

# Don't Mind the Gaps: Implicit Neural Representations for Resolution-Agnostic Retinal OCT Analysis

Bennet Kahrs<sup>1,2,\*</sup>, Julia Andresen<sup>2,\*</sup>, Fenja Falta<sup>2</sup>, Monty Santarossa<sup>3</sup>, Heinz Handels<sup>1,2</sup>, Timo Kepp<sup>1</sup>

<sup>1</sup> German Research Center for Artificial Intelligence, Luebeck, DE

<sup>2</sup> Institute of Medical Informatics, University of Luebeck, Luebeck, DE

<sup>3</sup> Multimedia Information Processing Group, Kiel University, Kiel, DE

\* These authors contributed equally to this work.

## Abstract

Routine clinical imaging of the retina using optical coherence tomography (OCT) is performed with large slice spacing, resulting in highly anisotropic images and a sparsely scanned retina. Most learning-based methods circumvent the problems arising from the anisotropy by using 2D approaches rather than performing volumetric analyses. These approaches inherently bear the risk of generating inconsistent results for neighboring B-scans. For example, 2D retinal layer segmentations can have irregular surfaces in 3D. Furthermore, the typically used convolutional neural networks are bound to the resolution of the training data, which prevents their usage for images acquired with a different imaging protocol. Implicit neural representations (INRs) have recently emerged as a tool to store voxelized data as a continuous representation. Using coordinates as input, INRs are resolution-agnostic, which allows them to be applied to anisotropic data. In this paper, we propose two frameworks that make use of this characteristic of INRs for dense 3D analyses of retinal OCT volumes. 1) We perform inter-B-scan interpolation by incorporating additional information from en-face modalities, that help retain relevant structures between B-scans. 2) We create a resolution-agnostic retinal atlas that enables general analysis without strict requirements for the data. Both methods leverage generalizable INRs, improving retinal shape representation through population-based training and allowing predictions for unseen cases. Our resolution-independent frameworks facilitate the analysis of OCT images with large B-scan distances, opening up possibilities for the volumetric evaluation of retinal structures and pathologies. Our code is available at <https://github.com/tkepp/ResA-OCT>.

## Keywords

Optical Coherence Tomography, Implicit Neural Representation, B-scan Interpolation, Retinal Atlas, Multi-modal Analysis, Image Registration

## Article informations

<https://doi.org/10.59275/j.melba.2026-38ba>

©2026 Kahrs, Andresen, Falta, Santarossa, Handels, Kepp. License: CC-BY 4.0

Volume 2026, Received: 2025-07-15, Published 2026-03



Corresponding author: [timo.kepp@dfki.de](mailto:timo.kepp@dfki.de)

Special issue: MELBA-BVM 2025 Special Issue

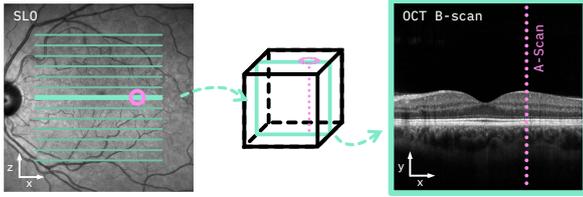
Guest editors: Andreas Maier, Thomas Deserno, Heinz Handels, Klaus Maier-Hein, Christoph Palm, Thomas Tolxdorff, Katharina Breining

## 1. Introduction

**O**ptical coherence tomography (OCT) is a standard imaging modality in ophthalmology (Huang et al., 1991). Due to its noninvasive nature, it is routinely used to measure alteration in layer thickness and visualize fluid accumulations or other structural changes of the central retina. With this, OCT supports the identification of a broad spectrum of ophthalmological diseases, such as age-related macular degeneration (AMD) (Schmidt-Erfurth

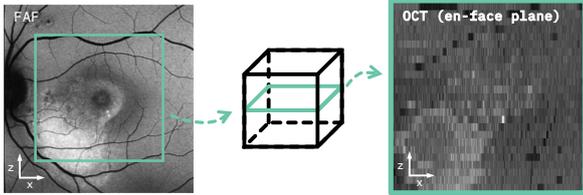
et al., 2017), central serous chorioretinopathy (CSCR) (Chhablani et al., 2020) and diabetic retinopathy (Virgili et al., 2015), as well as neurodegenerative conditions like multiple sclerosis (Britze and Frederiksen, 2018). Systematic learning-based analysis of OCT, whether of single scans, scan series, or population-based statistics, e.g., based on atlases (Chakravarty et al., 2018), can enhance diagnostics and therapy in clinical practice. However, the nature of its acquisition brings along some specific challenges: In favor

### Spatial correspondence between en-face SLO and OCT B-scans



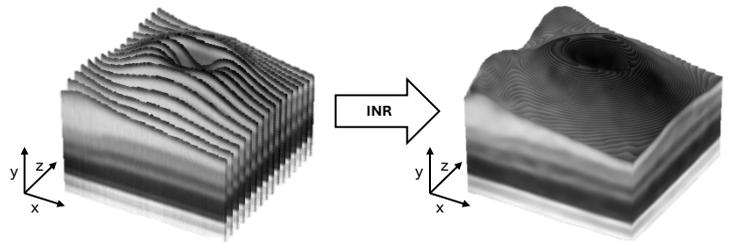
- Each scan line in the SLO corresponds to a single B-scan
- Each point on a scan line maps to an A-scan

### Resolution discrepancy between en-face image and OCT volume



- Difficult to match with FAF due to resolution gap

### From sparsely sampled volumes to continuous representations



- Large B-scan distances restrict volumetric OCT analysis
- INRs enable resolution-independent retinal analysis

Figure 1: Top left: Interrelation between en-face image (here: SLO) and OCT. The position marked in the SLO (●) corresponds to an A-scan (dotted line) in the OCT B-scan. Bottom left: Difference in resolution between an en-face image (here: FAF) and an analogous lateral slice of OCT. The limited amount of B-scans acquired in clinical practice leads to high anisotropy. Right: By leveraging a generalizable INR, we create resolution-independent representations of singular data instances or an atlas, respectively.

of shorter acquisition times, clinical routine OCT uses large inter-B-scan distances, leading to sparse sampling in the slow scanning direction and, consequently, to an anisotropic spatial resolution. Most learning-based approaches using convolutional neural networks (CNNs) therefore opt for two-dimensional (2D) solutions rather than performing full three-dimensional (3D) analyses, e.g., for the segmentation of retinal layers (He et al., 2019; Pekala et al., 2019; Li et al., 2020; Kepp et al., 2023). However, 2D segmentations do not take into account the information from neighboring B-scans, which may lead to inconsistencies between segmentations of individual slices. The resulting 3D surfaces typically are irregular and require extensive post-processing using graph-cut approaches (Li et al., 2006), additional refinement networks (Liu et al., 2024) or advanced regularization techniques (Kepp et al., 2019). Another problem that arises from the sparsely scanned retina is that small anatomical or pathological structures might be missed, reducing the diagnostic precision of OCT image analysis solutions (Velaga et al., 2017; von der Burchard et al., 2024). Volumetric measurements performed on these anisotropic scans are inherently imprecise since structures appear stretched or compressed. Some structural context for the space between the B-scans is maintained by en-face scanning laser ophthalmoscopy (SLO), which offers high-resolution imaging of the fundus and is used as a reference for eye tracking in certain OCT systems. However, SLO does not contain differentiated depth information. Similarly, other 2D modalities, such as color fundus photography or fundus autofluorescence (FAF), can also hold valuable

complementary information, especially about pathological structures visible on the retinal surface. The relation between these modalities and visible biomarkers is yet to be fully explored (Santarossa et al., 2022b; Velaga et al., 2022). The spatial correspondence and the resolution discrepancy between en-face SLO/FAF and OCT B-scans are illustrated in Fig. 1 (left, top and bottom).

The resolution-agnostic nature and structural flexibility of implicit neural representations (INRs) suggest that they have the potential to work well in the presence of these challenges. INRs aim to approximate individual data instances as continuous functions and have gained much traction in research recently. They have been shown to be able to efficiently represent and reconstruct natural images and scenes (Sitzmann et al., 2020; Mildenhall et al., 2021). With additional instance-specific latent codes or layer modifications, they become generalizable and can be trained population-based to represent multiple images at once (Chen and Zhang, 2019; Park et al., 2019; Kazerooni et al., 2024; Dupont et al., 2022; Kim et al., 2023). Consequently, a lot of research has been done to examine their applicability to medical tasks such as semantic segmentation (Stolt-Ansó et al., 2023), deformable registration (Wolterink et al., 2022) or atlas generation (Großbröhmer et al., 2024; Dannecker et al., 2024). However, these medical applications focus on data with far less anisotropy, such as CT or MRI, while highly anisotropic data like OCT remain rather underexplored.

**Our contribution** In this paper, we explore the usage of INRs for OCT-specific tasks focused on the challenging anisotropy of the data (Fig. 1, right). 1) We build upon previous work (Kepp et al., 2025) and make use of densely sampled but two-dimensional SLO and FAF information for the interpolation in-between B-scans. This increases the resolution in the slow scanning direction and is applied to images from healthy and diseased subjects and called anisotropic interpolation hereafter. We leverage a generalizable INR to incorporate the additional en-face modality and support this task through population-based retinal layer or pathology segmentation. 2) We perform implicit inter-subject registration and combine it with the introduced generalization approach to perform population-based registration and create an intensity- and shape-based atlas. An additional INR network enables the construction of a resolution-independent atlas, making it especially suited for anisotropic image data like OCT.

## 2. Related Works

### 2.1 Interpolation of Anisotropic OCT Images

Various interpolation methods exist to enhance anisotropic OCT image data resolution. Simple intensity-based methods are unable to capture shape and localization differences between B-scans, leading to artifacts and ambiguous results. Lindberg et al. (2018) propose a weighted combination of linear and transfinite interpolation to achieve an improved interpolation of OCT images, but rely on two OCT images taken with orthogonal scan directions. Registration-based interpolation methods (Ehrhardt et al., 2007) align structural features more accurately and reduce misalignment artifacts, but cannot infer information missing between slices, limiting reconstruction quality in sparsely sampled regions. Generative approaches for super-resolution are capable of producing realistic, detailed images by reconstructing fine textures and complex structures, e.g. (López-Varela et al., 2023), and thus constitute an approach to B-scan interpolation as well. However, most GAN-based approaches for OCT super-resolution only perform B-scan super-resolution but do not interpolate between them (Huang et al., 2019; Das et al., 2020; Yuan et al., 2023), or require pairs of high- and low-resolution images (Huang et al., 2019; Yuan et al., 2023). Furthermore, generative approaches require a large amount of training data and carry the risk of hallucinations, making them precarious in clinical context (Cohen et al., 2018).

A limited number of previous works have already used INRs for optical imaging. Li et al. (2025) built upon the capability for multi-view scene reconstruction demonstrated by neural radiance fields (NeRF) (Mildenhall et al., 2021) and proposed a ray tracing-based framework for multi-view

reconstruction for OCT of different animal anatomies. Similarly, Xiao et al. (2025) used an INR for limited view reconstruction of photoacoustic scans of vessels in mice. These works showed that INRs have the capability to work well with sparsely acquired optical imaging data, and the advances of NeRF can be applied well to the medical field. However, methods for multi-view reconstruction cannot directly be applied to interpolation, since slices have no varied angles or overlaps.

### 2.2 INR Super-Resolution

While CNN-based approaches are limited to the resolution of the training data, INRs are resolution-independent, as they deliver continuous representations of the imaged objects. Therefore, INRs inherently are a well-suited choice for the task of super-resolution – which is methodologically close to the here performed B-scan interpolation. Super-resolution of 3D medical images using INRs can, for example, be found in (McGinnis et al., 2023; Wu et al., 2023; Fang et al., 2024). Wu et al. (2023) use a CNN encoder and an INR decoder to generate high-resolution MR images from the respective low-resolution images. Similarly, Fang et al. use a CNN-based encoder to extract feature vectors from low-resolution CT images. These vectors are fed to an INR together with the 3D coordinates of the high-resolution image to predict the pixel intensities for the higher resolution (Fang et al., 2024). McGinnis et al. (2023) make use of multi-contrast MRI, in which different low-resolution contrasts are acquired in different directions, to produce densely sampled, high-resolution images for single subjects.

