Haoyu Dong

LinkedIn | Google Scholar | Email | Phone: +1 858 500-2825

Research Interests

My research focuses on foundation models and data-centric AI for medical imaging. I build large-scale vision and vision-language foundation models to enable universal and robust image analysis, and develop in-context learning approaches for image segmentation. I also design multimodal agents that integrate imaging with language and tool use, and am broadly interested in anomaly detection, test-time adaptation, and clinical decision support.

Education

Duke University *Durham, NC, USA*
Ph.D. in Electrical & Computer Engineering *Jun. 2022 – Jan. 2027 (Expected)*

- Advisor: Prof. Maciej A. Mazurowski

Duke University *Durham, NC, USA*
M.S. in Computer Science *Aug. 2019 – Jun. 2021*

University of California, San Diego *La Jolla, CA, USA*
B.S. in Mathematics–Computer Science & B.S. in Cognitive Science *Sep. 2015 – Jun. 2019*

Industry Experience

Siemens Healthineers *Jun. 2025 – Aug. 2025*
AI Research Intern (Mentor: Dr. Han Liu; Manager: Dr. Sasa Grbić) *Princeton, NJ, USA*

- Developed a training-free, in-context segmentation framework with fixed support–query prompting that reduces annotation overhead and scales to **14K** MRI/CT volumes (full preprocessing + evaluation). With only **5** supports per task—*without retraining*—it approaches a supervised nnU-Net trained on the same 5 supports, achieving **52.0%** vs **60.0%** Dice on our internal test set and **54.9%** vs **56.9%** Dice on AMOS CT val set; results prepared for CVPR submission.

Gradient Health *May 2020 – Sep. 2020*
Fellowship Recipient (Advisor: Ouwen Huang) *Durham, NC, USA*

- Built a private multi-modal dataset with **2.6M** patients' diagnostic images and paired radiology reports; designed a BYOL-based pretraining pipeline that incorporated language information for stronger representations, delivering a **+3%** accuracy improvement over a vanilla ResNet and enabling robust internal MRI/CT benchmarks.

Research Experience

Mazurowski's Lab, Duke University *Jul. 2022 – Present*
Graduate Research Associate (Mentor: Prof. Maciej A. Mazurowski) *Durham, NC, USA*

- **Built MRI foundation models that outperform strong baselines:** Pretrained **MRI-CORE** on **7M** MRI slices using modern self-supervised objectives, transferring to downstream segmentation with **+5% Dice** over SAM in few-shot settings (*Nature Biomedical Engineering submission*; [PDF]); currently building a **3D, multimodal breast MRI** foundation model that aligns volumetric imaging with paired radiology reports for language-guided analysis and zero/few-shot transfer.
- **Delivered practical adaptation recipes & a universal muscle model:** Comprehensive study on adapting foundation models for medical segmentation—comparing prompt design, PEFT (e.g., adapters/LoRA), and light finetuning to produce compute-aware guidelines (*MELBA 2025*; [PDF]); **SegmentAnyMuscle**, a universal muscle segmentation model emphasizing cross-location generalization and scalable evaluation across acquisition protocols (*The Lancet Digital Health submission*; [PDF]).
- **Validated promptable segmentation at scale (SAM-1/SAM-2):** Broad evaluation of promptable segmentation for clinical imaging, including a **SAM-1** study across **19** datasets, **5** modalities, and **18** body locations (*MedIA 2023*; [PDF]), and **SAM-2** extensions to **2D** and **3D** medical images (*TBME submission*; [PDF]), establishing data-efficient, interactive workflows that transfer across scanners/sites.

- **Increased real-world robustness with test-time adaptation (InTEnt):** Single-image TTA method that adjusts batch-norm statistics via integrated entropy weighting; improved out-of-domain segmentation by **+2.9% Dice** over the runner-up approach (*CVPR 2024 Workshop Oral*; [PDF]).

Mazurowski's Lab, Duke University *Jun. 2021 – Jun. 2022*
Associate in Research (Advisor: Prof. Maciej A. Mazurowski) *Durham, NC, USA*

- Revisited anomaly detection for high-resolution digital breast tomosynthesis and chest X-ray by proposing **SWSSL** with sliding-window training, removing resizing from augmentations, and adding a **continuity detection** head, surpassing prior SOTA with **+5.0% AUROC** on tomosynthesis and showing strong gains on chest X-ray (*IEEE TMI 2023*; [PDF]).

Sapiro's Lab, Duke University *Sep. 2019 – May 2021*
Graduate Research Assistant (Advisor: Prof. Guillermo Sapiro) *Durham, NC, USA*

- Built a joint visual–semantic space for text-conditioned image retrieval via **iterative point alignment** and **manifold-distance** search, improving recall by **+2%** on FashionIQ, **+5%** on CSS, and **+30%** on a new dataset (*CVPR 2021 Workshop*; [PDF]).

