



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.

# Fit Pixels, Get Labels: Meta-Learned Implicit Networks for Image Segmentation

Kushal Vyas<sup>[0009–0000–5952–697X]</sup>, Ashok Veeraraghavan<sup>[0000–0001–5043–7460]</sup>,  
and Guha Balakrishnan<sup>[0000–0001–8703–1368]</sup>

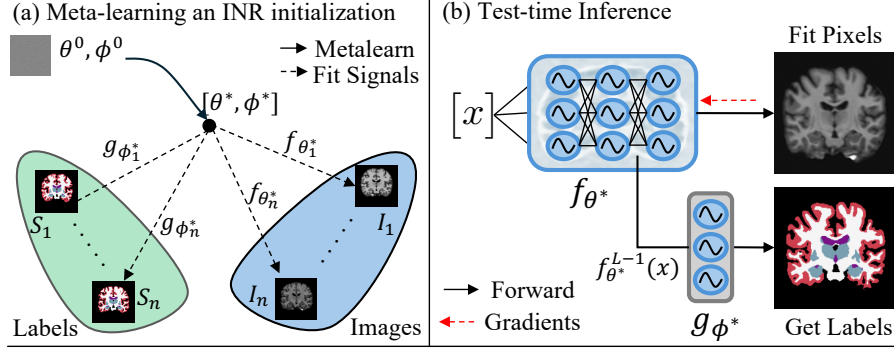
Rice University, Houston TX 77005, USA  
[guha@rice.edu](mailto:guha@rice.edu)

**Abstract.** Implicit neural representations (INRs) have achieved remarkable successes in learning expressive yet compact signal representations. However, they are not naturally amenable to predictive tasks such as segmentation, where they must learn semantic structures over a distribution of signals. In this study, we introduce MetaSeg, a meta-learning framework to train INRs for medical image segmentation. MetaSeg uses an underlying INR that simultaneously predicts per pixel intensity values and class labels. It then uses a meta-learning procedure to find optimal initial parameters for this INR over a training dataset of images and segmentation maps, such that the INR can simply be fine-tuned to fit pixels of an unseen test image, and automatically decode its class labels. We evaluated MetaSeg on 2D and 3D brain MRI segmentation tasks and report Dice scores comparable to commonly used U-Net models, but with 90% fewer parameters. MetaSeg offers a fresh, scalable alternative to traditional resource-heavy architectures such as U-Nets and vision transformers for medical image segmentation.

**Keywords:** Medical Image Segmentation, Implicit Neural Representations, Meta-Learning, MRI

## 1 Introduction

Implicit neural representations (INRs) have demonstrated impressive performance in encoding complex natural signals. A typical INR  $f_{\theta}(\cdot)$  is a multilayer perceptron (MLP) with parameters  $\theta$  that maps a coordinate  $\mathbf{x} \in \mathbb{R}^d$  to signal value  $I(\mathbf{x}) \in \mathbb{R}^D$  and is iteratively fit to a specific signal. Because they offer continuous signal representations, excellent reconstruction performance, and implicit signal priors, INRs have gained significant attention in the computer vision community for compactly modeling large signals [13, 17, 16], and in the medical imaging community for inverse imaging tasks such as accelerated MRI [1], and sparse-view CT reconstruction [18, 23]. Despite these successes, learned INR representations are highly specific to a given signal and to the way its parameters are initialized. As a result, unlike architectures such as U-Nets [15] and vision transformers [9], INR-produced features lack structural or semantic coherence, making them less used for tasks requiring learning over an image distribution.



**Fig. 1. Overview of MetaSeg.** (a) We use a meta-learning framework to learn optimal initial parameters  $\theta^*, \phi^*$  for an INR consisting of an  $L$ -layer reconstruction network  $f_{\theta}(\cdot)$  and shallow segmentation head  $g_{\phi}(\cdot)$ . (b) At test time, optimally initialized INR  $f_{\theta^*}$  is iteratively fit to the pixels of an unseen test scan. After convergence, the penultimate features  $f_{\theta^*}^{L-1}(x)$  are fed as input to the segmentation head  $g_{\phi^*}(\cdot)$  to predict per-pixel class labels.

