



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

Black-Box Adaptation for Medical Image Segmentation

Jay N. Paranjape¹✉, Shameema Sikder^{2,3}, S. Swaroop Vedula³, and Vishal M. Patel¹

¹ Department of Electrical and Computer Engineering, The Johns Hopkins University, Baltimore, USA

jparanj1@jhu.edu

² Wilmer Eye Institute, The Johns Hopkins University, Baltimore, USA

³ Malone Center for Engineering in Healthcare, The Johns Hopkins University, Baltimore, USA

Abstract. In recent years, various large foundation models have been proposed for image segmentation. These models are often trained on large amounts of data corresponding to general computer vision tasks. Hence, these models do not perform well on medical data. There have been some attempts in the literature to perform parameter-efficient finetuning of such foundation models for medical image segmentation. However, these approaches assume that all the parameters of the model are available for adaptation. But, in many cases, these models are released as APIs or Black-Boxes, with no or limited access to the model parameters and data. In addition, finetuning methods also require a significant amount of compute, which may not be available for the downstream task. At the same time, medical data can't be shared with third-party agents for finetuning due to privacy reasons. To tackle these challenges, we pioneer a Black-Box adaptation technique for prompted medical image segmentation, called BAPS. BAPS has two components - (i) An Image-Prompt decoder (IP decoder) module that generates visual prompts given an image and a prompt, and (ii) A Zero Order Optimization (ZOO) Method, called SPSA-GC that is used to update the IP decoder without the need for backpropagating through the foundation model. Thus, our method does not require any knowledge about the foundation model's weights or gradients. We test BAPS on four different modalities and show that our method can improve the original model's performance by around 4%. The code is available at <https://github.com/JayParanjape/Blackbox>.

Keywords: Black-Box Adaptation · Prompted Segmentation.

1 Introduction

Image segmentation is a fundamental problem in medical image analysis tasks. Many deep learning-based approaches have been proposed in the literature that segment out various regions of interest across different medical modalities [24]. In recent years, many foundation models have been proposed, which are large models with billions of parameters that show excellent performance on downstream

tasks like classification [18] and segmentation [13,16,28]. Following the success of these foundation models, various adaptation approaches have been proposed that transfer the knowledge learnt by these large-scale models to medical segmentation [14]. These adaptation methods are commonly termed as Parameter Efficient Finetuning (PEFT) methods and they aim to utilize a given foundation model while tuning only a fraction of its parameters. However, all PEFT methods make two over-optimistic assumptions. First, they assume that all parameters of the foundation model are available during training. However, it often occurs that companies release their AI models as APIs or Black-Boxes instead of releasing the entire parameter set, training dataset, or codebase due to proprietary concerns. In such conditions, all the PEFT methods would fail since they require gradient computation to work. Secondly, many PEFT methods assume the availability of high compute resources [25,27], which is not realistic. At the same time, the transfer of medical data to a third-party system with more resources would result in privacy concerns. In this work, we tackle these concerns by pioneering a Black-Box adaptation method for medical image segmentation called Black-Box Adapter for Prompted Segmentation (BAPS). BAPS uses a frozen image encoder and a trainable Image-Prompt Decoder (IP Decoder), to produce an input-dependent per-pixel prompt. This is added to the original image and provided to the foundation model Black-Box. We train the IP Decoder using a recently proposed ZOO method called Simultaneous Perturbation Stochastic Approximation with Gradient Correction (SPSA-GC) [17] which does not require computation of any gradient, thus, not needing any parameter information about the foundation model. The IP Decoder is a lightweight module and hence, our method requires minimal compute resources. In summary, our contributions can be listed as follows:

- 1) To the best of our knowledge, this is the first paper to explore Black-Box adaptation for medical image segmentation. For this, we propose BAPS which has a frozen pre-trained encoder and a lightweight decoder module that can be trained using derivative-free optimization methods. A visual comparison of our method with existing approaches can be seen in Figure 1.
- 2) We show that BAPS can improve the original foundation model’s performance on four widely used public datasets of different modalities, including endoscopic images, dermoscopic images, gastrointestinal polyp images, and retinal images. For this purpose, we conduct experiments using two recently proposed popular image segmentation foundation models - SAM [13] and MedSAM [16].

