

Anirudh Narasimha Bharadwaj

MEDICAL IMAGE ANALYSIS • IMPLICIT NEURAL REPRESENTATIONS • DEFORMABLE REGISTRATION

📍 Oviedo, Florida, USA 📞 (+1) 321-888-7707 ✉ an517162@ucf.edu ✉ anirudhbharadwaj13@gmail.com
🌐 github.com/anirudhnbharadwaj 🔗 linkedin.com/in/anirudhnbharadwaj 🌐 anirudhnbharadwaj.vercel.app

RESEARCH PROFILE

M.Sc. Computer Vision researcher at the University of Central Florida (GPA 3.889 / 4.00, graduating December 2026), specialising in **medical image analysis**, **implicit neural representations**, **deformable registration** and **uncertainty-aware deep learning**. First author of **QueryPhi** (under review, NeurIPS 2026), which develops continuous neural fields for diffeomorphic 3D registration under sparse and anisotropic sampling.

The through line is recovering continuous 3D anatomy from incomplete observations, building physical constraints into neural models, quantifying uncertainty, and validating under cross-cohort domain shift. Two years shipping medical-device software under ISO 13485, alongside clinicians, taught me what a method must survive before it reaches a clinical workflow.

Research interests: Medical image analysis • 3D reconstruction from incomplete observations • implicit neural and generative representations • physics-informed machine learning • uncertainty quantification • clinical validation.

WHAT I BRING

3D reconstruction from incomplete observations. My thesis studies precisely when a continuous implicit field recovers structure a voxel lattice cannot under sparse or anisotropic sampling, and I built the instruments to measure that rather than assert it. The same question governs any reconstruction from limited views: sparse angle tomography, undersampled acquisition, or a stack sampled far more coarsely through plane than in it.

Physics-informed modelling, as a foundation to extend. I work daily with differentiable field representations: a stationary velocity field integrated by scaling-and-squaring, a log-barrier constraint on the Jacobian determinant, bending and kinetic energy terms, and an analytic autograd Jacobian at continuous coordinates. That is the machinery to build on toward governing-equation constraints rather than generic smoothness priors. I have not yet worked on Navier-Stokes or computational fluid dynamics.

Uncertainty quantification. My correspondence objective is an explicit heteroscedastic estimator: a softmax matching distribution over a MIND cost volume with a learned precision $\kappa = 1/(1 + \sigma^2/\sigma_0)$, so the model down-weights coordinates where the image evidence is ambiguous.

Multi-cohort validation and clinical translation. A zero-shot cross-cohort transfer study with a statistical protocol built on subject-level resampling, Holm-corrected paired tests and a pre-specified interaction test; plus two years of ISO 13485 medical-device work designed alongside clinicians.

Generative modelling. Amortised conditional neural-field generation and LoRA-based foundation-model adaptation. I have not yet worked directly with diffusion or flow matching, and expanding there is a deliberate next step.

EDUCATION

University of Central Florida

Orlando, FL, USA

M.S. in Computer Vision | GPA 3.889 / 4.00

Jan. 2025 to Dec. 2026

- Coursework: Medical Image Computing, Advanced Computer Vision, Machine Learning, Mathematical Introduction to Deep Learning, 3D Computer Vision, Virtual Reality Engineering.
- Thesis: continuous diffeomorphic image registration with implicit neural velocity fields (QueryPhi).

Visvesvaraya Technological University

Bengaluru, India

B.E. in Electronics & Communication Engineering | GPA 3.50 / 4.00

2019 to 2023

- Digital signal processing, digital image processing, signals and systems. Sampling theory, reconstruction and projection geometry are first-language material for me.

