Agnibh Dasgupta

adg002@gmail.com
+1 (435) 754-9599
Omaha, NE, USA

Website: cent664.github.io
Github: github.com/cent664
Linkedin: linkedin.com/in/cent664
Google Scholar:
https://tinyurl.com/cent-googlescholar

Research Interests

Representation learning, Watermarking, Multimodality, Provenance

Education

University of Nebraska Omaha, Omaha, Nebraska, USA
Ph.D., Computing & Information Science 2021 – 2026
Dissertation: *Invariant Feature Learning in AI Models: Exploring Representation Spaces for Robust Watermarking and Beyond* (Dr. Xin Zhong)

Utah State University, Logan, Utah, USA
M.S., Computer Science 2018 – 2020
Thesis: *Deep Q-Learning Applied to Stock Trading* (Dr. Nicholas Flann)

West Bengal University of Technology, Kolkata, India
B.Tech., Computer Science 2013 – 2017

Professional Experience

University of Nebraska Omaha, Graduate Researcher 2021 – 2026

- Developed representation-level approaches for robust watermarking and model attribution across vision and language models, including semantically grounded invariant feature learning, black-box LLM watermarking, and geometric analysis of invariant subspaces in pretrained LLMs.
- Collaborated with medical practitioners on retinal vascular leakage detection from ultra-wide-field fluorescein angiography scans; trained patch-based classifiers and presented findings at the American Academy of Ophthalmology (AAO), 2025.

Independent Projects 2026

- Designed and deployed Siosa's Library, a production agentic RAG system for grounded Q&A; the pipeline integrates a LangGraph planner, multi-query hybrid retrieval, cross-encoder reranking, and LLM-as-judge evaluation, shipped with FastAPI.

Publications

Conference Publications

[C1] A. A. Tanvir, **A. Dasgupta**, X. Zhong. **TIACam: Text-Anchored Invariant Feature Learning with Auto-Augmentation for Camera-Robust Zero-Watermarking.** *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026. 2026

[C2] H. Yoon, X. Zhong, **A. Dasgupta**, G. Nugent, G. Trainin. **Leveraging Artificial Intelligence (AI) to Enhance Computer Science Instruction.** *IEEE Frontiers in Education Conference (FIE)*, 2024. 2024

[C3] **A. Dasgupta**, X. Zhong. **Robust Image Watermarking Based on Cross-Attention and Invariant Domain Learning.** *International Conference on Computational Science and Computational Intelligence (CSCI)*, 2023. 2023

[C4] N. Khatri, **A. Dasgupta**, Y. Shen, X. Zhong, F. Y. Shih. **Perspective Transformation Layer.** *International Conference on Computational Science and Computational Intelligence (CSCI)*, 2022. 2022

Journal Publications

[J1] **A. Dasgupta**, A. A. Tanvir, X. Zhong. **Watermarking Language Models through Language Models.** *IEEE Transactions on Artificial Intelligence*, pp. 1–10, 2025. DOI: 10.1109/TAI.2025.3605117 2025

Book Chapters

[B1] **A. Dasgupta**, X. Zhong. **Enhanced image watermarking through cross-attention and noise-invariant domain learning.** *Book Chapter, Imaging Science: Computer Vision, Image and Signal Processing, Pattern Recognition*, IEEE, 2024. 2024

Submitted / Under Review

[S1] **A. Dasgupta**, A. A. Tanvir, X. Zhong. **Invariant Features in Language Models: Geometric Characterization and Model Attribution.**

Theses and Other Research

[O1] **Invariant Feature Learning in AI Models: Exploring Representation Spaces for Robust Watermarking and Beyond** 2026
- Doctoral dissertation investigating invariant representation learning across vision and language domains, using watermarking as a quantitative probe of representational stability.

[O2] **Retinal Vascular Leakage Detection** 2025
- Trained classifiers on patches of ultra-wide-field fluorescein angiography scans to detect areas of leakage, in collaboration with medical practitioners. Presented at the American Academy of Ophthalmology (AAO), 2025.

[O3] **Sentiment Analysis on Electric Vehicle Discourse** 2022
- Analyzed public sentiment using lexical and Word2Vec-LSTM approaches on scraped Twitter data.

[O4] **Deep Q-Learning Applied to Stock Trading** 2020
- Designed and evaluated a reinforcement learning trading system using auto-scraped S&P 500 index data as part of the M.S. thesis.

Teaching Experience

Responsibilities across all courses included grading, holding office hours, and lab supervision where applicable. For Capstone classes, I provided guidance on project scope, feasibility, and implementation strategy for senior undergraduate projects.

University of Nebraska Omaha, Graduate Teaching Assistant 2023–2026

- **CSCI1620: Introduction to Computer Science II** (Undergraduate), Fall 2023 (class size: 59), Fall 2024 (class size: 32)
  Course overview: Covers object-oriented programming, data structures, and algorithm design in an introductory CS sequence. Designed quizzes with progressive difficulty structures in Fall 2024.
- **CSCI4970: Capstone Project** (Undergraduate), Spring 2024 (class size: 67), Fall 2025 (class size: 58), Spring 2026 (class size: 90)
  Course overview: Integrative senior project course in which students design and deliver a substantial service-learning oriented software or research project.

University of Nebraska Omaha, Intern Supervisor May – Aug 2021

- Interviewed and shortlisted high school intern candidates for AI research positions at the DLM Lab.