2 Related Work

The motivation for using Black-Box adaptation arose from the Natural Language Processing (NLP) field. Many custom AI models available in the NLP industry are deployed as a proprietary service or API [26]. Hence, several works attempt to adapt them for customized datasets using ZOO algorithms as the model weights and gradients are not available [22,21,4]. BBT [22] and BBTv2 [21] use a ZOO method called Covariance Matrix Adaptation Evolution Strategy

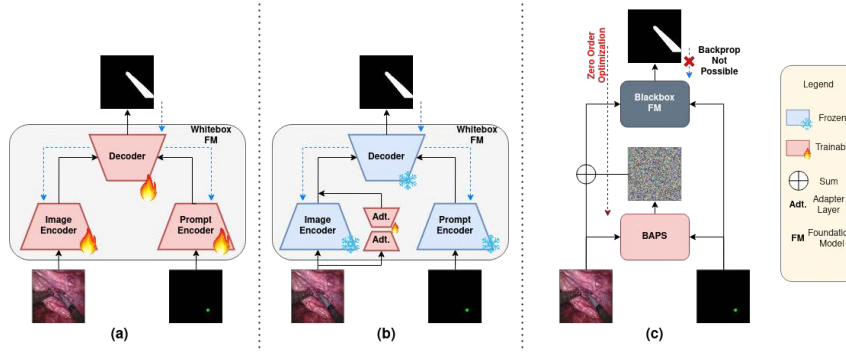


Fig. 1: Comparison of our method (c) against full finetuning of a Foundation Model (FM), (a) and common adaptation techniques (b).

(CMA-ES) [7,8], while RLPrompt [4] uses a reinforcement learning gradient-free based method to tune prompts, which are then added to the input to the foundation model. However, reinforcement learning and evolutionary strategy-based methods tend to show high variance and are unstable [15]. Further, they do not work as well for vision tasks [17]. There are very few methods that perform Black-Box adaptation for vision tasks [23,17]. BAR [23] uses a one-sided approximation for gradient estimation. However, this was found to be inaccurate empirically by a succeeding approach called BlackVIP [17]. BlackVIP uses a two-sided approximation for gradient estimation and proposes a different ZOO method for tuning weights. However, both BAR and BlackVIP are proposed only for the classification task with CLIP [18] as the Black-Box foundation model. In this work, we pioneer Black-Box adaptation for the recently proposed task of prompted segmentation.

3 Proposed Method

Model Architecture: Foundation models for segmentation usually perform the task of prompted segmentation. Given an input image and a prompt, the foundation model produces a segmentation mask that corresponds to the given prompt. Hence, for BAPS, we consider an image and a point prompt as the inputs. The overview of BAPS is shown in Figure 2. The input image is passed through a pre-trained image encoder that produces image embeddings. We use the Vision Transformer (ViT) encoder with Masked Autoencoder (MAE) pre-training [9] as our image encoder because of its strong innate understanding of images. Hence, it generates highly representative features. On the other hand, the point prompt is converted into positional embeddings based on its relative position in the image. This uses a sinusoid function for generating the embeddings, similar to ViT [5].

The image and prompt embeddings are then concatenated and passed to a module called the Image-Prompt Decoder (IP-Decoder). This module is a

deconvolution network that generates a visual prompt as its output. This is then added to the original image and passed to the Black-Box foundation model along with the original point prompt. Thus, the job of the IP-Decoder is to learn a residual visual prompt that when added to the original image, will make it easier for the foundation model to segment out the correct shape. Please note that the IP-Decoder is the only trainable module in BAPS.

Once the Black-Box model generates the prediction, it is compared with the label using a sum of the Binary Cross Entropy (BCE) loss and dice loss [20]. However, note that the gradients cannot be backpropagated since the weights in the Black-Box model are not available. Hence, we use a ZOO method called SPSA-GC [17], which we describe next.

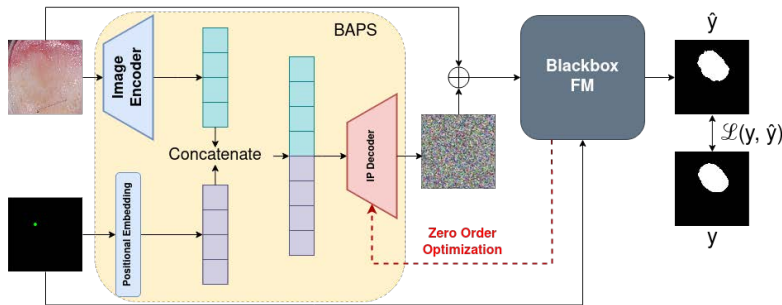


Fig. 2: Overview of BAPS.