While these methods either work with low-resolution but equidistantly sampled images or use more than one anisotropic image of the same scene, the resolution-agnostic nature of INRs also directly allows processing single anisotropic images. However, with very large inter-slice distances, as frequently used in retinal OCT imaging, fitting an INR to a single OCT volume results in overfitting to the positions given in the training data. That is, meaningful results are generated for the B-scan positions used during training, but there is no useful interpolation between the B-scans. In (Amiranashvili et al., 2022), this overfitting problem is overcome with a generalizable INR, which is trained on highly anisotropic segmentations. An MLP is trained that is shared between all training instances, while case-specific latent priors are adapted for each sample, allowing the generation of dense and smooth shapes. Building on this approach, Stolt-Ansó et al. (2023) train a generalizable INR to reconstruct and segment short-axis cardiac MRI and show smooth segmentations on long-axis images.

### 2.3 OCT Atlases

Equivalent to less anisotropic modalities, an OCT atlas can be constructed by combining deformable registration and optimization of an atlas image. In the context of optical imaging, an OCT atlas can be helpful to differentiate between physiological variability in normative tissue and pathological alterations. Based on these differences, many diseases have been shown to be classifiable using atlases. Chakravarty et al. (2018) created a normative atlas and calculated deviation to classify AMD. They proposed to use a geometric mean image as an initial template and employ discrete registration for computational efficiency due to the large size of OCT B-scans. Lee et al. (2017) employed geometry-based registration to create an atlas and visualized thinning of retinal nerve fiber layers due to glaucoma. Khansari et al. (2019) create an atlas based on pathological volumes to measure tissue contraction and expansion due to diabetic retinopathy.

In general, these previous approaches demonstrated the value of correctly aligning individual layers during the construction of the atlas, sometimes through explicit conditions on layer shapes. However, each of these approaches is limited by the resolution of the original OCT volumes. Consequently, the atlas is highly restrictive regarding the data it is applicable to. With different acquisition setups in clinical practice, this is an important factor to overcome. In this work, we will focus on the methodology to make the atlas generation resolution-agnostic and broadly applicable across varying data settings.

### 2.4 INR Registration and Atlases

Wolterink et al. (2022) introduced implicit deformable image registration, where a deformation field is represented by an INR and optimized for a single image pair. There have since been multiple adaptations proposed, e.g., B-spline regularization (Sideri-Lampretsa et al., 2024) or inverse consistency (Van Harten et al., 2023; Tian et al., 2024). To the best of our knowledge, the applicability of Wolterink et al. (2022) to OCT or similarly anisotropic data has not been explored in the academic literature yet.

INRs can further be used to represent an image atlas through different approaches. Großbröhmer et al. (2024) combine conventional registration with a single INR to represent the atlas image, whereas Dannecker et al. (2024) use a generalization approach to represent multiple images with one INR and thus omit any registration. It is still an open question how well the generalization capability of population-trained INRs translates to medical optical imaging with the associated challenges of sparsely sampled data, small datasets, and low signal-to-noise ratio.

## 3. Materials and Methods

### 3.1 Datasets

We use an in-house dataset consisting of paired OCT volumes and SLO images from the left and right eyes of 50 healthy volunteers (one OCT volume and one SLO image per eye), resulting in a total of 100 image pairs, to develop, train, and evaluate generalizable INRs for OCT interpolation, segmentation, and atlas generation. All OCT volumes were acquired with a Spectralis OCT scanner (Heidelberg Engineering) modified for research purposes to allow arbitrary inter-B-scan spacing, but image acquisition was performed using the angiography acquisition (OCT-A) mode consistent with standard clinical OCT systems. Acquisition in the angiography mode results in slightly reduced B-scan image quality, as fewer repeated scans per position are averaged compared to structural OCT acquired with typical clinical protocols, but provides densely sampled OCT volumes of size  $496 \times 512 \times 512$  voxels covering a field of view of  $2 \times 6 \times 6 \text{ mm}^3$ . In Spectralis scanners, an SLO (scanning laser ophthalmoscopy) is used for eye-tracking and acquired routinely alongside the OCT. The SLO images have a size of  $768 \times 768$  pixels, covering a larger field of view in the lateral plane (approx.  $8.6 \times 8.6 \text{ mm}^2$ ) than the OCT images. To match OCT and SLO, we use the localization information provided by the OCT device to subsample the SLO image at the A-scan positions corresponding to each B-scan. In the OCT images, twelve layers (eleven retinal layers plus choroid) were segmented automatically using the Iowa Reference Algorithms (Li et al., 2006). Subsequently, segmentation errors were corrected manually by an expert. For preprocessing, all B-scans are flattened in relation to the Bruch's membrane and cropped to  $230 \times 512$  pixels so that the retina is at the center height of the image. Images of right eyes are flipped for identical orientation of all images.

Additional experiments are performed on a clinical dataset containing images from 19 patients with central serous chorioretinopathy (CSCR). The CSCR dataset contains 146 OCT volumes consisting of 25 B-scans with a resolution of  $496 \times 512$  pixels. Like the healthy images, the CSCR images span a field of view of  $2 \times 6 \times 6 \text{ mm}^3$ . As an additional en-face modality, FAF (fundus autofluorescence) images are used for the CSCR data. They have the same size ( $496 \times 512 \times 512$  voxels) and field of view ( $8.6 \times 8.6 \text{ mm}^2$ ) as the SLO images. To match the FAF images to the OCT volumes, we use a two-step procedure utilizing the accompanying SLO. First, the FAF image is registered onto the SLO using the MedRegNet registration pipeline (Santarossa et al., 2022a), with improved RANSAC as proposed in (Santarossa et al., 2025). Next, the registered FAF is subsampled to the OCT resolution using the same SLO localization information as before. This entire two-step procedure, albeit in reverse (i.e., mapping OCT onto FAF),

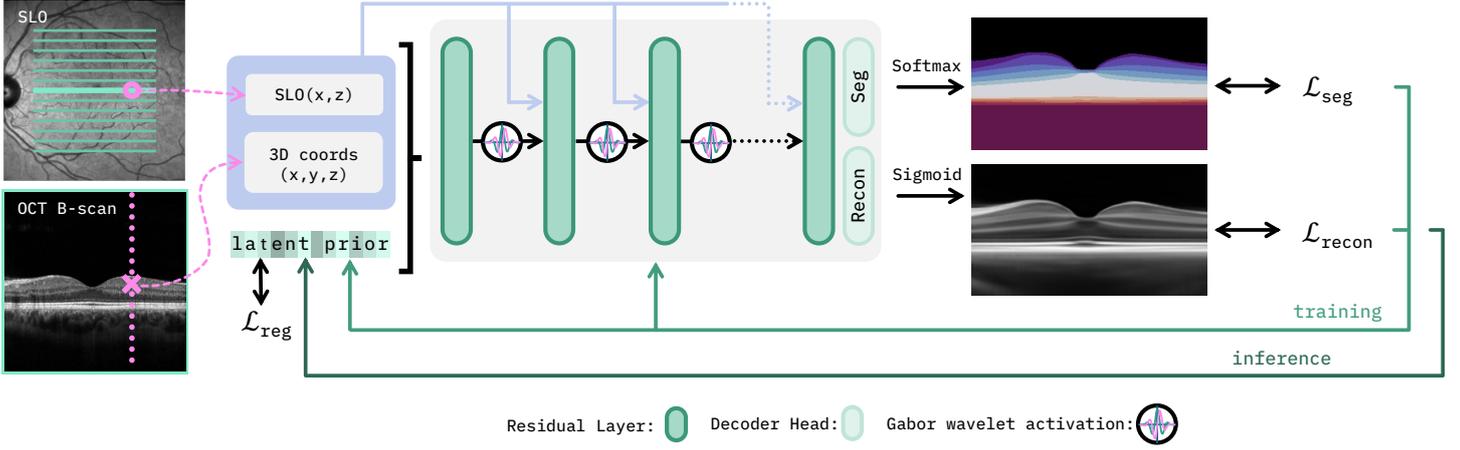


Figure 2: Schematic overview of the proposed interpolation framework. The input for the INR are 3D coordinates in the OCT ( $\times$ ) and intensities at the corresponding SLO position ( $\circ$ ). The INR is trained to reconstruct the intensity and segmentation label of the OCT volume at the respective coordinate. The reconstructed OCT can be evaluated at any arbitrary coordinate with a known SLO intensity. To reconstruct and segment an unseen case, the INR is kept frozen, and only the latent prior is adapted.

has been detailed in (Santarossa et al., 2022b). For the CSCR data, the retina as well as pathologies (intra- and subretinal fluid and pigment epithelial detachment) were segmented manually by experienced ophthalmologists.

### 3.2 Anisotropic Interpolation

To reduce the difficulties arising from the anisotropy of clinical OCT volumes, we build upon the framework for joint inter-B-scan interpolation and segmentation of retinal OCT first proposed in (Kepp et al., 2025). The method is shown exemplarily for healthy images in Fig. 2. A generalizable INR, extending the methods proposed by Amiranashvili et al. (2022) and Stolt-Ansó et al. (2023), is used that consists of a multi-layer perceptron (MLP) shared between all training samples and learnable subject-specific priors that condition the MLP on the respective subject. Since small structures between B-scans cannot be represented by the sparsely scanned OCT images alone, we propose to use additional en-face imaging modalities as input to the INR. Here, we exemplarily use SLO and FAF images that are densely scanned in the en-face plane. They therefore provide more detailed information about the localization of structures in the lateral direction than OCT. Although these images have no axial depth, we hypothesize that the lateral information enables the INR to represent structures that are not captured in the OCT scanning protocol and are therefore not visible in the B-scans. The INR  $f_\theta = (f_\theta^{recon}, f_\theta^{seg})$  with trainable parameters  $\theta$  has two separate output layers, one for the reconstruction of the OCT volume and one for the segmentation of anatomical or pathological retinal structures, and is trained with OCT and corresponding en-face images from  $n$  subjects. The network receives three inputs: a 3D coordinate  $(x, y, z) \in \mathbb{R}^3$  corresponding to a

voxel position in the OCT volume  $OCT_i$  of the subject  $i$ , the intensity of the en-face image  $EF_i(x, z)$  at the respective lateral position  $(x, z)$ , and the subject-specific prior  $\mathbf{p}_i \in \mathbb{R}^L$ .

For one forward pass, we use the coordinates  $(x, y, z)$  of an entire OCT B-scan and compute the loss

$$\mathcal{L} = \sum_{i=1}^n \sum_{(x,y,z)} [\mathcal{L}_{recon}(f_\theta^{recon}(x, y, z, EF_i(x, z), \mathbf{p}_i), OCT_i(x, y, z)) + \alpha \mathcal{L}_{seg}(f_\theta^{seg}(x, y, z, EF_i(x, z), \mathbf{p}_i), SEG_i(x, y, z)) + \beta \mathcal{L}_{reg}(\mathbf{p}_i)], \quad (1)$$

to adapt the network parameters and the latent priors simultaneously. Here,  $SEG_i(x, y, z)$  is the one-hot encoded ground truth segmentation label at position  $(x, y, z)$ . A combination of MSE and SSIM loss is used for the reconstruction loss  $\mathcal{L}_{recon} = 0.1 \cdot \mathcal{L}_{SSIM} + \mathcal{L}_{MSE}$  and the binary cross-entropy loss is used for the segmentation loss  $\mathcal{L}_{seg}$ . As proposed by Stolt-Ansó et al. (2023), an L2 regularization loss  $\mathcal{L}_{reg}$  is applied to penalize large values in the latent codes. The loss components are weighted by  $\alpha = \beta = 0.2$ .

After completion of the population-based training, the INR can be fit to an unseen subject  $n + 1$  by freezing the network  $f_\theta$  and adapting a new latent prior  $\mathbf{p}_{n+1}$ . Since for new subjects, typically the OCT image  $OCT_{n+1}(x, y, z)$  is given but not the corresponding segmentation, the loss for learning the new latent code is reduced to

$$\mathcal{L} = \sum_{(x,y,z)} \mathcal{L}_{recon}(f_\theta^{recon}(x, y, z, EF_{n+1}(x, z), \mathbf{p}_{n+1}), OCT_{n+1}(x, y, z)) + \beta \mathcal{L}_{reg}(\mathbf{p}_{n+1}). \quad (2)$$

As shown in our previous work (Kepp et al., 2025), learning the latent priors for unseen cases this way results in

plausible segmentation predictions. The learned latent prior  $\mathbf{p}_{n+1}$  enables both the interpolation of OCT intensities and segmentation labels for any coordinate  $(x, y, z)$  for which the intensity of the additional 2D modality  $\text{EF}_{n+1}(x, z)$  is known. To do so, a simple forward pass is used to calculate  $f_\theta(x, y, z, \text{EF}_{n+1}(x, z), \mathbf{p}_{n+1})$ .

