across Various Training Data Sizes.** $\Delta\text{ACC} = \text{ACC}_{\text{Linear Probing} + \text{Adapter}} - \text{ACC}_{\text{Linear Probing}}$. Experiments are conducted on the pretrained ViT-B, Swin-T, and Swin-B backbones for Tiny ImageNet and CIFAR-100 (In-Domain), COVID, BUSI, and ISIC-2019 (Out-of-Domain).

- *Data Size Significantly Impacts Adapter Performance.* We observe that the benefits of using Adapters diminish substantially as the size of the training dataset decreases. This effect is particularly pronounced for medical imaging (see Figure 1).
- *Adapters Can Harm Performance Under Severe Low Data Settings.* When the training data size is reduced to 1% or less, specifically in the context of adapting large pretrained models to medical imaging, Adapters perform worse than linear probing. This indicates that, under extreme data constraints, Adapters may negatively affect model adaptation (see Figure 1).
- *Effective Receptive Field (ERF) Decreases with Reduced Training Data.* Visualization of the ERF demonstrates that smaller training datasets result in reduced ERF, offering a plausible explanation for the observed performance degradation (see Figure 2).

Inspired by these insights, we propose a Dual-Kernel Adapter (DKA) that explicitly enlarges the ERF of standard Adapter modules. Each DKA module introduces a dual-branch convolution design: one branch leverages a depthwise large-kernel convolution to broaden the ERF, while the other employs a depthwise small-kernel convolution to preserve fine-grained local details. We evaluate DKA on a range of medical imaging tasks, including both classification and segmentation, across diverse datasets and large pretrained models. Experimental results demonstrate that DKA consistently outperforms existing methods in both data-constrained and data-rich settings.

Summary of Contributions

- We present the first systematic study of adapter-based fine-tuning for large pretrained models in low-data medical-imaging scenarios, showing that the conventional Adapter can actually degrade performance under severe data scarcity.
- We introduce the Dual-Kernel Adapter (DKA), a lightweight module that pairs large- and small-receptive-field convolutions in parallel, simultaneously broadening spatial context and preserving fine-grained detail.
- Extensive experiments on multiple segmentation and classification benchmarks demonstrate that DKA sets new state of the art, and ablation studies reveal that using asynchronous learning rates between adapters and linear head is critical to its gains.

2 UNDERSTANDING STANDARD ADAPTER IN CONSTRAINED-DATA SETTINGS

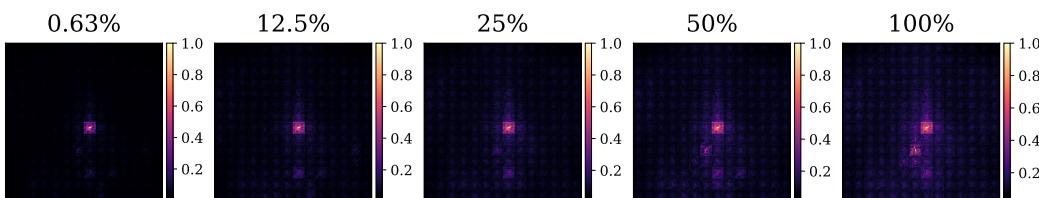
Adapters have gained popularity in medical image analysis due to their parameter efficiency and adaptability for fine-tuning large pretrained models. However, their performance under constrained data conditions remains underexplored. In this section, we investigate this critical aspect using a diverse set of datasets, including Tiny ImageNet (Le & Yang, 2015), CIFAR-100 (Sharma et al., 2018), COVID (Chowdhury et al., 2020), BUSI (Zhang et al., 2022), and ISIC-2019 (Gessert et al., 2020). Our experiments utilize three backbones—ViT-B (Dosovitskiy et al., 2021), Swin-T (Liu et al., 2021), and Swin-B (Liu et al., 2021)—all pretrained on the ImageNet (Deng et al., 2009), to evaluate

108 the impact of Adapters across varying data sizes, ranging from 0.63% to 100% of the training data.
 109 We report the difference in accuracy between the application of Adapters and the non-application of
 110 Adapters, which is denoted by $\Delta\text{ACC} = \text{ACC}_{\text{LinearProbing}+\text{Adapter}} - \text{ACC}_{\text{LinearProbing}}$. Our key
 111 observations are summarized as follows.

112 **① Degraded Adapter Performance with Less Training Data.** In Figure 1, we consistently observe
 113 that the performance gains provided by Adapters diminish across multiple tasks and pretrained
 114 models. Notably, this decline is significantly more pronounced in medical datasets (COVID, BUSI,
 115 ISIC-2019) compared to natural-vision datasets (TinyImageNet, CIFAR-100). A plausible explana-
 116 tion is that medical tasks represent out-of-domain scenarios, requiring feature representations that
 117 diverge considerably from those learned by the original pretrained models. This challenge is further
 118 exacerbated in low-data settings, where learning domain-specific features becomes more difficult. In
 119 contrast, natural-vision tasks remain largely in-domain, aligning more closely with the pretraining
 120 distribution, and thus benefiting more from adapter-based fine-tuning.

121 **② Negative Effects of Adapters in Extremely Low Training Data in Medical Imaging.** We
 122 further observe that in medical imaging tasks, when training data is limited to 1% or less, the
 123 performance gain from applying Adapters becomes negative, indicating a detrimental impact on
 124 model performance. Unlike natural images, medical images typically exhibit low contrast, ambiguous
 125 boundaries, and small or irregular pathological structures (Zhang et al., 2024), which usually demand
 126 a large effective receptive field (ERF) to capture long-range contextual dependencies. However,
 127 standard Adapter do not possess a strong inductive bias toward expanding the ERF. **We hypothesize**
 128 that *under limited supervision, this limitation restricts the Adapter’s capacity to learn spatially*
 129 *dispersed features and long-range contextual dependence, thereby contributing to the observed*
 130 *degradation in performance.*

131 **③ Reduced Effective Receptive Field Under Constrained Training Data.** To validate whether
 132 reduced supervision limits the Adapter’s ability to capture spatially dispersed features and long-
 133 range contextual dependencies, we visualize the ERF of Adapters trained on varying proportions
 134 of the training set, ranging from 0.63% to 100%, using the COVID dataset and the pretrained
 135 ViT-B model. Following the definition in (Araujo et al., 2019), the ERF of a neural network layer
 136 refers to the region encompassing all input pixels that exert a non-negligible influence on a given
 137 output unit. As shown in Figure 2, the ERF becomes progressively smaller as the amount of training
 138 data decreases. This observation supports our hypothesis that limited supervision restricts the
 139 Adapter’s ability to learn spatially diverse patterns and long-range contextual relationships, which
 140 are particularly crucial in medical imaging tasks. We provide complementary ERF analysis on an
 141 alternative backbone (Swin-T) in Appendix F.4, which leads to consistent conclusions, reinforcing
 142 the generality of our observations.



150
 151 **Figure 2: Effective Receptive Field of Standard Adapters Across Varying Training Set Ratios.**
 152 Experiments are conducted on the COVID dataset using the pretrained ViT-B.

153 These findings indicate that standard Adapter fail to enlarge the ERF in low-data settings, which in
 154 turn degrades model performance. Consequently, it is crucial to develop new Adapter architectures
 155 with a built-in inductive bias toward expanding the ERF.

157 3 DUAL-KERNEL ADAPTER

160 To mitigate the adverse effects of standard Adapter in low-data regimes, we propose the
 161 Dual-Kernel Adapter (DKA), which explicitly integrates a large-kernel convolution to ex-
 162 pand the ERF and a small-kernel convolution to preserve local spatial details. Prior studies have

162 demonstrated that large kernels introduce a strong inductive bias toward capturing broader contextual
163 information by expanding the ERF (Huang et al., 2023; Liu et al., 2022; Ding et al., 2022).
164

165 As illustrated in Figure 3, DKA first reduces the dimensionality of the input features through a linear
166 down-projection. After this step, the patch tokens are reshaped back to their 2D spatial layout so that
167 depthwise convolutions can be applied in the image domain. The reduced features are then processed
168 in parallel by two depthwise convolution branches—a computationally efficient form of convolution
169 that applies a distinct filter to each input channel (Chollet, 2017). One branch employs a large kernel
170 to significantly enlarge the receptive field, facilitating the modeling of long-range dependencies.
171 The other branch uses a smaller kernel to retain fine-grained spatial features essential for capturing
172 localized structures. The outputs from the two branches are aggregated via element-wise summation,
173 followed by a GELU activation, a linear up-projection, and a residual connection to the input.
174

175 Formally, the DKA operation can be expressed as:
176

$$f_{DKA}(\mathbf{x}) = \mathbf{x} + \text{Up}(\sigma(\text{DWConv}_{\text{large}}(\text{Down}(\mathbf{x})) + \text{DWConv}_{\text{small}}(\text{Down}(\mathbf{x})))) \quad (1)$$

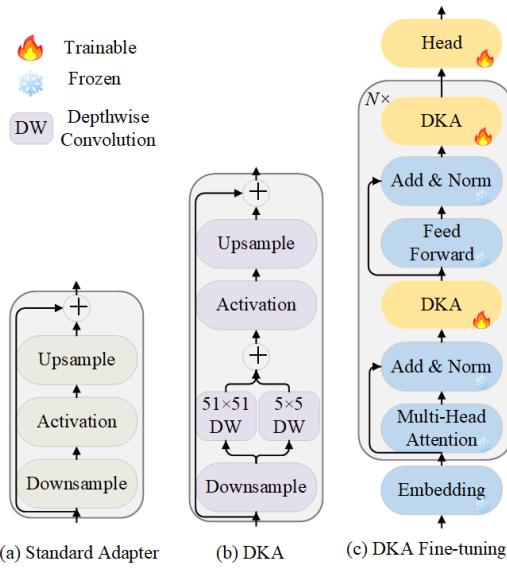
177 where $\text{Down}(\cdot)$ and $\text{Up}(\cdot)$ denote linear projection layers, $\text{DWConv}_{\text{large}}$ and $\text{DWConv}_{\text{small}}$ are depth-
178 wise convolutions with kernel sizes 51 and 5 respectively, and σ is the GELU activation.
179

180 4 EXPERIMENTS

181 To comprehensively evaluate the performance
182 of DKA, we conduct extensive experiments on
183 medical imaging tasks, spanning classification
184 and segmentation benchmarks across diverse
185 datasets, and additionally, models pretrained on
186 both natural and medical domains.
187

188 **Pretrained Models.** We consider two re-
189 presentative categories of pretrained models.
190 Natural-pretrained Models. For classification,
191 we conduct experiments using ViT-B (Dosovitskiy et al., 2021) and Swin-B (Liu et al., 2021),
192 both pretrained on the ImageNet. For segmentation,
193 we adopt Segmenteer-B (Strudel et al., 2021)
194 in the OpenMMLab MMSegmentation (Con-
195 tributors, 2020). Medical-pretrained Models.
196 To assess domain adaptability, we further eval-
197 uate DKA on medical backbones, including
198 RadImageNet-pretrained ResNet-50 (Mei et al.,
199 2022) for classification and MedSAM (Ma et al.,
200 2024) for segmentation.
201

202 **Datasets and Metrics.** We evaluate the pro-
203 posed DKA across both medical classification
204 and segmentation tasks to ensure broad applica-
205 bility. Medical Image Classification. We utilize
206 three widely-adopted medical image classification datasets: COVID (Chowdhury et al., 2020), BUSI
207 (Zhang et al., 2022), and ISIC-2019 (Gessert et al., 2020). The COVID dataset comprises chest X-ray
208 images for COVID-19 diagnosis. The BUSI dataset includes breast ultrasound images categorized as
209 benign, malignant, or normal. The ISIC-2019 dataset is a large-scale dermoscopic dataset designed
210 for multi-class skin lesion classification. For each dataset, we report *Top-1 Accuracy (ACC)*, *F1 Score*
211 (*F1*); and *Sensitivity (SEN)* under varying proportions of training data. Medical Image Segmentation.
212 We further assess DKA on three representative segmentation datasets: BRATS (Menze et al., 2014),
213 BUSI (Zhang et al., 2022), and ISIC-2018 (Codella et al., 2019). BRATS focuses on brain tumor
214 segmentation using multi-modal MRI scans. BUSI provides breast ultrasound images with annotated
215 tumor regions. ISIC-2018 contains dermoscopic images with pixel-level annotations for skin lesion
segmentation. We evaluate segmentation performance using standard metrics: *mean Intersection over*
Union (mIoU) and *Dice coefficient (Dice)*. More details are provided in Appendix B.1.



216 **Figure 3: Overview of the DKA Module.** (a) Stan-
217 dard Adapter. (b) The proposed DKA module. (c)
218 DKA Fine-tuning.

Baselines. We compare DKA against a broad range of baselines, categorized into three groups: Standard Fine-tuning Methods: (1) *Linear Probing*: freezing the backbone and tuning only the head; (2) *Full Fine-tuning*: updating all model parameters during downstream training. Adapter-tuning Methods: (1) *Adapter* (Houlsby et al., 2019): adding adapters within each transformer block; (2) *AdapterFormer* (Chen et al., 2022): adding parallel adapters with learnable scaling to each MLP layer; (3) *CIAT* (Zhu et al., 2021): adding adapters before transformer blocks and parallel adapters within blocks; (4) *Convpass* (Jie et al., 2024): enhancing adapters with parallel 3×3 convolution branches; (5) *AIM* (Yang et al., 2023): combining sequential and parallel adapters across spatial and temporal attention pathways. Other Parameter-efficient Fine-tuning (PEFT) Methods: (1) *Prompt Tuning* (Jia et al., 2022): fine-tuning extra learnable tokens; (2) *LoRA* (Hu et al., 2022): injecting low-rank trainable matrices into attention modules; (3) *BitFit* (Zaken et al., 2021): fine-tuning only the bias terms of pretrained models.

Implementation Details. Following the common training protocol, we freeze all pretrained model weights and update only the parameters of DKA and head during fine-tuning. Specifically, DKA modules are inserted within transformer blocks following the placement strategy described in (Yin et al., 2024). For the DKA module, the middle dimension d is set to 16 for classification tasks and 192 for segmentation tasks, as discussed in Section 4.5. The learning rates are set to 1e-4 for the task head and 1e-3 for the DKA modules, as shown in Section 4.3. Classification models are trained for 100 epochs, while segmentation models are trained for 300 epochs, balancing the need for convergence across task complexities. For experiments under constrained supervision (i.e., training with less than 100% of the training set), we perform a 5-fold cross-validation on the training set while keeping the test set fixed. All reported results are averaged across the cross-validation folds. More details can be found in Appendix B.2.