Results from recent studies suggest that this shortcoming of INRs may be fixable. In particular, by fitting an INR across multiple images starting from the same parameter initialization, the final parameters exhibit clear semantic, structural properties for datasets such as faces or MRI scans [6, 21, 22]. And as a result, these learned initializations allow the INR to fit unseen (test) images rapidly with far fewer gradient updates while learning more generalizable features [2, 21, 22]. In this study, we leverage this insight to develop an INR fitting strategy for medical image segmentation. In particular, we present the surprising discovery that an INR carefully optimized to fit many training pairs of images and associated segmentation maps can predict a segmentation map for an unseen test image *when simply fine-tuned to reconstruct the image’s pixels*.

We propose MetaSeg, a novel meta-learning [7] INR framework that performs this strategy. MetaSeg uses an INR that simultaneously predicts both image pixels and segmentation labels given a coordinate (Fig. 1b), and aims to find an optimal setting of the INR parameters ( $\theta$  and  $\phi$ ), such that the network can be quickly fine-tuned to fit any image-mask pair starting from these parameters. MetaSeg uses a nested, meta-learning procedure (Fig. 1a). In the inner optimization, the INR is fit to reconstruct a randomly selected batch of training image and segmentation mask pairs. The outer (meta) optimization applies a gradient descent step to the current values of  $\theta$  and  $\phi$  based on the converged inner parameters. These steps iterate to converge to an initialization that works well for both image reconstruction and segmentation objectives from the given image distribution. At test time, we simply fit the INR initialized with the learned optimal parameters to the test image (without needing any segmentation information),

generating viable segmentation masks in just 2 gradient descent updates and near state-of-the-art quality masks after 100 updates.

MetaSeg draws inspiration from Neural Implicit Segmentation Functions (NISFs) [20], which also use an INR to predict pixel and segmentation values. However, instead of forcing the INR to learn image-label correlations, NISF takes an additional latent vector as input to represent an image’s semantic content, which must then be optimized at test time to produce a mask. Due to this strategy, NISF must apply several regularization losses on INR parameters and the latent space to succeed. In contrast, MetaSeg offers a much simpler training methodology based on meta-learning, only requiring reconstruction and segmentation losses during training and only a reconstruction loss during testing.

We evaluated MetaSeg on 2D and 3D magnetic resonance imaging (MRI) segmentation tasks using the OASIS-MRI [8, 12] dataset. We first find that MetaSeg achieves similar performance for 2D (Dice score: 0.93) and 3D (Dice score: 0.91) MRI segmentation tasks to widely used U-Net baselines [4, 14] with 90% fewer parameters, and significantly outperforms NISF [20]. Second, we show that MetaSeg performs well for fine-grained segmentation tasks, accurately segmenting image regions which may be sparse and span only a few pixels. Third, we show that when trained on low-resolution scans, MetaSeg fares better on high-resolution scans than U-Nets, likely due to its continuous representation. Finally, we qualitatively demonstrate through principal component analysis (PCA) visualizations that MetaSeg’s learned initial parameters encode semantic information correlated with ground truth segmentations.

## 2 Methods

We assume a given dataset  $\mathcal{D} = \{I_j, S_j\}_{j=1}^N$  of  $N$  subjects, where  $I_j$  and  $S_j$  are a  $d$ -dimensional image scan and segmentation map, respectively ( $d = 2$  and  $3$  in our experiments). At location  $x \in \mathbb{R}^d$ , we have  $I_j(x) \in \mathbb{R}^D$  and  $S_j(x) \in \{0, 1\}^C$ , where  $D$  is the number of image channels, and  $C$  is the number of classes.

We propose the MetaSeg framework to estimate  $\hat{S}$  for an unseen scan  $I$  at test time. MetaSeg consists of an INR (Fig. 1b) with an  $L$ -layer MLP  $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^D$  that predicts  $\hat{I}(\mathbf{x})$  given coordinate  $\mathbf{x}$ , and a segmentation head  $g_\phi : \mathbb{R}^h \rightarrow \mathbb{R}^C$  that predicts  $C$  class probabilities given the penultimate features, of dimensionality  $h$ , computed by  $f_\theta: f_\theta^{L-1}(\mathbf{x})$ . At test time, we aim to fit  $f_\theta(\cdot)$  on  $I$  for  $T_f$  steps, and then simply compute  $\hat{S}(\mathbf{x}) = g_\phi(f_\theta^{L-1}(\mathbf{x}))$  in one feed-forward operation.

