In the far future, medical imaging is a solved problem. Noninvasive tools for diagnosis and study of the human body will be ubiquitous, affordable, fast, convenient, and provide high-quality data. Today, however, medical imaging systems are limited and specialized. Scanners are expensive to build and maintain, tradeoffs exist between image quality and scanning speed, and special training is required both to operate the scanners and to interpret the images. Artificial intelligence (AI) continues to improve automated medical image analysis, processing, and reconstruction algorithms, leading to improved quantification, quality, and access. Nevertheless, many current AI methods are still sensitive to the particular appearance of the image. Indeed, medical images exhibit vast heterogeneity due to a plethora of degrees of freedom, including device manufacturers and models, modalities, contrasts, acquisition parameters, signal-to-noise ratios, and reconstruction algorithms. Toward an idealized future of medical imaging, automated analysis algorithms must operate correctly for any kind of image.

I hypothesize that incorporating imaging physics into AI algorithms will improve their robustness, reliability, interpretability, consistency, and overall performance. My research aims to build the scientific foundations for a new class of AI imaging algorithms that adhere to modality-specific acquisition processes and provide accurate representations of the true patient anatomy while leveraging large datasets where appropriate. By grounding algorithms in the image acquisition process, anatomical geometry, and physical constraints, I intend to establish a rigorous foundation for AI-driven imaging that is stable, data-consistent, and clinically reliable. As a computer science researcher, I plan to develop grounded AI methods informed by knowledge of the physics, instrumentation, signal and image processing practices, and medical and scientific goals in medical imaging and health informatics in general.

The core vision of my lab is to develop general algorithms for medical imaging that work for every patient, on every scanner, anywhere, by tightly coupling data-driven models with the physics of image formation. This involves algorithms that function correctly at both individual and population scales. My previous work on super-resolution [1, 2, 3] and large-scale generative model training [4, 5] and deployment [6] supports my ability to continue development of techniques along this continuum.

This vision demands innovations in generative modeling, optimization, and algorithmic reliability that extend well beyond healthcare. My lab will contribute to core areas of computer science, including machine learning, optimization, computer vision, and explainable AI, while advancing our understanding of how learned models can interact with physical systems. The resulting frameworks have the potential to standardize medical imaging across scanners, improve access to low-cost and portable imaging, and enable diagnostic AI systems that are trustworthy, generalizable, and aligned with clinical needs.

EXISTING WORK

Physics-Grounded Modeling of Acquisition Processes Inverse problems cannot be solved reliably without modeling the acquisition process of the imaging system. While much of modern deep learning for inverse problems has been developed for natural images, there is a need to study and implement computational models for medical modalities. In MRI acquisition, for example, the relationship between a high-resolution (HR) and a low-resolution (LR) image volume may be modeled by differences in slice excitation profile shape, thickness, and separation and represented as a linear acquisition operator. Even in an example such as this, which disregards differences in magnetic field strength, acquisition protocol, scanner manufacturer, etc., the heterogeneity of these slice excitation profile parameters can confound data-driven super-resolution techniques trained for a particular resolution. For instance, in natural image super-resolution it is common to desire $2\times$ or $4\times$ super-resolution, whereas in medical image super-resolution the goal is often isotropy, involving non-integer scale factors.

In my dissertation [6], I developed methods that leverage knowledge of the acquisition operator for the super-resolution of anisotropic MRI volumes. ECLARE [3] is an internally trained deep learning method that defines the HR in-plane data as a reference dataset and simulates paired LR data according to an approximate acquisition operator learned from the data. This method requires no external training data.

To restrict super-resolution methods from excessive hallucination, I developed a formulation rooted in perfect reconstruction filter bank theory that embeds the acquisition operator into an analysis-synthesis filter bank. By doing so, the super-resolved solution is prevented from hallucinating components in the range space.

Beyond supervised super-resolution methods, I developed FOCVS [7, 6], a diffusion-based algorithm that enforces exact agreement between the reconstructed image and the acquired data via a range–null space decomposition of the acquisition operator. As a result, FOCVS achieves data consistency under anisotropic MRI acquisition. This formulation extends classical ideas from optimization and signal processing into the domain of generative modeling, showing how physics-grounded constraints can control and stabilize learned inverse problem solvers. This development was a powerful extension to previous super-resolution approaches in MRI, including my own early approaches, and has informed my planned approach to future work in super-resolution and other medical image analysis tasks.

Generative Models as Priors for Medical Images Modern generative models are powerful approximators of data distributions. They have been used successfully as priors for inverse problem solvers, including the previously mentioned FOCVS approach. However, achieving a generative model for 3D medical imaging is not as straightforward as obtaining all available data and training a model. Just as reinforcement learning from human feedback is a key element of aligning large language models, data curation is a key element of training medical imaging generative models. Data curation for medical imaging requires particular expertise.