PUBLICATIONS & PREPRINTS

QueryPhi: Continuous Diffeomorphic 3D Registration with Implicit Neural Velocity Fields

NeurIPS 2026, under review

- **A. N. Bharadwaj et al. (first author).** A 5.5M-parameter amortised model decoding a coordinate-continuous stationary velocity field with a modulated SIREN over a RAFT/PWC-style correlation pyramid, integrated by scaling-and-squaring. Highest mean Dice of the **ten methods evaluated** on IXI atlas-to-subject registration (**0.7726** over 30 FreeSurfer structures, 95% CI 0.769 to 0.776, against 0.7534 for TransMorph-bspl, $p = 1.5 \times 10^{-8}$), at **0.0001% folded voxels** and **8.5× fewer parameters** than the transformer baselines.
- The same paper carries the anisotropy and topology analysis: a controlled 150-run study on the **CHAOS** abdominal MR/CT benchmark, testing the pre-specified hypothesis that a continuous velocity parameterisation stays diffeomorphic under anisotropic sampling where a lattice does not.
- *Preprint submitted to arXiv and in processing, so no public identifier yet. Full manuscript with appendices and statistical protocol available on request.*

VISTA: Video Interaction Spatio-Temporal Analysis Benchmark

arXiv:2605.01391 | CVPR 2026, PVUW

- A. Aparcedo, A. Kumar, A. Garg, D. Pham, W.-K. Chen, **A. N. Bharadwaj**, A. Chadha, Y. Rawat. An interaction-aware diagnostic benchmark (~12K curated video and query pairs) for spatio-temporal reasoning in vision-language models. I contributed to benchmark construction and to the evaluation of 11 state-of-the-art VLMs.

RESEARCH EXPERIENCE

AI MIND Lab, University of Central Florida

Orlando, FL, USA

Graduate Research Assistant, Medical Image Registration

Jan. 2026 to present

- **Problem.** Grid-based registration stores the deformation on a fixed voxel lattice, so the field exists only at grid nodes, its derivatives are piecewise constant, and the discretisation itself can admit folded warps.
- **Method, designed and implemented end to end (5.5M parameters).** A 3D CNN encoder and three-level cost-volume correlation pyramid give a coarse-to-fine velocity estimate; a **modulated SIREN decoder** adds residual detail through per-coordinate FiLM modulation, zero-initialised so training begins at the identity warp; a **scaling-and-squaring integrator** yields $\varphi = \exp(\nu)$ with an analytic autograd Jacobian at arbitrary continuous coordinates. Fully unsupervised: MIND-SSC, multi-window LNCC and edge-NGF similarity plus a heteroscedastic correspondence term, over a log-barrier, diffusion, bending and kinetic-energy regulariser stack. The training Dice weight is exactly zero; labels only score the result.
- **Results, $n = 115$ paired test cases.** IXI atlas-to-subject: **first of ten methods**, Dice 0.7726, folding 0.0001%, best boundary distance (HD95 1.653 mm), best Friedman-Nemenyi average rank (1.77 against 1.81), all nine paired comparisons significant after Holm correction. **Zero-shot transfer** to **OASIS** with no retraining: best HD95 (**1.502 mm**) of nine methods, retaining **91%** of the published in-domain ceiling while staying near fold free.
- **A quantitative instrument for continuity.** Implicit-representation papers assert continuity; almost none measure it. I built the suite that does: field-continuity divergence, held-out-slice gain, super-Nyquist spectral energy, and a finite-difference-to-autograd Jacobian convergence certificate.
- **Self-falsification, the result I am most willing to be judged on.** To separate representation from objective I built *GridSVF*, a lattice control matched on the loss, the integrator, the data, the training schedule and an equal hyperparameter search budget. Across a 35-cell, 115-case ablation the control **beat** my model on 6 of 9 metrics on isotropic data. I reported it, traced the mechanism, re-scoped the claim from *representation* to *objective and topology*, and cut eight of my own components that improved results on removal.
- **The claim that survived: an anisotropy dose-response on topology.** On CHAOS, where MR is ~5:1 anisotropic and CT ~1.3:1, the continuous parameterisation is fold free in **88% (38/43)** of tuned configurations against **9% (3/34)** for the lattice (Fisher $p = 9.3 \times 10^{-13}$), while on near-isotropic CT the two are indistinguishable. I declined to report the accompanying accuracy gap, which is a sampling artefact, and specified the confound-elimination arm that must run before publication.
- **Research infrastructure.** A unified harness placing 13 competing methods (ANTs-SyN, NiftyReg, LDDMM, uniGradICON, TransMorph family, nnFormer, CoTr, VoxelMorph family, MIDIR, ViT-V-Net) behind one interface with identical metrics, splits and output schema; a 150-run matrix on SLURM arrays; four parallel Optuna studies of 160 trials each under a composite objective penalising HD95 and field complexity rather than Dice. GPU accounting (42.8 hours per ablation cell) showed the queued 43-cell plan would take 92% of a shared 2,000-hour cap, so I cut it to the 14 cells with a demonstrated effect. I also found and fixed silent correctness bugs that would have contaminated published numbers, among them a metric scoring absent structures as perfect.