SPSA-GC: Simultaneous Perturbation Stochastic Approximation (SPSA) [19] is a Zero Order Optimization (ZOO) technique that estimates gradients effectively using a two-sided approximation of the derivative. Given a set of parameters ϕ and a loss function \mathbb{L} , the estimated gradients \hat{g} are given as follows:

$$\hat{g}_i = \frac{\mathbb{L}(\phi_i + c\Delta_i) - \mathbb{L}(\phi_i - c\Delta_i)}{2c} \Delta_i^{-1}, \quad (1)$$

where $c \in [0, 1]$ is a hyperparameter and Δ represents a random perturbation vector of the same shape as ϕ . Each element of Δ is sampled uniformly from $[-1, -0.5] \cup [0.5, 1]$. Here, i represents the iteration number. This estimate of the gradient can be used to update the weights of the model as $\phi_{i+1} \leftarrow \phi_i - \alpha \hat{g}_i$, where α is the learning rate. While SPSA is a good approximator, it has been shown that it can lead to slower convergence [17]. Hence, in SPSA-GC a momentum term m is added while updating the weights as follows:

$$m_{i+1} = \beta m_i - \alpha \hat{g}_i(\phi_i + \beta m_i), \quad \phi_{i+1} = \phi_i + m_{i+1}, \quad (2)$$

where β denotes the weight of the momentum. This allows the IP Decoder to train quicker and more stably. However, we found that given a point prompt, the foundation model produces a reasonable mask without any sort of optimization (zero-shot performance). Thus, the approximated gradients are small in magnitude and the system gets stuck at local minimas more often. To alleviate this, as

an implementation detail, we add a mechanism to detect the local minima and increase the learning rate for some iterations during training to help the model get out of the local minima. This is detailed in the supplementary document.

4 Experiments and Results

We choose the recently proposed Segment Anything Model (SAM) [13] and MedSAM [16] as the Black-Box foundation models for our experiments. SAM and MedSAM are foundation models for the task of promptable segmentation. Given an image and a prompt in the form of point, text, mask or bounding box, they can segment out the object of interest corresponding to the prompt. SAM is trained on a large corpus of natural images while MedSAM is tuned on medical data including 3D and 2D images. Majority of the training data of MedSAM includes CT and MRI images. We use the point-based prompt for our experiments. Thus, for a given image and a point prompt, we pass it through BAPS and generate a visual prompt. This is added to the original image and sent to the Black-Box foundation model along with the point prompt. We evaluate our method on four different datasets and calculate the Dice Score (DSC) and Hausdorff Distance at the 95th percentile (HD95). For our baselines, we consider the zero-shot performance of the foundation model without any adaptation, and Visual Prompting (VP) [12] with SPSA-GC. We measure the upper bound of the segmentation performance using a white box adaptation of SAM and MedSAM using LoRA [10], similar to various SAM-based adaptation methods [25,27].

Datasets. We use four widely used publicly available datasets for our experiments. Kvasir-Seg [11] consists of gastrointestinal polyp images, divided into 600 training, 100 validation, and 300 test images. ISIC2018 [3] contains dermoscopic images for skin lesion segmentation. It is divided into 2594 images for training, 100 for validation, and 100 for testing. The third dataset is REFUGE [6], which has retinal images for optic disk and optic cup segmentation. It is divided into 800 training, 800 validation, and 800 testing images. The fourth dataset is Endovis 17 [1], which has images of endoscopic surgery. It is further divided into 2878 testing and 3231 training images, out of which we use 366 for validation. The dataloaders will be made available along with the code after the review process.

Experimental Setup. For all the datasets, we use the same set of data augmentations. These include random rotation up to 10 degrees, random brightness changes with scale 2, and random saturation changes with scale 2. The images are scaled to the resolution 512×512 . Based on the validation set performance of these datasets, we set hyperparameters of the model as: $c = 0.01$, $\alpha = 0.005$ for Endovis17 and $\alpha = 0.01$ for other datasets, $\beta = 0.9$. For all datasets, we use a batch size of 32. Training is done on a single Nvidia RTX A5000 GPU and uses only 6GB of memory. Note that in our case, we also run the Black-Box on the same GPU. However, in practice, training would be cheaper since the forward

propagation in the Black-Box will occur at the server. The Black-Boxes SAM and MedSAM both are initialized using their ViT-base checkpoints.