In Sec. 2.1, we describe MetaSeg’s strategy for learning optimal values  $\theta^*, \phi^*$  over  $\mathcal{D}$ , such that  $f_{\theta^*}(\cdot)$  and  $g_{\phi^*}(\cdot)$  (networks initialized with  $\theta^*$  and  $\phi^*$ ) may be easily fine-tuned to reconstruct and segment a scan from the distribution. In Sec. 2.2, we describe how the INR initialized with these parameters may be used at inference time to estimate a segmentation map for a scan.

## 2.1 Training: Meta-Learning and Segmentation Head Optimization

MetaSeg uses a MAML[7] meta-learning strategy to learn optimal parameters  $\theta^*, \phi^*$ , consisting of a nested optimization with *inner* and *outer* (meta) routines (Fig. 1a). Starting with initial random values  $\theta^0$  and  $\phi^0$  at iteration  $t = 0$ , the outer routine takes  $T_o$  gradient descent steps to move the parameters to generalizable values. The gradient at each outer step is based on the results of the inner routine, which fits  $f_{\theta^t}(\cdot)$  and  $g_{\phi^t}(\cdot)$  for  $T_i$  steps on one data pair  $(I_j, S_j)$ . The outer routine uses the final inner parameter values as a gradient signal to update  $\theta^t, \phi^t$ . To avoid a noisy gradient due to overfitting to example  $j$ , we set  $T_i$  to a small value in practice (e.g.,  $T_i = 2$ ).

In the inner routine, at time step  $t$ , we optimize both  $f_{\theta^t}(\cdot)$  and  $g_{\phi^t}(\cdot)$  for a single subject  $j$  with a loss function:

$$\mathcal{L}_{inner}(I_j, \hat{I}_j, S_j, \hat{S}_j) = \sum_{\mathbf{x}} \mathcal{L}_{recon}(I_j(\mathbf{x}), \hat{I}_j(\mathbf{x})) + \mathcal{L}_{cls}(S_j(\mathbf{x}), \hat{S}_j(\mathbf{x})), \quad (1)$$

where  $\mathcal{L}_{recon}(\mathbf{a}, \mathbf{b}) = \|\mathbf{a} - \mathbf{b}\|_2^2$  is a per-pixel reconstruction loss, and

$$\mathcal{L}_{cls}(\mathbf{a}, \mathbf{b}) = \sum_{c=1}^C -(1 - \mathbf{b}(c))^\gamma \cdot \delta_{c,\mathbf{a}} \log(\mathbf{b}(c)) \quad (2)$$

is a per-pixel multiclass focal classification loss [11], where  $\gamma$  is a hyperparameter. We use a focal loss to account for heavy class imbalance across pixels. In the outer step, we make a gradient descent update:

$$[\theta^{t+1}, \phi^{t+1}] \leftarrow [\theta^t, \phi^t] - \beta \nabla_{[\theta^t, \phi^t]} (\mathcal{L}_{inner}(I_j, \hat{I}_j, S_j, \hat{S}_j)), \quad (3)$$

where gradient  $\nabla_{[\theta^t, \phi^t]}(\cdot)$  computes the difference between the converged parameters  $(\theta_j^t, \phi_j^t)$  of the inner optimization and current parameters  $\theta^t, \phi^t$ . After  $T_o$  outer gradient descent steps, we obtain parameters  $\theta^{T_o}$  and  $\phi^{T_o}$ .


We freeze  $\theta^* = \theta^{T_o}$ , but further optimize  $\phi^{T_o}$  such that when, at test time,  $f_{\theta^*}(\cdot)$  is fit for  $T_f$  iterations on any random scan, we can obtain accurate segmentation predictions, in one step. To do so, we first fit  $f_{\theta^*}(\cdot)$  separately on each training scan  $I_j$  for  $T_f$  iterations and populate a dataset of learned scan-specific features and segmentation maps,  $\{f_{\theta_j^*}^{L-1}(\mathbf{x}), S_j(\mathbf{x})\}_{j=1}^N$ . We then globally optimize  $g_{\phi}(\cdot)$  by minimizing the loss function  $\mathcal{L}_{seg}$ :

$$\mathcal{L}_{seg}(\mathcal{D}) = \sum_{\mathbf{x}} \sum_j \mathcal{L}_{cls}(S_j(\mathbf{x}), g_{\phi}(f_{\theta_j^*}^{L-1}(\mathbf{x}))). \quad (4)$$

We freeze the converged parameters as  $\phi^*$ .