239

240 4.1 SUPERIOR PERFORMANCE ON CONSTRAINED DATA

241

242 We evaluate DKA under constrained-data settings (0.63% and 1.25%) and the full-data setting. 243 Experiments are organized by the type of pretrained backbone: natural-image models (ViT-B, 244 Segmenter-B) and medical-image models (RadImageNet-ResNet-50, MedSAM). This design enables 245 a clear examination of DKA’s generalization across different pretraining sources.

246

Natural-pretrained Models. We first evaluate DKA using natural-pretrained backbones, as summarized in Tables 1 and 2, DKA consistently outperforms all baselines across classification and segmentation tasks. Notably, under low-data regimes (0.63% and 1.25%), DKA even surpasses full fine-tuning and linear probing, while other PEFT methods exhibit clear performance degradation. These observations confirm the effectiveness of DKA when adapting natural-pretrained models to medical domains. Additional results on the natural-pretrained ViT-B and Segmenter-B are provided in Appendix D and C. Furthermore, similar performance improvements are observed when DKA is applied to other natural-pretrained models such as Swin-B (see Appendix F.1 for details).

254

255

256 **Table 1: Comparison of Baselines and DKA on Three Classification Datasets Across Varying**
 257 **Data Sizes.** Results are reported in terms of ACC (%). Experiments are based on the pretrained
 258 ViT-B. The best results are highlighted in bold.

259 260 Methods	261 COVID			262 BUSI			263 ISIC-2019		
	264 0.63%	265 1.25%	266 100%	267 0.63%	268 1.25%	269 100%	270 0.63%	271 1.25%	272 100%
273 Full Fine-tuning	274 87.43	275 88.00	276 98.43	277 71.17	278 76.73	279 94.62	280 60.04	281 61.21	282 82.05
283 Linear Probing	284 86.84	285 87.50	286 94.85	287 73.48	288 77.64	289 89.78	290 59.15	291 59.44	292 71.83
293 BitFit (Zaken et al., 2021)	294 73.91	295 79.65	296 96.95	297 57.19	298 60.20	299 88.27	300 50.80	301 53.22	302 79.84
303 Prompt (Jia et al., 2022)	304 77.91	305 83.75	306 98.45	307 61.34	308 64.30	309 93.07	310 52.55	311 53.59	312 81.02
313 LoRA (Hu et al., 2022)	314 80.43	315 85.91	316 98.73	317 63.64	318 67.41	319 94.75	320 51.08	321 53.96	322 81.75
323 Adapter (Houlsby et al., 2019)	324 83.29	325 86.26	326 98.33	327 63.18	328 73.68	329 93.33	330 52.77	331 54.88	332 79.54
333 Adapterformer (Chen et al., 2022)	334 82.46	335 84.32	336 98.18	337 63.42	338 72.75	339 92.65	340 51.62	341 52.71	342 78.19
343 Convpass (Jie et al., 2024)	344 84.72	345 86.94	346 98.45	347 64.83	348 74.63	349 93.97	350 54.72	351 56.25	352 80.45
353 CIAT (Zhu et al., 2021)	354 77.34	355 82.85	356 96.54	357 60.28	358 65.07	359 89.86	360 48.35	361 49.96	362 72.18
363 AIM (Yang et al., 2023)	364 80.92	365 83.55	366 97.23	367 62.72	368 70.34	369 90.12	370 50.12	371 52.72	372 77.39
373 DKA	374 89.01	375 91.06	376 99.21	377 74.23	378 79.46	379 95.89	380 60.52	381 62.32	382 83.09

270
 271 **Table 2: Comparison of Baselines and DKA on Three Segmentation Datasets Under Varying**
 272 **Data Sizes.** Experiments are based on the pretrained Segmenteer-B. mIoU is reported as percentages.
 273 The best results are highlighted in bold.

274 275 Methods	276 BRATS			277 BUSI			278 ISIC-2018		
	0.63%	1.25%	100%	0.63%	1.25%	100%	0.63%	1.25%	100%
279 Full Fine-tuning	9.25	22.39	73.08	26.67	32.31	57.41	62.27	73.63	77.58
280 Linear Probing	7.95	20.20	69.86	25.53	31.96	54.07	60.90	71.06	74.10
281 BitFit (Zaken et al., 2021)	1.20	14.33	63.52	7.13	17.08	52.56	53.21	65.84	73.10
282 Prompt (Jia et al., 2022)	1.22	15.21	64.53	9.19	18.77	53.57	56.56	67.40	73.61
283 LoRA (Hu et al., 2022)	3.84	16.19	68.48	14.10	22.39	53.96	58.70	69.03	73.84
284 Adapter (Houlsby et al., 2019)	6.16	18.95	72.02	18.18	25.90	55.01	59.78	72.80	76.71
285 Adapterformer (Chen et al., 2022)	5.99	18.77	72.54	17.41	25.68	55.14	59.67	72.85	76.58
286 Convpass (Jie et al., 2024)	7.13	19.64	73.32	19.84	28.52	56.09	60.19	73.56	77.54
287 CIAT (Zhu et al., 2021)	3.80	16.81	70.41	13.40	20.20	54.76	58.26	69.52	75.09
288 AIM (Yang et al., 2023)	4.58	17.61	71.98	15.40	23.54	54.78	59.24	72.05	76.65
289 DKA	9.47	23.02	74.96	26.85	34.52	58.90	63.13	74.27	78.53

288 **Medical-pretrained Models.** We further evaluate DKA on medical-pretrained backbones, including
 289 RadImageNet-pretrained ResNet-50 (Mei et al., 2022) for classification and MedSAM (Ma et al.,
 290 2024) for segmentation. As reported in Table 3, DKA consistently surpasses full fine-tuning, linear
 291 probing, and other PEFT baselines across ISIC-2019 and BUSI under all data scales. These im-
 292 provements demonstrate that the gains of DKA are not limited to natural image pretraining, but also
 293 generalize effectively to domain-specific medical large pretrained models. More results are reported
 294 in Appendix E.

295
 296 **Table 3: Performance of DKA with medical-pretrained models.** (a) Classification results on the
 297 ISIC-2019 using RadImageNet-pretrained ResNet-50 by reporting ACC (%). (b) Segmentation results
 298 on the BUSI using MedSAM by reporting mIoU (%). The best results are highlighted in bold.

(a) Classification Results.				(b) Segmentation Results.			
300 Methods	0.63%	1.25%	100%	301 Methods	0.63%	1.25%	100%
302 Full Fine-tuning	52.70	55.17	76.63	303 Full Fine-tuning	36.21	45.04	70.62
303 Linear Probing	51.27	53.71	67.39	304 Linear Probing	34.72	42.76	66.54
304 BitFit	48.84	51.62	70.38	305 BitFit	27.85	36.92	64.43
305 Prompt	50.12	53.29	73.31	306 Prompt	29.57	38.71	64.62
306 LoRA	50.03	52.51	72.92	307 LoRA	32.39	40.35	67.04
307 Adapter	51.32	54.04	74.26	308 Adapter	35.40	43.62	68.06
309 DKA	53.69	56.58	78.56	310 DKA	37.13	46.27	72.53

311 4.2 ERF VISUALIZATION

312 To assess whether DKA effectively expands the effective receptive field (ERF), we visualize the
 313 ERFs of DKA alongside those of other adapter-tuning baselines, including Adapter (Houlsby et al.,
 314 2019), AdapterFormer (Chen et al., 2022), Convpass (Jie et al., 2024), CIAT (Zhu et al., 2021),
 315 and AIM (Yang et al., 2023), under both the constrained (0.63%) and the full (100%) training data
 316 settings. As shown in Figure 4, DKA consistently exhibits the broadest ERF across both settings,
 317 demonstrating its superior ability to capture extensive spatial context even under constrained-data
 318 conditions. In contrast, other baselines yield more localized ERFs, particularly in the constrained-data
 319 setting, which likely contributes to their comparatively weaker performance.

320 4.3 LARGE KERNEL MATTERS INSTEAD OF TRAINABLE PARAMETERS

321 Given that DKA introduces a slightly higher number of trainable parameters, a natural question arises:
 322 *Are the observed performance gains primarily attributable to the increased parameter count or to the
 323 effect of the large kernel?* To investigate this, we perform a controlled comparison where the increase
 in trainable parameters arises either from enlarging the kernel size or from expanding the intermediate

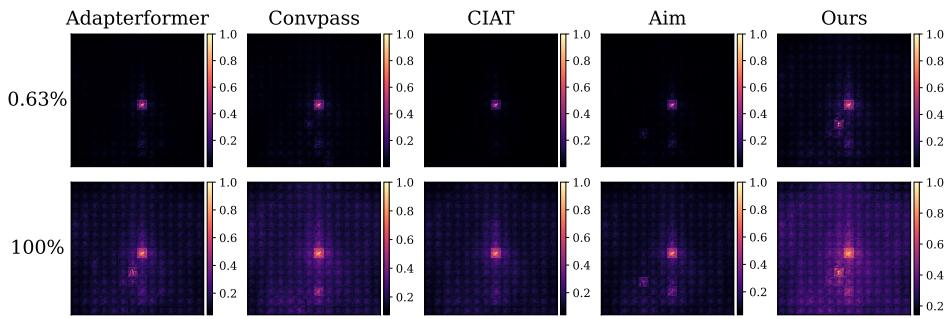


Figure 4: **Effective Receptive Field of DKA and Other Adapter-based Methods.** Experiments are conducted on the COVID dataset based on the pretrained ViT-B under both constrained (0.63%) and full (100%) training settings.

dimension \hat{d} . Specifically, we adjust the middle dimension \hat{d} of other adapter-based baselines so that their total trainable parameter counts align with that of DKA with increasing kernel sizes. We fix the middle dimension $\hat{d} = 16$ in DKA throughout all classification experiments, ensuring that any parameter increase stems solely from the enlarged kernel. As shown in Figure 5, the results reveal that: ① DKA consistently outperforms all baselines across the 0.63%, 1.25%, and 100% training data settings under similar parameter constraints; ② the performance improvements resulting from increasing the kernel size (from 11 to 51) are significantly steeper than those obtained by merely enlarging the hidden dimension, highlighting the effectiveness of large kernels in enhancing DKA’s representational capacity.

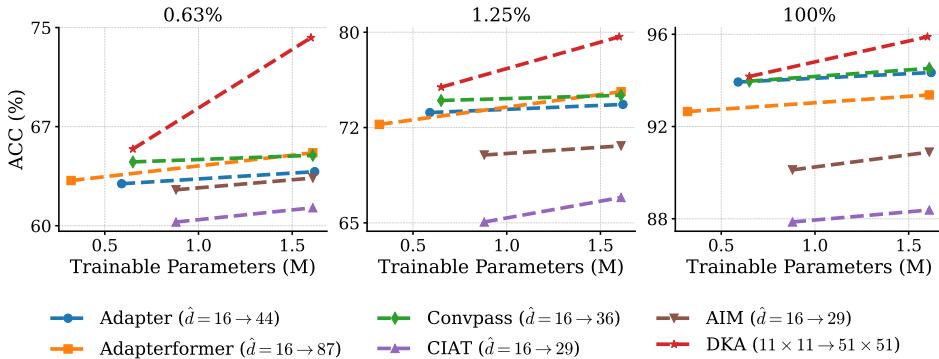


Figure 5: **Comparison of Baselines and DKA with Comparable Numbers of Trainable Parameters.** Experiments are conducted on 0.63%, 1.25%, and 100% subsets of the BUSI dataset using the pretrained ViT-B. The symbol \hat{d} denotes the intermediate dimensionality in adapter-based methods.

4.4 ASYNCHRONOUS LEARNING RATES MATTER

In standard adapter-based fine-tuning, it is common practice to use the same learning rate for both the adapter and the head. However, given their different roles and characteristics, this one-size-fits-all strategy may not be optimal. To investigate the effect of asynchronous learning rates for the adapter and head, we conduct experiments on the COVID dataset using a pretrained ViT-B. We systematically vary the learning rates assigned to the adapter and head, and assess the results under both the 0.63% and 1.25% training data. As shown in Figure 6, asynchronous learning rate schedules—where the adapter and head use different learning rates—often outperform symmetric configurations. We observe that the best results are not achieved when both components share the same learning rate, suggesting that the adapter and head benefit from distinct optimization dynamics. This trend holds

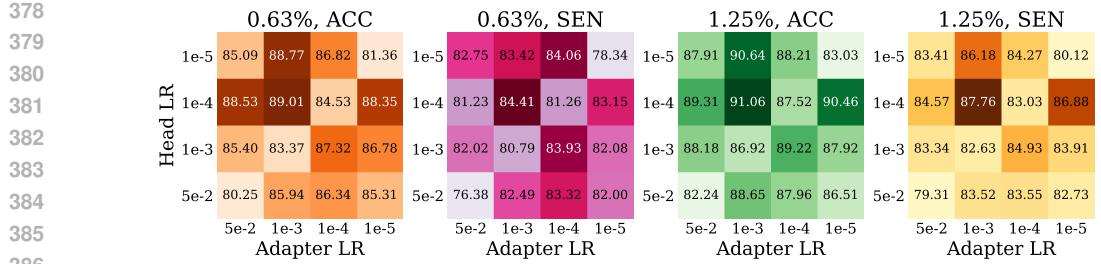


Figure 6: **Performance Comparison Across Varying Learning Rates for DKA and Classification Head.** Experiments are conducted on the COVID dataset using the pretrained ViT-B under 0.63% and 1.25% training data. Results are reported in terms of ACC (%) and SEN (%).

across data scales, highlighting the importance of tuning these components independently. More results are available in Appendix F.2.

