# Deep Probabilistic Cerebrovascular Atlas

Xiaoming Zhang[1], Maik Dannecker[2], Jon Cleary[3], Daniele Falcetta[1], Kaiyuan Yang[4], Bjoern Menze[4], Yu Shi[5,6], Susanne Wegener[7], Sebastien Ourselin[3], Daniel Rueckert[2,8], and Maria A. Zuluaga[1,3]

[1] EURECOM, Sophia Antipolis, France
[2] Chair for AI in Healthcare and Medicine, Technical University of Munich (TUM) and TUM University Hospital, Munich, Germany
[3] School of Biomedical Engineering & Imaging Sciences, King's College London, UK
[4] Department for Quantitative Biomedicine, University of Zurich, Switzerland
[5] Department of Radiology, Shengjing Hospital of China Medical University, Shenyang, China
[6] Department of Radiology, the Second Affiliated Hospital, Zhejiang University School of Medicine, Hangzhou, China
[7] Department of Neurology, University Hospital Zurich, University of Zurich, Switzerland
[8] Department of Computing, Imperial College London, London, UK
zxiaoming360@gmail.com, maria.zuluaga@eurecom.fr

**Abstract.** Cerebrovascular modeling remains a foundational challenge in medical image analysis due to the intricate, multi-scale topology and extreme anatomical variability of the brain's vascular tree. Traditional discrete representations, such as voxel grids, face challenges in capturing the continuous nature of these highly variable structures, often leading to spatial wash-out during population averaging. In this study, we introduce VITAL, an implicit neural atlas framework for brain vessel modeling. To address the multi-scale nature of the cerebrovasculature, spanning from major arteries to distal branches, VITAL utilizes a dynamically routed Mixture-of-Experts architecture that decouples multi-scale features, enabling effective and continuous vascular reconstruction. To our knowledge, we present the first continuous conditioned probabilistic neural atlas of the human cerebrovasculature, enabling cross-sectional modeling of age-related morphological changes and sex-stratified vascular differences across the adult lifespan. Furthermore, we evaluate VITAL on subjects with stroke-related neurological disorders, demonstrating its potential as an exploratory normative reference for anomaly analysis.

**Keywords:** Cerebrovasculature · Conditioned atlas · MR

## 1 Introduction

The cerebrovascular system is essential to brain function, yet our understanding of it remains limited compared with that of other organs that are more extensively studied. Despite its critical role in oxygen delivery and waste removal,

the complexity of the cerebrovascular tree, characterized by an intricate hierarchy ranging from the large, high-flow arteries of the Circle of Willis (CoW) to the microscopic, tortuous capillary beds, has historically defied comprehensive mapping. Methodological challenges hinder the development of neurovascular atlases at large scale, limiting insights into the role of the neurovascular tree in pathology and aging. Consequently, the mechanisms underlying widespread conditions such as cerebral small vessel disease, vascular cognitive impairment, and the progression of neurodegenerative disorders remain partially obscured by the lack of high-fidelity, population-level anatomical references.

For decades, the construction of vessel atlases has been pursued as a critical tool to bridge this gap [13,18]. Ideally, such atlases would serve as canonical representations of vascular anatomy, facilitating cross-population analysis, supporting the training of automated segmentation algorithms [27], and enabling the quantification of pathological deviations (e.g., ischemic stroke, intracranial aneurysms). However, traditional cerebrovascular atlases have relied on explicit, discrete geometric representations that scale poorly to the biological reality. Discretization of continuous vascular CT or MRI signals onto Cartesian grids introduces partial volume effects and aliasing, especially for vessels approaching the scanner resolution limit [28]. Similarly, mesh-based representations and centerline methods, despite their sub-voxel precision, struggle to capture the complex and variable topology of vascular networks, where bifurcations, loops, and anastomoses vary significantly across patients. These discretization bottlenecks fundamentally limit the ability to model the continuous, multi-scale nature of the brain's vascular system.

The paradigm of Implicit Neural Representations (INRs) has emerged recently as a promising alternative [19,25], moving away from discrete data structures toward continuous, functionally parameterized deep representations. Specifically, INRs encode anatomical structures not as lists of coordinates, but as the weights of a neural network that learns a continuous mapping from spatial position to a signal of interest [2] (e.g., occupancy probability or signed distance fields). This continuous parameterization decouples anatomical representation from input data resolution, theoretically enabling structure and topology recovery at arbitrary spatial scales without the stair-casing artifacts inherent to discrete grids. Furthermore, by conditioning the network on subject-specific latent codes, INRs provide a strong framework for deformable atlas construction, enabling the joint learning of a shared canonical anatomy and individual variations crucial for precision medicine [3,23].