Machine Learning, Perception, and Cognition Lab, UC San Diego *Jul. 2018 – May 2019*
Research Assistant (Advisor: Prof. Zhuowen Tu) *La Jolla, CA, USA*

- Mitigated exposure bias in autoregressive language models by introducing **multi-entropy sampling** and **multi-range reinforcing**, outperforming GAN-styled text generation by **+3.5%** in F-score BLEU [PDF].

Publications (*Alphabetical; †Mentor)

Refereed Journal & Conference Papers

[P1] Yaqian Chen, Lin Li, Hanxue Gu, **Haoyu Dong**, Derek L. Nguyen, Allan D. Kirk, Maciej A. Mazurowski, E. Shelley Hwang: *Breast density in MRI: an AI-based quantification and relationship to assessment in mammography.* In **npj Breast Cancer**. 2025.

[P2] Hanxue Gu, Roy Colglazier, **Haoyu Dong**, Jikai Zhang, Yaqian Chen, Zafer Yildiz, Yuwen Chen, ..., Maciej A. Mazurowski: *SegmentAnyBone: A universal model that segments any bone at any location on MRI.* In **Medical Image Analysis**. April 2025.

[P3] Pedro R. A. S. Bassi, Wenxuan Li, Yucheng Tang, Fabian Isensee, ..., **Haoyu Dong**, ..., Zongwei Zhou, Alan L. Yuille: *Touchstone benchmark: Are we on the right way for evaluating AI algorithms for medical segmentation?.* In **Advances in Neural Information Processing Systems (NeurIPS)**, December 2024.

[P4] Nicholas Konz, Yuwen Chen, **Haoyu Dong**, Maciej A. Mazurowski: *Anatomically-controllable medical image generation with segmentation-guided diffusion models.* In **International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)**, October 2024.

[P5] Binxu Li, Tiankai Yan, Yuanting Pan, Zhe Xu, Jie Luo, Ruiyang Ji, Shilong Liu, **Haoyu Dong**, Zihao Lin, Yixin Wang: *MMedAgent: Learning to use medical tools with multi-modal agent.* In **Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)**, July 2024.

[P6] Hanxue Gu*, **Haoyu Dong***, Jichen Yang, Maciej A. Mazurowski: *How to build the best medical image segmentation algorithm using foundation models: a comprehensive empirical study with Segment Anything Model.* In **Machine Learning for Biomedical Imaging (MELBA) Journal**, Volume 3, Article 2025:006, May 2025.

[P7] Yixin Wang, Zihao Lin, Zhe Xu, **Haoyu Dong**, Jie Luo, Jiang Tian, Zhongchao Shi, Lifu Huang, Yang Zhang, Jianping Fan, Zhiqiang He: *Trust it or not: Confidence-guided automatic radiology report generation.* In **Neurocomputing**, April 2024.

[P8] Christopher O. Lew, Majid Harouni, Ella R. Kirksey, Elianne J. Kang, **Haoyu Dong**, Hanxue Gu, Lars J. Grimm, Ruth Walsh, Dorothy A. Lowell, Maciej A. Mazurowski: *A publicly available deep learning model and dataset for segmentation of breast, fibroglandular tissue, and vessels in breast MRI.* In **Scientific Reports**, Volume 14, Issue 1, Page 5383. March 2024.

[P9] **Haoyu Dong**, Nicholas Konz, Hanxue Gu, Maciej A. Mazurowski: *Medical image segmentation with intent: Integrated entropy weighting for single image test-time adaptation.* In **Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops – DEF-AI-MIA**, 2024.

[P10] Maciej A. Mazurowski†, **Haoyu Dong**, Hanxue Gu, Jichen Yang, Nicholas Konz, Yixin Zhang: *Segment anything model for medical image analysis: an experimental study.* In **Medical Image Analysis**, Volume 89, October 2023.

[P11] **Haoyu Dong**, Yifan Zhang, Hanxue Gu, Nicholas Konz, Yixin Zhang, Maciej A. Mazurowski: *SWSSL: Sliding window-based self-supervised learning for anomaly detection in high-resolution images.* In **IEEE Transactions on Medical Imaging**, Volume 42, Issue 12, Pages 3860–3870. September 2023.

[P12] Nicholas Konz, **Haoyu Dong**, Maciej A. Mazurowski: *Unsupervised anomaly localization in high-resolution breast scans using deep pluralistic image completion.* In **Medical Image Analysis**, Volume 87, Page 102836. July 2023.

[P13] Yifan Zhang, **Haoyu Dong**, Nicholas Konz, Hanxue Gu, Maciej A. Mazurowski: *Lightweight transformer backbone for medical object detection.* In **MICCAI Workshop on Cancer Prevention through Early Detection**, Pages 47–56. September 2022. Springer Nature Switzerland.

[P14] Nicholas Konz, Hanxue Gu, **Haoyu Dong**, Maciej A. Mazurowski: *The intrinsic manifolds of radiological images and their role in deep learning.* In **International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)**, Pages 684–694. September 2022. Springer Nature Switzerland.

[P15] **Haoyu Dong**, Ze Wang, Qiang Qiu, Guillermo Sapiro: *Using text to teach image retrieval.* In **Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)**, 2021.