4.5 ABLATION STUDY

Kernel Size Selection. We investigate the impact of different kernel size combinations in DKA’s dual-branch convolution design. Specifically, we sweep over five candidate sizes (3×3 , 5×5 , 7×7 , 9×9 , 11×11 , 31×31 , 51×51 , 71×71) for both small- and large-depthwise convolution branches. As shown in Figure 7, the combination of 5×5 (small) and 51×51 (large) consistently yields the best accuracy on both the data-constrained (0.63% and 1.25%) setting and the full data setting (100%). Notably, kernel sizes that are too small or too large result in performance degradation, particularly under low-data conditions. These results support our dual-branch design choice, balancing fine-grained detail extraction and large receptive field expanding. Extra experimental results are included in Appendix F.5.

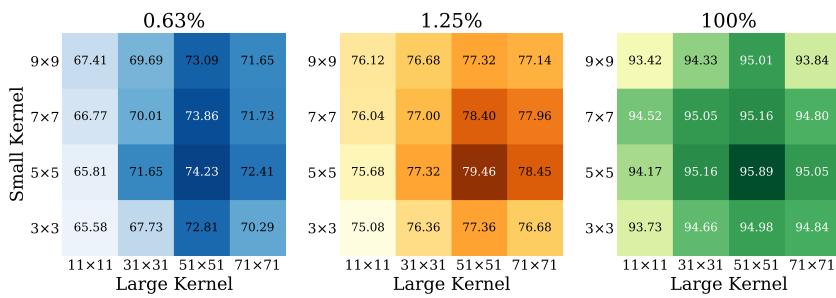


Figure 7: **Performance of Different Kernel Size Combination.** Experiments are conducted on the BUSI dataset using the pretrained ViT-B across three training setting (0.63%, 1.25%, and 100%). ACC (%) is reported.

Single vs. Dual Convolutions. To evaluate the effectiveness of using both large- and small-depthwise convolutions in DKA, we compare the full dual-branch design ($5 \times 5 + 51 \times 51$) against single-branch variants that use only one of the two kernels. Experiments are conducted on the COVID, BUSI, and ISIC-2019 datasets across varying training set sizes. As shown in Figure 8, the dual-branch design consistently outperforms both single-branch variants, particularly in low-data regimes. While the 5×5 branch better captures localized detail and the 51×51 branch improves global coverage, neither alone matches the full-dual structure. See Appendix F.6 for additional results.

Middle Dimension. We investigate the impact of the middle dimension \hat{d} in the DKA module for both classification and segmentation tasks. As shown in Table 4, increasing \hat{d} generally leads to improved performance. For classification on the BUSI dataset, performance steadily improves from $\hat{d} = 1$ to $\hat{d} = 16$, reaching a peak at $\hat{d} = 16$, after which performance slightly declines or saturates, indicating potential overfitting or redundancy. For segmentation on the ISIC-2018 dataset, the optimal middle dimension is also moderate, with $\hat{d} = 192$ offering the best trade-off between parameter efficiency and accuracy. This observation aligns with the findings in (Chen et al., 2022), which also highlight the task-specific nature of optimal middle dimensions, emphasizing that different tasks require different parameter settings for optimal performance.

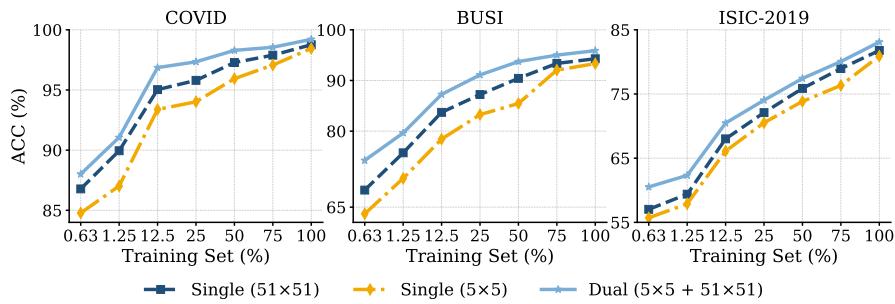


Figure 8: **Ablation for Dual-Convolution Design.** Experiments are based on the pretrained ViT-B across three classification datasets. ACC (%) is reported.

Table 4: **Performance Comparison of Different Middle Dimensions \hat{d} in DKA.** (a) Classification performance using the pretrained ViT-B on the BUSI dataset by reporting ACC (%). (b) Segmentation performance using the pretrained Segmenter-B on the ISIC-2018 dataset by reporting mIoU (%).

(a) Middle Dimension for Classification. (b) Middle Dimension for Segmentation.

\hat{d}	0.63%	1.25%	100%	\hat{d}	0.63%	1.25%	100%
1	65.02	70.69	87.86	64	57.93	68.68	74.10
4	69.65	74.38	91.42	96	60.25	71.95	76.18
8	72.18	76.68	93.61	128	62.30	73.56	77.48
16	74.23	79.64	95.89	192	63.13	74.27	78.06
32	73.75	79.23	95.29	256	62.72	73.88	77.95

5 RELATED WORK

Adapter-based Fine-tuning adapts large pretrained models to downstream tasks by inserting and training lightweight modules while keeping the original model parameters frozen, offering significant computational and memory advantages over full fine-tuning (Xu et al., 2023; Ding et al., 2023; Han et al., 2024b). Early Adapter methods introduced task-specific bottleneck layers between transformer blocks (Houlsby et al., 2019), with subsequent innovations improving architectural designs (Wang et al., 2020b; Zhang et al., 2021) and multi-modal applications (Sung et al., 2022; Pan et al., 2022; Gao et al., 2024). Recent medical imaging adaptations demonstrate how Adapter modules can effectively transfer pretrained knowledge to diagnostic tasks while maintaining model integrity (Chen et al., 2023; Dutt et al., 2023b; Lian et al., 2024). Recent large-kernel designs such as LKA (Zhu et al., 2025) enhance global receptive fields within convolutional backbones. However, adapter-based methods typically assume access to a moderate amount of labeled data (Dutt et al., 2023a; Liu et al., 2024). Their effectiveness in data-constrained settings remains underexplored, raising critical questions about their performance in such scenarios.

Limited Data in Medical Imaging remains a major challenge, as labeled samples are scarce due to privacy, high annotation cost, and limited experts (Shaikhina & Khovanova, 2017; Chlap et al., 2021). This constraint is especially critical in clinical practice, where collecting large and diverse datasets is difficult (Ewing, 2017; Lee & Yoon, 2017; Li et al., 2010). To address scarcity, strategies include data augmentation (Garcea et al., 2023; Goceri, 2023; Zhao et al., 2019; Islam et al., 2024), transfer learning (Raghu et al., 2019; Kim et al., 2022; Kora et al., 2022), semi-supervised learning (Huynh et al., 2022; Chebli et al., 2018; Wang et al., 2020a; Han et al., 2024a; Zhou et al., 2019), and self-supervised representation learning (Ericsson et al., 2022; Jiao et al., 2020; Ye et al., 2024; Krishnan et al., 2022). These approaches leverage unlabeled data or external sources to improve generalization under low-resource settings (Zheng et al., 2024; Lin et al., 2021), and recent work highlights large pretrained models to further reduce annotation needs (Moor et al., 2023; Zhang & Metaxas, 2024; Khan et al., 2025). Nevertheless, most rely on full fine-tuning or moderately sized datasets (Davila et al., 2024; Khan & Fang, 2023), while their integration with adapter-based methods under severe scarcity remains underexplored but highly relevant in practice.

486

6 CONCLUSION

488 In this paper, we revisit adapter-based fine-tuning for medical image analysis and uncovered key
 489 limitations in low-data settings: existing Adapters often struggle to capture relevant features under
 490 data scarcity, partly due to their constrained receptive field. To address this, we introduced DKA, a
 491 dual-branch adapter module that integrates large- and small-kernel depthwise convolutions to enhance
 492 the receptive field while preserving local detail. Experiments on both medical image classification
 493 and segmentation tasks, evaluated across various datasets and backbones show that DKA consistently
 494 outperform both full fine-tuning and other PEFT baselines by a good margin, particularly in the
 495 constrained-data setting, without significantly increasing parameter count.

496

7 ETHICS STATEMENT

499 We confirm that this study adheres to the ICLR Code of Ethics, in every respect. All datasets used are
 500 publicly available and have been widely adopted in prior research.

502

8 REPRODUCIBILITY STATEMENT

504 We provide complete implementation details, including dataset splits, training settings, and evaluation
 505 protocols in Section 4, Appendix B.1, and Appendix B.2. Code and data will be released to facilitate
 506 reproducibility and further research.

508

REFERENCES

510 André Araujo, Wade Norris, and Jack Sim. Computing receptive fields of convolutional neural
 511 networks. *Distill*, 4(11):e21, 2019.

513 Asma Chebli, Akila Djebbar, and Hayet Farida Marouani. Semi-supervised learning for medical
 514 application: A survey. In *2018 international conference on applied smart systems (ICASS)*, pp.
 515 1–9. IEEE, 2018.

516 Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo.
 517 Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural
 518 Information Processing Systems*, 35:16664–16678, 2022.

520 Tianrun Chen, Lanyun Zhu, Chaotao Deng, Runlong Cao, Yan Wang, Shangzhan Zhang, Zejian
 521 Li, Lingyun Sun, Ying Zang, and Papa Mao. Sam-adapter: Adapting segment anything in
 522 underperformed scenes. In *Proceedings of the IEEE/CVF International Conference on Computer
 523 Vision*, pp. 3367–3375, 2023.

524 Tianrun Chen, Ankang Lu, Lanyun Zhu, Chaotao Ding, Chunan Yu, Deyi Ji, Zejian Li, Lingyun
 525 Sun, Papa Mao, and Ying Zang. Sam2-adapter: Evaluating & adapting segment anything 2 in
 526 downstream tasks: Camouflage, shadow, medical image segmentation, and more. *arXiv preprint
 527 arXiv:2408.04579*, 2024.

528 Phillip Chlap, Hang Min, Nym Vandenberg, Jason Dowling, Lois Holloway, and Annette Haworth. A
 529 review of medical image data augmentation techniques for deep learning applications. *Journal of
 530 medical imaging and radiation oncology*, 65(5):545–563, 2021.

532 François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of
 533 the IEEE conference on computer vision and pattern recognition*, pp. 1251–1258, 2017.

534 KR1442 Chowdhary and KR Chowdhary. Natural language processing. *Fundamentals of artificial
 535 intelligence*, pp. 603–649, 2020.

537 Muhammad EH Chowdhury, Tawsifur Rahman, Amith Khandakar, Rashid Mazhar, Muhammad Ab-
 538 dul Kadir, Zaid Bin Mahbub, Khandakar Reajul Islam, Muhammad Salman Khan, Atif Iqbal,
 539 Nasser Al Emadi, et al. Can ai help in screening viral and covid-19 pneumonia? *Ieee Access*, 8:
 132665–132676, 2020.

540 Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David
 541 Gutman, Brian Helba, Aadi Kalloo, Konstantinos Liopyris, Michael Marchetti, et al. Skin lesion
 542 analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging
 543 collaboration (isic). *arXiv preprint arXiv:1902.03368*, 2019.

544

545 MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and
 546 benchmark. <https://github.com/open-mmlab/mmsegmentation>, 2020.

547

548 Ana Davila, Jacinto Colan, and Yasuhisa Hasegawa. Comparison of fine-tuning strategies for transfer
 549 learning in medical image classification. *Image and Vision Computing*, 146:105012, 2024.

550

551 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale
 552 hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,
 553 pp. 248–255. Ieee, 2009.

554

555 Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin
 556 Chen, Chi-Min Chan, Weize Chen, et al. Parameter-efficient fine-tuning of large-scale pre-trained
 557 language models. *Nature Machine Intelligence*, 5(3):220–235, 2023.

558

559 Xiaohan Ding, Xiangyu Zhang, Jungong Han, and Guiguang Ding. Scaling up your kernels to 31x31:
 560 Revisiting large kernel design in cnns. In *Proceedings of the IEEE/CVF conference on computer
 561 vision and pattern recognition*, pp. 11963–11975, 2022.

562

563 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
 564 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit,
 565 and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale.
 566 *arXiv preprint arXiv:2010.11929*, 2021.

567

568 Raman Dutt, Ondrej Bohdal, Sotirios A Tsaftaris, and Timothy Hospedales. Fairtune: Opti-
 569 mizing parameter efficient fine tuning for fairness in medical image analysis. *arXiv preprint
 570 arXiv:2310.05055*, 2023a.

571

572 Raman Dutt, Linus Ericsson, Pedro Sanchez, Sotirios A Tsaftaris, and Timothy Hospedales.
 573 Parameter-efficient fine-tuning for medical image analysis: The missed opportunity. *arXiv preprint
 574 arXiv:2305.08252*, 2023b.

575

576 Linus Ericsson, Henry Gouk, Chen Change Loy, and Timothy M Hospedales. Self-supervised
 577 representation learning: Introduction, advances, and challenges. *IEEE Signal Processing Magazine*,
 578 39(3):42–62, 2022.

579

580 European Union. General data protection regulation (gdpr) – regulation (eu) 2016/679, 2016. URL
 581 <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

582

583 GW Ewing. The limitations of big data in healthcare. *MOJ Proteomics Bioinform*, 5(2):00152, 2017.

584

585 Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and
 586 Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *International Journal
 587 of Computer Vision*, 132(2):581–595, 2024.

588

589 Fabio Garcea, Alessio Serra, Fabrizio Lamberti, and Lia Morra. Data augmentation for medical
 590 imaging: A systematic literature review. *Computers in Biology and Medicine*, 152:106391, 2023.

591

592 Nils Gessert, Maximilian Nielsen, Mohsin Shaikh, René Werner, and Alexander Schlaefer. Skin
 593 lesion classification using ensembles of multi-resolution efficientnets with meta data. *MethodsX*, 7:
 100864, 2020.

594

595 Evgin Goceri. Medical image data augmentation: techniques, comparisons and interpretations.
 596 *Artificial Intelligence Review*, 56(11):12561–12605, 2023.

597

598 Shizhan Gong, Yuan Zhong, Wenao Ma, Jinpeng Li, Zhao Wang, Jingyang Zhang, Pheng-Ann Heng,
 599 and Qi Dou. 3dsam-adapter: Holistic adaptation of sam from 2d to 3d for promptable medical
 600 image segmentation. *arXiv e-prints*, pp. arXiv–2306, 2023.

