Keanu Nichols

kmn5409@bu.edu | cs-people.bu.edu/kmn5409

Education

Boston University — Boston, MA
Ph.D. in Computer Science — *September 2023 – Expected Fall 2027*

- **Research Advisors:** Prof. Deepti Ghadiyaram and Prof. Bryan A. Plummer

Boston University — Boston, MA
M.S. in Artificial Intelligence; GPA: 3.95/4.00 — *September 2021 – December 2023*

The University of the West Indies — Trinidad and Tobago
B.Sc. in Computer Science (Special); GPA: 4.03/4.30 — *August 2017 – May 2021*

Publications

Multi-axis Analysis of Image Manipulation Localization
K. Nichols, D. Appapogu, G. Biamby, D. Bashkirova, A. Rohrbach, and B. A. Plummer.
Transactions on Machine Learning Research (TMLR), 2026. ***Accepted****.* OpenReview | arXiv.

Seeing Isn't Orienting: A Cognitively Grounded Hierarchical Benchmark for Object Orientation in MLLMs
N. Tasnim*, **K. Nichols***, Y. Yan, N. Ikechukwu, E. Zou, D. Ghadiyaram, and B. A. Plummer.
*European Conference on Computer Vision (ECCV), 2026. *Equal contribution.* ECCV paper | arXiv.

FuTCR: Future-Targeted Contrast and Repulsion for Continual Panoptic Segmentation
N. Ikechukwu*, **K. Nichols***, D. Ghadiyaram, and B. A. Plummer.
*arXiv preprint arXiv:2605.12451, 2026. *Equal contribution.* arXiv.

Estimating Deforestation Using Machine Learning Algorithms
K. Nichols and P. Hosein.
International Conference on Intelligent Data Science Technologies and Applications (IDSTA), 2021. Paper.

Recommendations for Long-Term Profit Optimization
P. Hosein, I. Rahaman, **K. Nichols**, and K. Maharaj.
Workshop on the Impact of Recommender Systems at ACM RecSys, 2019. Best Paper Award. Paper.

Selected Research Projects

DORI: Object-Orientation Reasoning in MLLMs — *ECCV 2026*

- Co-developed DORI, a hierarchical benchmark of 33,656 multiple-choice questions over 13,652 images, decomposing object orientation into four dimensions evaluated at coarse and granular levels; evaluated 26 state-of-the-art MLLMs.
- Best-performing models reached only 64.2% on coarse and 42.9% on granular orientation judgments, with major failures on compound rotations and reference-frame reasoning.

AUDITS: Image Manipulation Localization — *TMLR 2026 (Accepted)*

- Introduced AUDITS, a ∼530K image-mask benchmark spanning two image domains (news and COCO), 11 diffusion-based manipulation techniques, and axes of manipulation type, image source, manipulation size, and perceptual quality.
- Evaluations revealed a 23-point AUC decline for MMFusion under unseen manipulation types ($96.8 \rightarrow 73.8$), exposing generalization failures in image manipulation localization.

Temporal Reasoning and Information Flow in Video-LLMs — *In Progress*

- Built layer-wise probing pipelines for InternVideo2, InternVideo2.5-Chat-8B, and Qwen2.5-VL-7B, evaluating linear, MLP, and attention-based readouts across Arrow-of-Time, Something-Something-v2, and synthetic temporal-ordering tasks; developing lightweight interventions to route temporally informative visual features into language-model layers.

FuTCR: Continual Panoptic Segmentation *Under Review*

- Co-developed (equal contribution) a Mask2Former-based continual panoptic segmentation framework that applies pixel-to-region contrastive learning and feature repulsion to structure unlabeled background regions and reserve representational space for future categories.
- Evaluated across six ADE20K class-incremental settings; improved relative new-class Panoptic Quality by up to 28% ($17.3 \rightarrow 22.3$ on ADE20K 100-5) while preserving or improving base-class PQ by up to 4%.

Additional Research Experience

Machine Learning Undergraduate Researcher April 2018 – August 2021
TTLAB *Trinidad and Tobago*

- Contributed to an AI for Earth computer-vision proposal awarded $70,800; evaluated segmentation models for drone-based weed detection and developed a Best Paper-winning profit-aware recommender system.

Deep Learning Research Intern June 2019 – August 2019
NASA Ames Research Center *Mountain View, CA*

- Developed a fully convolutional neural network in TensorFlow/Keras to segment clouds in Landsat satellite imagery.

Technical Skills

Programming/Frameworks: Python, C++, Bash, PyTorch, TensorFlow/Keras, Hugging Face Transformers, scikit-learn
Research/Infrastructure: Vision-language models, video understanding, panoptic segmentation, image forensics, continual learning, representation probing, Linux, Git, Slurm, multi-GPU training