Despite the theoretical advantages of INRs, standard coordinate-based networks (e.g., MLPs, SIREN [26]) struggle to scale to the complexity of whole-brain vasculature. Although networks such as SIREN [26] mitigate spectral bias through periodic activations, a single global parameter set must still represent both low-frequency geometry of major arteries and high-frequency details of distal branches. This capacity competition can lead to over-smoothed representations of fine vascular structures, limiting the fidelity of continuous cerebrovascular atlas construction [15,29].

In this work, we propose the Vascular Implicit Atlas (VITAL). To overcome the representational constraints of monolithic architectures, VITAL introduces a dynamically routed Mixture-of-Experts (MoE) mechanism [5,31] that disentangles the learning of scale-specific patterns of the vascular tree. This sparse activation mechanism enables the network to simultaneously capture the low-frequency geometry of major arteries and the high-frequency topology of distal branches without incurring prohibitive computational costs.

In summary, we present the first continuous, probabilistic spatio-temporal neural atlas of the human cerebrovascular system, built using an adaptive MoE-INR architecture that addresses spectral bias in multi-scale vascular modeling. Our approach enables generation of demographic-conditioned atlases through latent space interpolation and provides a normative reference for localizing cerebrovascular anomalies in subjects with stroke-related neurological disorders.

## 2 Method

### 2.1 Vascular Implicit Atlas

Building upon the CINeMA framework [9], we formulate the construction of the demographic-conditioned vessel atlas as learning a continuous implicit function $f_\Theta$, parameterized by a deep neural network with weights $\Theta$ (Fig. 1). Unlike discrete voxel grids, $f_\Theta$ maps a continuous spatial coordinate $\mathbf{x} \in \mathbb{R}^3$ to a MR intensity value $I(\mathbf{x})$ and a vessel probability $P(\mathbf{x})$. To account for the large anatomical variability across the population, we condition this representation on a subject-specific latent code $\mathbf{z}_i \in \mathbb{R}^D$ and demographic covariates $\mathbf{c}_i$ (e.g., age, sex).

To align individual anatomies into a shared canonical space without explicit pre-registration, we define a learnable spatial transformation module. Similarly to [9], for a subject $i$, a coordinate $\mathbf{x}_s$ sampled from the subject space $\Omega_{\text{Subject}}$ is mapped to the canonical atlas space $\Omega_{\text{Atlas}}$ via an affine transformation via $\mathbf{x}a = R_i\mathbf{x}_s + \mathbf{t}i$, where $\{R_i, \mathbf{t}_i\}$ represents the learnable rotation and translation parameters unique to subject $i$. This creates a deformable atlas framework where the network learns the canonical structure in $\Omega_{\text{Atlas}}$.

Representing vascular structures across multiple scales poses a fundamental challenge for standard coordinate-based MLPs due to spectral bias: from low-frequency global geometry (e.g., large arteries) to high-frequency local details (e.g., distal branches). VITAL addresses this by employing the MoE architecture, where each expert is implemented as a SIREN [26] sub-network. To synthesize subject-specific anatomical details, we condition the SIREN experts using FiLM [20] modulation. Following [10], the concatenated conditioning vector $\tilde{\mathbf{z}}_i = [\mathbf{z}_i, \mathbf{c}_i]$ is mapped to layer-specific scale ($\boldsymbol{\gamma}$) and shift ($\boldsymbol{\beta}$) parameters via a linear projection. The activation of the $l$-th hidden layer in an expert network is defined as $\mathbf{h}_l(\mathbf{x}) = \sin\left(\omega_0 \cdot \boldsymbol{\gamma}_l(\tilde{\mathbf{z}}_i) \odot (\mathbf{W}_l\mathbf{x} + \mathbf{b}_l) + \boldsymbol{\beta}_l(\tilde{\mathbf{z}}_i)\right)$, where $\mathbf{W}_l$ and $\mathbf{b}_l$ are shared weights and biases, and $\omega_0$ is the frequency scaling factor. This design choice ensures that latent code variations translate into smooth semantic