594 Kai Han, Victor S Sheng, Yuqing Song, Yi Liu, Chengjian Qiu, Siqi Ma, and Zhe Liu. Deep semi-
 595 supervised learning for medical image segmentation: A review. *Expert Systems with Applications*,
 596 245:123052, 2024a.

597 Zeyu Han, Chao Gao, Jinyang Liu, Jeff Zhang, and Sai Qian Zhang. Parameter-efficient fine-tuning
 598 for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*, 2024b.

600 Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe,
 601 Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for
 602 nlp. In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.

603 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 604 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

606 Jihong Hu, Yinhao Li, Hao Sun, Yu Song, Chujie Zhang, Lanfen Lin, and Yen-Wei Chen. Lga: A
 607 language guide adapter for advancing the sam model’s capabilities in medical image segmentation.
 608 In *International Conference on Medical Image Computing and Computer-Assisted Intervention*,
 609 pp. 610–620. Springer, 2024.

610 Tianjin Huang, Lu Yin, Zhenyu Zhang, Li Shen, Meng Fang, Mykola Pechenizkiy, Zhangyang Wang,
 611 and Shiwei Liu. Are large kernels better teachers than transformers for convnets? In *International
 612 Conference on Machine Learning*, pp. 14023–14038. PMLR, 2023.

614 Tri Huynh, Aiden Nibali, and Zhen He. Semi-supervised learning for medical image classification
 615 using imbalanced training data. *Computer methods and programs in biomedicine*, 216:106628,
 616 2022.

617 Tauhidul Islam, Md Sadman Hafiz, Jamin Rahman Jim, Md Mohsin Kabir, and MF Mridha. A
 618 systematic review of deep learning data augmentation in medical imaging: Recent advances and
 619 future research directions. *Healthcare Analytics*, pp. 100340, 2024.

621 Cheng Ji, Fan Wu, Zongwei Zhu, Li-Pin Chang, Huanghe Liu, and Wenjie Zhai. Memory-efficient
 622 deep learning inference with incremental weight loading and data layout reorganization on edge
 623 systems. *Journal of Systems Architecture*, 118:102183, 2021.

624 Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and
 625 Ser-Nam Lim. Visual prompt tuning. In *European conference on computer vision*, pp. 709–727.
 626 Springer, 2022.

628 Wenyu Jiang, Zhenlong Liu, Zejian Xie, Songxin Zhang, Bingyi Jing, and Hongxin Wei. Exploring
 629 learning complexity for efficient downstream dataset pruning. *arXiv preprint arXiv:2402.05356*,
 630 2024.

631 Jianbo Jiao, Richard Droste, Lior Drukker, Aris T Papageorghiou, and J Alison Noble. Self-supervised
 632 representation learning for ultrasound video. In *2020 IEEE 17th international symposium on
 633 biomedical imaging (ISBI)*, pp. 1847–1850. IEEE, 2020.

635 Shibo Jie, Zhi-Hong Deng, Shixuan Chen, and Zhijuan Jin. Convolutional bypasses are better vision
 636 transformer adapters. In *ECAI 2024*, pp. 202–209. IOS Press, 2024.

637 Muhammad Osama Khan and Yi Fang. Revisiting fine-tuning strategies for self-supervised medical
 638 imaging analysis. *arXiv preprint arXiv:2307.10915*, 2023.

640 Wasif Khan, Seowung Leem, Kyle B See, Joshua K Wong, Shaoting Zhang, and Ruogu Fang. A
 641 comprehensive survey of foundation models in medicine. *IEEE Reviews in Biomedical Engineering*,
 642 2025.

643 Hee E Kim, Alejandro Cosa-Linan, Nandhini Santhanam, Mahboubeh Jannesari, Mate E Maros, and
 644 Thomas Ganslandt. Transfer learning for medical image classification: a literature review. *BMC
 645 medical imaging*, 22(1):69, 2022.

647 Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*,
 2014.

648 Padmavathi Kora, Chui Ping Ooi, Oliver Faust, U Raghavendra, Anjan Gudigar, Wai Yee Chan,
 649 K Meenakshi, K Swaraja, Pawel Plawiak, and U Rajendra Acharya. Transfer learning techniques
 650 for medical image analysis: A review. *Biocybernetics and biomedical engineering*, 42(1):79–107,
 651 2022.

652 Rayan Krishnan, Pranav Rajpurkar, and Eric J Topol. Self-supervised learning in medicine and
 653 healthcare. *Nature Biomedical Engineering*, 6(12):1346–1352, 2022.

654 Ya Le and Xuan S. Yang. Tiny imagenet visual recognition challenge. *CS231n*, 2015.

655 Choong Ho Lee and Hyung-Jin Yoon. Medical big data: promise and challenges. *Kidney research
 656 and clinical practice*, 36(1):3, 2017.

657 Der-Chiang Li, Chiao-Wen Liu, and Susan C Hu. A learning method for the class imbalance problem
 658 with medical data sets. *Computers in biology and medicine*, 40(5):509–518, 2010.

659 Chenyu Lian, Hong-Yu Zhou, Yizhou Yu, and Liansheng Wang. Less could be better: Parameter-
 660 efficient fine-tuning advances medical vision foundation models. *arXiv preprint arXiv:2401.12215*,
 661 2024.

662 Shuai Lin, Pan Zhou, Xiaodan Liang, Jianheng Tang, Ruihui Zhao, Ziliang Chen, and Liang Lin.
 663 Graph-evolving meta-learning for low-resource medical dialogue generation. In *proceedings of the
 664 AAAI conference on artificial intelligence*, volume 35, pp. 13362–13370, 2021.

665 Gang Liu, Jinlong He, Pengfei Li, Genrong He, Zhaolin Chen, and Shenjun Zhong. Pefomed:
 666 Parameter efficient fine-tuning of multimodal large language models for medical imaging. *arXiv
 667 preprint arXiv:2401.02797*, 2024.

668 Shiwei Liu, Tianlong Chen, Xiaohan Chen, Xuxi Chen, Qiao Xiao, Boqian Wu, Tommi Kärkkäinen,
 669 Mykola Pechenizkiy, Decebal Mocanu, and Zhangyang Wang. More convnets in the 2020s: Scaling
 670 up kernels beyond 51x51 using sparsity. *arXiv preprint arXiv:2207.03620*, 2022.

671 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.
 672 Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the
 673 IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.

674 Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. Segment anything in medical
 675 images. *Nature Communications*, 15(1):654, 2024.

676 N Mahendran. Analysis of memory consumption by neural networks based on hyperparameters.
 677 *arXiv*, 2021.

678 Xueyan Mei, Zelong Liu, Philip M Robson, Brett Marinelli, Mingqian Huang, Amish Doshi, Adam
 679 Jacobi, Chendi Cao, Katherine E Link, Thomas Yang, et al. Radimagenet: an open radiologic deep
 680 learning research dataset for effective transfer learning. *Radiology: Artificial Intelligence*, 4(5):
 681 e210315, 2022.

682 Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin
 683 Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain
 684 tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):
 685 1993–2024, 2014.

686 Michael Moor, Oishi Banerjee, Zahra Shakeri Hossein Abad, Harlan M Krumholz, Jure Leskovec,
 687 Eric J Topol, and Pranav Rajpurkar. Foundation models for generalist medical artificial intelligence.
 688 *Nature*, 616(7956):259–265, 2023.

689 Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient
 690 image-to-video transfer learning. *Advances in Neural Information Processing Systems*, 35:26462–
 691 26477, 2022.

692 A Paszke. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint
 693 arXiv:1912.01703*, 2019.

702 Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding
 703 transfer learning for medical imaging. *Advances in neural information processing systems*, 32,
 704 2019.

705 Felix Ritter, Tobias Boskamp, André Homeyer, Hendrik Laue, Michael Schwier, Florian Link, and
 706 H-O Peitgen. Medical image analysis. *IEEE pulse*, 2(6):60–70, 2011.

708 Raphael Schäfer, Till Nicke, et al. Overcoming data scarcity in biomedical imaging with a founda-
 709 tional multi-task model. *Nature Machine Intelligence*, 5:1–10, 2023.

710 Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh,
 711 and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based local-
 712 ization. In *Proceedings of the IEEE international conference on computer vision*, pp. 618–626,
 713 2017.

714 Torgyn Shaikhina and Natalia A Khovanova. Handling limited datasets with neural networks in
 715 medical applications: A small-data approach. *Artificial intelligence in medicine*, 75:51–63, 2017.

716 Fereshteh Shakeri, Yunshi Huang, Julio Silva-Rodríguez, Houda Bahig, An Tang, Jose Dolz, and
 717 Ismail Ben Ayed. Few-shot adaptation of medical vision-language models. In *International
 718 Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 553–563.
 719 Springer, 2024.

720 Neha Sharma, Vibhor Jain, and Anju Mishra. An analysis of convolutional neural networks for image
 721 classification. *Procedia computer science*, 132:377–384, 2018.

722 Julio Silva-Rodriguez, Sina Hajimiri, Ismail Ben Ayed, and Jose Dolz. A closer look at the few-shot
 723 adaptation of large vision-language models. In *Proceedings of the IEEE/CVF Conference on
 724 Computer Vision and Pattern Recognition*, pp. 23681–23690, 2024.

725 Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for
 726 modern deep learning research. In *Proceedings of the AAAI conference on artificial intelligence*,
 727 volume 34, pp. 13693–13696, 2020.

728 Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for
 729 semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer
 730 vision*, pp. 7262–7272, 2021.

731 Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. Vi-adapter: Parameter-efficient transfer learning for
 732 vision-and-language tasks. In *Proceedings of the IEEE/CVF conference on computer vision and
 733 pattern recognition*, pp. 5227–5237, 2022.

734 Kenji Suzuki. Overview of deep learning in medical imaging. *Radiological physics and technology*,
 735 10(3):257–273, 2017.

736 Richard Szeliski. *Computer vision: algorithms and applications*. Springer Nature, 2022.

737 Neil C Thompson, Kristjan Greenewald, Keeheon Lee, Gabriel F Manso, et al. The computational
 738 limits of deep learning. *arXiv preprint arXiv:2007.05558*, 10, 2020.

739 U.S. Department of Health & Human Services. Summary of the hipaa privacy
 740 rule, 2003. URL <https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html>.

741 Dong Wang, Yuan Zhang, Kexin Zhang, and Liwei Wang. Focalmix: Semi-supervised learning for
 742 3d medical image detection. In *Proceedings of the IEEE/CVF conference on computer vision and
 743 pattern recognition*, pp. 3951–3960, 2020a.

744 Jiaji Wang, Shuihua Wang, and Yudong Zhang. Deep learning on medical image analysis. *CAAI
 745 Transactions on Intelligence Technology*, 10(1):1–35, 2025.

746 Rongguang Wang, Pratik Chaudhari, and Christos Davatzikos. Embracing the disharmony in medical
 747 imaging: A simple and effective framework for domain adaptation. *Medical Image Analysis*, 72:
 748 102–110, 2021.

756 Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Guihong Cao, Daxin Jiang,
 757 Ming Zhou, et al. K-adapter: Infusing knowledge into pre-trained models with adapters. *arXiv*
 758 *preprint arXiv:2002.01808*, 2020b.

759

760 Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. Medclip: Contrastive learning from
 761 unpaired medical images and text. In *Proceedings of the Conference on Empirical Methods in*
 762 *Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*,
 763 volume 2022, pp. 3876, 2022.

764 Junde Wu, Ziyue Wang, Mingxuan Hong, Wei Ji, Huazhu Fu, Yanwu Xu, Min Xu, and Yueming
 765 Jin. Medical sam adapter: Adapting segment anything model for medical image segmentation.
 766 *Medical Image Analysis*, pp. 103547, 2025.

767 Lingling Xu et al. Parameter-efficient fine-tuning methods for pretrained language models: A critical
 768 review and assessment. *arXiv preprint arXiv:2312.12148*, 2023.

769

770 Taojannan Yang, Yi Zhu, Yusheng Xie, Aston Zhang, Chen Chen, and Mu Li. Aim: Adapting image
 771 models for efficient video action recognition. *arXiv preprint arXiv:2302.03024*, 2023.

772

773 Yiwen Ye, Yutong Xie, Jianpeng Zhang, Ziyang Chen, Qi Wu, and Yong Xia. Continual self-
 774 supervised learning: Towards universal multi-modal medical data representation learning. In
 775 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 776 11114–11124, 2024.

777

778 Dongshuo Yin, Leiyi Hu, Bin Li, Youqun Zhang, and Xue Yang. 5%*i*, 100%: Breaking performance
 shackles of full fine-tuning on visual recognition tasks. *arXiv preprint arXiv:2408.08345*, 2024.

779

780 Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. Bitfit: Simple parameter-efficient fine-tuning
 781 for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*, 2021.

782

783 Renrui Zhang, Rongyao Fang, Wei Zhang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and
 784 Hongsheng Li. Tip-adapter: Training-free clip-adapter for better vision-language modeling. *arXiv*
 785 *preprint arXiv:2111.03930*, 2021.

786

787 Shaoting Zhang and Dimitris Metaxas. On the challenges and perspectives of foundation models for
 788 medical image analysis. *Medical image analysis*, 91:102996, 2024.

789

790 Yichi Zhang, Zhenrong Shen, and Rushi Jiao. Segment anything model for medical image segmenta-
 791 tion: Current applications and future directions. *Computers in Biology and Medicine*, pp. 108238,
 792 2024.

793

794 Yingtao Zhang, Min Xian, Heng-Da Cheng, Bryar Shareef, Jianrui Ding, Fei Xu, Kuan Huang,
 795 Boyu Zhang, Chunping Ning, and Ying Wang. Busis: a benchmark for breast ultrasound image
 796 segmentation. In *Healthcare*, volume 10, pp. 729. MDPI, 2022.

797

798 Amy Zhao, Guha Balakrishnan, Fredo Durand, John V Guttag, and Adrian V Dalca. Data augmenta-
 799 tion using learned transformations for one-shot medical image segmentation. In *Proceedings of*
 800 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8543–8553, 2019.

801

802 Fudan Zheng, Jindong Cao, Weijiang Yu, Zhiguang Chen, Nong Xiao, and Yutong Lu. Exploring
 803 low-resource medical image classification with weakly supervised prompt learning. *Pattern*
 804 *Recognition*, 149:110250, 2024.