## 2.2 Inference: Fit Pixels, Get Labels

At inference time, we assume a given (unseen) scan  $I$ . We fit  $f_{\theta^*}(\cdot)$  to  $I$  for  $T_f$  iterations, solely optimizing the pixelwise reconstruction loss  $\mathcal{L}_{recon}$ . Upon completion, we compute  $\hat{S}(\mathbf{x}) = g_{\phi^*}(f_{\theta^*}^{L-1}(\mathbf{x}))$ , and apply softmax and argmax operations to  $\hat{S}(\mathbf{x})$  to obtain the predicted class per pixel.



**Fig. 2. Example progression of reconstruction and segmentation performance as a MetaSeg INR fits the pixels of a test image.** We report PSNR (top) and Dice (bottom) scores. MetaSeg’s initialization leads to rapid convergence, with a Dice score of 0.85 emerge as early as  $T_f = 2$  iterations. It achieves optimal segmentation (Dice=0.95, PSNR=41.3) at  $T_f = 100$  iterations, and eventually declines as it overfits to the image pixels at  $T_f = 5000$  iterations.


### 3 Experiments and Results

We evaluated MetaSeg on the tasks of 2D and 3D MRI image segmentation. We used 414 T1 brain MRI scans from the OASIS-MRI [8, 12] dataset, randomly split into 214 training, 100 validation, and 100 test scans. OASIS-MRI provides 5 coarse (background, white matter, gray matter, cortex, and cerebrospinal fluid), 24 fine-grained segmentation labels for 2D slices, and 35 segmentation labels for full 3D MRI volumes. We used aligned coronal sections provided in OASIS-MRI for 2D segmentation and aligned 3D volumes for 3D segmentation. We normalized all images and volumes to have intensity in  $[0, 1]$ . We padded and resized images to resolution  $192 \times 192$ , and cropped 3D volumes to  $160 \times 160 \times 200$  followed by  $2 \times$  downsampling to yield  $80 \times 80 \times 100$  resolution scans.

**Implementation.** For all experiments, we used a SIREN [19] INR with  $L = [6, 5, 5]$  layers, width  $h = [128, 512, 256]$ , and  $w_0 = 30$ . We set the segmentation head  $g_\phi(\cdot)$  as one fully connected layer with Leaky-ReLU activation, followed by a linear layer with  $C$  output values. We used the Adam [10] optimizer with a learning rate of  $10^{-4}$  for inner and outer meta-learning optimizations. When tuning the segmentation head for 5 classes, we reduced the learning rate to  $5e - 5$ . We set  $T_i = 2$  and  $T_f = 100$ , and ran the outer loop for  $T_o = 10$  epochs over the training data while validating the performance of optimized parameters  $\theta^t, \phi^t$  every 50 steps. For 2D segmentation with 5 classes, we set the focal loss hyperparameter  $\gamma = 1.0$ , and for 24 output classes we set  $\gamma = 2.0$ . For 3D MRI volumes which exhibit high background class imbalance, we set  $\gamma = 3.0$  and further scaled the reconstruction loss by 0.1 for background pixels.

**Table 1. Quantitative performance of MetaSeg and baseline models on 2D and 3D MRI segmentation tasks.** MetaSeg achieves comparable Dice scores to popular U-Net baselines with 90% fewer parameters.

Task	Num. classes	Model	Dice Score $\uparrow$	Num. parameters $\downarrow$
2D MRI Segmentation	5	U-Net [4]	$0.96 \pm 0.008$	7.7M
		MetaSeg	$0.93 \pm 0.012$	83K
	24	U-Net [4]	$0.84 \pm 0.097$	7.7M
		MetaSeg	$0.86 \pm 0.032$	1.06M
3D MRI Segmentation	5	SegResNet [14]	$0.95 \pm 0.006$	4.7M
		NISF [20]	$0.81 \pm 0.007$	293K
		MetaSeg	$0.91 \pm 0.011$	330K