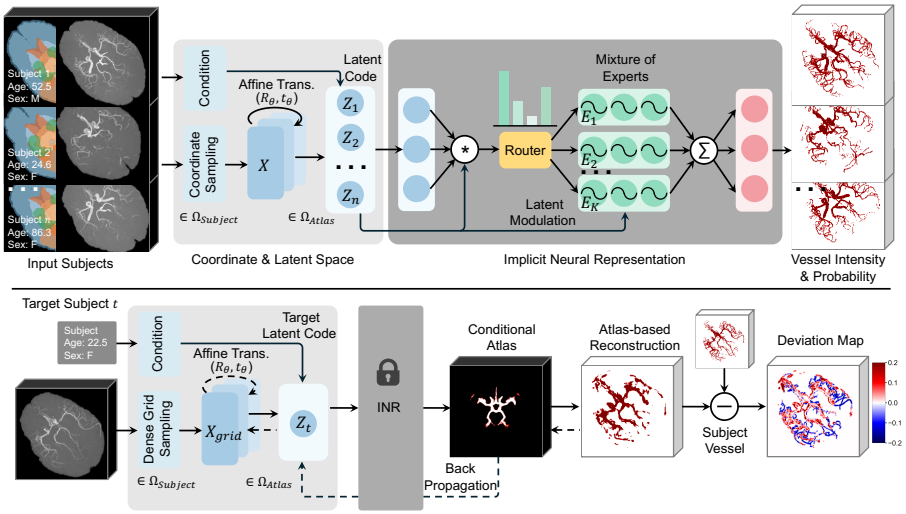


Fig. 1: Overview of the VITAL framework. *Top:* Demographic-conditioned atlas learning with a continuous probabilistic vessel atlas learned from a multi-subject cohort. Coordinates sampled from the subject space ($\Omega_{\text{Subject}}$) are mapped to a canonical atlas space ($\Omega_{\text{Atlas}}$) via learnable affine transformations ($R_\theta, t_\theta$). A MoE-based INR encodes the vascular topology, modulated by subject-specific latent codes and demographic covariates. *Bottom:* Inference via test-time adaptation. For an unseen target subject, the pre-trained INR weights are frozen. The framework iteratively optimizes the target latent code ($Z_t$) and alignment parameters via backpropagation to fit the atlas prior. An exploratory atlas-to-subject deviation map can be computed when based on observed vessel.

deformations rather than high-frequency artifacts, facilitating a compact and interpolatable latent space.

The network aggregates these modulated experts to form the final representation. Specifically, each expert sub-network $E_k$ is parameterized as a composition of such modulated layers $\mathbf{h}_l(\mathbf{x})$. It consists of $K$ parallel experts $\{E_1, \ldots, E_K\}$ and a gating network $G(\mathbf{x}_a)$. For a query coordinate $\mathbf{x}_a$, only the top-$k$ experts are evaluated. The output is formulated as $f_\Theta(\mathbf{x}_a, \mathbf{z}_i) = \sum_{j \in \mathcal{T}(\mathbf{x}_a)} \tilde{G}(\mathbf{x}_a)_j \cdot E_j(\mathbf{x}_a, \mathbf{z}_i)$, where $\mathcal{T}(\mathbf{x}_a)$ contains the indices of the top-$k$ experts, and $\tilde{G}$ is the routing weight normalized over the active experts, allowing spatial and spectral specializations without prohibitive computational overhead.

### 2.2 Training Paradigm

The VITAL framework is trained end-to-end by minimizing a composite objective function. The total loss $\mathcal{L}_{\text{total}}$ is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rec}} + \mathcal{L}_{\text{seg}} + \lambda_{\text{moe}} \mathcal{L}_{\text{moe}}, \tag{1}$$

where $\mathcal{L}_{\text{rec}}$ is the L1 reconstruction loss on MR intensities, and $\mathcal{L}_{\text{seg}}$ is the weighted cross-entropy loss for semantic segmentation. Following [9], we optimize the rigid alignment parameters implicitly by minimizing the reconstruction objective $\mathcal{L}_{\text{rec}}$ through the differentiable STN [14], without explicit regularization on $\theta_{\text{stn}}$. To ensure efficient utilization of the expert ensemble and prevent expert collapse, where only a few experts are active while others remain idle [7,24], we incorporate an auxiliary load balancing loss $\mathcal{L}_{\text{moe}}$. For a given MoE layer $l$ with $K$ experts and a training batch of size $N$, we define the importance of the $k$-th expert as $I_k^{(l)}$, and the sum of its gating probabilities across batches as $I_k^{(l)} = \sum_{n=1}^{N} G^{(l)}(\mathbf{x}_n)_k$, where $G^{(l)}(\mathbf{x}_n)_k$ denotes the softmax output of the gating network for input $\mathbf{x}_n$. With ideal uniform load distribution, each expert should process an equal workload share, yielding a target importance $\tau = N/K$. The load balancing loss is formulated as the normalized variance of expert importance, averaged across all $L$ MoE layers as:

$$\mathcal{L}_{\text{moe}} = \frac{1}{L} \sum_{l=1}^{L} \sum_{k=1}^{K} \left( \frac{I_k^{(l)} - \tau}{\tau} \right)^2 . \tag{2}$$

By penalizing uneven expert utilization, $\mathcal{L}_{\text{moe}}$ promotes uniform distribution of queries across available experts, thereby maximizing neural representation.

VITAL is built with 5-layer SIREN backbone and $K = 4$ experts, with $k = 2$ active experts per query. Feature modulation layers are applied after the intermediate layers to inject the learned spatial latent codes, which are parameterized with dimensions $256 \times 3 \times 3 \times 3$ to capture anatomical variations. To address the inherent sparsity of the cerebrovascular tree, we employ a label-stratified coordinate sampling strategy during training. Specifically, foreground coordinates (i.e., vessels and brain tissues) are exhaustively sampled, while background coordinates are randomly sampled to match the foreground count.

## 3 Experiments

**Experimental Setup & Implementation Details.** In this study, we used MRA scans and corresponding T1 images from healthy adults aged 19–87 years in the IXI dataset [6]. The cohort was randomly partitioned into 497 subjects for training and 55 for testing. To establish anatomical and vascular reference labels, we processed the data through a standardized pipeline. Vascular structures were annotated from MRA scans following the VesselVerse protocol [11]. For anatomical context, T1 images were segmented using FastSurfer [12] and rigidly registered to the MRA space via ANTs [4]. The final label maps were harmonized into ten classes (e.g., vessel, cerebrospinal fluid, cortical gray matter, white matter, basal ganglia). Preprocessing involved non-brain tissue removal and in-plane resampling to a resolution of $256 \times 256$.

The model was trained for 200 epochs using the AdamW optimizer [16] with a cosine annealing scheduler, setting the initial learning rates to 1e-4 for the INR

Table 1: Quantitative comparison and ablation study of atlas-to-subject reconstruction and topological alignment on the healthy IXI dataset [6]. We report global anatomical and vascular-specific topological metrics evaluated on both the full brain and the macroscopic region (CoW).

| Metric | CiNeMA [9] | MoE ($K = 2$) | MoE ($K = 4$) | MoE ($K = 4$)+$\mathcal{L}_{\text{moe}}$ |
|---|---|---|---|---|
| PSNR ↑ (dB) | $21.76 \pm 0.96$ | $23.72 \pm 2.01$ | $24.94 \pm 2.02$ | $\mathbf{24.98} \pm 2.29$ |
| SSIM ↑ | $0.79 \pm 0.03$ | $0.82 \pm 0.03$ | $0.82 \pm 0.03$ | $\mathbf{0.84} \pm 0.03$ |
| Dice ↑ | $\mathbf{0.72} \pm 0.02$ | $0.70 \pm 0.03$ | $0.71 \pm 0.02$ | $0.71 \pm 0.03$ |
| ASD↓ (mm) | $5.16 \pm 1.63$ | $4.30 \pm 0.82$ | $4.42 \pm 0.75$ | $\mathbf{4.18} \pm 0.60$ |
| ASD (CoW)↓ (mm) | $3.45 \pm 1.14$ | $3.22 \pm 1.31$ | $3.11 \pm 1.37$ | $\mathbf{2.93} \pm 0.98$ |
| Conn. Error↓ | $296.08 \pm 180.29$ | $204.50 \pm 151.24$ | $198.62 \pm 150.27$ | $\mathbf{183.69} \pm 130.67$ |
| Conn. Error (CoW)↓ | $28.23 \pm 20.63$ | $18.04 \pm 15.08$ | $18.27 \pm 12.79$ | $\mathbf{15.65} \pm 12.17$ |
| Center. Dist ↓ (mm) | $5.52 \pm 0.70$ | $4.75 \pm 0.50$ | $4.66 \pm 0.42$ | $\mathbf{4.57} \pm 0.34$ |
| Center. Dist (CoW)↓ (mm) | $3.78 \pm 0.82$ | $3.77 \pm 0.82$ | $3.71 \pm 0.85$ | $\mathbf{3.58} \pm 0.79$ |