### 3.2.1 Implementation Details

The generalizable INR for OCT interpolation consists of a six-layer MLP using complex Garbor wavelet activation functions (Saragadam et al., 2023) plus two separate linear layers as decoder heads for reconstruction output  $f_\theta^{\text{recon}}$  and segmentation output  $f_\theta^{\text{seg}}$ . All hidden layers have 128 nodes, and the reconstruction head outputs scalar grayscale values, while the segmentation head generates  $C$  output values for one-hot encoded class labels, where  $C$  is the number of classes considered in the respective task. That is, for the segmentation of retinal layers, twelve labels corresponding to the different tissue layers – including the choroid – are predicted, and for the segmentation of pathological images, three labels corresponding to the background, retina, and pathologies are predicted. The reconstruction output layer uses sigmoid activation, whereas the segmentation head uses softmax activation. All hidden layers use residual connections, and to enhance focus on input coordinates as well as the additional imaging modality, we pass these inputs repeatedly to the hidden layers. The length  $L$  of the latent vectors is set to 128 empirically. Network training is performed for 1,500 epochs, where one epoch corresponds to a run over all B-scans of the training images. We use an exponential learning rate decay of 0.99 and an initial learning rate of  $10^{-4}$ . Early stopping is enabled after 100 epochs using a patience of 20 epochs and a minimum loss improvement of  $10^{-5}$ . The generalizable INR was trained for up to three days on a single NVIDIA GeForce GTX 1080 Ti.

### 3.3 Registration and Atlas Generation

To overcome the resolution dependence of retinal atlases, we continue to use an INR for efficient handling of the anisotropic data. We combine INR-based inter-subject registration with atlas generation – enabling joint optimization of registration and atlas without meta-learning. Our proposed method consists of two key components that jointly model the deformation from the atlas to the subject space, as well as the atlas representation itself. These components are 1) a learnable atlas volume encoded with an INR and 2) a generalizable INR that models the deformation fields with latent priors that modulate the generalizable INR to capture subject-specific variations. The joint framework is shown in Fig. 3 and is described in detail below. The resulting framework takes 3D coordinates and a latent prior

as input to predict the deformed atlas values with auxiliary information, in our case retina layer labels.

**The learnable atlas** is a key component to enable resolution independence. It builds upon the idea proposed by Großbröhmer et al. (2024) of representing the atlas using an MLP  $f_\theta$  with adjustable parameters  $\theta$ . As such, the atlas is trained using warped input coordinates to predict the most similar subject representation. In this work, we extend this approach by incorporating a combined representation of intensity values and segmentation labels, similar to our previously proposed interpolation method. Thus, we introduce an INR  $f_\theta = (f_\theta^{\text{recon}}, f_\theta^{\text{seg}})$ , which takes warped 3D coordinates  $(x', y', z') \in \mathbb{R}^3$  to reconstruct the intensity value of the OCT image and the class label of the retinal layers. Since we aim to construct a single atlas volume that represents the whole dataset, this INR is (different from our previous approach) non-generalizable. This continuous representation of the atlas enables the intensity and segmentation values to be sampled instantly, without the need for explicit interpolation in the registration process.

**The displacement model** is a generalizable INR, denoted as  $g_\vartheta$  with trainable parameters  $\vartheta$ . In contrast to the proposed method for interpolation, the displacement model uses a hypernetwork to modulate its interim outputs, similar to Dannecker et al. (2024) and Dupont et al. (2022). To do so, we use the hypernetwork (depicted in purple in Fig. 3) – consisting of one linear layer for each layer of the INR (depicted in turquoise in Fig. 3) – to predict individual scaling and shifting vectors based on the latent prior input  $\mathbf{p}_i \in \mathbb{R}^L$ . Let  $H$  be the hidden size of the INR. The individual scaling vector  $\phi \in \mathbb{R}^H$  and shifting vector  $\psi \in \mathbb{R}^H$  are the result of the linear mapping  $(\phi, \psi)^T = W_{\text{hyper}}\mathbf{p}_i + \mathbf{b}_{\text{hyper}}$  with weight matrix  $W_{\text{hyper}} \in \mathbb{R}^{2H \times L}$ , and bias  $\mathbf{b}_{\text{hyper}} \in \mathbb{R}^{2H}$ . By modulating  $g_\vartheta$  with scaling vector  $\phi$  and shifting vector  $\psi$ , the output of a linear layer  $\sigma(W\mathbf{x} + \mathbf{b})$  with input  $\mathbf{x}$ , weight matrix  $W$ , bias  $\mathbf{b}$  and activation function  $\sigma$  becomes  $\sigma(\phi(W\mathbf{x} + \mathbf{b}) + \psi)$ .

The resulting latent code-driven architecture is then used to represent a continuous displacement field  $u_i : \mathbb{R}^3 \rightarrow \mathbb{R}^3$  for each subject  $i$  as already introduced for classic INRs by Wolterink et al. (2022). Thus, this model gets the original 3D input coordinates  $(x, y, z) \in \mathbb{R}^3$  and a latent prior  $\mathbf{p}_i \in \mathbb{R}^L$ . For each coordinate  $(x, y, z)$ , we thereby obtain a subject-specific displacement

$$u_i(x, y, z) = g_\vartheta(x, y, z, \mathbf{p}_i) \quad (3)$$

and warped coordinates

$$\varphi_i(x, y, z) = u_i(x, y, z) + (x, y, z) \quad (4)$$

that act as the input  $(x', y', z')$  to the previously described learnable atlas.

In contrast to the architecture of Dannecker et al. (2024), we separate the parts  $\mathbf{p}_{i_1}$  and  $\mathbf{p}_{i_2}$  of the latent

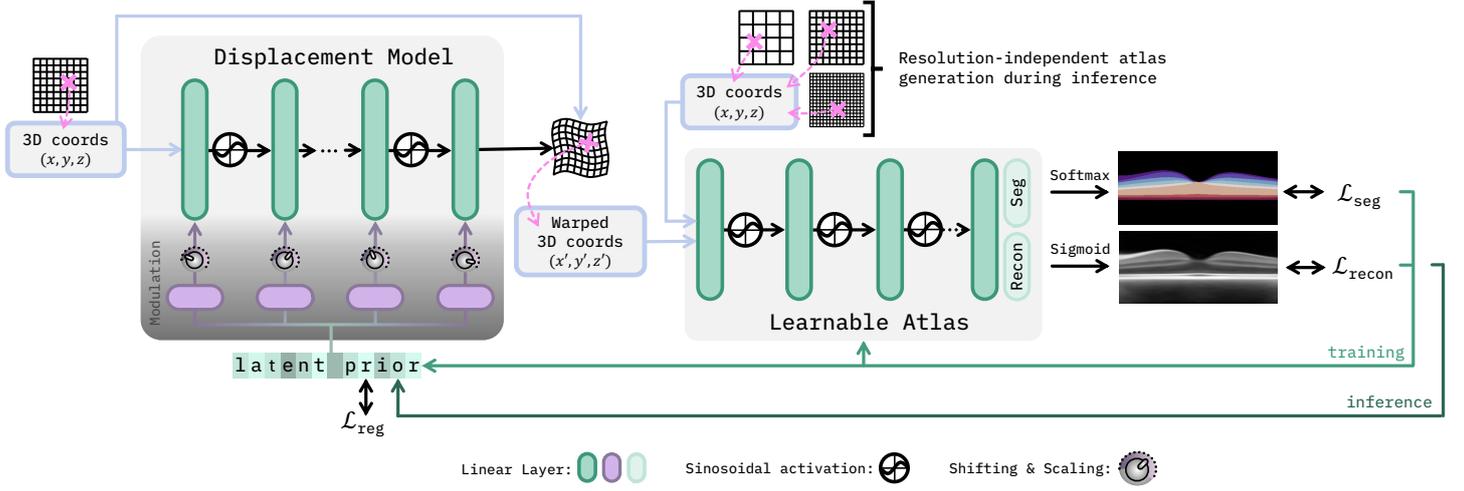


Figure 3: Schematic overview of the proposed atlas generation framework. The generalizable INR predicts a deformation field based on 3D coordinates ( $\times$ ) and an instance-specific latent prior. The INR layers are modulated through scale and shift vectors predicted from the latent prior by a hypernetwork (purple layers). The deformed coordinates are then input into another INR that represents the atlas with intensities and segmentation labels. The learned atlas is resolution-independent and can be evaluated on any arbitrary number of coordinates.

prior  $\mathbf{p}_i = (\mathbf{p}_{i_1}, \mathbf{p}_{i_2})$  that predict  $\phi$  and  $\psi$  and use two normal distributions  $\mathbf{p}_{i_1} \sim \mathcal{N}(1, 0.1)$  and  $\mathbf{p}_{i_2} \sim \mathcal{N}(0, 0.1)$  to initialize them. With  $\phi = 1$  and  $\psi = 0$  the resulting generalizable INR corresponds to a classic INR. This initialization therefore reduces the disturbance from a non-trained latent prior right after initialization.

Training and inference for atlas generation work analogously to those we used for B-scan interpolation. Assembled into one framework, a subset of coordinates  $(x, y, z)$  of an OCT volume is used to compute the reconstruction and alignment error as

$$\begin{aligned} \mathcal{L} = & \sum_{i=1}^n \sum_{(x,y,z)} [\mathcal{L}_{\text{recon}}(f_{\theta}^{\text{recon}}(\varphi_i(x, y, z)), \text{OCT}_i(x, y, z)) \\ & + \alpha \mathcal{L}_{\text{seg}}(f_{\theta}^{\text{seg}}(\varphi_i(x, y, z)), \text{SEG}_i(x, y, z)) \\ & + \beta \mathcal{L}_{\text{reg}}(u_i(x, y, z))]. \end{aligned} \quad (5)$$

Note that this loss is dependent on the network parameters  $\vartheta$  and the latent priors  $\mathbf{p}_1, \dots, \mathbf{p}_n$  through Eq. (3) and Eq. (4).  $\text{OCT}_i(x, y, z)$  and  $\text{SEG}_i(x, y, z)$  are the ground truth intensity or segmentation labels at position  $(x, y, z)$ . The reconstruction loss  $\mathcal{L}_{\text{recon}}$  is based on the image intensity and uses MSE loss. Binary cross-entropy loss is used as the segmentation loss  $\mathcal{L}_{\text{seg}}$ . To regularize the deformation, L1 loss is used as  $\mathcal{L}_{\text{reg}}$ . The weighting factors are experimentally chosen and set to  $\alpha = 1.0$  and  $\beta = 0.01$ .

To apply the model to previously unseen subjects, only the latent prior  $\mathbf{p}_{n+1}$  needs to be optimized, while the deformation model  $g_{\vartheta}$  and the atlas representation  $f_{\theta}$  stay frozen. Since the available information during inference is limited to the OCT intensity values of one subject, the loss

function is modified to

$$\begin{aligned} \mathcal{L} = & \sum_{(x,y,z)} [\mathcal{L}_{\text{recon}}(f_{\theta}^{\text{recon}}(\varphi_i(x, y, z)), \text{OCT}_{n+1}(x, y, z)) \\ & + \gamma \mathcal{L}_{\text{reg}}(u_i(x, y, z))] \end{aligned} \quad (6)$$

where  $\gamma$  is a weighting factor and set to 0.01.