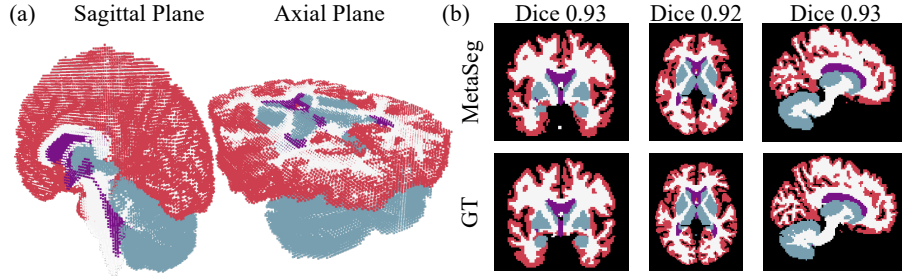


**Fig. 3. Coarse and fine-grained 2D brain MRI segmentation results with 24 classes, for two subjects.** MetaSeg accurately segments structures and is robust to high variations across subjects for structures such as ventricles (depicted in purple), brain stem (gray), and hippocampus (yellow). MetaSeg also adapts well to structures such as the cerebral cortex (red) which are not compact and localized in space.

**Baselines:** We adopt a widely used U-Net [15] baseline for 2D MRIs [4] and SegResNet [14] from the Monai [5] package for 3D MRIs. We trained all models using a learning rate of  $1e-3$  until the validation loss saturated. We also compare against NISF [20], the most closely related INR method in the literature.

### 3.1 Results

Table 1 presents average Dice[3] scores with standard deviations for MetaSeg and baseline methods for both 2D and 3D cases. For 2D, MetaSeg achieves an excellent Dice score of  $0.93 \pm 0.012$  for 5 classes, comparable to that achieved by the U-Net, but with 90% fewer parameters. Interestingly, as Fig. 2 illustrates, MetaSeg fits a signal remarkably fast, yielding a viable segmentation in just 2 update steps to  $f_\theta(\cdot)$  (Dice: 0.85), and reaching peak performance at  $T_f$  steps. If fit further, MetaSeg’s segmentation accuracy eventually falls off as parameters overfit to the particular scan. MetaSeg also performs well on the finer-grained segmentation task with 24 classes, achieving a Dice score of  $0.86 \pm 0.032$  and outperforming the baseline U-Net. As shown in Fig. 3, MetaSeg provides consistently accurate segmentation results for the ventricles and brain stem, which have drastic variations in size and area across examples.



**Fig. 4. MetaSeg is capable of accurately segmenting 3D MRI volumes.** (a) Various cross sections of one learned volume. (b) Renderings of intermediate coronal, axial, and sagittal planes, with high agreement ( $\text{Dice} \approx 0.93$ ) with ground truth.

MetaSeg also performs well for 3D segmentation, achieving an average Dice score of  $0.91 \pm 0.011$  with 90% fewer parameters than SegResNet, and significantly outperforms NISF. Fig. 4 depicts rendered coronal, sagittal, and axial planes from the segmented volume showing high agreement with ground truth segmentation planes, also demonstrating that MetaSeg-encoded volumes can be readily queried at any viewing plane. We also explored whether MetaSeg can generate high resolution ( $2\times$ ) 3D segmentation maps when only trained on low resolution scans, due to its underlying continuous representation. MetaSeg achieved a Dice score of  $0.78 \pm 0.011$  on this task, while SegResNet [14] achieved a Dice score of  $0.73 \pm 0.019$ . Hence, while there is a significant drop in performance in this scenario, MetaSeg still outperforms U-Nets.


Finally, we conducted an ablation study on parameter initialization strategy for 2D segmentation, with results reported in Table 2. MetaSeg performs significantly better than INRs initialized with random, fixed, and image-only meta-learning strategies. The latter result, in particular, demonstrates the benefit of learning structural priors with joint segmentation supervision.

**Table 2. Ablation studies on INR initialization strategies for 2D segmentation.** MetaSeg’s approach of jointly learning initial INR and segmentation parameters outperforms random, fixed, and image-only meta-learning initialization strategies.