805

806 Yi Zhou, Xiaodong He, Lei Huang, Li Liu, Fan Zhu, Shanshan Cui, and Ling Shao. Collaborative
 807 learning of semi-supervised segmentation and classification for medical images. In *Proceedings of*
 808 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2079–2088, 2019.

809

810 Yaoming Zhu, Jiangtao Feng, Chengqi Zhao, Mingxuan Wang, and Lei Li. Counter-interference
 811 adapter for multilingual machine translation. *arXiv preprint arXiv:2104.08154*, 2021.

812

813 Ziquan Zhu, Si-Yuan Lu, Tianjin Huang, Lu Liu, and Zhe Liu. Lka: Large kernel adapter for
 814 enhanced medical image classification. In *International Conference on Medical Image Computing*
 815 *and Computer-Assisted Intervention*, pp. 394–404. Springer, 2025.

810	APPENDIX	
811		
812	A The Use of Large Language Models (LLMs)	17
813		
814		
815	B Experiment Settings	17
816	B.1 Datasets	17
817	B.2 Implementation Details	17
818		
819		
820	C Additional Segmentation Results on Natural-pretrained Models	18
821	C.1 Segmentation Results with Additional Metrics	18
822	C.2 Segmentation Results with More Data Scales	18
823		
824		
825	D Additional Classification Results on Natural-pretrained Models	20
826	D.1 Classification Results with Additional Metrics	20
827	D.2 Classification Results with More Data Scales	21
828		
829		
830	E Additional Results on Medical-pretrained Models	22
831		
832	F Additional Ablations	23
833	F.1 DKA in Different Backbone	23
834	F.2 Performance of Dual-Kernel Convolution and Asynchronous Learning Rates	23
835	F.3 DKA Position	24
836	F.4 ERF of Adapter in Different Backbone	24
837	F.5 Effect of Dilated vs. Standard Kernels	25
838	F.6 Effect of Diverse Kernel Combinations	26
839	F.7 Effect under Extreme Low-data Regimes	27
840	F.8 Effect on Medical-pretrained Vision-Language Models.	28
841	F.9 Frequency-domain Analysis of Kernel Sizes	28
842	F.10 Inference Latency and Memory Usage	28
843	F.11 Parameter Efficiency Analysis	29
844	F.12 Complementary CAM Visualization	29
845	F.13 Training Loss Analysis	29
846		
847		
848		
849		
850		
851		
852	G Pseudocode	31
853		
854		
855		
856		
857		
858		
859		
860		
861		
862		
863		

864 **A THE USE OF LARGE LANGUAGE MODELS (LLMs)**
865866 Large language models (LLMs) were not used as part of the core methodology, experiments, or
867 original research contributions. They were only used as an assistive tool for language editing and
868 formatting.
869870 **B EXPERIMENT SETTINGS**
871872 **B.1 DATASETS**
873874 **B.1.1 CLASSIFICATION DATASETS**
875876 **COVID** (Chowdhury et al., 2020): The COVID-19 Radiography Database, developed through
877 a collaborative effort involving researchers from Qatar University, the University of Dhaka, and
878 medical experts from Pakistan and Malaysia, includes approximately 3616 COVID-19 cases, 10,192
879 normal cases, 6012 lung opacity cases (representing non-COVID lung infections), and 1345 viral
880 pneumonia cases. This dataset provides a comprehensive collection of chest X-ray (CXR) images for
881 the diagnosis of COVID-19 and other lung conditions.
882883 **BUSI** (Zhang et al., 2022): The BUSI (Breast Ultrasound Images) dataset, collected in 2018,
884 comprises approximately 780 ultrasound images from 600 female patients aged 25 to 75. It includes
885 around 437 normal, 210 benign, and 133 malignant breast lesion images, each with an average
886 resolution of approximately 500×500 pixels. The dataset is organized into three primary categories
887 based on the clinical classification of breast lesions: normal, benign, and malignant.
888889 **ISIC-2019** (Gessert et al., 2020): The ISIC-2019 dataset, part of the annual ISIC (International Skin
890 Imaging Collaboration) challenges, includes 25,331 dermoscopic images compiled from previous
891 ISIC challenges (2017 and 2018). It spans nine diagnostic categories: melanoma, melanocytic nevus,
892 basal cell carcinoma, actinic keratosis, benign keratosis (including solar lentigo, seborrheic keratosis,
893 and lichen planus-like keratosis), dermatofibroma, vascular lesion, squamous cell carcinoma, and a
category for images that do not fit into any of the other classes, providing a comprehensive benchmark
for multi-class skin lesion classification.
894895 **B.1.2 SEGMENTATION DATASETS**
896897 **BRATS** (Menze et al., 2014): The BRATS (Brain Tumor Segmentation) dataset provides multi-
898 institutional, clinically acquired multi-parametric MRI (mpMRI) scans of gliomas, including T1,
899 post-contrast T1-weighted (T1Gd), T2, and T2-FLAIR volumes. It includes pathologically confirmed
900 cases with expert-annotated tumor sub-regions, including the enhancing tumor (ET), tumor core (TC),
and whole tumor (WT), providing a comprehensive benchmark for brain tumor segmentation.
901902 **BUSI** (Zhang et al., 2022): The BUSI (Breast Ultrasound Images) dataset includes ground truth
903 masks for precise lesion segmentation, facilitating the evaluation of automated lesion detection and
904 boundary delineation methods. These masks correspond to the original ultrasound images, capturing
905 the exact regions of interest for each lesion type. The dataset primarily supports the segmentation of
benign and malignant breast lesions, providing a detailed representation of lesion morphology.
906907 **ISIC-2018** (Codella et al., 2019): The ISIC-2018 dataset, released as part of the ISIC 2018 Task 1
908 challenge, contains a total of 3,694 dermoscopic images, each paired with a binary segmentation mask
909 outlining the precise lesion boundaries. This dataset serves as a critical benchmark for evaluating
automated skin lesion segmentation methods.
910911 **B.2 IMPLEMENTATION DETAILS**
912913 Our experiments are conducted on NVIDIA RTX 3090 GPUs. The code is based on PyTorch (Paszke,
914 2019). Due to hardware memory constraints, we use different batch sizes for classification and
915 segmentation tasks. Specifically, classification tasks use a batch size of 8, while segmentation tasks
916 are restricted to a batch size of 1, as larger batch sizes lead to out-of-memory errors. The test set is
917 fixed at 20% of the total dataset, ensuring consistent evaluation across all experimental settings. We
employ Adam optimizer (Kingma, 2014) for both the head and DKA modules, albeit with different
918

learning rates for each. When reducing the training data sizes, we proportionally decrease samples per class to maintain balance. Even in extreme low-data scenarios, we ensure that each class retains some images to prevent complete absence of any category. The weights of the down-projection layers and the biases and weights of the up-projection layers are initialized to zero, providing a stable starting point for training.

C ADDITIONAL SEGMENTATION RESULTS ON NATURAL-PRETRAINED MODELS

C.1 SEGMENTATION RESULTS WITH ADDITIONAL METRICS

To complement the segmentation results in the main text, we provide additional evaluations in Table 5, where performance is reported in terms of Dice scores across three datasets (BRATS, BUSI, ISIC-2018) under different training ratios (0.63%, 1.25%, 100%). The results confirm that DKA consistently achieves higher Dice scores than all baseline methods across datasets and data scales, further demonstrating its robustness in capturing both local details and global context for medical image segmentation.

Table 5: Comparison of Baselines and DKA on Three Segmentation Datasets Under Varying Data Sizes. Experiments are based on the pretrained Segmenter-B. Results are reported in terms of Dice (%). The best results are highlighted in bold.

Methods	BRATS			BUSI			ISIC-2018		
	0.63%	1.25%	100%	0.63%	1.25%	100%	0.63%	1.25%	100%
Full Fine-tuning	16.94	36.59	84.45	42.11	48.83	72.94	76.75	84.81	87.37
Linear Probing	14.72	33.62	82.26	40.67	48.44	70.19	75.70	83.08	85.12
BitFit (Zaken et al., 2021)	2.40	25.06	77.69	13.31	29.18	68.90	69.46	79.41	84.46
Prompt (Jia et al., 2022)	2.42	26.41	78.44	16.84	31.61	69.76	72.25	80.52	84.80
LoRA (Hu et al., 2022)	7.40	28.86	81.29	24.72	36.59	70.09	73.97	81.68	84.95
Adapter (Houlsby et al., 2019)	11.60	31.86	83.73	30.77	41.14	70.98	74.83	84.26	86.82
Adapterformer (Chen et al., 2022)	11.30	31.61	84.08	29.66	40.86	71.08	74.74	84.29	86.73
Convpass (Jie et al., 2024)	13.31	32.83	84.60	33.11	44.39	71.87	75.14	84.77	87.35
CIAT (Zhu et al., 2021)	7.33	28.72	82.64	23.63	33.62	70.77	73.63	82.02	85.77
AIM (Yang et al., 2023)	8.75	30.49	83.71	26.69	38.11	70.79	74.41	83.75	86.78
DKA	17.29	37.43	85.69	42.34	51.33	74.13	77.39	85.24	87.97

C.2 SEGMENTATION RESULTS WITH MORE DATA SCALES

To further validate the performance of DKA in segmentation tasks, we present additional results across a wider range of data scales (12.5% to 75%) in Table 6. These results, covering three segmentation medical imaging datasets, provide a comprehensive assessment of our model’s performance in mid-scale training regimes, complementing the extreme low-data (0.63% and 1.25%) and full-data (100%) findings discussed in Section 4.1). Consistent with our previous observations, DKA outperforms all baselines across all datasets and training set ratios, achieving the highest mIoU and Dice metrics. This demonstrates the effectiveness of DKA, highlighting its ability to effectively capture both local detail and global context across varying supervision levels, providing a reliable solution for medical image segmentation.

972
 973 **Table 6: Additional Segmentation Results on Three Datasets with Varying Training Set Ratios.**
 974 Experiments are based on the pretrained Segmenter-B. Results are reported as percentages for mIoU
 975 and Dice. The best results are highlighted in bold.

976 Methods	977 BRATS				978 BUSI				979 ISIC-2018			
	980 12.5%	981 25%	982 50%	983 75%	984 12.5%	985 25%	986 50%	987 75%	988 12.5%	989 25%	990 50%	991 75%
988 mIoU (%)												
989 Full Fine-tuning	990 55.28	991 61.35	992 70.43	993 71.50	994 51.44	995 53.70	996 55.35	997 56.8	998 74.60	999 75.36	1000 75.50	1001 77.10
1002 Linear Probing	1003 53.22	1004 59.43	1005 67.75	1006 68.11	1007 48.93	1008 50.76	1009 51.97	1010 53.41	1011 71.65	1012 71.99	1013 72.18	1014 73.63
1015 BitFit (Zaken et al., 2021)	1016 47.21	1017 53.46	1018 60.30	1019 61.59	1020 36.30	1021 41.07	1022 42.45	1023 48.09	1024 65.27	1025 66.36	1026 68.00	1027 71.55
1028 Prompt (Jia et al., 2022)	1029 48.38	1030 54.08	1031 62.73	1032 63.01	1033 38.49	1034 43.38	1035 46.91	1036 50.30	1037 68.54	1038 69.16	1039 69.63	1040 72.31
1041 LoRA (Hu et al., 2022)	1042 50.77	1043 57.19	1044 64.80	1045 67.75	1046 41.22	1047 47.72	1048 50.42	1049 52.97	1050 69.64	1051 70.83	1052 71.61	1053 72.95
1054 Adapter (Houlsby et al., 2019)	1055 52.52	1056 59.90	1057 69.22	1058 70.65	1059 45.25	1060 49.66	1061 52.11	1062 54.19	1063 73.54	1064 73.87	1065 74.48	1066 76.29
1067 Adapterformer (Chen et al., 2022)	1068 52.12	1069 59.36	1070 69.11	1071 70.37	1072 45.22	1073 49.25	1074 52.40	1075 54.44	1076 73.45	1077 73.79	1078 74.15	1079 75.66
1080 Conpass (Jie et al., 2024)	1081 54.08	1082 61.08	1083 69.67	1084 71.26	1085 49.30	1086 51.77	1087 53.61	1088 55.27	1089 74.31	1090 74.98	1091 75.13	1092 76.94
1093 CIAT (Zhu et al., 2021)	1094 51.70	1095 58.23	1096 67.59	1097 68.45	1098 40.74	1099 47.72	1100 51.71	1101 53.85	1102 70.83	1103 71.42	1104 72.34	1105 74.57
1106 AIM (Yang et al., 2023)	1107 51.09	1108 58.43	1109 68.11	1110 69.34	1111 43.44	1112 48.15	1113 51.20	1114 53.96	1115 72.71	1116 73.36	1117 73.55	1118 75.30
1119 DKA	1120 56.54	1121 62.92	1122 71.95	1123 72.92	1124 52.30	1125 54.76	1126 56.43	1127 58.08	1128 75.29	1129 75.83	1130 76.02	1131 78.06
1132 Dice (%)												
1133 Full Fine-tuning	1134 71.20	1135 76.05	1136 82.65	1137 83.38	1138 67.93	1139 69.88	1140 71.26	1141 72.45	1142 85.46	1143 85.95	1144 86.04	1145 87.07
1146 Linear Probing	1147 69.47	1148 73.33	1149 80.77	1150 81.03	1151 65.71	1152 67.34	1153 68.39	1154 69.63	1155 83.48	1156 83.71	1157 83.85	1158 84.81
1159 BitFit (Zaken et al., 2021)	1160 64.86	1161 69.43	1162 76.86	1163 76.14	1164 53.57	1165 58.29	1166 61.45	1167 66.42	1168 78.20	1169 79.21	1170 80.13	1171 81.78
1172 Prompt (Jia et al., 2022)	1173 65.21	1174 70.20	1175 77.10	1176 77.31	1177 55.59	1178 60.51	1179 63.86	1180 67.07	1181 81.33	1182 81.77	1183 82.10	1184 83.93
1185 LoRA (Hu et al., 2022)	1186 67.35	1187 72.77	1188 78.64	1189 80.77	1190 58.37	1191 64.61	1192 67.04	1193 69.26	1194 82.10	1195 82.92	1196 83.45	1197 84.36
1198 Adapter (Houlsby et al., 2019)	1199 68.87	1200 74.92	1201 81.81	1202 82.80	1203 62.30	1204 66.36	1205 68.52	1206 70.29	1207 84.75	1208 84.97	1209 85.37	1210 86.55
1211 Adapterformer (Chen et al., 2022)	1212 68.53	1213 74.49	1214 81.73	1215 82.61	1216 62.28	1217 66.00	1218 68.76	1219 70.50	1220 84.69	1221 84.92	1222 85.16	1223 86.15
1224 Conpass (Jie et al., 2024)	1225 70.20	1226 75.84	1227 82.13	1228 83.22	1229 66.04	1230 68.22	1231 69.80	1232 71.19	1233 85.26	1234 85.70	1235 85.80	1236 86.97
1237 CIAT (Zhu et al., 2021)	1238 68.17	1239 73.60	1240 80.66	1241 81.27	1242 57.89	1243 64.61	1244 68.17	1245 70.01	1246 82.92	1247 83.33	1248 83.95	1249 85.43
1250 AIM (Yang et al., 2023)	1251 67.63	1252 73.76	1253 81.03	1254 81.90	1255 60.56	1256 65.00	1257 67.73	1258 70.10	1259 84.20	1260 84.63	1261 84.76	1262 85.91
1263 DKA	1264 72.24	1265 77.24	1266 83.69	1267 84.34	1268 68.68	1269 70.77	1270 72.15	1271 73.48	1272 85.91	1273 86.25	1274 86.38	1275 87.68