Equidor Medtech LLP

Bengaluru, India

Computer Vision Engineer, Medical Device Software

2023 to 2024

- Built **EquiCOG**, a production real-time binocular eye-tracking engine processing **100 Hz** synchronised sensor streams from embedded hardware (Python, OpenCV) for oculomotor and vestibular analysis, with live filtering, artifact and dropout rejection, and derived motion metrics.
- Shipped **ISO 13485-compliant** PyQt calibration and quality-control tooling whose automated signal-quality checks flag a bad acquisition while the patient is still present, **cutting device setup time by 25%**; built the hardware-accelerated FFmpeg pipeline for synchronised video-plus-signal logging.
- Authored traceable design and test documentation to regulatory standard and worked directly with clinicians on acquisition protocol, experience I expect to matter in a clinically partnered research setting.

Center for Research in Computer Vision (CRCV), UCF

Orlando, FL, USA

Research Assistant, Vision-Language Model Evaluation

Aug. 2025 to Jan. 2026

- Co-developed the **VISTA** benchmark (~12K curated video and query pairs) for interaction-aware spatio-temporal reasoning, including schema design and the curation protocol.
- Built the evaluation harness running **11 state-of-the-art vision-language models** end to end, normalising free-form generative output and scoring it automatically. Slice-level failure analysis exposed systematic spatio-temporal biases hidden by aggregate accuracy, the same diagnostic habit I now apply to per-structure and per-cohort registration error.

Cyclops Medtech Pvt. Ltd.

Bengaluru, India

Research Intern, Signal & Data Analysis

2022 to 2023

- Analysed large-scale clinical ocular-movement time series and built cleaning, feature-extraction and batch pipelines (NumPy, Pandas, OpenCV) that replaced manual spreadsheet workflows and refined vestibular-diagnostic algorithms.

SELECTED PROJECTS

The Geometry of Generalisation: optimisation dynamics in paediatric brain-tumour segmentation

PyTorch, PyHessian, H100

BraTS-PED

2025

- Four matched training arms varying only the loss and the optimiser, read through the Hessian spectrum rather than the leaderboard, under **3500:1** class imbalance. A signed-distance boundary loss traded one point of Dice (85.9% to 84.9%) for **1.21 mm** of HD95 (8.05 to 6.84 mm); Lion never left initialisation, because sign-based updates erase the gradient magnitude that encodes class rarity; and sharpness-aware minimisation could not stabilise it, because $\lambda_{\max} \approx 501$ made the perturbation radius wider than the valley. Diagnosing *why* a regime fails, not merely that it does.

From Nuclei to Tissue: panoptic melanoma histopathology

Mask R-CNN, DynUNet

PUMA Challenge

2025

- Two specialised models rather than one shared backbone, to place nuclei in the tissue context that gives them clinical meaning for tumour-infiltrating-lymphocyte assessment. Nuclei stage across five folds: **AJI 0.62**, Panoptic Quality 0.54, macro F1 0.65, slightly ahead of the HoVer-Net figures reported on comparable data. The tissue stage and the panoptic merge are specified and partly built, and stand as the next step rather than a finished result.