Init. Strategy	Random	Fixed	Meta-learn, Image Only	MetaSeg
Dice scores	$0.30 \pm 0.057$	$0.53 \pm 0.1$	$0.81 \pm 0.033$	<b><math>0.93 \pm 0.012</math></b>

## 4 Discussion and Conclusion

Results demonstrate that MetaSeg achieves strong performance on both 2D and 3D MRI segmentation tasks, with comparable Dice scores to common U-Net



**Fig. 5. PCA visualization of penultimate features learned by MetaSeg and a standard INR for a 2D test scan.** We performed PCA separately on the features returned by each INR across spatial coordinates. We see a strong correlation for MetaSeg features with anatomical structures. For example, component #2 approximately resembles inner brain regions, #3 resembles regions like hippocampus and basal ganglia, #4 resembles ventricles, and #5 captures the cerebral cortex. On the contrary, a standard INR yields seemingly random features.

baselines [4, 14] but with 90% fewer parameters. In addition, MetaSeg significantly outperforms NISF [20], a related INR approach for segmentation. Ablation results demonstrate that our meta-learning strategy to initialize parameters with joint reconstruction and segmentation supervision is crucial for success.

MetaSeg offer a new perspective on the capabilities of INRs for imaging tasks, particularly beyond signal representation, which has been their predominant application in the literature. In particular, the results demonstrate that the meta-learning strategy can be useful for image translation tasks such as segmentation, where there is strong cross-correlation across signals. Indeed, as demonstrated by PCA decomposition in Fig. 5, MetaSeg learns low-dimensional embeddings that seem to jointly model both pixel intensity and semantic structure.

Because an INR explicitly conditions its features on input coordinates, it is likely that an INR will learn highly signal-specific features localized in space. While this is a good property for signal representation, it may lead to poor performance on tasks requiring generalization, such as segmentation. We conducted a pilot test to measure the sensitivity of MetaSeg to the spatial alignment of test scans to the training set distribution by applying random orientation and translation augmentations. We found that Dice score drops by 2%-6% for test images rotated randomly in  $[5^\circ, 15^\circ]$ , and drops by 3%-9% for test images with a random translation in  $[5 - 10]$  pixels. Hence, MetaSeg can be somewhat sensitive to spatial alignment, though further experimentation is needed to understand whether this can be mitigated itself with proper training augmentation.

MetaSeg learns fast, is resource-friendly, and scales well with more dimensions, unlike typical vision and transformer models. It also is easy to train, using only standard reconstruction and classification losses. In an age where AI re-



quires increasing resource and financial investments, MetaSeg provides a fresh new perspective to solving image segmentation tasks.

**Acknowledgments.** The authors acknowledge support from NIH award R01DE032051, ARPA-H award D24AC00296, NSF awards IIS-2107313 and EEC-1648449.

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

## References

1. Arefeen, Y., Levac, B., Stoebner, Z., Tamir, J.: Infusion: Diffusion regularized implicit neural representations for 2d and 3d accelerated mri reconstruction. arXiv preprint arXiv:2406.13895 (2024)
2. Bauer, M., Dupont, E., Brock, A., Rosenbaum, D., Schwarz, J.R., Kim, H.: Spatial functa: Scaling functa to imagenet classification and generation (2023), <https://arxiv.org/abs/2302.03130>
3. Bertels, J., Eelbode, T., Berman, M., Vandermeulen, D., Maes, F., Bisschops, R., Blaschko, M.B.: Optimizing the dice score and jaccard index for medical image segmentation: Theory and practice. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2019. pp. 92–100. Springer International Publishing, Cham (2019)
4. Buda, M., Saha, A., Mazurowski, M.A.: Association of genomic subtypes of lower-grade gliomas with shape features automatically extracted by a deep learning algorithm. *Computers in Biology and Medicine* **109** (2019). <https://doi.org/10.1016/j.compbiomed.2019.05.002>
5. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murray, B., Myronenko, A., Zhao, C., Yang, D., Nath, V., He, Y., Xu, Z., Hatamizadeh, A., Zhu, W., Liu, Y., Zheng, M., Tang, Y., Yang, I., Zephyr, M., Hashemian, B., Alle, S., Zalbagi Darestani, M., Budd, C., Modat, M., Vercauteren, T., Wang, G., Li, Y., Hu, Y., Fu, Y., Gorman, B., Johnson, H., Genereaux, B., Erdal, B.S., Gupta, V., Diaz-Pinto, A., Dourson, A., Maier-Hein, L., Jaeger, P.F., Baumgartner, M., Kalpathy-Cramer, J., Flores, M., Kirby, J., Cooper, L.A., Roth, H.R., Xu, D., Bericat, D., Floca, R., Zhou, S.K., Shuaib, H., Farahani, K., Maier-Hein, K.H., Aylward, S., Dogra, P., Ourselin, S., Feng, A.: MONAI: An open-source framework for deep learning in healthcare (Nov 2022). <https://doi.org/https://doi.org/10.48550/arXiv.2211.02701>
6. Dupont, E., Kim, H., Eslami, S.M.A., Rezende, D.J., Rosenbaum, D.: From data to functa: Your data point is a function and you can treat it like one. In: 39th International Conference on Machine Learning (ICML) (2022)
7. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: Precup, D., Teh, Y.W. (eds.) *Proceedings of the 34th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 70, pp. 1126–1135. PMLR (06–11 Aug 2017), <https://proceedings.mlr.press/v70/finn17a.html>
8. Hoopes, A., Hoffmann, M., Greve, D.N., Fischl, B., Gutttag, J., Dalca, A.V.: Learning the effect of registration hyperparameters with hypermorph. *Machine Learning for Biomedical Imaging 1(IPMI 2021)*, 1–30 (Apr 2022). <https://doi.org/10.59275/j.melba.2022-74f1>, <http://dx.doi.org/10.59275/j.melba.2022-74f1>