### 3.3.1 Implementation Details

The atlas representation  $f_{\theta}$  is a four-layer MLP with two additional separate layers for reconstruction  $f_{\theta}^{\text{recon}}$  and segmentation  $f_{\theta}^{\text{seg}}$ . Each hidden layer has 256 nodes and uses a sinusoidal activation function as proposed by Sitzmann et al. (2020). The prediction heads are based on sigmoid activation for the image intensity values and softmax activation to predict the class values as one-hot encoding. Different from the interpolation task, here, the class labels are redefined to ten retinal layers by merging very thin layers and discarding the choroid label. While Großbröhmer et al. (2024) suggest using a warm-start procedure by increasing the learning rate, we propose to pre-train the atlas INR  $f_{\theta}$  for 500 epochs to represent the voxel-wise median of the dataset, resulting in a less blurred initialization. The generalizable INR used to model the deformation consists of four layers with 128 nodes and uses sinusoidal activations throughout, including a single sine activation in the final layer to predict the three-dimensional displacement vector. The hypernetwork's layers are three separate linear layers with 256 nodes each. The model weights are initialized by sampling from a uniform distribution  $W \sim \mathcal{U}(-0.001, 0.001)$  for small and smooth initial deformations. The vector length of 128 for the latent space is chosen empirically. All components –

Table 1: OCT interpolation and segmentation results for the healthy OCT data. Results are reported for our proposed generalizable INR with (GenINR<sub>SLO</sub>) and without (GenINR) additional SLO input for train and test data separately. For comparison, interpolation results are reported for linear and registration-based interpolation, as well as for standard INRs adapted to single images with and without SLO guidance (SingleINR<sub>SLO</sub> / SingleINR), evaluated on the training data only since these instance-based methods require ground truth annotations for some of the B-scans. Metrics are stated as the mean  $\pm$ std. The best results are highlighted. All INR methods show significant differences in Dice, ASSD, and SSIM in comparison to the registration-based interpolation ( $p < 0.001$ , Wilcoxon signed-rank test with Holm-Bonferroni correction).

	Method	Image Reconstruction				Image Segmentation		
		$\downarrow$ MAE[%]	$\uparrow$ PSNR	$\uparrow$ SSIM[%]	$\downarrow$ LPIPS	$\uparrow$ Dice[%]	$\downarrow$ ASSD[ $\mu$ m]	$\downarrow$ HD[ $\mu$ m]
train	linear	5.92 $\pm$ 0.55	21.7 $\pm$ 0.5	39.7 $\pm$ 5.6	<b>0.12</b> $\pm$ 0.03	88.8 $\pm$ 1.1	9.3 $\pm$ 0.7	37.2 $\pm$ 3.3
	reg.	5.76 $\pm$ 0.55	22.0 $\pm$ 0.5	41.3 $\pm$ 5.6	<b>0.13</b> $\pm$ 0.03	89.1 $\pm$ 1.2	7.5 $\pm$ 0.8	39.2 $\pm$ 4.4
	SingleINR	6.65 $\pm$ 0.64	20.6 $\pm$ 0.5	36.7 $\pm$ 5.9	0.26 $\pm$ 0.02	90.8 $\pm$ 1.1	7.7 $\pm$ 0.7	37.6 $\pm$ 3.6
	SingleINR <sub>SLO</sub>	9.42 $\pm$ 1.51	17.5 $\pm$ 1.0	26.8 $\pm$ 6.4	0.45 $\pm$ 0.07	83.6 $\pm$ 5.3	18.7 $\pm$ 12.0	603.4 $\pm$ 179.8
	GenINR	<b>5.24</b> $\pm$ 0.51	<b>22.6</b> $\pm$ 0.5	<b>48.9</b> $\pm$ 5.7	0.53 $\pm$ 0.03	<b>91.9</b> $\pm$ 1.0	<b>7.0</b> $\pm$ 0.7	<b>32.5</b> $\pm$ 3.7
	GenINR <sub>SLO</sub>	<b>5.24</b> $\pm$ 0.51	<b>22.6</b> $\pm$ 0.5	<b>48.9</b> $\pm$ 5.7	0.52 $\pm$ 0.03	<b>91.9</b> $\pm$ 1.0	<b>7.0</b> $\pm$ 0.6	<b>32.5</b> $\pm$ 3.7
test	GenINR	6.04 $\pm$ 0.57	21.4 $\pm$ 0.5	44.3 $\pm$ 6.0	0.55 $\pm$ 0.03	86.5 $\pm$ 3.7	11.4 $\pm$ 2.1	45.1 $\pm$ 4.9
	GenINR <sub>SLO</sub>	6.08 $\pm$ 0.59	21.4 $\pm$ 0.5	44.2 $\pm$ 6.0	0.54 $\pm$ 0.03	86.2 $\pm$ 3.8	11.7 $\pm$ 2.2	45.4 $\pm$ 5.0

the deformation model, the atlas model, and the latent priors – are trained simultaneously for 2,000 iterations using the Adam optimizer with learning rates of  $10^{-5}$ ,  $10^{-6}$ , and  $10^{-5}$ , respectively. The optimizer is combined with an exponential learning rate scheduler with a multiplication factor of 0.999. The batch size is set to one subject, with 916,608 coordinates each. This number corresponds to an equidistant downsampling to a sub-volume containing  $231 \times 62 \times 64$  voxels. This volume is randomly selected from all shifted variations of downsampled sub-volumes. Training was conducted on two NVIDIA RTX A5000s and completed in approximately eight hours. During inference, one latent prior is trained for each subject for 100 epochs with a learning rate of  $5 \times 10^{-4}$  and the same exponential learning rate scheduler as for training.

## 4. Experiments and Results

### 4.1 Anisotropic Interpolation

The interpolation capabilities of the proposed INR are first assessed using the healthy dataset. To do so, the network is trained with 80 OCT volumes and SLO images of 40 volunteers using only 16 B-scans per volume as input to the INR, resulting in even higher inter-slice distances than for typical clinical images. The input B-scans are sampled equidistantly from the entire depth of the image volumes, enabling evaluation on intermediate B-scans. For comparison, we use simple linear interpolation, registration-based interpolation (Ehrhardt et al., 2007) and non-generalizable INRs adapted to single OCT images. To segment intermediate B-scans for the interpolation methods, we proceed as

follows: For the linear interpolation, we interpolate the label images linearly and round the resulting continuous values to integers. For the registration-based method, we use the deformation fields resulting from the B-scan interpolation and apply them to the label images using nearest neighbor interpolation. The non-generalizable INRs (SingleINRs) are implemented with a slightly changed architecture, since we observed improved performance for the modified architecture. The SingleINR, thus, consists of three linear layers of size 512 with sinusoidal activation functions (Sitzmann et al., 2020). Furthermore, to get the SingleINR to interpolate between the given B-scans instead of overfitting to the given positions, we intentionally use smaller B-scan distances than the real ones so that the input coordinates lie approximately on an isotropic grid.

Since the proposed generalizable INR (GenINR) is the only one of the considered methods that can be used to interpolate and segment unseen test cases, we additionally evaluate the GenINR for the 20 unseen cases from the healthy dataset. For these cases, new latent priors are learned, while the INR weights remain frozen. We repeat this experiment for the pathological data, using 120 OCT and FAF images from 14 patients for training and 26 images from five patients for testing. Here, we use every second B-scan for training, i.e., 13 B-scans per volume, and evaluate performance for the remaining twelve B-scans. Results for the healthy dataset can be found in Tab. 1 and in Tab. 2 for the CSCR dataset. For both datasets, several image similarity metrics, including mean absolute error (MAE), peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), and learned perceptual

Table 2: OCT interpolation and segmentation results for the CSCR data. We compare the proposed generalizable INR with and without FAF input (GenINR<sub>FAF</sub>, GenINR) to linear and registration-based interpolation as well as INRs adapted to single images (SingleINR / GenINR). Results are again reported separately for train and test images. Metrics are stated as the mean $\pm$ std. The best results are highlighted. All INR methods show significant differences in Dice<sub>retina</sub>, Dice<sub>fluid</sub>, and SSIM in comparison to the registration-based interpolation ( $p < 0.001$ , Wilcoxon signed-rank test with Holm-Bonferroni correction).

		Image Reconstruction				Image Segmentation	
Method		$\downarrow$ MAE[%]	$\uparrow$ PSNR	$\uparrow$ SSIM[%]	$\downarrow$ LPIPS	$\uparrow$ Dice <sub>retina</sub> [%]	$\uparrow$ Dice <sub>fluid</sub> [%]
train	linear	6.12 $\pm$ 1.35	21.30 $\pm$ 1.37	39.90 $\pm$ 9.06	<b>0.12<math>\pm</math>0.03</b>	98.05 $\pm$ 0.59	49.54 $\pm$ 23.15
	reg.	5.82 $\pm$ 1.28	21.68 $\pm$ 1.34	42.55 $\pm$ 9.13	0.13 $\pm$ 0.03	98.45 $\pm$ 0.53	44.04 $\pm$ 24.00
	SingleINR	6.86 $\pm$ 1.55	20.23 $\pm$ 1.38	36.83 $\pm$ 9.80	0.29 $\pm$ 0.05	98.49 $\pm$ 0.64	47.52 $\pm$ 23.05
	SingleINR <sub>FAF</sub>	7.54 $\pm$ 1.64	19.32 $\pm$ 1.34	33.69 $\pm$ 9.51	0.31 $\pm$ 0.06	98.08 $\pm$ 0.73	45.94 $\pm$ 23.39
	GenINR	<b>5.33<math>\pm</math>1.24</b>	<b>22.50<math>\pm</math>1.42</b>	<b>49.16<math>\pm</math>9.74</b>	0.55 $\pm$ 0.07	<b>98.99<math>\pm</math>0.48</b>	50.71 $\pm$ 27.92
	GenINR <sub>FAF</sub>	<b>5.34<math>\pm</math>1.25</b>	22.48 $\pm$ 1.42	49.11 $\pm$ 9.74	0.54 $\pm$ 0.07	98.98 $\pm$ 0.47	<b>51.03<math>\pm</math>27.51</b>
test	GenINR	6.28 $\pm$ 1.25	20.25 $\pm$ 1.27	44.45 $\pm$ 9.41	0.57 $\pm$ 0.07	97.53 $\pm$ 1.57	16.84 $\pm$ 25.86
	GenINR <sub>FAF</sub>	6.59 $\pm$ 1.21	20.29 $\pm$ 1.21	44.67 $\pm$ 9.50	0.56 $\pm$ 0.07	97.56 $\pm$ 1.57	17.59 $\pm$ 26.75

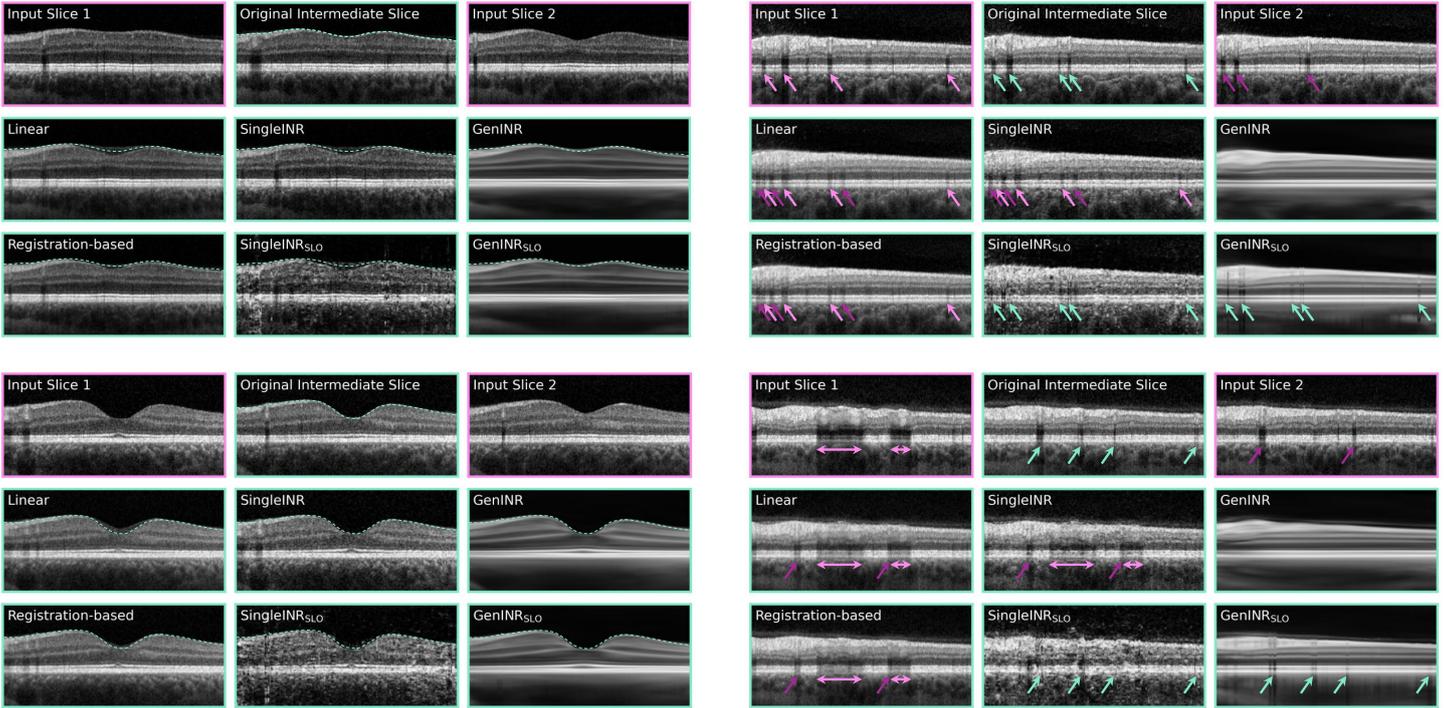


Figure 4: Four examples of interpolated B-scans using linear, registration-based (Ehrhardt et al., 2007), and INR-based interpolation. The input slices for the different methods are shown with pink frames, while the original intermediate and interpolated B-scans are shown with turquoise frames. On the left side, two examples from the foveal region of the retina are depicted, showing interpolation artifacts for instance-based methods, while our GenINR manages to correctly continue the shape of the retina. On the right side, examples from the outer border of the retina are shown. Here, blood vessels are correctly localized by the GenINR with SLO integration (turquoise arrows). For the other methods, vessels can only be propagated from the input slices, again leading to interpolation artifacts.