Table 1: Our approach improves over the zero-shot performance of the Black-Box model SAM (row 1). Having the IP Decoder improves performance over direct visual prompt tuning (row 2). The performance of our method is upper bounded by the white box adaptation methods like LoRA (row 4).

<i>Foundation Model - SAM</i>								
Method	Kvasir-Seg		ISIC2018		REFUGE		Endovis 17	
	DSC (\uparrow)	HD95 (\downarrow)	DSC (\uparrow)	HD95 (\downarrow)	DSC (\uparrow)	HD95 (\downarrow)	DSC (\uparrow)	HD95 (\downarrow)
SAM (ZS) [13]	0.68	99.3	0.66	85.98	0.38	214.98	0.60	89.46
VP [12]	0.69	99.3	0.70	70.96	0.39	201.67	0.63	83.77
BAPS (Our Approach)	0.72	83.55	0.74	70.3	0.44	176.67	0.65	81.56
Whitebox LoRA [10]	<i>0.88</i>	<i>22.23</i>	<i>0.85</i>	<i>30.95</i>	<i>0.85</i>	<i>19.22</i>	<i>0.68</i>	<i>45.39</i>

Results. We tabulate the quantitative results on each of the datasets in Table 1 for SAM and in Table 2 for MedSAM. For all four datasets, we see an average of 5% improvement in the Dice Score with our method over the foundation model SAM and an average of 7% over MedSAM. For all the results, evaluation is done five times with randomly selected point prompts and the mean value is listed in the table. The standard deviation in each case is less than 0.01. The increase in performance can be attributed to the strong pre-trained encoder and IP Decoder of BAPS, which generates a visual prompt as a function of the input image and the point prompt. Some samples of the modified images after adding the visual prompt are shown in Figure 4. In row 2, we compare our method with simply adding the same visual prompt for each image-point pair (no encoder or I-P Decoder). We see significant improvement, showing the effectiveness of the encoder-decoder structure. This can also be seen in supplementary Figure 1, where we plot the training progress of BAPS in comparison to VP, with MedSAM as the Black-Box for ISIC2018 and REFUGE. For both these cases, we see that the average error for BAPS decreases consistently with the number of iterations. All results with BAPS have a p-value of at most 10^{-8} . Qualitative results are shown in Figure 3.

Table 2: Our approach improves over the zero-shot performance of the Black-Box model MedSAM (row 1). Having the IP Decoder improves performance over direct visual prompt tuning (row 2). The performance of our method is upper bounded by the white box adaptation methods like LoRA (row 4).

<i>Foundation Model - MedSAM</i>								
Method	Kvasir-Seg		ISIC2018		REFUGE		Endovis 17	
	DSC (\uparrow)	HD95 (\downarrow)	DSC (\uparrow)	HD95 (\downarrow)	DSC (\uparrow)	HD95 (\downarrow)	DSC (\uparrow)	HD95 (\downarrow)
MedSAM (ZS) [16]	0.68	98.4	0.70	87.23	0.36	214.29	0.63	87.52
VP [12]	0.68	90.36	0.71	99.2	0.36	214.48	0.63	85.58
BAPS (Our Approach)	0.72	80.01	0.79	66.6	0.44	168.07	0.65	82.4
Whitebox LoRA [10]	<i>0.90</i>	<i>20.05</i>	<i>0.86</i>	<i>30.1</i>	<i>0.85</i>	<i>19.15</i>	<i>0.68</i>	<i>42.92</i>




Fig. 3: Qualitative Results on all the datasets. GT - ground truth, VP - visual prompting [12]. The green dot in the image denotes the point prompt given to the Black-Box foundation model

5 Ablation Analysis

Ablation Over Pretrained Image Encoder. We use the MAE pre-trained ViT image encoder from Meta (ViT-MAE) [9] since it gives a stronger featurization of images over other encoders. To verify this claim, we make an ablation analysis by changing the pre-trained encoder and measuring the performance on the ISIC 2018 dataset. The results are shown in Table 3, where we compare the results of ViT-MAE with three popular image encoders. We find that ViT-MAE outperforms CLIP [18] and DINO-Resnet50 [2] significantly. ViT [5] is a strong encoder. However, the MAE pretraining further improves the downstream performance by 2%

Ablation Over SPSA. We empirically test the effectiveness of SPSA-GC as a ZOO method by removing each of the components of the algorithm and measuring performance on the ISIC 2018 dataset. As seen in Table 4, we start with the zero-shot performance of SAM [13], which has no required optimization. Using just SPSA [19] with zero momentum increases this performance by only 1%. This is further improved by adding a momentum of 0.9 as suggested by SPSA-GC [17], finally giving a DSC of 0.74.