- Supervised interns on active research projects in representation learning and computer vision, adapting technical content for students with varying levels of prior knowledge.

Utah State University, Graduate Teaching Assistant 2018–2020

- **CS-4700: Programming Languages** (Undergraduate), Fall 2018 (class size: 49), Spring 2019 (class size: 54), Fall 2019 (class size: 64), Spring 2020 (class size: 69)
  Course overview: Covers paradigms, syntax, semantics, and implementation of programming languages including functional and logic programming. Designed assignments with a progressive difficulty structure in Spring 2020.
- **ECE-1410: Computer Programming for ECE II** (Undergraduate), Fall 2020 (class size: 26)
  Course overview: Continuation of introductory programming for electrical and computer engineering students, covering data structures and algorithmic problem solving.
- **CS-5890/6890: Temporal Data Mining** (Cross-listed), Fall 2020 (class size: 13)
  Course overview: Cross-listed course covering techniques for mining patterns and knowledge from time-series and sequential data.

Grants and Awards

Graduate Research and Creative Activity (GRACA) Fellowship, $5,000 2025
University of Nebraska Omaha, Office of Research and Creative Activity

- Supported summer research on training a neural network augmenter for automatic image augmentations; presented at UNO Research & Creative Activity Fair

Graduate Research and Creative Activity (GRACA) Fellowship, $5,000 2021
University of Nebraska Omaha, Office of Research and Creative Activity

- Supported summer research on invariant domain learning for image watermarking; presented at UNO Research & Creative Activity Fair

Service

External Reviewer

- ACCV, 2022
- Scientific Reports (Nature), 2023
- Springer Nature, 2025

Software and Code

Siosa's Library (2026) — `https://github.com/cent664/Siosa`

A production wiki-grounded RAG application for Path of Exile 1 mechanics Q&A, deployed at poesiosa.net. The system retrieves live content from poewiki.net via multi-query fusion and title probes, reranks passages with a cross-encoder, and generates cited answers through swappable Claude/GPT-4 backends. A LangGraph planner adds targeted search terms for multi-topic questions before a single fused retrieval pass. Built with FastAPI and React; shipped on Railway with Docker, GitHub Actions CI, optional LLM-as-judge scoring, and a dev trace UI for pipeline debugging.

LLMID (2026) — `https://github.com/cent664/LLMID`

A geometric framework for identifying invariant subspaces in pretrained LLMs across 9 open-source model families, without additional training. The framework applies contrastive sensitivity decomposition to separate semantic and nuisance directions in hidden representations, validated through causal patching experiments. The discovered invariant zones achieve 90%+ zero-shot model attribution accuracy across base, fine-tuned, and distilled model variants.

TIACam (2026) — https://github.com/tanvir097/TIACam

A camera-robust zero-watermarking framework that combines CLIP image–text features with a learned adversarial augmentation module to ground invariant representations in semantic content. The auto-augmentor discovers camera-like distortions through differentiable operators, removing the need for manually designed augmentation families. Achieves 95%+ watermark recovery on ImageNet and Visual Genome datasets under real-world camera distortions.

LLMWM (2025) — https://github.com/cent664/LLMWM

A model-agnostic prompt-based watermarking framework for LLM-generated text that operates without access to model weights or decoding internals. The pipeline was evaluated across 25 open-source LLMs including fine-tuned and distilled variants trained using LoRA and DeepSpeed. Achieves 88%+ detection accuracy under paraphrasing, fine-tuning, and adversarial prompt attacks.

SSRIW (2023) — https://github.com/cent664/SSRIW

A ViT-based cross-attention watermarking system that directly trains an invariant domain alongside the watermarking task using a triplet loss objective. The cross-attention mechanism localizes semantically relevant regions for watermark embedding, and the invariant domain provides robustness under compound distortions. Achieves state-of-the-art comparable accuracy under common noise types with improved performance under combined augmentations.

PTL (2022) — https://github.com/kcnishan/Perspective_Transformation_Layer

A lightweight differentiable layer that learns 2D projections of 3D views for multiple viewpoints, implemented as a drop-in module within standard deep learning pipelines. Rather than relying on fixed augmentation pipelines, PTL internalizes perspective warping as a learnable operator, allowing the model to adaptively select viewpoint transformations that improve downstream task accuracy.

Technical Skills

ML concepts:	Representation Learning, Watermarking, Computer Vision, NLP, LLMs
Frameworks:	PyTorch, TensorFlow, HuggingFace Transformers, scikit-learn
Programming:	Python, Bash, CUDA, Git, Linux
Distributed:	DeepSpeed, LoRA, PyTorch DDP, Slurm/HPC
Agentic AI:	LangGraph, ChromaDB, hybrid retrieval, cross-encoder reranking
Deployment:	FastAPI, OpenAI & Anthropic APIs

References

Dr. Xin Zhong
Associate Professor
University of Nebraska Omaha
xzhong@unomaha.edu
+1 (402) 554-6637
PKI 281A, Omaha, NE 68182

Dr. Qiuming Zhu
Professor
University of Nebraska Omaha
qzhu@unomaha.edu
+1 (402) 554-3685
PKI 281D, Omaha, NE 68182

Dr. Dario Ghersi
Associate Professor
University of Nebraska Omaha
dghersi@unomaha.edu
+1 (402) 554-6004
PKI 173B, Omaha, NE 68182