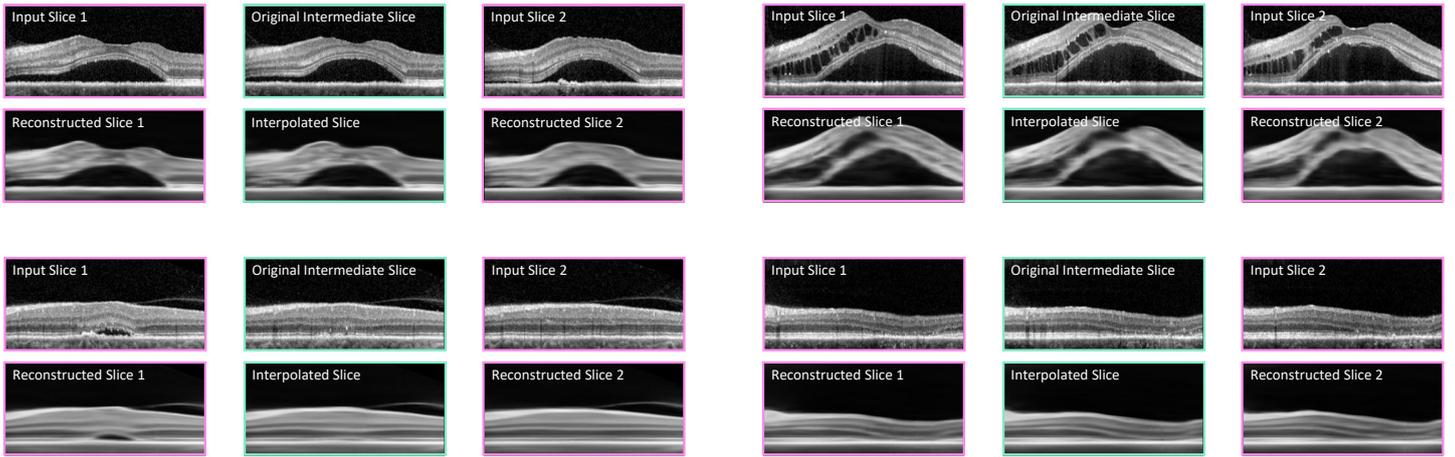


Figure 5: Four examples of interpolated B-scans from a CSCR patient using the proposed GenINR with FAF integration. Subretinal fluid (examples above and below left) is common in CSCR, whereas intraretinal fluid (above right) is seldom). The example below right shows a case with photoreceptor atrophy. In all cases, the retinal shape is reproduced well, but the images appear blurred in the areas of displaced retinal layers and no blood vessel shadows are visible. The color coding is the same as in Fig. 4.

image patch similarity (LPIPS), are reported to assess the image reconstruction and interpolation capability of the considered methods. For a consistent evaluation of both proposed approaches, we used the redefined class labels in which the choroid is discarded and very thin layers are merged. To evaluate the segmentation performance, the Dice score averaged over the nine retinal layers, as well as the average symmetric surface distance (ASSD) and Hausdorff distance (HD) of the layer borders, are reported for the healthy dataset. For the CSCR dataset, we evaluate the segmentation performance with Dice scores of the retina ( $Dice_{retina}$ ) and of the three disease-related fluid accumulations combined into one binary pathology label ( $Dice_{fluid}$ ).

The quantitative results for the healthy dataset show that our proposed GenINR achieves the best results for three out of four image similarity metrics. Only for the LPIPS metric, the GenINRs underperform, while the simple interpolation methods (linear and registration-based) achieve the highest values. Simple interpolation hardly changes the contrast and noise level, whereas the GenINR interpolation results look comparably smooth, missing the noisy and high-frequency components of the images (cf. Fig. 4), which might explain the LPIPS results. Regarding image segmentation, the proposed method performs better than the instance-based methods, indicating a better interpolation of the retinal shape achieved by generalization using population-based training. Especially in the fovea region, where large shape differences can be observed for neighboring B-scans, the GenINR is the only method that can meaningfully represent the shape, as shown in the examples on the left side in Fig. 4. The SingleINR, in turn,

shows a strong overfitting on the slices used for training, hindering it from interpolating the retinal shape correctly. This overfitting also prevents the SingleINR from benefiting from the additional SLO information; in fact, the additional information, which is different at the interpolated positions than at those used for training, causes SingleINR to fail. The GenINR, in turn, is enabled to reconstruct blood vessels, as shown in the examples on the right in Fig. 4. The localization information provided by the en-face SLO enables the correct positioning of blood vessels and the reconstruction of blood vessel shadows in the outer retina. Reconstruction in the fovea region, where fewer blood vessels are located, appears to be more challenging, as no reconstructions of vessels or vessel shadows are notable. Interpolation artifacts from neighboring B-scans in the form of misplaced shadows, as observed in competing methods, do not occur for the GenINR – all observed vessel shadows are correctly placed. Furthermore, the GenINR enables the interpolation and segmentation of unseen test cases, which is not possible with the instance-specific methods. Although a decline in performance for the test cases compared to the training images can be observed, the GenINR still manages to reconstruct the OCT volumes reliably and to represent the retinal layer boundaries with an average ASSD of  $7.0\ \mu\text{m}$ .

Similar results can be observed for the CSCR data. Again, the GenINR demonstrates better performance than the baseline methods both in terms of image reconstruction and segmentation, except for the LPIPS metric. As for the healthy data, the simple linear interpolation delivers the best results for this metric. As observed before, the retinal shape is best aligned with the generalizable INR (average Dice score of 98.98 % for the retina), and also the

pathologies are better interpolated. For the test data, all metrics considered show similar performance to the training data, except for the Dice score of pathologies. Without FAF integration, an average Dice score of 16.84 % is achieved for the interpolated slices of the test images and improved slightly to 17.59 % with the additional FAF modality. Although there is no known clear relationship between FAF and OCT appearance, the INR seems to benefit from the additional modality, indicating a direct correlation between FAF and OCT. The reduced segmentation performance for the test cases indicates that the highly variable appearance of retinal pathologies in CSCR cannot yet be completely covered by the individualization through the latent codes. Some exemplary interpolation results for the CSCR data are shown in Fig. 5. Here, it can be seen that the retinal layers are reconstructed well for B-scans without severe pathological changes. However, no reconstruction of blood vessels can be observed, which we attribute to the distorting effect of pathologies on training. Also, in pathological areas the retinal layers appear more blurred compared to the regions unaffected by fluids. Still, the shape of both the retina and the pathologies is accurately reproduced by the  $\text{GenINR}_{\text{FAF}}$ .

Across all experiments, the runtime for each subject was below one second for linear interpolation, around 19 seconds for registration-based interpolation, approximately three minutes for the  $\text{GenINR}$  (inference adaptation to a new subject), and about five minutes for the  $\text{SingleINR}$ .

## 4.2 Atlas Registration

This section assesses the usage of generalizable INRs for image registration tasks – evaluating how effectively they capture structural consistency within a population. Registration performance and atlas generation capabilities are evaluated using the healthy dataset and compared with two different approaches. First, we compared our  $\text{GenINR}$  framework with the learnable implicit neural atlas approach of Großbröhmer et al. (2024) (SINA), which uses B-spline control points for registration. For the purpose of supervised segmentation, SINA is extended by a second INR that acts as a label atlas, which introduces weak supervision into the atlas generation process. In addition, we modified our proposed method with a non-generalizable INR ( $\text{SingleINR}$ ). Since  $\text{SingleINRs}$  are trained on individual instances, joint optimization of atlas and deformation presents new challenges. Thus, we used the computed voxel-wise median as a fixed atlas that is registered using the inter-subject approach of Wolterink et al. (2022). The  $\text{SingleINR}$  consists of a SIREN with four hidden linear layers of size 128 (Sitzmann et al., 2020). As a reference, we include the initial alignment to the atlas ( $\text{GenINR}_{\text{init}}$ ), and the alignment after affine registration ( $\text{GenINR}_{\text{aff}}$ ).

Intensity similarity is evaluated using the SSIM, while

anatomical alignment is evaluated using the Dice and the ASSD. Deformation’s integrity and plausibility are evaluated using the percentage of negative Jacobian determinants ( $|J_\varphi| \leq 0$ ) and the size of the displacements, as measured by the L1 norm ( $\|u(x, y, z)\|_1$ ). All similarity metrics are evaluated on single B-scans and averaged over each volume. The methods are trained using the healthy dataset with 64 equidistantly sampled B-scans. The experiments are performed using five-fold cross-validation with 80 OCT and segmentation volumes for training and 20 volumes for inference. We conducted four experiments to evaluate the performance of our proposed registration framework. We 1) evaluate the atlas generation capabilities, 2) investigate the registration performance, 3) show that the framework can be used independently of the inference resolution, and 4) perform an ablation study by replacing the implicit atlas representation with an explicit discrete parameter representation.

Our proposed registration framework creates a well-aligned mean representation, i.e. atlas, as indicated by  $\text{GenINR}_{\text{init}}$  in Tab. 3. It achieves scores of 19.4  $\mu\text{m}$ , 76.4 %, and 38.5 % for ASSD, Dice, and SSIM, respectively. These findings are supported by the qualitative results shown in Fig. 7. We consider an atlas to be good when it is sharp – especially on the layer edges, artifact-free, and anatomically plausible. Our framework shows qualitatively sharper and more precise layer separation as well as hardly any artifacts compared to the median representation or SINA. Moreover, SINA exhibits artifacts at the center of the fovea and, in general, a blurrier representation. While the median atlas already shows minimal artifacts and reasonable layer separation, our method achieves higher sharpness and better separation in the upper retinal layers. Although both generated atlases show a consistent alignment between class labels and intensity representations, our method yields more anatomically plausible results, particularly in the innermost retinal layers, with a continuous representation of the RNFL across the retina.

Focusing on the results of the atlas-to-subject registration experiment in Tab. 3, our approach achieved the best results by about 11.3  $\mu\text{m}$ , 86.4 %, and 42.4 % for ASSD, Dice, and SSIM, respectively. Significant improvements are shown in Fig. 6 for  $\text{GenINR}$  over the initial alignment, affine alignment, and the two approaches that were compared: SINA and  $\text{SingleINR}$ . Meanwhile, the time taken for inference execution is reduced by around four minutes per subject compared to SINA and around one minute per subject compared to the  $\text{SingleINR}$  – resulting in faster processing at around 23.1 s per subject. However, SINA presents a significant deterioration in the segmentation metrics compared to the base alignment ( $\text{GenINR}_{\text{init}}$ ), indicating an insufficient number of control points of the interpolated displacement field. SINA and  $\text{GenINR}$  produce smooth

Table 3: Quantitative performance comparison of the atlas registration methodology. Similarity metrics, deformation metrics, and the inference time per subject are reported. GenINR<sub>init</sub> shows the initial alignment of the generated atlas without registration. Metrics are stated as the mean<sub>±std</sub>. The best results are highlighted.