Visualizing the Modified Images: To test the effectiveness of BAPS, we compare the visual results of the Black-Box foundation model MedSAM with and without adding the visual prompt generated by BAPS. These results are shown in Figure 4. Here, the zero-shot prediction of MedSAM generates inaccurate masks as seen in column 3. However, after adding the learnt visual prompt to the original image, MedSAM can correctly generate the masks. The modified image can be seen in column 4 of the figure.

Table 3: Ablation analysis of different pre-trained image encoders on ISIC 2018 dataset

Image encoder	DSC
CLIP [18]	0.65
DINO-Resnet50 [2]	0.69
ViT [5]	0.72
ViT-MAE [9] (used in BAPS)	0.74

Table 4: Ablation analysis of component-wise importance of the zero-order optimization method.

Optimization method	DSC
SAM (ZS) [13] (No ZOO)	0.66
SPSA [19]	0.68
SPSA-GC (Ours)	0.74

Memory Analysis. We perform an analysis of the number of parameters required to be trained by the client and the memory requirement of the generated checkpoint. Here, the client refers to the machine used at the downstream application level. Note that for our method, SAM is a Black-Box and does not have to be run on the client machine, resulting in significantly lower memory consumption. This is one of the advantages of the proposed Black-Box paradigm. BAPS requires just one hundred thousand parameters to be tuned, thus making the memory required to store the model checkpoint only 0.4 MB as compared to hundreds to thousands of MBs required by whitebox adaptation methods, which have at least ten times as many parameters.




Fig. 4: Visualizing the effect of adding the visual prompt on ISIC 2018 dataset. The Black-Box used for these results is MedSAM.

6 Conclusion

In this work, we proposed one of the first Black-Box adaptation methods, called BAPS, for the adaptation of foundation models for prompted segmentation. BAPS consists of a pretrained image encoder and a trainable IP decoder, that generates a visual prompt as a function of the input image and given prompt. This visual prompt is added to the original image and given to the foundation

model. The IP decoder is trained using a novel Zero Order Optimization (ZOO) method called SPSA-GC. We test BAPS on four different public datasets in medical segmentation and verify its effectiveness for the recently proposed and popular foundation models SAM and MedSAM. Finally, we carry out an ablation study to gauge the effectiveness of different design decisions of BAPS. Thus, our proposed method can efficiently perform Black-Box adaptation of SAM without the requirement of gradients.

7 Disclosure of Interest

This research was supported by a grant from the National Institutes of Health, USA; R01EY033065. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The authors have no competing interests in the paper.

References

1. Allan, M., Shvets, A., Kurmann, T., Zhang, Z., Duggal, R., Su, Y.H., Rieke, N., Laina, I., Kalavakonda, N., Bodenstedt, S., Herrera, L., Li, W., Igloukov, V., Luo, H., Yang, J., Stoyanov, D., Maier-Hein, L., Speidel, S., Azizian, M.: 2017 robotic instrument segmentation challenge (2019)
2. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. CoRR **abs/2104.14294** (2021), <https://arxiv.org/abs/2104.14294>
3. Codella, N., Rotemberg, V., Tschandl, P., Celebi, M.E., Dusza, S., Gutman, D., Helba, B., Kalloo, A., Liopyris, K., Marchetti, M., Kittler, H., Halpern, A.: Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic) (2019)
4. Deng, M., Wang, J., Hsieh, C.P., Wang, Y., Guo, H., Shu, T., Song, M., Xing, E.P., Hu, Z.: Rlprompt: Optimizing discrete text prompts with reinforcement learning (2022)
5. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=YicbFdNTTy>
6. Fu, H., Li, F., Orlando, J.I., Bogunović, H., Sun, X., Liao, J., Xu, Y., Zhang, S., Zhang, X.: Refuge: Retinal fundus glaucoma challenge (2019), <https://dx.doi.org/10.21227/tz6e-r977>
7. Hansen, N., Müller, S.D., Koumoutsakos, P.: Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es). *Evolutionary Computation* **11**(1), 1–18 (2003)
8. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation* **9**(2), 159–195 (2001). <https://doi.org/10.1162/106365601750190398>
9. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.B.: Masked autoencoders are scalable vision learners. CoRR **abs/2111.06377** (2021), <https://arxiv.org/abs/2111.06377>