994
 995
 996
 997
 998
 999
 1000
 1001
 1002
 1003
 1004
 1005
 1006
 1007
 1008
 1009
 1010
 1011
 1012
 1013
 1014
 1015
 1016
 1017
 1018
 1019
 1020
 1021
 1022
 1023
 1024
 1025

1026 **D ADDITIONAL CLASSIFICATION RESULTS ON NATURAL-PRETRAINED**
 1027 **MODELS**
 1028

1029 **D.1 CLASSIFICATION RESULTS WITH ADDITIONAL METRICS**
 1030

1031 In addition to the accuracy results reported in the Section 4.1, we further provide a comprehensive
 1032 evaluation of DKA and baselines using F1 and sensitivity (SEN) scores on three classification datasets:
 1033 COVID, BUSI, and ISIC-2019. As summarized in Table 7, DKA consistently achieves the best
 1034 or near-best performance across all data scales (0.63%, 1.25%, and 100%). In particular, under
 1035 low-data regimes, DKA exhibits notable improvements over existing PEFT approaches such as BitFit,
 1036 LoRA, and Adapter-based methods, highlighting its robustness in capturing discriminative features
 1037 when supervision is scarce. These results further validate the effectiveness of DKA beyond standard
 1038 accuracy and confirm its ability to improve both predictive balance (F1) and clinical reliability (SEN)
 1039 in medical image classification.

1040
 1041 **Table 7: Comparison of Baselines and DKA on Three Classification Datasets Across Varying**
 1042 **Data Sizes.** Experiments are based on the pretrained ViT-B. SEN is reported as percentages, while
 1043 F1 is presented as raw value. The best results are highlighted in bold.

Methods	COVID			BUSI			ISIC-2019		
	0.63%	1.25%	100%	0.63%	1.25%	100%	0.63%	1.25%	100%
F1									
Full Fine-tuning	0.845	0.855	0.970	0.697	0.752	0.939	0.293	0.317	0.539
Linear Probing	0.831	0.843	0.932	0.710	0.756	0.897	0.287	0.307	0.442
BitFit (Zaken et al., 2021)	0.718	0.762	0.950	0.542	0.573	0.879	0.184	0.235	0.501
Prompt (Jia et al., 2022)	0.753	0.819	0.972	0.588	0.618	0.930	0.218	0.235	0.534
LoRA (Hu et al., 2022)	0.786	0.814	0.973	0.603	0.641	0.945	0.214	0.235	0.528
Adapter (Houlsby et al., 2019)	0.802	0.833	0.978	0.607	0.704	0.931	0.231	0.248	0.501
Adapterformer (Chen et al., 2022)	0.793	0.807	0.974	0.614	0.693	0.918	0.228	0.231	0.495
Convpass (Jie et al., 2024)	0.807	0.833	0.969	0.622	0.718	0.938	0.253	0.275	0.512
CIAT (Zhu et al., 2021)	0.748	0.794	0.947	0.574	0.635	0.899	0.162	0.195	0.429
AIM (Yang et al., 2023)	0.776	0.804	0.953	0.602	0.682	0.907	0.182	0.227	0.433
DKA	0.865	0.881	0.981	0.720	0.774	0.951	0.296	0.326	0.542
SEN (%)									
Full Fine-tuning	83.57	84.15	97.24	69.27	74.26	93.46	27.76	30.64	52.35
Linear Probing	82.64	83.49	92.74	70.25	74.43	87.74	26.87	28.68	42.71
BitFit (Zaken et al., 2021)	70.36	75.34	94.45	51.33	54.41	85.52	18.50	21.74	47.55
Prompt (Jia et al., 2022)	74.31	81.76	96.57	56.71	59.35	91.34	20.65	23.92	49.28
LoRA (Hu et al., 2022)	77.42	80.76	96.76	59.37	61.41	92.07	20.11	24.13	50.75
Adapter (Houlsby et al., 2019)	81.38	82.40	97.41	59.40	66.69	92.58	21.84	24.93	49.50
Adapterformer (Chen et al., 2022)	79.34	81.51	97.40	60.96	65.44	90.72	20.36	22.78	48.31
Convpass (Jie et al., 2024)	81.23	82.71	95.76	61.79	67.90	92.74	24.63	27.36	51.12
CIAT (Zhu et al., 2021)	74.93	78.34	94.35	54.52	61.57	87.85	16.24	20.17	42.18
AIM (Yang et al., 2023)	76.04	80.76	94.50	58.75	65.46	89.45	18.06	21.07	47.16
DKA	84.41	87.76	98.12	71.82	76.78	95.46	28.20	31.50	53.69

1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

1080
1081

D.2 CLASSIFICATION RESULTS WITH MORE DATA SCALES

1082
1083
1084
1085
1086
1087
1088
1089

To further substantiate the advantages of DKA identified in the main text, we present additional classification results across a wider range of data scales (12.5% to 75%) in Table 8. These results reinforce the key findings, demonstrating that DKA consistently outperforms other baselines across all datasets (COVID, BUSI, ISIC-2019). This superior performance holds not only in extreme low-data regimes (0.63% and 1.25%) and full-data regimes (100%) but also across mid-scale settings, confirming that DKA maintains its advantage regardless of data availability. This performance reflects the effectiveness of its dual-branch design, which simultaneously captures local detail and broader context, providing a critical edge in medical image classification.

1090

1091
1092
1093

Table 8: Additional Classification Results on Three Datasets with Varying Training Set Ratios. Experiments are based on the pretrained ViT-B. ACC and SEN are reported as percentages, while F1 is presented as raw value. The best results are highlighted in bold.

1094
1095
1096
1097
1098
1099
1100
1101
1102
1103

Methods	COVID				BUSI				ISIC-2019			
	12.5%	25%	50%	75%	12.5%	25%	50%	75%	12.5%	25%	50%	75%
ACC (%)												
Full Fine-tuning	95.90	97.10	97.92	98.28	85.42	90.01	91.32	93.00	69.81	72.70	76.57	79.06
Linear Probing	92.84	93.43	94.16	94.34	79.23	84.03	86.54	87.22	63.66	65.75	66.60	68.28
BitFit (Zaken et al., 2021)	83.02	90.51	93.57	93.50	69.01	75.00	84.12	90.98	60.96	63.17	66.33	72.10
Prompt (Jia et al., 2022)	87.32	94.06	96.38	97.53	72.52	80.40	87.91	91.15	62.46	67.24	69.67	75.09
LoRA (Hu et al., 2022)	88.50	95.88	97.46	98.43	77.45	85.50	90.77	93.25	62.62	68.15	73.47	77.56
Adapter (Houlsby et al., 2019)	94.32	96.04	97.14	97.67	80.51	86.18	89.78	90.73	65.54	68.51	70.55	73.41
Adapterformer (Chen et al., 2022)	90.00	92.19	94.67	97.39	78.00	85.19	87.75	89.37	63.94	69.45	69.73	71.88
Conpass (Jie et al., 2024)	94.89	96.48	97.32	97.73	81.79	86.89	89.65	90.94	66.96	69.56	71.27	74.25
CIAT (Zhu et al., 2021)	88.98	92.73	94.23	95.45	72.47	79.32	83.56	87.55	59.84	65.16	67.55	70.25
AIM (Yang et al., 2023)	89.38	93.88	95.76	96.17	76.61	82.00	85.63	86.92	65.62	69.24	70.82	73.58
DKA	96.86	97.34	98.29	98.55	87.26	91.10	93.73	95.01	70.47	74.04	77.42	80.06
F1												
Full Fine-tuning	0.930	0.951	0.956	0.968	0.833	0.874	0.902	0.925	0.402	0.435	0.481	0.502
Linear Probing	0.898	0.916	0.922	0.929	0.782	0.828	0.856	0.861	0.335	0.351	0.379	0.392
BitFit (Zaken et al., 2021)	0.812	0.860	0.903	0.917	0.636	0.743	0.821	0.875	0.276	0.344	0.363	0.411
Prompt (Jia et al., 2022)	0.848	0.921	0.938	0.953	0.683	0.795	0.862	0.905	0.319	0.384	0.397	0.446
LoRA (Hu et al., 2022)	0.866	0.934	0.956	0.965	0.747	0.847	0.886	0.932	0.331	0.387	0.439	0.462
Adapter (Houlsby et al., 2019)	0.920	0.946	0.958	0.967	0.783	0.842	0.871	0.897	0.354	0.382	0.423	0.457
Adapterformer (Chen et al., 2022)	0.883	0.905	0.920	0.953	0.767	0.826	0.856	0.873	0.340	0.408	0.416	0.427
Conpass (Jie et al., 2024)	0.923	0.941	0.958	0.962	0.794	0.840	0.878	0.907	0.372	0.409	0.421	0.458
CIAT (Zhu et al., 2021)	0.843	0.905	0.911	0.925	0.708	0.788	0.813	0.844	0.293	0.364	0.375	0.417
AIM (Yang et al., 2023)	0.869	0.917	0.928	0.944	0.731	0.813	0.839	0.842	0.334	0.404	0.415	0.425
DKA	0.949	0.958	0.970	0.978	0.856	0.891	0.931	0.943	0.407	0.456	0.488	0.514
SEN (%)												
Full Fine-tuning	93.24	95.16	95.52	96.68	82.24	87.69	89.92	91.05	38.84	40.01	41.56	49.36
Linear Probing	88.72	90.13	90.85	92.06	76.50	81.25	84.31	85.28	32.67	34.81	36.12	38.21
BitFit (Zaken et al., 2021)	79.33	85.27	89.38	91.15	63.85	72.62	80.77	86.31	28.23	35.85	41.62	44.38
Prompt (Jia et al., 2022)	83.56	91.52	93.24	94.95	68.27	77.25	85.47	90.54	31.54	37.63	43.10	46.75
LoRA (Hu et al., 2022)	85.15	92.90	95.65	96.22	71.12	81.34	87.93	90.23	31.88	37.31	42.61	46.90
Adapter (Houlsby et al., 2019)	91.93	93.65	95.36	96.29	76.53	85.78	87.39	88.03	34.65	37.31	41.16	44.12
Adapterformer (Chen et al., 2022)	87.35	89.80	91.04	95.54	73.68	82.74	85.79	87.16	32.52	39.10	39.75	41.56
Conpass (Jie et al., 2024)	92.00	94.51	94.92	95.36	76.78	80.73	85.03	88.21	35.36	39.84	41.58	43.99
CIAT (Zhu et al., 2021)	84.52	89.49	91.72	93.61	69.64	76.49	80.83	84.21	21.79	24.31	36.25	36.50
AIM (Yang et al., 2023)	86.51	90.22	92.53	93.27	73.86	79.80	83.77	84.27	34.23	37.80	40.65	42.62
DKA	95.24	96.27	97.18	97.59	83.88	89.62	93.46	93.89	39.72	44.07	47.50	50.39

1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

1134 **E ADDITIONAL RESULTS ON MEDICAL-PRETRAINED MODELS**
1135

1136 To provide more comprehensive evidence beyond the main paper, we further report detailed results
 1137 using medical-pretrained models. Table 9 presents classification performance on the ISIC-2019
 1138 dataset with RadImageNet-pretrained ResNet-50 (Mei et al., 2022), where DKA consistently improves
 1139 over linear probing and standard adapter tuning across different training ratios in terms of ACC,
 1140 F1, and SEN. Table 10 reports segmentation results on the BUSI dataset with MedSAM (Ma et al.,
 1141 2024), showing that DKA achieves clear gains over both baselines in terms of mIoU and Dice. These
 1142 extended results corroborate our main findings in Section 4.5, demonstrating that the advantages of
 1143 DKA generalize robustly to medical-pretrained models and are not restricted to natural-pretrained
 1144 backbones.

1145
 1146 **Table 9: Additional Classification Results based on the RadImageNet-pretrained ResNet-50**
 1147 **backbone on the ISIC-2019 dataset under varying training ratios.** SEN is reported as percentages,
 1148 while F1 is presented as raw value. The best results are highlighted in bold.