Method	Similarity			Deformation		
	↓ ASSD[ $\mu\text{m}$ ]	↑ Dice[%]	↑ SSIM[%]	↓ $ J_\varphi  \leq 0$ [%]	↓ $\ u(x, y, z)\ _1$	↓ Time[s]
GenINR <sub>init</sub>	19.4 $\pm$ 5.8	76.4 $\pm$ 7.9	38.5 $\pm$ 4.6	—	—	—
GenINR <sub>aff</sub>	14.3 $\pm$ 3.3	82.9 $\pm$ 4.3	41.7 $\pm$ 3.7	—	—	—
SINA	29.8 $\pm$ 11.0	65.1 $\pm$ 10.3	35.1 $\pm$ 4.9	<b>0.0</b> $\pm$ 0.0	0.009 $\pm$ 0.003	279.3 $\pm$ 32.52
SingleINR	12.8 $\pm$ 2.3	84.3 $\pm$ 2.9	41.9 $\pm$ 4.0	0.001 $\pm$ 0.007	<b>0.005</b> $\pm$ 0.003	97.3 $\pm$ 4.6
GenINR	<b>11.3</b> $\pm$ 1.6	<b>86.4</b> $\pm$ 2.3	<b>42.4</b> $\pm$ 3.6	<b>0.0</b> $\pm$ 0.0	0.04 $\pm$ 0.011	<b>23.1</b> $\pm$ 0.85

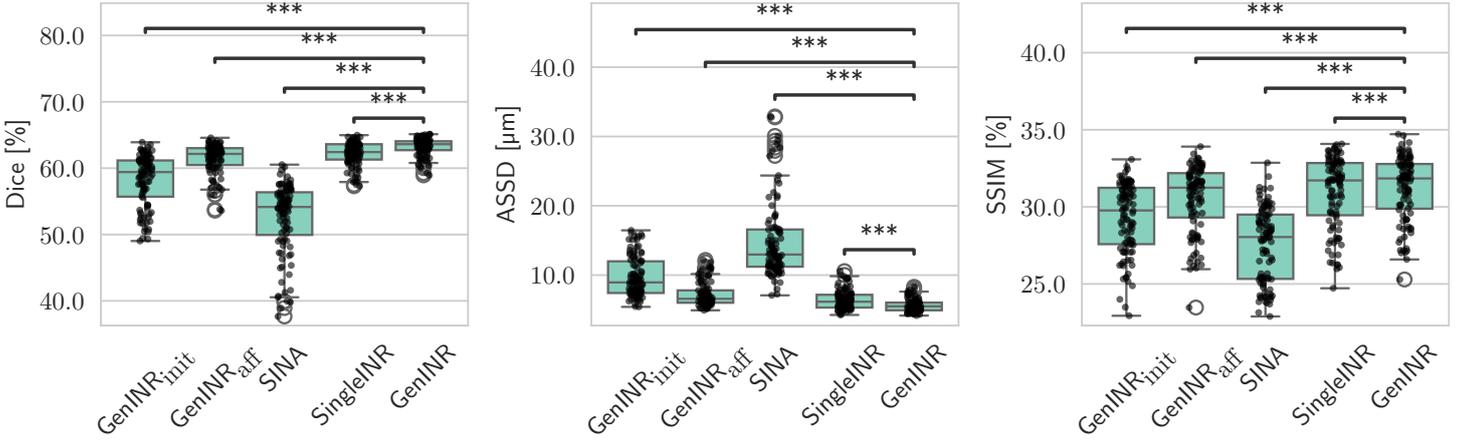


Figure 6: Comparison of the population-based atlas registration methods across all similarity metrics. Statistical significance is indicated using stars (\*\*\*:  $p < 0.001$ ), determined via the Wilcoxon signed-rank test and Holm-Bonferroni correction using the Python package by Charlier et al. (2022).

and fold-free deformations without the need for additional regularization. SingleINR exhibits a slight percentage of folding. Fig. 8 presents the anatomical alignment results using the separated ASSD metric for all layers. Due to pre-processing, the bottom layers (RPE, OPR & SVS, and IS/OS) are already well aligned initially and challenge the methods to generate a proper transition from small to larger displacements. Despite this, our framework achieved a uniform performance distributed across all layers, whereas the SingleINR struggled with the upper middle layers GCL, IPL, and INL.

Due to the continuous nature of the INR and generalizable INR, the proposed framework is independent of the discrete grid representation and the associated resolution. Therefore, during inference, the latent codes of the trained atlas framework are learned and evaluated on 16, 32, 64, and 512 B-scans. Fig. 9 illustrates no significant changes ( $p > 0.01$ ) in the alignment performance when compared with the resolution of 64 B-scans and, in general, similar performance scores. All setups result in around 85.2%, 12.1  $\mu\text{m}$ , and 42% for Dice, ASSD, and SSIM, respectively.

We conducted an ablation study to illustrate the performance gains of the implicit atlas representation. To do so,

we replaced the INR with two learnable parameter images for intensities and class labels, which are then transformed using the deformation field and linear interpolation to be compared with the target volume. This explicit atlas representation scales  $\mathcal{O}(n^3)$  with the size of the three dimensions  $n$ , which limited the experiment to a resolution of 16 B-scans. The results are shown in Fig. 4 and support the advantage of implicit atlas representation by even improving alignment performance across all segmentation-based metrics. It should be noted that the training time refers to the latent prior training during inference.

Table 4: Results of the ablation study conducted on 16 B-scans. Similarity metrics, deformation metrics, and the inference time per subject are reported. Metrics are stated as the mean<sub>±std</sub>. The best results are highlighted.

Atlas	↓ ASSD[ $\mu\text{m}$ ]	↑ Dice[%]	↑ SSIM[%]	↑ Time[s]
Explicit	11.2 $\pm$ 1.8	86.1 $\pm$ 2.6	<b>42.7</b> $\pm$ 3.9	15.5 $\pm$ 0.63
Implicit	<b>10.9</b> $\pm$ 2.1	<b>86.7</b> $\pm$ 3.1	41.8 $\pm$ 4.7	<b>10.6</b> $\pm$ 0.85

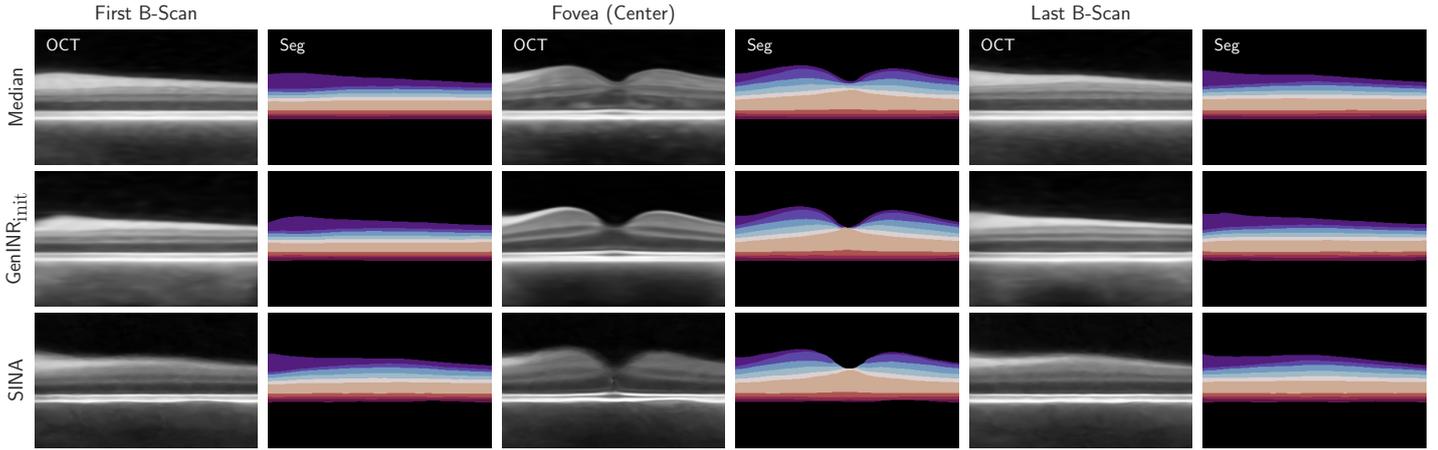


Figure 7: Three examples of B-scans located at the center and each side of the volume. Compared are the voxel-wise median, our learned atlas, and the learned atlas of SINA. For each B-scan, the intensity values and the segmentation map are displayed.

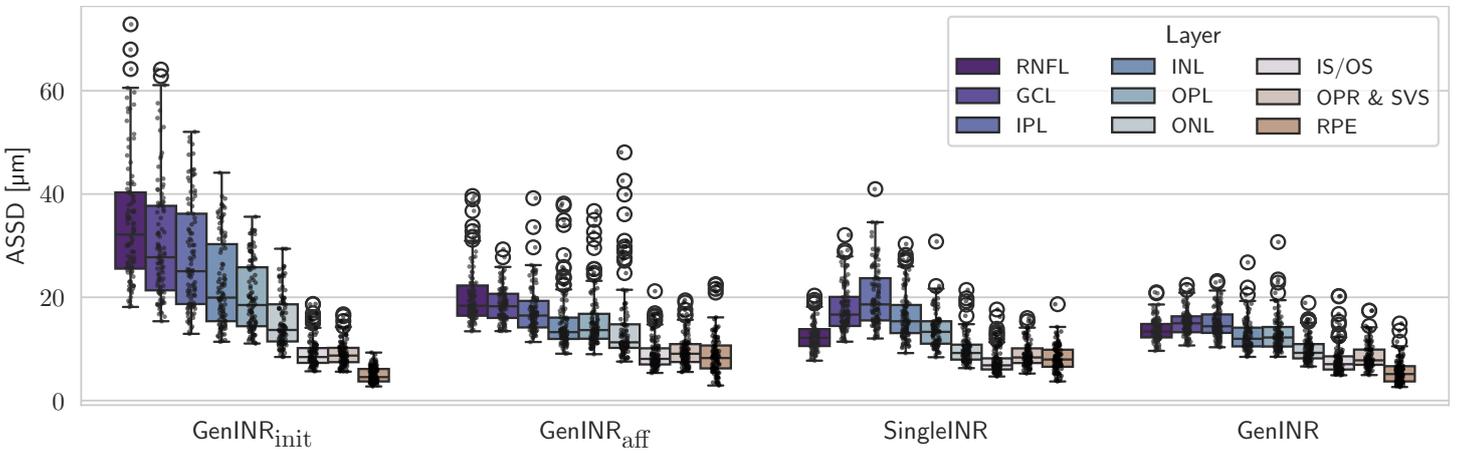


Figure 8: Detailed performance overview of the SingleINR and GenINR on the example of ASSD for each retinal layer: retinal nerve fiber layer (RNFL), ganglion cell layer (GCL), inner plexiform layer (IPL), inner nuclear layer (INL), outer plexiform layer (OPL), outer nuclear layer (ONL), IS/OS junction (IS/OS), outer photoreceptor segment (OPR), subretinal virtual space (SVS), and retinal pigment epithelium (RPE).

### 4.3 Latent Space Analysis

To gain a better insight into the INRs, we perform a principal component analysis (PCA) on the latent priors for both approaches. Fig. 10 shows the result of the PCA analysis of the first two principal components for the inter-B-scan interpolation – explaining about 20 % of the variance – as well as a t-SNE visualization of the training latent space. While the latent priors for the training data are more compactly centered around 0, the latent priors for the test cases deviate quite heavily from the training cases, showing that individual solutions are found for the test data rather than the most similar training case. The deviation is even stronger for the CSCR data, possibly due to a stronger variability in the data. Furthermore, training cases of the same patient appear to have grouped latent priors, indicating that individual structural differences are coded efficiently in

the vector.

For the registration approach, we analyzed the latent priors and their relation to the deformation. We again performed a PCA on the latent priors and related the first two principal components (explaining around 13 % of the variance) with the mean Jacobian determinant that is computed using the retina as a mask. Fig. 11 shows the results, indicating a significant correlation between the mean Jacobian determinant – which quantifies the volume changes caused by the deformation – and the second principal component of the latent space.