10. Hu, E.J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: LoRA: Low-rank adaptation of large language models. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=nZeVKeeFYf9>
11. Jha, D., Smedsrud, P.H., Riegler, M.A., Halvorsen, P., de Lange, T., Johansen, D., Johansen, H.D.: Kvasir-seg: A segmented polyp dataset. In: International Conference on Multimedia Modeling. pp. 451–462. Springer (2020)
12. Jia, M., Tang, L., Chen, B.C., Cardie, C., Belongie, S., Hariharan, B., Lim, S.N.: Visual prompt tuning. In: Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII. p. 709–727. Springer-Verlag, Berlin, Heidelberg (2022), https://doi.org/10.1007/978-3-031-19827-4_41
13. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything (2023)
14. Lee, H.H., Gu, Y., Zhao, T., Xu, Y., Yang, J., Usuyama, N., Wong, C., Wei, M., Landman, B.A., Huo, Y., Santamaria-Pang, A., Poon, H.: Foundation models for biomedical image segmentation: A survey (2024)
15. Liu, S., Chen, P.Y., Kailkhura, B., Zhang, G., Hero III, A.O., Varshney, P.K.: A primer on zeroth-order optimization in signal processing and machine learning: Principals, recent advances, and applications. *IEEE Signal Processing Magazine* **37**(5), 43–54 (2020)
16. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images (2023)
17. Oh, C., Hwang, H., Lee, H.y., Lim, Y., Jung, G., Jung, J., Choi, H., Song, K.: Blackvip: Black-box visual prompting for robust transfer learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 24224–24235 (June 2023)
18. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision (2021)
19. Spall, J.: Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control* **37**(3), 332–341 (1992)
20. Sudre, C.H., Li, W., Vercauteren, T.K.M., Ourselin, S., Cardoso, M.J.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. Deep learning in medical image analysis and multimodal learning for clinical decision support : Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, held in conjunction with MICCAI 2017 Quebec City, QC,... **2017**, 240–248 (2017), <https://api.semanticscholar.org/CorpusID:21957663>
21. Sun, T., He, Z., Qian, H., Zhou, Y., Huang, X., Qiu, X.: BBTv2: Towards a gradient-free future with large language models. In: Goldberg, Y., Kozareva, Z., Zhang, Y. (eds.) Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing. pp. 3916–3930. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Dec 2022), <https://aclanthology.org/2022.emnlp-main.259>
22. Sun, T., Shao, Y., Qian, H., Huang, X., Qiu, X.: Black-box tuning for language-model-as-a-service. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S. (eds.) Proceedings of the 39th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 162, pp. 20841–20855. PMLR (17–23 Jul 2022), <https://proceedings.mlr.press/v162/sun22e.html>

23. Tsai, Y.Y., Chen, P.Y., Ho, T.Y.: Transfer learning without knowing: Reprogramming black-box machine learning models with scarce data and limited resources. In: III, H.D., Singh, A. (eds.) Proceedings of the 37th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 119, pp. 9614–9624. PMLR (13–18 Jul 2020), <https://proceedings.mlr.press/v119/tsai20a.html>
24. Wang, R., Lei, T., Cui, R., Zhang, B., Meng, H., Nandi, A.K.: Medical image segmentation using deep learning: A survey. IET Image Processing **16**(5), 1243–1267 (2022), <https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/ipr2.12419>
25. Wu, J., Ji, W., Liu, Y., Fu, H., Xu, M., Xu, Y., Jin, Y.: Medical sam adapter: Adapting segment anything model for medical image segmentation (2023)
26. Ye, J., Chen, X., Xu, N., Zu, C., Shao, Z., Liu, S., Cui, Y., Zhou, Z., Gong, C., Shen, Y., Zhou, J., Chen, S., Gui, T., Zhang, Q., Huang, X.: A comprehensive capability analysis of gpt-3 and gpt-3.5 series models (2023)
27. Zhang, K., Liu, D.: Customized segment anything model for medical image segmentation (2023)
28. Zou, X., Yang, J., Zhang, H., Li, F., Li, L., Wang, J., Wang, L., Gao, J., Lee, Y.J.: Segment everything everywhere all at once (2023)