Parameter-efficient adaptation of the Segment Anything Model

PyTorch, LoRA

PUMA

2025

- Adapted **SAM** to nuclei segmentation by training **1.04%** of its weights (6.69M of 643.7M) through low-rank adapters in the attention projections: Dice 0.8233, AJI 0.6197, Panoptic Quality 0.5724 across five folds. The gap between Dice and AJI is the informative part: the model finds nuclei reliably and separates adjacent ones less reliably.

Volumetric brain-tumour segmentation with five-fold cross-validation

PyTorch, MONAI, SLURM

BraTS

2025

- 3D U-Net with DiceFocalLoss for small-lesion recall under extreme class imbalance; whole-tumour Dice **0.6181** across folds (range 0.5824 to 0.6434, HD95 8.85 mm). The useful finding was the spread rather than the mean: per-case enhancing-tumour Dice is bimodal (0.08 to 0.84), which the average conceals entirely.

Classical image processing and reconstruction

MATLAB

Course projects, UCF

2025

- **Wiener deconvolution** scored by PSNR and SSIM, with a mismatch study quantifying how a two-degree error in the assumed motion angle degrades restoration. **Fourier against Haar compression**, where Haar holds a 5 to 15 dB SNR advantage at every ratio from 2 to 50. **Switching median filtering** at 30 to 40% impulse noise, **covariance-based colour transfer**, and **ResNet-34 pneumonia classification** where transfer learning reached 82% accuracy against 70% from scratch, the per-class recall split mattering more than the mean.

TECHNICAL SKILLS

Programming & infrastructure: Python (expert), MATLAB, Bash, LaTeX, Git; Linux, HPC and SLURM (job arrays, accounting, fair-share budgeting), Docker, H100 and A100 nodes, reproducible experiment tracking and regression checks on scientific results

Deep learning: PyTorch, MONAI, scikit-learn, Optuna, Weights & Biases, multi-GPU and distributed training, mixed precision, LoRA and parameter-efficient fine-tuning

Medical image analysis: Deformable and diffeomorphic registration (stationary velocity fields, scaling-and-squaring, LDDMM, B-spline), implicit neural representations (SIREN, FiLM-modulated INRs), segmentation (U-Net family, SwinUNETR, DynUNet, Mask R-CNN, SAM), similarity metrics (LNCC, MIND-SSC, NGF, mutual information), FreeSurfer, ANTs, NiftyReg, SimpleITK, Nifti and DICOM

Modelling & mathematics: Differential geometry of deformations (Jacobian determinants, log-Jacobian dispersion, inverse consistency), variational regularisation (diffusion, bending, kinetic energy, log-barrier), optical-flow and cost-volume methods, Fourier and sampling theory, Hessian-spectrum and loss-landscape analysis, heteroscedastic uncertainty modelling

Validation & statistics: Subject-level rather than pair-level bootstrap, since pair-level intervals are anticonservative when pairs share subjects; paired Wilcoxon with Holm correction over a pre-specified family; Friedman with Nemenyi critical-difference diagrams; Wilson intervals for rates reaching 0% and 100%; trend and interaction tests, so a confound constant across the dose cannot manufacture the effect

Devices, data & languages: Real-time 100 Hz embedded sensor pipelines, calibration and QC tooling, OpenCV, FFmpeg, PyQt, ISO 13485 documentation. Datasets: IXI, OASIS, CHAOS, AbdomenMRCT, BraTS and BraTS-PED, PUMA, proprietary clinical oculomotor data. English (full professional and academic proficiency), Hindi, Kannada

HONOURS & ADDITIONAL INFORMATION

- **Clinical Innovation Award**, Department of Otorhinolaryngology, Yenepoya Medical College, India, for contributions to vertigo and balance rehabilitation technology.
- **Availability:** M.S. completes December 2026 and I am available to start from January 2027. Open to relocation, including outside the United States.
- **References:** available on request, including my thesis advisor at the AI MIND Lab and my supervising researcher at CRCV.