## 5. Discussion

In this work, we examined the usage of INRs for highly anisotropic OCT data. By representing the OCT volume as a continuous function, we obtain a volume of arbitrary

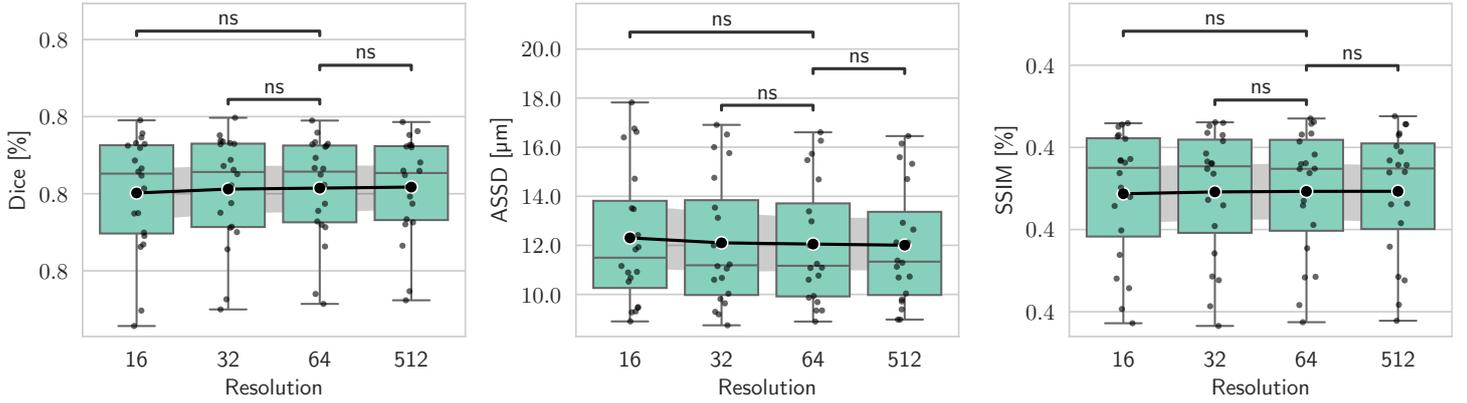


Figure 9: Results of our GenINR during inference – trained and evaluated on different numbers of B-scans. No statistical significance is indicated using (ns: not significant,  $p \geq 0.01$ ), determined via the Wilcoxon signed-rank test and Holm-Bonferroni correction using the Python package by Charlier et al. (2022).

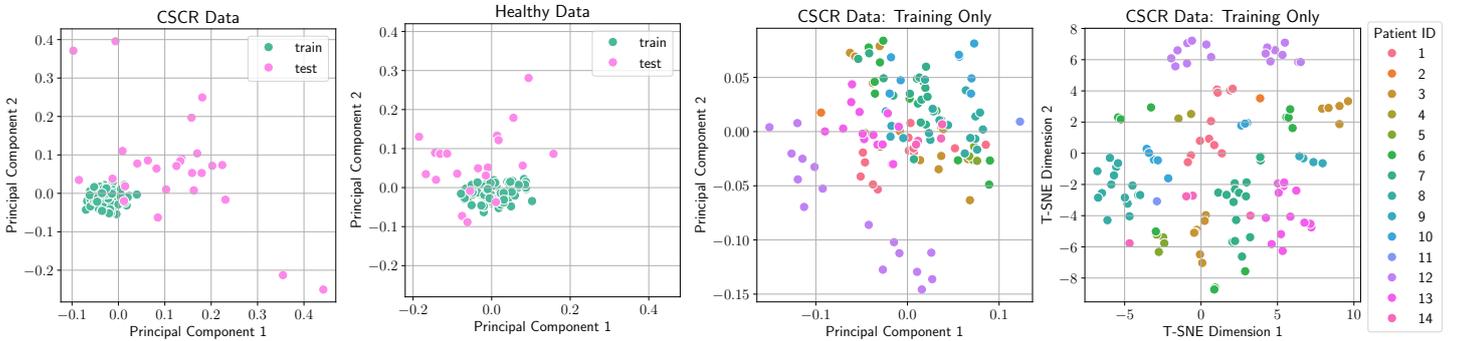


Figure 10: PCA analysis of the learned latent codes of the interpolation INR (GenINR<sub>SLO</sub> for the healthy data and GenINR<sub>FAF</sub> for the CSCR data). On the left-hand side, the PCA for the latent codes of training and test cases is performed together. The two plots on the right show a PCA and t-SNE performed exclusively on the training images of GenINR<sub>FAF</sub> and are color-coded by subject.

resolution. Our results prove that the reconstructed space between B-scans carries meaningful information. This way, we overcome the limitation of established approaches.

Our proposed interpolation method achieves better performance than the considered competitive approaches. For healthy cases, it can interpolate the retinal layer shape most accurately while using even larger slice distances than clinical OCT volumes. On pathological data, we were able to repeat good OCT B-scan interpolation performance while retaining pathologies within these interpolated B-scans. Our approach is mainly limited by the low-frequency reconstruction by the generalizable INR, leaving a disparity in intensity reconstruction. Despite that, the generalizable INR enables the incorporation of en-face modality information and segmentation prediction of unseen OCT volumes. This in turn results in better reconstruction of retina shape, vessel structures, and fluid positions in interpolated slices. Further improvement might be gained by incorporating sub-pixel-level layer boundaries instead of the voxel-wise segmentation used in this work. Shape-based interpolation methods similarly make use of segmentations and boundaries and are

well-suited for interpolating label maps, but do not address the B-scan interpolation task because intensity information is not modeled. In contrast, our approach generates super-sampled OCT data with corresponding labels, which can benefit various downstream applications, such as volumetric segmentation, 3D reconstruction, longitudinal analyses or image registration, as in Sec. 4.2. Demonstrating direct clinical benefit remains an important direction for future work and will be addressed in subsequent studies.

To provide further analysis and understanding of the population, we built upon the previous findings and proposed a similar framework for registration and atlas generation. We combine methods from Großbröhmer et al. (2024) and Wolterink et al. (2022) with a generalization approach to obtain a pipeline that enables the joint prediction of a resolution-agnostic atlas and corresponding deformation fields. Großbröhmer et al. (2024) proposed an INR as an atlas representation but used a simple registration based on B-spline control points. This instance-based B-spline registration showed good performance on other datasets but – as shown in this study – scales poorly to

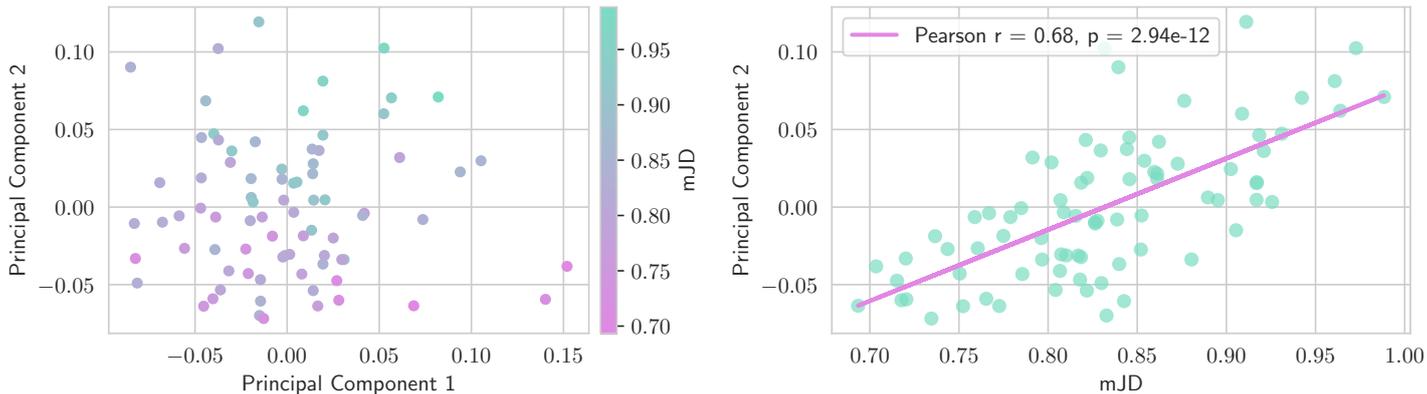


Figure 11: Exploration of the latent space using PCA and the masked mean of the Jacobian Determinant (mJD) as a metric to describe the general volume change. On the left side, the first two principal components are displayed with the mJD color-coded. On the right, the correlation between mJD and the second principal component is plotted.

the challenges associated with OCT. Implicit deformable image registration (Wolterink et al., 2022) on the other hand, exhibited good alignment performance with close to no foldings on our OCT data. To resolve the restriction of the SingleINR to prior knowledge, we proposed a generalizable INR registration framework to include population-based knowledge and auxiliary information (i.e., retinal layer segmentation) in a joint atlas and deformation optimization process. Thus, we enable the creation of a sharp and accurate atlas representation as well as a continuous representation of the deformation. The resulting method is, by construction, resolution-agnostic and can be used on volumes containing varying numbers of B-scans. Our combined framework improved the alignment and representation performance, reduced the execution time during inference, and reduced the deformation field representation to a single lower-dimensional vector. In line with the findings of Tian et al. (2024), we show that generalizable INRs present an efficient approach to model the registration process. Since smoothness of the deformation field is intended, the drawback of low-frequency reconstruction observed for interpolation does not transfer to registration.

Altogether, INRs simplify the handling and processing of high-dimensional image data. Especially for image registration, INRs seem to cause a trend for new approaches (Wolterink et al., 2022; Sideri-Lampretsa et al., 2024; Van Harten et al., 2023; Tian et al., 2024; Großbröhmer et al., 2024). When the image is represented by an INR, it eliminates the need for linear interpolation in the registration process while maintaining differentiability. Furthermore, implicitly representing data as functions reduces memory and runtime requirements when compared to other iterative approaches. This is particularly the case once the INR is combined with a generalization approach and the representation is compressed into the latent priors. With this, new potentials and analysis options arise, as we demonstrated

for our proposed methods.

## 6. Conclusion

This work explored the potential of INRs in a highly anisotropic domain such as retinal OCT imaging. We proposed two methods for interpolation and atlas registration using generalizable INRs. Our findings highlight the power of INRs as a resolution-agnostic interpolation or registration tool. With both approaches, we were able to generate more accurate reconstructions than the considered baselines. Our approach overcomes limitations regarding the resolution introduced by device-specific sampling and enables the fusion of different modalities.

Generalizable INRs offer the advantage of low-dimensional data representations due to latent priors. These latent vectors provide a more efficient way to store the data and also enable the analysis of the underlying structures. That way, we visualized relations in the latent space – suggesting new use cases such as classification or anomaly detection. In addition, further modalities such as OCT-angiography or fluorescence lifetime imaging ophthalmoscopy may be evaluated. In summary, INRs, and generalizable INRs in particular, are a powerful tool for medical image analysis. They highlight the value of continuous representation by enabling a resolution-agnostic interpolation and atlas registration framework.

## Acknowledgments

This work was supported by funding from the federal state of Schleswig-Holstein, Germany (grant 22024003).

## Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of animals or human subjects.

## Conflicts of Interest

We declare we don't have conflicts of interest.

## Data availability

The CSCR dataset was acquired during clinical routine and is not publicly available due to data protection regulations. The in-house dataset of healthy volunteers is currently being prepared for public release. Further details will be provided in the GitHub repository after release.