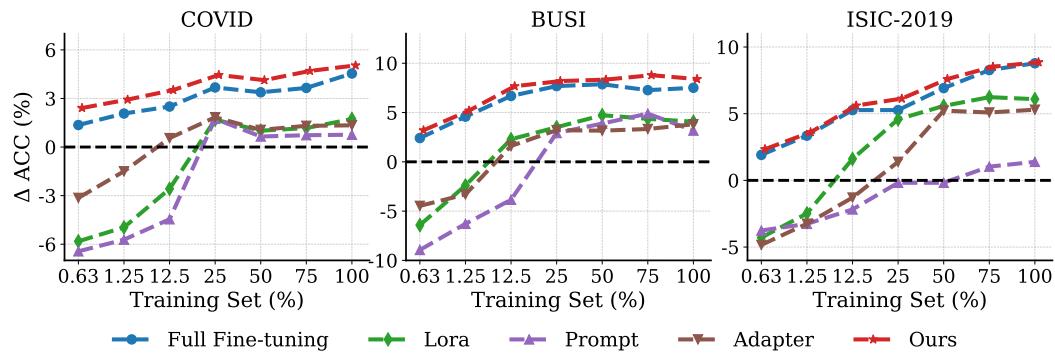
Methods	0.63%	1.25%	100%
F1			
Full Fine-tuning	0.223	0.262	0.483
Linear Probing	0.216	0.236	0.387
BitFit	0.172	0.208	0.411
Prompt	0.200	0.232	0.439
LoRA	0.203	0.230	0.431
Adapter	0.218	0.242	0.440
DKA	0.230	0.278	0.498
SEN (%)			
Full Fine-tuning	22.57	26.86	43.14
Linear Probing	20.28	24.52	37.61
BitFit	17.69	22.86	40.20
Prompt	19.83	24.42	44.05
LoRA	19.79	25.77	44.38
Adapter	20.34	25.13	44.57
DKA	23.40	27.89	45.39

1165
 1166 **Table 10: Additional Segmentation Results based on MedSAM backbone on the BUSI segmenta-**
 1167 **tion task under varying training ratios.** Results are reported in terms of Dice (%). The best results
 1168 are highlighted in bold.

Methods	0.63%	1.25%	100%
Dice (%)			
Full Fine-tuning	54.54	63.08	81.15
Linear Probing	50.09	60.90	79.68
BitFit	23.46	47.31	74.15
Prompt	25.19	50.65	78.82
LoRA	32.58	55.03	80.42
Adapter	52.28	61.81	80.35
DKA	55.39	65.24	83.97

1188 **F ADDITIONAL ABLATIONS**
11891190 **F.1 DKA IN DIFFERENT BACKBONE**
1191

1192 To evaluate whether the advantages of DKA transfer to other architectures, we repeat classification
1193 experiments on the pretrained Swin-B using the same datasets: COVID, BUSI, and ISIC-2019. As
1194 shown in Figure 9, DKA consistently outperforms full fine-tuning and all PEFT baselines across
1195 different data scales. Critically, while other PEFT methods struggle to match linear probing, and
1196 often degrade substantially in constrained-data regimes, DKA maintains strong performance in both
1197 settings.



1211 **Figure 9: Performance of Baselines and DKA across Various Training Data Sizes.** $\Delta ACC =$
1212 $ACC_{Baselines} - ACC_{LinearProbing}$. Experiments are based on the pretrained Swin-B for COVID,
1213 BUSI, and ISIC-2019.

1214
1215 **F.2 PERFORMANCE OF DUAL-KERNEL CONVOLUTION AND ASYNCHRONOUS LEARNING
1216 RATES**
1217

1218 To further evaluate the impact of our proposed enhancements, including the introduction of dual-
1219 convolution design and a learning rate split strategy within the DKA module, we present additional
1220 results in Table 11. This table compares DKA against two baselines (Adapter + Dual-Conv and Adapter
1221 + LR Split) across multiple datasets (COVID, BUSI, ISIC-2019) and training set ratios. Consistently,
1222 DKA outperforms both baselines, achieving the highest overall accuracy in each dataset, with notable
1223 gains observed even in low-data regimes. This confirms the effectiveness of our combined strategy in
1224 capturing richer local-global features and improving learning stability.

1225
1226 **Table 11: Performance Comparison of Enhanced Adapter Designs.** Experiments are conducted
1227 on the pretrained ViT-B across three medical imaging classification datasets and varying training set
1228 ratios by reporting ACC (%).

Datasets	Methods	Training Set Size						
		0.63%	1.25%	12.5%	25%	50%	75%	100%
COVID	Adapter + Dual-Conv	86.78	88.89	96.01	96.98	98.13	97.95	98.61
	Adapter + LR Split	88.04	88.95	96.20	97.04	97.64	98.07	98.79
	Ours	89.01	91.06	96.86	97.34	98.29	98.55	99.21
BUSI	Adapter + Dual-Conv	67.73	77.32	83.39	89.14	90.73	91.05	94.25
	Adapter + LR Split	70.37	77.64	84.71	89.46	91.10	93.37	94.93
	Ours	74.23	79.64	87.26	91.10	93.73	95.01	95.89
ISIC-2019	Adapter + Dual-Conv	59.49	60.51	70.12	73.27	76.60	78.52	82.25
	Adapter + LR Split	59.06	60.09	69.99	72.24	76.07	78.52	81.67
	Ours	60.52	62.32	70.47	74.04	77.42	80.06	83.09

1242 F.3 DKA POSITION
1243

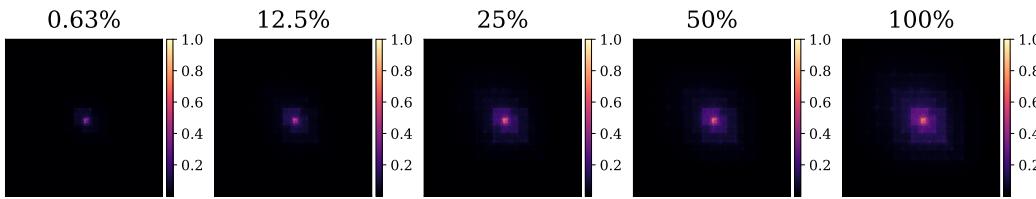
1244 To understand how the position of inserted DKA influences performance, we explore three placement
1245 strategies using the pretrained ViT-B with 12 transformer blocks: inserting DKA into the bottom
1246 4 blocks (Blocks 0–3), middle 4 blocks (Blocks 4–7), and top 4 blocks (Blocks 8–11). We also
1247 include a reference setting where DKA is inserted into all layers. We extend our analysis to three
1248 medical imaging classification datasets: COVID, BUSI, and ISIC-2019. As shown in Table 12, the
1249 position of DKA significantly affects model performance across various training set sizes. Among
1250 partial configurations, placing DKA in the middle layers consistently outperforms the top and bottom
1251 placements across datasets. Notably, under low-data regimes (e.g., 0.63%), placing DKA in the
1252 top layers offers stronger performance than bottom or middle placement, highlighting the value
1253 of adapting higher-level representations when supervision is scarce. This trend is consistent with
1254 observations from prior work (Yang et al., 2023), which reported that inserting adapters in bottom
1255 blocks yields limited performance.

1256
1257 Table 12: **Effect of Position.** Experiments are conducted on the pretrained ViT-B across three medical
1258 imaging classification datasets and varying training set ratios by reporting ACC (%) with standard
1259 deviations.

Datasets	Positions	Training Set Size						
		0.63%	1.25%	12.5%	25%	50%	75%	100%
COVID	Bottom	86.04 ± 1.61	88.68 ± 1.38	93.69 ± 0.89	94.29 ± 0.75	95.52 ± 0.60	96.64 ± 0.45	97.34 ± 0.35
	Middle	86.80 ± 1.45	89.18 ± 1.25	95.43 ± 0.53	96.80 ± 0.41	97.76 ± 0.29	98.43 ± 0.25	98.79 ± 0.21
	Top	87.78 ± 1.21	90.50 ± 1.13	94.72 ± 0.69	96.12 ± 0.55	97.22 ± 0.38	98.23 ± 0.32	98.53 ± 0.26
	ALL	89.01 ± 0.90	91.06 ± 0.84	96.86 ± 0.40	97.34 ± 0.32	98.29 ± 0.26	98.55 ± 0.24	99.21 ± 0.13
BUSI	Bottom	71.47 ± 2.06	77.16 ± 1.47	84.86 ± 1.11	88.00 ± 0.65	91.19 ± 0.43	92.79 ± 0.33	93.94 ± 0.27
	Middle	73.48 ± 1.79	78.22 ± 1.28	86.62 ± 0.75	90.43 ± 0.44	93.43 ± 0.25	94.94 ± 0.19	95.46 ± 0.11
	Top	73.86 ± 1.64	78.16 ± 1.22	86.48 ± 0.84	90.08 ± 0.45	93.07 ± 0.28	94.62 ± 0.20	95.00 ± 0.13
	ALL	74.23 ± 1.53	79.64 ± 1.17	87.26 ± 0.64	91.10 ± 0.39	93.73 ± 0.23	95.01 ± 0.12	95.89 ± 0.09
ISIC-2019	Bottom	56.87 ± 2.47	58.68 ± 2.11	67.28 ± 1.52	70.37 ± 1.26	73.43 ± 1.05	76.16 ± 0.73	79.52 ± 0.46
	Middle	57.64 ± 2.33	60.01 ± 2.05	69.88 ± 1.34	73.95 ± 1.06	76.74 ± 0.74	79.43 ± 0.43	82.59 ± 0.20
	Top	59.89 ± 2.17	61.79 ± 1.95	69.32 ± 1.38	72.59 ± 1.18	76.04 ± 0.79	78.42 ± 0.52	81.90 ± 0.25
	ALL	60.52 ± 2.02	62.32 ± 1.85	70.47 ± 1.26	74.04 ± 0.99	77.42 ± 0.67	80.06 ± 0.36	83.09 ± 0.14

1270
1271 F.4 ERF OF ADAPTER IN DIFFERENT BACKBONE
1272

1273 To further validate our findings on the impact of data scarcity on Adapter performance, we extend the
1274 effective receptive field (ERF) analysis to the pretrained Swin-T model (Liu et al., 2021), as shown in
1275 Figure 10. This complementary analysis reinforces our observations from the pretrained ViT-B model
1276 (Figure 2), revealing a similar contraction of ERFs as the training set size decreases. Specifically,
1277 under extreme data scarcity (e.g., 0.63% and 1.25% training data), the pretrained Swin-T exhibits
1278 a sharply reduced ERF, aligning with the earlier findings (see Section 2) that Adapters can disrupt
1279 pretrained feature representations under severe data constraints. This result suggests that the negative
1280 impacts observed in the main text are not limited to a single architecture but are likely a more general
1281 phenomenon affecting a wide range of vision backbones.



1290 Figure 10: **Effective Receptive Field of Adapter under Different Training Set Ratios.** Experiments
1291 are conducted on the COVID dataset using the pretrained Swin-T.

1296
1297

F.5 EFFECT OF DILATED VS. STANDARD KERNELS

1298
1299
1300
1301
1302

We further compare dilated convolutional kernels with our standard large-kernel design. In particular, we evaluate several dilated kernel settings, including 3×3 with dilation rates of $d = 3, 25, 11 \times 11$ with $d = 5$, and 26×26 with $d = 2$, against the dual-kernel configuration of $5 \times 5 + 51 \times 51$. The classification results on the BUSI dataset are shown in Table 13, and the segmentation results on the ISIC-2018 dataset are reported in Table 14.

1303
1304
1305
1306
1307
1308

Across both datasets, the $5 \times 5 + 51 \times 51$ design consistently achieves the best performance in terms of ACC, F1, and SEN for classification, as well as mIoU and Dice for segmentation. Although dilated kernels provide a larger effective receptive field, they underperform compared to explicitly using a large standard kernel. This suggests that our dual-kernel design captures both local and global information more effectively than dilated alternatives, highlighting the efficiency of standard large kernels in the DKA framework.

1309
1310
1311
1312
1313

Table 13: **Comparison of different kernel designs on the BUSI dataset for classification under varying training ratios.** Experiments are based on the pretrained ViT-B. ACC and SEN are reported as percentages, while F1 is presented as raw value. The best results are highlighted in bold. “ d ” represents the dilation rate in dilated convolution.

Kernel Designs	0.63%	1.25%	100%
ACC (%)			
$3 \times 3 (d=3) + 3 \times 3 (d=25)$	66.77	77.32	93.61
$3 \times 3 (d=3) + 11 \times 11 (d=5)$	67.41	77.64	94.25
$3 \times 3 (d=3) + 26 \times 26 (d=2)$	68.69	77.96	94.37
$5 \times 5 + 51 \times 51$	74.23	79.46	95.89
F1			
$3 \times 3 (d=3) + 3 \times 3 (d=25)$	0.641	0.751	0.931
$3 \times 3 (d=3) + 11 \times 11 (d=5)$	0.650	0.756	0.936
$3 \times 3 (d=3) + 26 \times 26 (d=2)$	0.663	0.758	0.942
$5 \times 5 + 51 \times 51$	0.720	0.774	0.951
SEN (%)			
$3 \times 3 (d=3) + 3 \times 3 (d=25)$	62.81	74.35	92.36
$3 \times 3 (d=3) + 11 \times 11 (d=5)$	63.99	74.43	93.44
$3 \times 3 (d=3) + 26 \times 26 (d=2)$	64.72	74.71	94.03
$5 \times 5 + 51 \times 51$	71.82	76.78	95.46

1320
1321
1322
1323
1324

Table 14: **Comparison of different kernel designs on the ISIC-2018 dataset for segmentation under varying training ratios.** Experiments are based on the pretrained Segmenter-B. mIoU and Dice are reported as percentages. The best results are highlighted in bold. “ d ” represents the dilation rate in dilated convolution.

Kernel Designs	0.63%	1.25%	100%
mIoU (%)			
$3 \times 3 (d=3) + 3 \times 3 (d=25)$	60.25	72.71	76.43
$3 \times 3 (d=3) + 11 \times 11 (d=5)$	61.46	73.11	76.97
$3 \times 3 (d=3) + 26 \times 26 (d=2)$	62.27	73.43	77.48
$5 \times 5 + 51 \times 51$	63.13	74.27	78.53
Dice (%)			
$3 \times 3 (d=3) + 3 \times 3 (d=25)$	75.19	84.20	86.64
$3 \times 3 (d=3) + 11 \times 11 (d=5)$	76.13	84.46	86.99
$3 \times 3 (d=3) + 26 \times 26 (d=2)$	76.75	84.68	87.31
$5 \times 5 + 51 \times 51$	77.39	85.24	87.97

1346
1347
1348
1349

1350
1351

F.6 EFFECT OF DIVERSE KERNEL COMBINATIONS

1352
1353
1354
1355
1356

To further analyze the contribution of different kernel configurations in DKA, we evaluate multiple kernel combinations for both classification and segmentation tasks. Specifically, we compare three settings: (i) $5 \times 5 + 11 \times 11 + 51 \times 51$, (ii) $5 \times 5 + 31 \times 31 + 51 \times 51$, and (iii) $5 \times 5 + 51 \times 51$. Table 15 reports results on the BUSI classification dataset, while Table 16 presents results on the ISIC-2018 segmentation dataset.