decoder, 5e-4 for the latent codes, and $\lambda_{\text{moe}} = 0.01$. During inference, we generated conditioned atlases across continuous age domains by computing target latent codes via Gaussian kernel regression over the trained population following CINA [8]. For atlas-to-subject reconstruction, the atlas was projected to target subjects by optimizing the subject-specific latent codes and transformation parameters for 50 epochs using $\mathcal{L}_{\text{rec}}$, while freezing the INR weights. The framework was implemented in PyTorch and trained on Nvidia A100 GPUs.

**Main Experimental Results.** We divide our evaluation into the architecture representation and the constructed atlas quality. Table 1 shows the quantitative ablation study of the model's alignment accuracy and topology reconstruction. We compare VITAL with CiNeMA [9], the closest INR-based generative atlas baseline. The integration of the MoE architecture improves image reconstruction and topological consistency while maintaining comparable anatomical overlap. Specifically, using $K = 4$ experts increases the global Peak Signal-to-Noise Ratio (PSNR) from 21.76 dB to 24.94 dB and the Structural Similarity index (SSIM) from 0.79 to 0.84. The application of the load balancing loss ($\mathcal{L}_{moe}$) further refines the vascular topology. For the whole-brain vasculature, this configuration reduces the Average Surface Distance (ASD) to 4.18 mm and the symmetric centerline distance (Center. Dist) to 4.57 mm. It also lowers the whole-brain connectivity error (Conn. Error) from 296.08 to 183.69, which measures the absolute difference in Betti-0 connected components, thus indicating a reduction in vascular breaks. To evaluate hemodynamically critical structures, we isolated the CoW via a bounding volume encompassing primary cerebral arteries. The ASD and centerline distance further decreased to 2.93 mm and 3.58 mm, respectively, while the connectivity error dropped to 15.65. These results suggest the MoE mechanism enhances multi-scale representation.

**Comparison of Atlas Construction.** Fig. 2 compares VITAL with conventional voxel-based atlas construction methods. While voxel-based registration methods are effective for relatively consistent brain tissues [1], vascular struc-

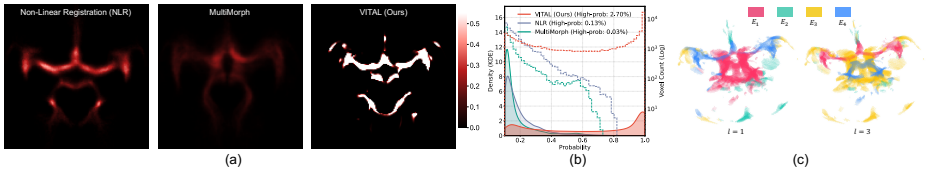


Fig. 2: Comparison of atlas construction. (a) Vascular probability maps from NLR [18], MultiMorph [1], and VITAL. (b) Probability density distributions. Solid curves represent overall probability distributions, and dashed lines represent voxel counts. (c) 3D visualization of MoE routing distribution in the atlas.

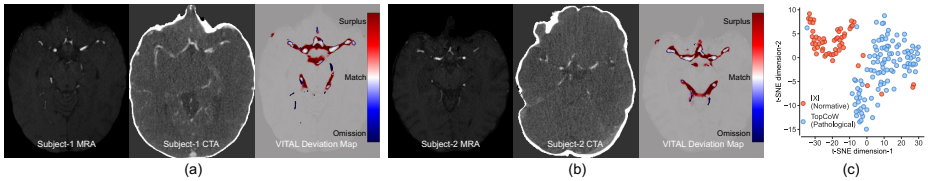


Fig. 3: (a, b) Qualitative visualization of VITAL results on stroke-related subjects on TopCoW [30]. (c) t-SNE [17] visualization of the learned latent codes.

tures exhibit substantial inter-subject variability. Averaging these highly variable structures can lead to spatial wash-out, where non-overlapping vessels dilute each other and reduce the visibility of small or distal vessel structures (Fig. 2(a)). This probability dilution is reflected in Fig. 2(b), where the density distributions for discrete methods are dominated by low-confidence voxels. By assigning a higher proportion of probability mass to high-probability vessel regions (2.70%) and achieving the lowest Shannon entropy (0.56), VITAL produces sharper high-probability core vessel structures while preserving a probabilistic representation of anatomical variability. Fig. 2(c) visualizes MoE spatial concentration via top-1 expert routing ($\arg\max_k G^{(l)}(\mathbf{x}_a)_k$) in $\Omega_{\text{Atlas}}$, showing that the MoE dynamically partitions vessel anatomy across scales, with some experts capturing low-frequency arteries and others reconstructing high-frequency branches.