## References

- Tamaz Amiranashvili, David Lüdke, Hongwei Bran Li, Björn Menze, and Stefan Zachow. Learning shape reconstruction from sparse measurements with neural implicit functions. In *International Conference on Medical Imaging with Deep Learning*, pages 22–34. PMLR, 2022.
- Josefine Britze and Jette Laurrup Frederiksen. Optical coherence tomography in multiple sclerosis. *Eye*, 32(5): 884–888, 2018.
- Arunava Chakravarty, Divya Jyothi Gaddipati, and Jayanthi Sivaswamy. Construction of a retinal atlas for macular oct volumes. In *Image Analysis and Recognition: 15th International Conference, ICIAR 2018, Póvoa de Varzim, Portugal, June 27–29, 2018, Proceedings 15*, pages 650–658. Springer, 2018.
- Florian Charlier, Marc Weber, Dariusz Izak, Emerson Harkin, Marcin Magnus, Joseph Lalli, Louison Fresnais, Matt Chan, Nikolay Markov, Oren Amsalem, Sebastian Proost, Agamemnon Krasoulis, getzze, and Stefan Replinger. Statannotations, October 2022.
- Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5939–5948, 2019.
- Jay Chhablani, Francine Behar Cohen, Pauline Aymard, Talal Beydoun, Elodie Bousquet, Alejandra Daruich-Matet, Alexandre Matet, Min Zhao, Chui Ming Gemmy Cheung, K Bailey Freund, et al. Multimodal imaging-based central serous chorioretinopathy classification. *Ophthalmology Retina*, 4(11):1043–1046, 2020.
- Joseph Paul Cohen, Margaux Luck, and Sina Honari. Distribution matching losses can hallucinate features in medical image translation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part I*, pages 529–536. Springer, 2018.
- Maik Dannecker, Vanessa Kyriakopoulou, Lucilio Cordero-Grande, Anthony N Price, Joseph V Hajnal, and Daniel Rueckert. CINA: Conditional implicit neural atlas for spatio-temporal representation of fetal brains. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 181–191. Springer, 2024.
- Vineeta Das, Samarendra Dandapat, and Prabin Kumar Bora. Unsupervised super-resolution of OCT images using generative adversarial network for improved age-related macular degeneration diagnosis. *IEEE Sensors Journal*, 20(15):8746–8756, 2020.
- Emilien Dupont, Hyunjik Kim, SM Ali Eslami, Danilo Jimenez Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you can treat it like one. In *International Conference on Machine Learning*, pages 5694–5725. PMLR, 2022.
- Jan Ehrhardt, Dennis Säring, and Heinz Handels. Structure-preserving interpolation of temporal and spatial image sequences using an optical flow-based method. *Methods of Information in Medicine*, 46(03):300–307, 2007.
- Wei Fang, Yuxing Tang, Heng Guo, Mingze Yuan, Tony CW Mok, Ke Yan, Jiawen Yao, Xin Chen, Zaiyi Liu, Le Lu, et al. CycleNR: Cycle implicit neural representation for arbitrary-scale volumetric super-resolution of medical data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11631–11641, 2024.
- Christoph Großbröhmer, Ziad Al-Haj Hemidi, Fenja Falta, and Mattias P Heinrich. SINA: Sharp implicit neural atlases by joint optimisation of representation and deformation. In *International Workshop on Biomedical Image Registration*, pages 165–180. Springer, 2024.
- Yufan He, Aaron Carass, Yihao Liu, Bruno M. Jedynek, Sharon D. Solomon, Shiv Saidha, Peter A. Calabresi, and Jerry L. Prince. Fully convolutional boundary regression for retina oct segmentation. In Dinggang Shen, Tianming Liu, Terry M. Peters, Lawrence H. Staib, Caroline Essert, Sean Zhou, Pew-Thian Yap, and Ali Khan, editors,

- Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, pages 120–128, Cham, 2019. Springer International Publishing.
- David Huang, Eric A Swanson, Charles P Lin, Joel S Schuman, William G Stinson, Warren Chang, Michael R Hee, Thomas Flotte, Kenton Gregory, Carmen A Puliafito, et al. Optical coherence tomography. *Science*, 254(5035): 1178–1181, 1991.
- Yongqiang Huang, Zexin Lu, Zhimin Shao, Maosong Ran, Jiliu Zhou, Leyuan Fang, and Yi Zhang. Simultaneous denoising and super-resolution of optical coherence tomography images based on generative adversarial network. *Optics Express*, 27(9):12289–12307, 2019.
- Amirhossein Kazerouni, Reza Azad, Alireza Hosseini, Dorit Merhof, and Ulas Bagci. INCODE: Implicit neural conditioning with prior knowledge embeddings. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1298–1307, 2024.
- Timo Kepp, Jan Ehrhardt, Mattias P. Heinrich, Gereon Hüttmann, and Heinz Handels. Topology-preserving shape-based regression of retinal layers in oct image data using convolutional neural networks. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 1437–1440. IEEE, 2019.
- Timo Kepp, Julia Andresen, Claus von der Burchard, Johann Roider, Gereon Hüttmann, and Heinz Handels. Shape-based segmentation of retinal layers and fluids in oct image data. In *Medical Imaging 2023: Computer-Aided Diagnosis*, volume 12465, pages 207–216. SPIE, 2023.
- Timo Kepp, Julia Andresen, Fenja Falta, and Heinz Handels. Bridging gaps in retinal imaging: Fusing OCT and SLO information with implicit neural representations for improved interpolation and segmentation. In *BVM Workshop*, pages 107–112. Springer, 2025.
- Maziyar M Khansari, Jiong Zhang, Yuchuan Qiao, Jin Kyu Gahm, Mona Sharifi Sarabi, Amir H Kashani, and Yonggang Shi. Automated deformation-based analysis of 3D optical coherence tomography in diabetic retinopathy. *IEEE transactions on medical imaging*, 39(1):236–245, 2019.
- Chiheon Kim, Doyup Lee, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Generalizable implicit neural representations via instance pattern composers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11808–11817, 2023.
- Sieun Lee, Nicolas Charon, Benjamin Charlier, Karteek Popuri, Evgeniy Lebed, Marinko V Sarunic, Alain Trouvé, and Mirza Faisal Beg. Atlas-based shape analysis and classification of retinal optical coherence tomography images using the functional shape (fshape) framework. *Medical image analysis*, 35:570–581, 2017.
- Kang Li, Xiaodong Wu, Danny Z Chen, and Milan Sonka. Optimal surface segmentation in volumetric images—a graph-theoretic approach. *IEEE transactions on pattern analysis and machine intelligence*, 28(1):119–134, 2006.
- Qiaoliang Li, Shiyu Li, Zhuoying He, Huimin Guan, Runmin Chen, Ying Xu, Tao Wang, Suwen Qi, Jun Mei, and Wei Wang. DeepRetina: layer segmentation of retina in OCT images using deep learning. *Translational vision science & technology*, 9(2):61–61, 2020.
- Yangxi Li, Chuanhao Zhang, Tianqi Huang, Yingwei Fan, Guochen Ning, and Hongen Liao. Computational multi-angle optical coherence tomography using implicit neural representation. *Optics & Laser Technology*, 184:112551, 2025.
- Anne-Sofie Wessel Lindberg, Thomas Martini Jørgensen, and Vedrana Andersen Dahl. Linear, transfinite and weighted method for interpolation from grid lines applied to OCT images. *Applied Soft Computing*, 68:293–302, 2018.
- Hong Liu, Dong Wei, Donghuan Lu, Xiaoying Tang, Liansheng Wang, and Yefeng Zheng. Simultaneous alignment and surface regression using hybrid 2D–3D networks for 3D coherent layer segmentation of retinal oct images with full and sparse annotations. *Medical Image Analysis*, 91:103019, 2024.
- Emilio López-Varela, Noelia Barreira, Nuria Olivier Pascual, Maria Rosa Arroyo Castillo, and Manuel G. Penedo. Generation of synthetic intermediate slices in 3d oct cubes for improving pathology detection and monitoring. *Computers in Biology and Medicine*, 163:107214, 2023. ISSN 0010-4825. . URL <https://www.sciencedirect.com/science/article/pii/S0010482523006790>.
- Julian McGinnis, Suprosanna Shit, Hongwei Bran Li, Vasiliki Sideri-Lampretsa, Robert Graf, Maik Dannecker, Jiazhen Pan, Nil Stolt-Ansó, Mark Mühlau, Jan S. Kirschke, Daniel Rueckert, and Benedikt Wiestler. Single-subject multi-contrast mri super-resolution via implicit neural representations. In Hayit Greenspan, Anant Madabhushi, Parvin Mousavi, Septimiu Salcudean, James Duncan, Tanveer Syeda-Mahmood, and Russell Taylor, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, pages 173–183, Cham, 2023. Springer Nature Switzerland.

- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106, 2021.
- Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019.
- M. Pekala, N. Joshi, T.Y. Alvin Liu, N.M. Bressler, D. Cabrera DeBuc, and P. Burlina. Deep learning based retinal OCT segmentation. *Computers in Biology and Medicine*, 114:103445, 2019. ISSN 0010-4825. . URL <https://www.sciencedirect.com/science/article/pii/S0010482519303221>.
- Monty Santarossa, Ayse Kilic, Claus von der Burchard, Lars Schmarje, Claudius Zelenka, Stefan Reinhold, Reinhard Koch, and Johann Roider. MedRegNet: Unsupervised multimodal retinal-image registration with GANs and ranking loss. In *Medical Imaging 2022: Image Processing*, volume 12032, pages 321–333. SPIE, 2022a.
- Monty Santarossa, Ayse Tatli, Claus von der Burchard, Julia Andresen, Johann Roider, Heinz Handels, and Reinhard Koch. Chronological registration of oct and autofluorescence findings in CSCR: Two distinct patterns in disease course. *Diagnostics*, 12(8):1780, 2022b.
- Monty Santarossa, Reinhard Koch, and Meysam Tavakoli. Robust multimodal retinal image registration in diabetic retinopathy using a light-weight neural network and improved RANSAC algorithm. *IEEE Sensors Journal*, 2025.
- Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G Baraniuk. Wire: Wavelet implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18507–18516, 2023.
- U Schmidt-Erfurth, S Klimescha, SM Waldstein, and H Bogunović. A view of the current and future role of optical coherence tomography in the management of age-related macular degeneration. *Eye*, 31(1):26–44, 2017.
- Vasiliki Sideri-Lampretsa, Julian McGinnis, Huaqi Qiu, Magdalini Paschali, Walter Simson, and Daniel Rueckert. SINR: Spline-enhanced implicit neural representation for multi-modal registration. *Medical Imaging with Deep Learning*, pages 1462–1474, 2024.
- Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.
- Nil Stolt-Ansó, Julian McGinnis, Jiazhen Pan, Kerstin Hammernik, and Daniel Rueckert. NISF: Neural implicit segmentation functions. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 734–744. Springer, 2023.
- Lin Tian, Hastings Greer, Raúl San José Estépar, Roni Sengupta, and Marc Niethammer. NePhi: Neural deformation fields for approximately diffeomorphic medical image registration. In *European Conference on Computer Vision*, pages 213–237. Springer, 2024.
- Louis D Van Harten, Jaap Stoker, and Ivana Išgum. Robust deformable image registration using cycle-consistent implicit representations. *IEEE transactions on medical imaging*, 43(2):784–793, 2023.
- SB Velaga, MG Nittala, RK Konduru, F Heussen, PA Keane, and SR Sadda. Impact of optical coherence tomography scanning density on quantitative analyses in neovascular age-related macular degeneration. *Eye*, 31(1):53–61, 2017.
- Swetha Bindu Velaga, Muneeswar G Nittala, Amir Hariri, and Srinivas R Sadda. Correlation between fundus autofluorescence and en face oct measurements of geographic atrophy. *Ophthalmology Retina*, 6(8):676–683, 2022.
- Gianni Virgili, Francesca Menchini, Giovanni Casazza, Ruth Hogg, Radha R Das, Xue Wang, and Manuele Michelessi. Optical coherence tomography (OCT) for detection of macular oedema in patients with diabetic retinopathy. *Cochrane Database of Systematic Reviews*, 1, 2015.
- Claus von der Burchard, Johann Roider, and Timo Kepp. Analysis of oct scanning parameters in AMD and RVO. *Diagnostics*, 14(5):516, 2024.
- Jelmer M Wolterink, Jesse C Zwienenberg, and Christoph Brune. Implicit neural representations for deformable image registration. In *International Conference on medical imaging with deep learning*, pages 1349–1359. PMLR, 2022.
- Qing Wu, Yuwei Li, Yawen Sun, Yan Zhou, Hongjiang Wei, Jingyi Yu, and Yuyao Zhang. An arbitrary scale super-resolution approach for 3D MR images via implicit neural representation. *IEEE Journal of Biomedical and Health Informatics*, 27(2):1004–1015, 2023.

Youshen Xiao, Yuting Shen, Sheng Liao, Bowei Yao, Xiran Cai, Yuyao Zhang, and Fei Gao. Limited-view photoacoustic imaging reconstruction via high-quality self-supervised neural representation. *Photoacoustics*, page 100685, 2025.

Zhuoqun Yuan, Di Yang, Weike Wang, Jingzhu Zhao, and Yanmei Liang. Self super-resolution of optical coherence tomography images based on deep learning. *Optics Express*, 31(17):27566–27581, 2023.