When developing MedForj [5], my colleagues and I curated a large aggregation of multiple open-source datasets of human brain MR volumes. Automated quality control alongside manual verification of every single image volume resulted in a dataset of more than 70,000 image volumes for training. I then trained several diffusion model estimation types to characterize their behavior on very high-dimensional 3D data and analyzed the anatomical plausibility with statistical and downstream segmentation-based similarity tests.

This model has successfully served as the generative prior backbone for several inverse problem solvers with my colleagues, including single image super-resolution, multi-image super-resolution, MRI tag tracking, and normal pressure hydrocephalus shunt inpainting.

FUTURE RESEARCH AGENDA

My future work builds on these foundations to pursue three research thrusts that together define a new paradigm of physics-grounded generative image processing and analysis. Each direction introduces fundamental problems at the interface of generative modeling, optimization, and structured acquisition processes, and each direction has direct clinical and societal impact.

Multi-Image and Multi-Modal Computational Imaging Patients are usually not imaged only once. They typically undergo multiple scans, often across scanners, institutions, modalities, and time points. Each observation is a distinct partial rendering of their anatomy at that location and time. There is a computational challenge in reliably leveraging all available single-subject multi-modal imaging data.

I plan to formalize multi-image inference as a structured generative inverse problem. In a structured graphical model of the acquisition process, each scan results in a node (corresponding to the image measurement) and an edge (corresponding to the acquisition physics) connecting from a source: the true patient anatomical structure. Building on my operator-grounded single-image methods [7], I plan to develop algorithms that exploit this structure through factorizations, spectral properties, and operator-aware regularization. This unifies multi-view geometry, inverse problems, and generative modeling into a scalable framework for real-world imaging.

This framework enables consistent reconstructions across visits and modalities, supports longitudinal disease monitoring, and stabilizes downstream algorithms for segmentation, registration, and diagnosis. More broadly, it advances multi-view inference and structured operator learning, which are problems central to geometry, vision, and probabilistic modeling.

Co-Design of Acquisition and Reconstruction Algorithms Current scanners were designed with physics and signal processing principles because it was possible to provide guarantees on patient safety and image accuracy. However, all scanners must make a tradeoff between acquisition time, signal-to-noise ratio, and resolution. The Shannon-Nyquist sampling theorem [8] provides guarantees if the underlying signal is sufficiently bandlimited. The developments in compressed sensing [9, 10, 11] provide guarantees if the data is sufficiently sparse under some transform. The modern era must make use of deep generative priors to provide guarantees when the data is sufficiently close to a learned manifold.

My lab will develop joint optimization of acquisition and reconstruction using differentiable physics and generative priors. This includes studying, modeling, and training data-driven systems to learn acquisition parameters that maximize information transfer under generative constraints, designing reconstruction and inverse-problem solvers that are robust to pathology, contrast, and modality, and enabling ultra-low-field MRI systems and other portable, low-cost imaging systems to produce high-quality and clinically relevant images.

Studying how acquisition operators allow generative priors to better reconstruct images may lead to promising new research endeavors. Conventional image acquisition relies on Shannon-Nyquist or compressed sensing principles. Grounded generative prior reconstruction may lead to new insights and flexible acquisition protocols that previously were infeasible. For example, in MRI, current slice excitation profiles often take a shape similar to a Gaussian, rect, and are implemented via the Shinnar-Le Roux algorithm [12, 13]. However, with powerful generative priors, alternative slice excitation profiles may lead to faster imaging, more accurate reconstructions, and higher signal-to-noise ratios.

This work builds directly on my operator-grounded models [6] and generative priors [4, 5, 7]. Co-design can reduce scan time, expand imaging access globally, and enable diagnostic-quality imaging on low-cost systems. For computer science, it opens new research on differentiable operators, optimal experimental design, and generative models interacting with real-world sensing systems.

Generative Priors with Properties Required for Medicine Medical imaging requires generative models that are not just expressive, but also stable, controllable, and privacy-preserving. Generated image volumes should not only have excellent perceptual quality, but also be anatomically plausible. I will develop generative models for 3D medical imaging with, geometry-constrained image generation with structural anatomical priors, pathology-aware and pathology-preserving priors for robust analysis and downstream tasks, privacy-preserving training to enable larger effective datasets for training, and theoretical guarantees regarding data consistency, posterior correctness, and bounds on synthesized data.

Early work from my collaborations provide initial footholds for this thrust. Previously, I explored the feasibility of StyleGAN3 for medical image generation [4]. Towards robust disentanglement of imaging protocols and underlying anatomy, HACA3 trains a contrastive learning-based encoder and conditional decoder [14]. Towards anatomical consistency, latent metric space Schrödinger Bridges train invertible neural networks towards a common space and learn the diffusion bridge matching in the latent metric space [15].

Grounding is important; data is not just data. This research establishes the reliability theory for generative imaging. There are implications for safety-critical systems beyond medicine as well, including robotics, scientific measurement, and autonomous sensing.