**Atlas-Based Deviation Analysis.** To assess VITAL on pathological out-of-distribution cases, we incorporated 109 stroke-related scans from the TopCoW dataset [30] (ages 18–80 years). Given the imaging protocol differences between datasets and the morphological variability of stroke subjects, relying solely on intensity reconstruction can lead to misalignment. Therefore, we employed mask-guided test-time adaptation using $\mathcal{L}_{\text{rec}}+\mathcal{L}_{\text{seg}}$. As shown in Fig. 3(a-b), projecting the atlas anatomy into subject space enables visualization of atlas-to-subject mismatches, such as vessel occlusions. This suggests that the atlas may provide an exploratory topological reference for visualizing atlas-to-subject deviations. Latent space clustering (Fig. 3(c)) also reveals that the learned latent codes tend to

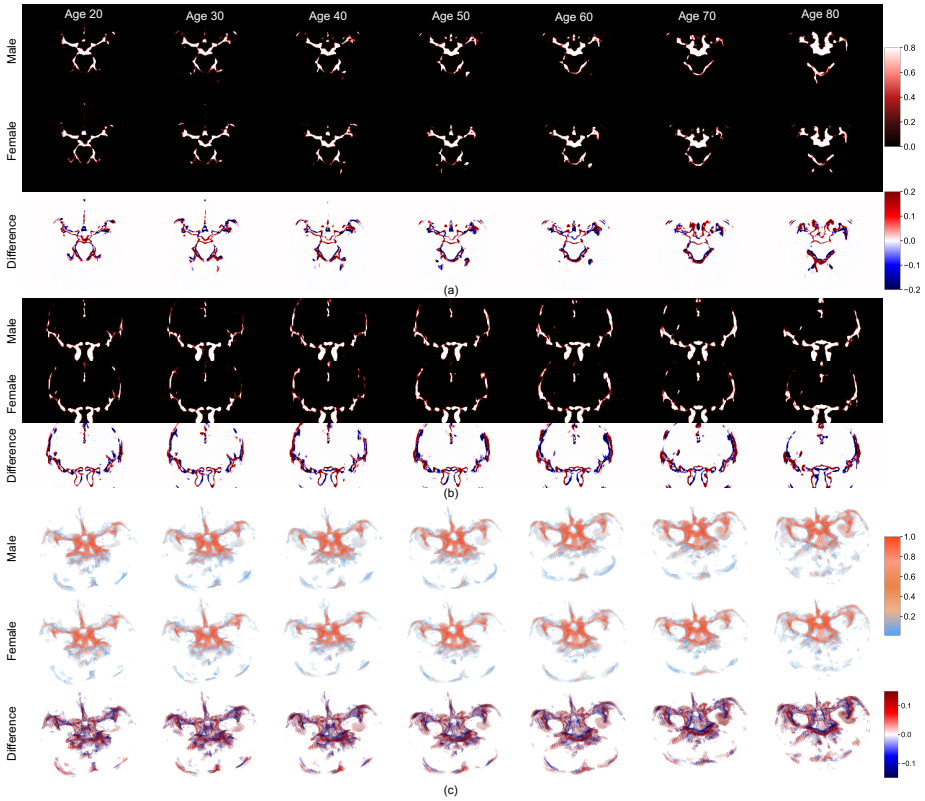


Fig. 4: Age- and sex-conditioned probabilistic cerebrovascular atlas generated by VITAL from cross-sectional population data. (a-b) Axial (CoW) and coronal (Internal Carotid Artery and Middle Cerebral Artery) slices across ages 20 to 80. (c) 3D rendering in the CoW plane. (Bottom) Male-minus-female difference maps highlighting sex-stratified vascular variation.

distinguish between pathological and healthy subjects, indicating that the latent representation appears to capture out-of-distribution anatomical variation.