Submitting & Under Review

[R1] Roy Colglazier*, Jisoo Lee*, **Haoyu Dong***, Hanxue Gu, Yaqian Chen, Joseph Cao, Zafer Yildiz, Zhonghao Liu, Nicholas Konz, Jichen Yang, Jikai Zhang, Yuwen Chen, Lin Li, Adrian Camarena, Maciej A. Mazurowski: *SegmentAnyMuscle: A universal muscle segmentation model across different locations in MRI.* Submitting to **The Lancet Digital Health**.

[R2] **Haoyu Dong**, Yuwen Chen, Hanxue Gu, Nicholas Konz, Yaqian Chen, Qihang Li, Maciej A. Mazurowski: *MRI-CORE: A Foundation Model for Magnetic Resonance Imaging.* Submitting to **Nature Biomedical Engineering**.

[R3] Yaqian Chen, Hanxue Gu, **Haoyu Dong**, Qihang Li, Yuwen Chen, Nicholas Konz, Lin Li, Maciej A. Mazurowski: *GuidedMorph: Two-Stage Deformable Registration for Breast MRI.* Submitting to **IEEE Journal of Biomedical and Health Informatics (JBHI)**.

[R4] Yuwen Chen, Zafer Yildiz, Qihang Li, Yaqian Chen, **Haoyu Dong**, Hanxue Gu, Nicholas Konz, Maciej A. Mazurowski: *Accelerating Volumetric Medical Image Annotation via Short-Long Memory SAM 2.* Submitting to **IEEE Transactions on Medical Imaging (TMI)**.

[R5] Yaqian Chen, Hanxue Gu, Yuwen Chen, Jichen Yang, **Haoyu Dong**, Joseph Y. Cao, Adrian Camarena, Christopher Mantyh, Roy Colglazier, Maciej A. Mazurowski: *Automated muscle and fat segmentation in computed tomography for comprehensive body composition analysis.* Submitting to **???**.

[R6] Nicholas Konz, Richard Osuala, Preeti Verma, Yuwen Chen, Hanxue Gu, **Haoyu Dong**, Yaqian Chen, Andrew Marshall, Lidia Garrucho, Kaisar Kushibar, Daniel M. Lang, Gene S. Kim, Lars J. Grimm, John M. Lewin, James S. Duncan, Julia A. Schnabel, Oliver Diaz, Karim Lekadir, Maciej A. Mazurowski: *Fréchet Radiomic Distance (FRD): A Versatile Metric for Comparing Medical Imaging Datasets.* Submitting to **Medical Image Analysis (MIA)**.

[R7] **Haoyu Dong**, Hanxue Gu, Yaqian Chen, Jichen Yang, Yuwen Chen, Maciej A. Mazurowski: *Segment anything model 2: an application to 2D and 3D medical images.* Submitting to **IEEE Transactions on Biomedical Engineering (TBME)**.

[R8] Yuwen Chen, Nicholas Konz, Hanxue Gu, **Haoyu Dong**, Yaqian Chen, Lin Li, Jisoo Lee, Maciej A. Mazurowski: *ContourDiff: Unpaired image translation with contour-guided diffusion models.* Submitting to **Machine Learning for Biomedical Imaging (MELBA) Journal**.

Non-Archival Workshops & Symposiums Papers

[N1] Nicholas Konz, Yuwen Chen, Hanxue Gu, **Haoyu Dong**, Yaqian Chen, Maciej A. Mazurowski: *RaD: A Metric for Medical Image Distribution Comparison in Out-of-Domain Detection and Other Applications.* In **arXiv e-prints**, December 2024.

[N2] Nicholas Konz, Yuwen Chen, Hanxue Gu, **Haoyu Dong**, Maciej A. Mazurowski: *Rethinking Perceptual Metrics*

for Medical Image Translation. In **Medical Imaging with Deep Learning (MIDL 2024) Short Papers**, April 2024.

[N3] Jikai Zhang, Zafer Yildiz, Hanxue Gu, **Haoyu Dong**, Maciej A. Mazurowski: *SAM-Geo3D: A Geometrical Method to Extend SAM to 3D*. In **Medical Imaging with Deep Learning (MIDL 2024) Short Papers**, April 2024.

[N4] Yixin Wang, Zihao Lin, **Haoyu Dong**: *Rethinking Medical Report Generation: Disease Revealing Enhancement with Knowledge Graph*. In **Interpretable Machine Learning in Healthcare (IMLH), International Conference on Machine Learning (ICML Workshop)**, July 2023.

[N5] Hanxue Gu, **Haoyu Dong**, Nicholas Konz, Maciej A. Mazurowski: *A systematic study of the foreground-background imbalance problem in deep learning for object detection*. In **arXiv preprint arXiv:2306.16539**, June 2023.

Academic Services

Reviewer

- Conferences: ICLR, NeurIPS, ICML, CVPR, ICCV, ECCV, MICCAI, MIDL
- Journals: IEEE TMI (Distinguished Reviewer, Silver Level 2023-24); IEEE JBHI; IEEE TNNLS; IEEE TBME; Medical Image Analysis; Artificial Intelligence Review