1357
1358
1359
1360
1361
1362

Across both tasks, we observe that the dual-kernel configuration ($5 \times 5 + 51 \times 51$) consistently achieves the best trade-off, outperforming the three-branch alternatives in terms of ACC, F1, and SEN for classification, as well as mIoU and Dice for segmentation. These results indicate that adding intermediate kernels (e.g., 11×11 or 31×31) does not provide additional benefits, and a simpler dual-kernel design is sufficient to capture both local and global dependencies. This further validates the efficiency of our proposed kernel selection strategy within DKA.

1363
1364
1365
1366
1367

Table 15: **Comparison of diverse kernel combinations on the BUSI dataset for classification under varying training ratios.** Experiments are based on the pretrained ViT-B. ACC and SEN are reported as percentages, while F1 is presented as raw value. The best results are highlighted in bold.

Kernel Combinations	0.63%	1.25%	100%
ACC (%)			
$5 \times 5 + 11 \times 11 + 51 \times 51$	72.68	77.00	94.89
$5 \times 5 + 31 \times 31 + 51 \times 51$	71.04	76.68	94.57
$5 \times 5 + 51 \times 51$	74.23	79.46	95.89
F1			
$5 \times 5 + 11 \times 11 + 51 \times 51$	0.709	0.743	0.942
$5 \times 5 + 31 \times 31 + 51 \times 51$	0.694	0.738	0.936
$5 \times 5 + 51 \times 51$	0.720	0.774	0.951
SEN (%)			
$5 \times 5 + 11 \times 11 + 51 \times 51$	69.17	73.56	94.11
$5 \times 5 + 31 \times 31 + 51 \times 51$	68.39	72.85	93.62
$5 \times 5 + 51 \times 51$	71.82	76.78	95.46

1380
13811382
1383
1384
1385

Table 16: **Comparison of diverse kernel combinations on the ISIC-2018 dataset for segmentation under varying training ratios.** Experiments are based on the pretrained Segmenteer-B. mIoU and Dice are reported as percentages. The best results are highlighted in bold.

Kernel Combinations	0.63%	1.25%	100%
mIoU (%)			
$5 \times 5 + 11 \times 11 + 51 \times 51$	62.72	73.76	78.06
$5 \times 5 + 31 \times 31 + 51 \times 51$	62.40	73.63	77.54
$5 \times 5 + 51 \times 51$	63.13	74.27	78.53
Dice (%)			
$5 \times 5 + 11 \times 11 + 51 \times 51$	77.09	84.90	87.68
$5 \times 5 + 31 \times 31 + 51 \times 51$	76.84	84.81	87.35
$5 \times 5 + 51 \times 51$	77.39	85.24	87.97

1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

1404
1405 F.7 EFFECT UNDER EXTREME LOW-DATA REGIMES1406 To further investigate the limits of DKA, we conduct experiments under an extreme low-data regime
1407 with only 0.125% of the training set available. This setup approximately corresponds to a 5-shot
1408 setting for COVID classification and genuine 1-shot settings for BUSI and ISIC-2019 (classification),
1409 as well as BRATS, BUSI, and ISIC-2018 (segmentation). The results are reported in Table 17 and
1410 Table 18.1411 Across all datasets, DKA consistently outperforms both Linear Probing and standard Adapter, even
1412 under such highly constrained supervision. For example, on ISIC-2019 classification, DKA improves
1413 ACC from 52.54% (Linear Probing) to 55.95%, and F1 from 0.227 to 0.263. Similarly, on BUSI
1414 segmentation, DKA achieves a Dice of 37.43%, substantially higher than Linear Probing (23.02%) and
1415 Adapter (24.93%). These results confirm that the proposed large-kernel design remains robust and
1416 effective even in the most challenging few-shot scenarios, highlighting its practicality for data-scarce
1417 medical applications.1418
1419 Table 17: **Comparison of Linear Probing, Adapter, and DKA on three classification datasets**
1420 **under 0.125% training data.** Experiments are based on the pretrained ViT-B. ACC and SEN are
1421 reported as percentages, while F1 is presented as raw value. The best results are highlighted in bold.

Methods	COVID	BUSI	ISIC-2019
ACC (%)			
Linear Probing	76.51	37.06	52.54
Adapter	71.38	35.46	47.25
DKA	78.74	39.62	55.95
F1			
Linear Probing	0.735	0.273	0.227
Adapter	0.704	0.247	0.153
DKA	0.787	0.340	0.263
SEN (%)			
Linear Probing	72.53	33.87	21.58
Adapter	68.40	30.41	15.07
DKA	75.64	37.13	24.39

1435
1436
1437 Table 18: **Comparison of Linear Probing, Adapter, and DKA on three segmentation datasets**
1438 **under 0.125% training data.** Experiments are based on the pretrained Segmenter-B. mIoU and
1439 Dice are reported as percentages. The best results are highlighted in bold.

Methods	BRATS	BUSI	ISIC-2018
mIoU (%)			
Linear Probing	3.25	19.64	53.21
Adapter	1.20	14.24	47.05
DKA	6.36	32.83	60.25
Dice (%)			
Linear Probing	6.94	23.02	69.46
Adapter	2.40	24.93	63.99
DKA	11.96	37.43	75.19

1458 F.8 EFFECT ON MEDICAL-PRETRAINED VISION-LANGUAGE MODELS.
1459

1460 To further validate the generalization ability of DKA, we extend our evaluation to medical vision-
 1461 language models. Specifically, we adopt MedCLIP (Wang et al., 2022) as the backbone and follow
 1462 the few-shot image classification protocol commonly used in prior works on vision-language adapta-
 1463 tion (Shakeri et al., 2024; Silva-Rodriguez et al., 2024). We report results on the BUSI dataset under
 1464 1-shot, 4-shot, and 8-shot settings, where each configuration is averaged over five random seeds. As
 1465 summarized in Table 19, DKA consistently surpasses linear probing and standard adapter tuning across
 1466 all support sizes. These results demonstrate that the benefits of DKA are not limited to vision-only
 1467 large pretrained models, but also extend to multimodal vision-language models, highlighting its
 1468 robustness in broader medical AI scenarios.

1469
1470 **Table 19: Comparison of different methods under few-shot settings on the BUSI dataset based**
1471 **on MedCLIP.** Results are reported in terms of ACC, F1, and SEN.

Tuning Strategies	1-shot	4-shot	8-shot
ACC (%)			
Linear Probing	66.25	74.21	78.41
Adapter	65.86	74.08	79.32
DKA	71.20	79.82	82.27
F1			
Linear Probing	0.631	0.712	0.776
Adapter	0.625	0.706	0.782
DKA	0.697	0.762	0.815
SEN (%)			
Linear Probing	62.55	69.72	77.45
Adapter	61.87	69.46	78.96
DKA	68.17	75.85	81.87

1485
1486 F.9 FREQUENCY-DOMAIN ANALYSIS OF KERNEL SIZES
1487

1488 To better understand the complementary roles of small and large kernels in DKA, we analyze their
 1489 frequency responses using the radial power spectral density (PSD). Specifically, we compute the
 1490 spectral centroid f_c and the normalized frequency f_{90} that captures 90% of the cumulative energy.
 1491 Table 20 shows the results for 5×5 and 51×51 kernels. The smaller kernel exhibits a higher spectral
 1492 centroid ($f_c = 0.618$ vs. 0.503), indicating stronger sensitivity to high-frequency details. In contrast,
 1493 the larger kernel shifts energy towards lower frequencies, facilitating global context modeling. This
 1494 analysis provides further evidence for the effectiveness of combining diverse kernel sizes in DKA.
 1495

1496
1497 **Table 20: Frequency-domain analysis of kernel sizes.** Results are reported as mean \pm std over different
1498 trained models.

Kernel Size	f_c	f_{90}
5×5	0.618 ± 0.018	0.924 ± 0.011
51×51	0.503 ± 0.004	0.904 ± 0.003

1503
1504 F.10 INFERENCE LATENCY AND MEMORY USAGE
1505

1506 We also report the inference efficiency of different methods in terms of latency and memory consump-
 1507 tion on the BUSI dataset using the pretrained ViT-B. As summarized in Table 21, DKA introduces only
 1508 marginal overhead compared to the standard adapter framework, with inference latency increasing by
 1509 less than 0.5 ms and memory usage by less than 6 MB. These differences are negligible in practice,
 1510 indicating that the proposed large-kernel design achieves substantial performance gains with minimal
 1511 computational cost.

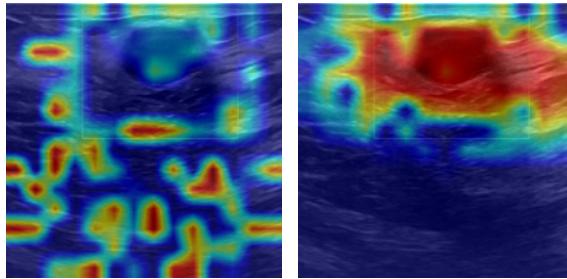


Figure 11: **Grad-CAM comparison on BUSI.** Left: standard Adapter; right: DKA with more focused, lesion-related activations.

Table 21: **Comparison of inference latency and memory usage on the BUSI dataset based on the pretrained ViT-B.** DKA introduces only negligible overhead compared to standard adapter variants.

Methods	Inference Latency (ms)	Memory (MB)
Linear Probing	6.78	352.81
Adapter	11.54	433.70
Adapter + 5×5 Conv	11.66	433.90
Adapter + 51×51 Conv	11.69	439.56
DKA	11.97	439.63

F.11 PARAMETER EFFICIENCY ANALYSIS

The DKA module integrates two depthwise convolution branches with kernel sizes k_1 and k_2 within each adapter. These convolutions are applied independently on each of the \hat{d} channels, contributing $\hat{d}(k_1^2 + k_2^2)$ parameters per module. Since DKA is inserted twice in every Transformer block (after the attention and feedforward layers), the total number of additional parameters grows linearly with the number of blocks. Even with two kernels (e.g., $k_1 = 51$, $k_2 = 5$), the total number of trainable parameters introduced by all DKA modules remains less than 2% of the pretrained backbone, maintaining strong parameter efficiency while providing strong performance gains, especially in low-data regimes.

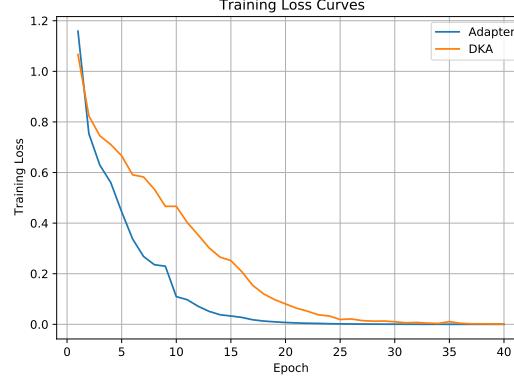
F.12 COMPLEMENTARY CAM VISUALIZATION

To complement the ERF analysis, we additionally compare Grad-CAM (Selvaraju et al., 2017) responses of the standard Adapter and our DKA on BUSI. As shown in Figure F.11, the standard Adapter exhibits diffuse and background-driven activations, whereas DKA produces coherent, lesion-centered responses. This confirms that the enlarged ERF of DKA captures meaningful contextual information rather than merely expanding activation range.

F.13 TRAINING LOSS ANALYSIS

To further assess whether DKA introduces any optimization difficulties, we provide the training loss curves of the standard Adapter and our DKA on the COVID classification task using the pretrained ViT-B. As shown in Figure F.11, both methods exhibit smooth and stable convergence across all epochs without oscillation, divergence, or plateaus, indicating that neither method suffers from optimization instability. Adapter exhibits a faster early loss drop due to its lighter structure, whereas DKA converges more gradually because of its large-kernel branch. Both ultimately reach near-zero loss, indicating that DKA does not introduce optimization difficulties and its gains are not due to training dynamics..

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586 Training Loss Curves
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599 Figure 12: Training loss curves of Adapter and DKA on the COVID dataset using the pretrained
1600 ViT-B.
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619



1620 G PSEUDOCODE
16211622

1623 **Algorithm 1** Pseudo-code of a Transformer block with DKA

```

1624 class DKA:
1625     def __init__(self, dim, middle_dim, kernel_large, kernel_small):
1626         self.downsample = Linear(dim, middle_dim)
1627         self.conv_large = DepthwiseConv(middle_dim, kernel_large)
1628         self.conv_small = DepthwiseConv(middle_dim, kernel_small)
1629         self.activation = GELU()
1630         self.upsample = Linear(middle_dim, dim)
1631
1632     def forward(self, x):
1633         # Store the input for the residual connection
1634         residual = x
1635         x = self.downsample(x)
1636
1637         # Dual-Path Convolutions (Large + Small)
1638         x_large = self.conv_large(x)
1639         x_small = self.conv_small(x)
1640         x = x_large + x_small
1641
1642         x = self.activation(x)
1643         x = self.upsample(x)
1644         x = x + residual
1645
1646     return x
1647
1648 class TransformerBlock_with_DKA:
1649     def __init__(self, dim, num_heads, mlp_ratio, middle_dim, kernel_large, kernel_small):
1650         # Original ViT components
1651         self.attn = MultiheadAttention(dim, num_heads)
1652         self.norm1 = LayerNorm(dim)
1653         self.mlp = MLP(dim, mlp_ratio)
1654         self.norm2 = LayerNorm(dim)
1655
1656         # DKA Adapter
1657         self.dka = DKA(dim, middle_dim, kernel_large, kernel_small)
1658
1659     def forward(self, x):
1660         residual = x
1661         x = self.norm1(x)
1662         x = self.attn(x)
1663         x = x + residual
1664
1665         x = self.dka(x)
1666
1667         residual = x
1668         x = self.norm2(x)
1669         x = self.mlp(x)
1670         x = x + residual
1671
1672         x = self.dka(x)
1673
1674     return x

```

1666
1667
1668
1669
1670
1671
1672
1673