**Conditioned Atlas Generation.** Since VITAL parameterizes the atlas using INR, it provides a continuous functional mapping and enables dynamic generation of conditional atlases and observation of anatomical variation across demographic variables. Fig. 4 presents cross-sectional age-conditioned cerebrovascular probability maps spanning the adult lifespan from 20 to 80 years. Major arterial trunks (e.g., Middle and Anterior Cerebral Artery) are delineated by high-probability regions that reconstruct the complex topology of the CoW. Across age, increased spatial dispersion appears in older groups, reflecting greater inter-subject geometric variability and age-associated vascular morphological variation. This dispersion suggests that the atlas captures population-level age-

associated changes. The sex-stratified difference maps reveal distinct sexual dimorphism, with morphological deviations primarily localized at the vascular periphery: a pronounced high-probability rim in the male cohort reflects the effect of global cranial scaling on vascular distributions, consistent with known anatomical differences in cranial size and vascular geometry between sexes [21,22], and demonstrating VITAL's ability to disentangle demographic differences from anatomical registration bias.

## 4 Conclusion

We proposed VITAL, a continuous probabilistic neural atlas for multi-scale modeling of the human cerebrovasculature. Through a dynamically routed MoE representation, VITAL improves atlas-to-subject reconstruction and topological consistency while enabling demographic-conditioned atlas generation. Experiments on healthy subjects and exploratory analyses on stroke-related cases demonstrate its potential as a normative reference for vascular analysis. Nevertheless, VITAL is currently intended as a probabilistic population-level reference rather than a surgical navigation tool. Future work will extend VITAL to large-scale, multi-institutional cohorts and further validate its utility for population-level cerebrovascular studies.

**Acknowledgments.** This work is funded by the European Union (ERC CARAVEL, 101171357). J.C. acknowledges the seed grant from the Royal College of Radiologists. S.W. acknowledges support from the Swiss National Science Foundation (SNSF, so10006824). We acknowledge computational resources provided by EuroHPC JU on LEONARDO, hosted by CINECA, Italy (EHPC-DEV-2026D01-008), and by GENCI-IDRIS on Jean Zay A100 partition (2026-A0200617550).

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the EU or the European Research Council. Neither the EU nor the granting authority can be held responsible for them.

## References

1. Abulnaga, S.M., Hoopes, A., Dey, N., Hoffmann, M., Fischl, B., Guttag, J., Dalca, A.: Multimorph: On-demand atlas construction. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 30906–30917 (2025)
2. Alblas, D., Suk, J., Brune, C., Yeung, K.K., Wolterink, J.M.: Sire: scale-invariant, rotation-equivariant estimation of artery orientations using graph neural networks. Medical Image Analysis **101**, 103467 (2023)
3. Amiranashvili, T., Lüdke, D., Li, H.B., Zachow, S., Menze, B.H.: Learning continuous shape priors from sparse data with neural implicit functions. Medical Image Analysis **94**, 103099 (2024)
4. Avants, B.B., Tustison, N., Song, G., et al.: Advanced normalization tools (ANTS). The Insight Journal **2**(365), 1–35 (2009)