9. Khan, A., Rauf, Z., Khan, A.R., Rathore, S., Khan, S.H., Shah, N.S., Farooq, U., Asif, H., Asif, A., Zahoor, U., Khalil, R.U., Qamar, S., Asif, U.H., Khan, F.B., Majid, A., Gwak, J.: A recent survey of vision transformers for medical image segmentation (2023), <https://arxiv.org/abs/2312.00634>
10. Kingma, D.P.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
11. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **42**(2), 318–327 (2020). <https://doi.org/10.1109/TPAMI.2018.2858826>
12. Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L.: Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI Data in Young, Middle Aged, Nondemented, and Demented Older Adults. *Journal of Cognitive Neuroscience* **19**(9), 1498–1507 (09 2007). <https://doi.org/10.1162/jocn.2007.19.9.1498>, <https://doi.org/10.1162/jocn.2007.19.9.1498>
13. Martel, J.N.P., Lindell, D.B., Lin, C.Z., Chan, E.R., Monteiro, M., Wetzstein, G.: Acorn: Adaptive coordinate networks for neural scene representation. *ACM Trans. Graph. (SIGGRAPH)* **40**(4) (2021)
14. Myronenko, A.: 3d mri brain tumor segmentation using autoencoder regularization. In: Crimi, A., Bakas, S., Kuijf, H., Keyvan, F., Reyes, M., van Walsum, T. (eds.) *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*. pp. 311–320. Springer International Publishing, Cham (2019)
15. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. pp. 234–241. Springer International Publishing, Cham (2015)
16. Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., Baraniuk, R.G.: Wire: Wavelet implicit neural representations. In: *Conf. Computer Vision and Pattern Recognition* (2023)
17. Saragadam, V., Tan, J., Balakrishnan, G., Baraniuk, R., Veeraraghavan, A.: Miner: Multiscale implicit neural representations. In: *European Conf. Computer Vision* (2022)
18. Shi, J., Zhu, J., Pelt, D.M., Batenburg, K.J., Blaschko, M.B.: Implicit neural representations for robust joint sparse-view ct reconstruction. arXiv preprint arXiv:2405.02509 (2024)
19. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. *Advances in neural information processing systems* **33**, 7462–7473 (2020)
20. Stolt-Ansö, N., McGinnis, J., Pan, J., Hammernik, K., Rueckert, D.: Nisf: Neural implicit segmentation functions. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 734–744. Springer Nature Switzerland, Cham (2023)
21. Tancik, M., Mildenhall, B., Wang, T., Schmidt, D., Srinivasan, P.P., Barron, J.T., Ng, R.: Learned initializations for optimizing coordinate-based neural representations. In: *CVPR* (2021)
22. Vyas, K.K., Humayun, I., Dashpute, A., Baraniuk, R., Veeraraghavan, A., Balakrishnan, G.: Learning transferable features for implicit neural representations. In: *Advances in Neural Information Processing Systems*. vol. 37 (2024)

23. Zha, R., Zhang, Y., Li, H.: Naf: Neural attenuation fields for sparse-view cbct reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 442–452. Springer (2022)