VISION OF IMPACT

Machine learning has been pervasive in medical imaging for decades. With the past decade of deep learning and AI development, these systems have begun to find their place in clinical and industrial settings. To ensure scientific and patient trust in these systems, and to enable a future of new medical devices, grounding AI in the imaging physics is paramount.

My research program will build the foundations for these grounded systems. By unifying physics-grounded modeling, generative priors, and algorithmic guarantees, my lab will create image analysis and processing algorithms that behave reliably outside curated datasets, develop approaches to standardize image quality across centers and populations, facilitate imaging through low-cost and portable systems, and develop broadly applicable principles for generative inference in scientific and physical domains.

My long-term ambition is for my lab to define the theory and practice of generative imaging, shaping a future where learning algorithms and physical measurement systems are co-designed, interpretable, and fundamentally trustworthy.

As a faculty member, I will build an interdisciplinary research program that connects theory with practice, bridges expertise between modalities from MRI to ultrasound to CT and beyond, shapes the next generation of robust computational imaging systems, and contributes foundational ideas to the broader machine learning and computer science communities.

References

[1] **S. W. Remedios**, S. Han, Y. Xue, A. Carass, T. D. Tran, D. L. Pham, and J. L. Prince. "Deep filter bank regression for super-resolution of anisotropic MR brain images". In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VI*. Springer. 2022, pp. 613–622.

[2] **S. W. Remedios**, S. Han, L. Zuo, A. Carass, D. L. Pham, J. L. Prince, and B. E. Dewey. "Self-supervised super-resolution for anisotropic MR images with and without slice gap". In: *International Workshop on Simulation and Synthesis in Medical Imaging*. Springer. 2023, pp. 118–128.

[3] **S. W. Remedios**, S. Wei, S. Han, J. Zhang, A. Carass, K. G. Schilling, D. L. Pham, J. L. Prince, and B. E. Dewey. "ECLARE: efficient cross-planar learning for anisotropic resolution enhancement". In: *Journal of Medical Imaging* 13.2 (2026), pp. 024001–024001.

[4] **S. W. Remedios**, B. E. Dewey, A. Carass, D. L. Pham, and J. L. Prince. "A deep generative prior for high-resolution isotropic MR head slices". In: *Medical Imaging 2023: Image Processing*. International Society for Optics and Photonics. SPIE, 2023.

[5] **S. W. Remedios**, A. Carass, J. L. Prince, and B. E. Dewey. "Diffusion-Driven Generation of Minimally Preprocessed Brain MRI". In preparation. 2025.

[6] **S. W. Remedios**. "Through-plane super-resolution of anisotropic multi-slice magnetic resonance images". Ph.D. Dissertation. Johns Hopkins University, Oct. 2025.

[7] **S. W. Remedios**, S. Wei, A. Carass, B. E. Dewey, and J. L. Prince. "Cycle-Consistent Zero-Shot Through-Plane Super-Resolution for Anisotropic Head MRI". In: *International Conference on Information Processing in Medical Imaging*. Best paper award. Springer. 2025, pp. 249–264.

[8] C. Shannon. "Communication in the Presence of Noise". In: *Proceedings of the IRE* 37.1 (1949), pp. 10–21.

[9] D. L. Donoho. "Compressed sensing". In: *IEEE Transactions on Information Theory* 52.4 (2006), pp. 1289–1306.

[10] E. J. Candès, J. Romberg, and T. Tao. “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information”. In: *IEEE Transactions on Information Theory* 52.2 (2006), pp. 489–509.

[11] M. Lustig, D. Donoho, and J. M. Pauly. “Sparse MRI: The application of compressed sensing for rapid MR imaging”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 58.6 (2007), pp. 1182–1195.

[12] M. Shinnar, S. Eleff, H. Subramanian, and J. S. Leigh. “The synthesis of pulse sequences yielding arbitrary magnetization vectors”. In: *Magnetic resonance in medicine* 12.1 (1989), pp. 74–80.

[13] J. Pauly, P. Le Roux, D. Nishimura, and A. Macovski. “Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm (NMR imaging)”. In: *IEEE transactions on medical imaging* 10.1 (1991), pp. 53–65.

[14] L. Zuo, Y. Liu, Y. Xue, B. E. Dewey, **S. W. Remedios**, S. P. Hays, M. Bilgel, E. M. Mowry, S. D. Newsome, P. A. Calabresi, S. M. Resnick, J. L. Prince, and A. Carass. “HACA3: A unified approach for multi-site MR image harmonization”. In: *Computerized Medical Imaging and Graphics* 109 (2023), p. 102285.

[15] S. Wei, **S. W. Remedios**, B. E. Dewey, Z. Bian, S. Wang, J. Chen, B. M. Jedynak, shiv saidha, P. A. Calabresi, A. Carass, and J. L. Prince. “Optical Coherence Tomography Harmonization with Anatomy-Guided Latent Metric Schrödinger Bridges”. In: *The Thirty-ninth Annual Conference on Neural Information Processing Systems*. 2025.