5. Ben-Shabat, Y., Hewa Koneputugodage, C., Ramasinghe, S., Gould, S.: Neural experts: Mixture of experts for implicit neural representations. Advances in Neural Information Processing Systems **37**, 101641–101670 (2024)
6. Brain Development: IXI Dataset. https://brain-development.org/ixi-dataset/ (2008)
7. Chi, Z., Dong, L., Huang, S., Dai, D., Ma, S., Patra, B., Singhal, S., Bajaj, P., Song, X., Mao, X.L., et al.: On the representation collapse of sparse mixture of experts. Advances in Neural Information Processing Systems **35**, 34600–34613 (2022)
8. Dannecker, M., Kyriakopoulou, V., Cordero-Grande, L., Price, A.N., Hajnal, J.V., Rueckert, D.: Cina: Conditional implicit neural atlas for spatio-temporal representation of fetal brains. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 181–191 (2024)
9. Dannecker, M., Sideri-Lampretsa, V., Starck, S., Mihailov, A., Milh, M., Girard, N., Auzias, G., Rueckert, D.: Cinema: Conditional implicit neural multi-modal atlas for a spatio-temporal representation of the perinatal brain. IEEE Transactions on Medical Imaging **45**, 738–750 (2025)
10. 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: Proceedings of the International Conference on Machine Learning. vol. 162, pp. 5694–5725 (2022)
11. Falcetta, D., Marciano, V., Yang, K., Cleary, J., Legris, L., Rizzaro, M.D., Pitsiorlas, I., Chaptoukaev, H., Lemasson, B., Menze, B., Zuluaga, M.A.: Vesselverse: A dataset and collaborative framework for vessel annotation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 655–665 (2025)
12. Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., Reuter, M.: Fastsurfera fast and accurate deep learning based neuroimaging pipeline. NeuroImage **219**, 117012 (2020)
13. Huck, J., Wanner, Y., Fan, A.P., Jäger, A.T.P., Grahl, S., Schneider, U., Villringer, A., Steele, C.J., Tardif, C.L., Bazin, P.L., Gauthier, C.J.: High resolution atlas of the venous brain vasculature from 7 T quantitative susceptibility maps. Brain Structure and Function **224**, 2467 – 2485 (2019)
14. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. Advances in Neural Information Processing Systems **28** (2015)
15. Lee, D., Park, C., Lee, S., Lee, S., Kim, M.: Convolutional implicit neural representation of pathology whole-slide images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 191–200 (2024)
16. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: Proceedings of the International Conference on Learning Representations (2019)
17. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of Machine Learning Research **9**(11) (2008)
18. Mouches, P., Forkert, N.D.: A statistical atlas of cerebral arteries generated using multi-center mra datasets from healthy subjects. Scientific data **6**(1), 29 (2019)
19. Muffoletto, M., Xu, H., Xu, Y., Williams, S.E., Williams, M.C., Kunze, K.P., Neji, R., Niederer, S.A., Rueckert, D., Young, A.A.: Neural implicit functions for 3d shape reconstruction from standard cardiovascular magnetic resonance views. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 130–139. Springer (2023)
20. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Proceedings of the AAAI conference on Artificial Intelligence. vol. 32 (2018)

21. Ritchie, S.J., Cox, S.R., Shen, X., Lombardo, M.V., Reus, L.M., Alloza, C., Harris, M.A., Alderson, H.L., Hunter, S., Neilson, E., et al.: Sex differences in the adult human brain: evidence from 5216 uk biobank participants. Cerebral Cortex **28**(8), 2959–2975 (2018)
22. Ruigrok, A.N., Salimi-Khorshidi, G., Lai, M.C., Baron-Cohen, S., Lombardo, M.V., Tait, R.J., Suckling, J.: A meta-analysis of sex differences in human brain structure. Neuroscience & Biobehavioral Reviews **39**, 34–50 (2014)
23. Sander, J., de Vos, B.D., Bruns, S., Planken, N., Viergever, M.A., Leiner, T., Išgum, I.: Reconstruction and completion of high-resolution 3D cardiac shapes using anisotropic CMRI segmentations and continuous implicit neural representations. Computers in Biology and Medicine **164**, 107266 (2023)
24. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., Dean, J.: Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In: International Conference on Learning Representations (2017)
25. Sideri-Lampretsa, V., McGinnis, J., Qiu, H., Paschali, M., Simson, W., Rueckert, D.: Sinr: Spline-enhanced implicit neural representation for multi-modal registration. In: Proceedings of the Medical Imaging with Deep Learning (2024)
26. 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)
27. Ward, P.G.D., Ferris, N.J., Raniga, P., Dowe, D.L., Ng, A.C.L., Barnes, D.G., Egan, G.F.: Combining images and anatomical knowledge to improve automated vein segmentation in MRI. NeuroImage **165**, 294–305 (2017)
28. Wu, Q., Li, Y., Sun, Y., Zhou, Y., Wei, H., Yu, J., Zhang, Y.: 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 (2022)
29. Wu, Z., Jin, Y., Yi, K.M.: Neural fourier filter bank. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14153–14163 (2023)
30. Yang, K., Musio, F., Ma, Y., Juchler, N., Paetzold, J.C., Al-Maskari, R., Höher, L., Li, H.B., Hamamci, I.E., Sekuboyina, A., et al.: Benchmarking the cow with the topcow challenge: Topology-aware anatomical segmentation of the circle of willis for CTA and MRA. ArXiv pp. arXiv–2312 (2025)
31. Zhenxing, M., Xu, D.: Switch-nerf: Learning scene decomposition with mixture of experts for large-scale neural radiance fields. In: Proceedings of the International Conference on Learning Representations (2022)