Hongyu Tu

240-706-6288 | tu141@purdue.edu | linkedin.com/in/h-tu/ | h-tu.github.io

Education

Purdue University — West Lafayette, IN
Doctor of Philosophy in Construction Management Technology — *Aug. 2026 – Present*

University of Massachusetts Amherst — Amherst, MA
Master of Science in Computer Science — *Aug. 2021 – May 2023*

University of Maryland, College Park — College Park, MD
Bachelor of Science in Computer Engineering — *Aug. 2017 – May 2021*

Publication

- **H. Tu**, S. Shorewala, T. Ma and V. Thost. **"Retrosynthesis Prediction Revisited."** *NeurIPS 2022 AI for Science: Progress and Promises, **2022***
  - **Project Overview:** Studied how three types of machine learning approachs perform on retrosynthesis, a multi-step multi-route planning problem in organic chemistry, including Transformer-based **template-free** (e.g., Chemformer) and **Semi-template** GNNs (e.g., GraphRetro) against traditional **template-based** baselines (e.g., NeuralSym, MHN). Gained insight toward evaluating result's creativity and validness in multi-step search scenarios beyond merely single-step.
  - **System Development:** Built wrappers for each models with different environment requirement and constrains, resulting in a standardized framework for execution and evaluation tasks. Proposed and curated new datasets based on time-splits (USPTO-rt) and route diversity (USPTO-rd) to fully examine the model robustness in cases of distribution shifts.
  - **Key Finding:** Discovered that prior popular single-step accuracy metrics often do not guarantee good performance in actual planning tasks, as Transformer models usually produce results with poor diversity needed for effective search comparing to traditional template-based methods. This further proves the gap between single-step prediction and the uncertainty management required for multi-step reasoning.
  - **Impact:** Proposed a new multi-step score (**mss**) as metric to estimate a model's planning potential in the long run. This testing framework was also made open-source so the community with can use the tools to better measure generative models' generalization and diversity caoability.

Experience

R&D Machine Learning Engineer — Feb. 2024 – Apr. 2026
LCE Optics (Lianchuang Electronics US Research Lab) — *Remote*

- **Deep Learning-Based SLAM:** Replaced traditional feature descriptors in ORB-SLAM with deep learning-based SuperPoint, SuperGlue, aiming to improve localization robustness in texture-less environments. Used **TensorRT** for speedups and achieved real-time performance on computationally-limited edge devices.

- **Multimodal Sensor Fusion:** Developed sensor fusion algorithms that combine onboard camera RGB feed with LiDAR depth data, these data with additional dimensions can help improve the state estimation accuracy for tasks like autonomous vehicle navigation and optical flow-based map building.
- **Sim-to-Real Robotics Control:** Trained control policies for manipulation tasks like lens assembly in factory automation settings on UR5 Robotic Arm using **NVIDIA Isaac Sim**. Compared performance between different machine learning approaches including Reinforcement Learning and SOTA Vision-Language-Action models (OpenVLA).
- **DRL for Optical Design:** Applied Deep Reinforcement Learning (**LensRL**) for automatic optical lens optimization. Used differentiable ray-tracing engines (**DiffOptics**) as the ML backbone to optimize high-dimensional non-convex parameters, solving the inverse design problem and addressed Memory bottleneck, which allowed more rays being traced and resulting in less noise and quicker convergence.

ARVR System Research Assistant June – Dec. 2023
Futurewei Technologies, IC Lab *Austin, TX*

- **3D High Precision Motion Tracking:** Built a multi-modal tracking pipeline that combines the accurate distance data that captures torso position from 3D depth sensors with high-resolution data from vision-based module to capture fine-grained details (fingers and facial expressions). The two modules run in parallel on GPU and CPU without putting too much strain on either single hardware, achieved higher accuracy without sacrificing real-time performance.
- **Custom Data Format & Toolkit:** Created a dataset by combining in-house recordings and Adobe Mixamo assets into a custom defined tree-based data structure to represent skeleton hierarchy and built utility tools to convert between different formats, allowing the dataset later to be used for different models' fine-tuning stage for better downstream task performance.
- **LLM-Based Motion Correction and Completion:** Proposed to use Large Language Models to repair faulty tracking data and predict future frames to reduce latency. By modeling the **spatial** (joint-wise) and **temporal** (frame-wise) dependencies, the model can predict and patch missing motion frames or certain missing joints within a frame from noisy inputs.

Undergraduate Research Assistant Jun. 2020 – Dec. 2020
University of Maryland, College Park *College Park, MD*

- **Physics-Informed RL Environments:** Built a **Reinforcement Learning** gym environment and used Unity ML-Agents to train robotic arm manipulation policies. Proposed to fix the "Sim-to-Real" gap by creating custom low-resolution mesh colliders, resolving convergence issues caused by high-fidelity geometry artifacts.
- **Geometric Simulation Optimization:** Fine-tuned physics parameters to balance simulation accuracy with computational efficiency. Solved the "collision overlay" problem where complex mesh interactions introduced uncertainty and prevented the agent from learning stable grasping policies.
- **Control Policies:** Partially achieved the objective to train agents to maximize object visibility and minimize displacement penalties, showing that policies for complex geometric tasks can be trained with only simplified physical proxies.

Undergraduate Research Assistant Aug. 2017 - Dec. 2018
University of Maryland, First-Year Innovation & Research Experience *College Park, MD*

- **Bio-Acoustic Signal Processing:** Built a preprocessing pipeline to transform raw, noisy echolocation data into time-frequency spectrograms using **Fast Fourier Transform (FFT)**. Implemented spectral denoising and temporal windowing to split 15ms call segments, and handling hardware-induced noise and environmental interference.
- **Deep Metric Learning:** Developed a **Triplet Network** architecture with a ResNet backbone to learn a discriminative embedding space for audio samples. Optimized the model using triplet loss to minimize intra-class variance while maximizing inter-class separation, allowing for the robust identification of individual bats from limited samples.

- **Applications:** Feed conservation of the endangered Mexican Fishing Bat to the triplet network and confirmed the model can handle tasks like distinguishing individuals based on subtle frequency modulation patterns and later be used for population monitoring.

Course Projects

Bilimemenet (COMPSCI 685) | *Python, Hugging Face, ByT5, Chinese Roberta, GPT-2*

- **Dataset Creation and Processing:** Processed and tokenized a large-scale unstructured dataset of over 1 million comments, handling extreme class imbalance and noise to create a robust training corpus for above mentioned NLP tasks.
- **Conditional Text Generation:** Fine-tuned a pre-trained **GPT-2** model to generate domain-specific text (Danmu) conditioned on video categories, effectively modeling the distribution of "in-joke" slang.
- **Discriminator-Guided Learning:** Used **Chinese-RoBERTa** to evaluate, filter generated outputs, guiding the generative model towards high-fidelity, category-aligned samples and remove out-of-sample results.

Auto Crossy Road (COMPSCI 682) | *Python, PyTorch*

- **Perception-Control Pipeline:** Designed a modular autonomous navigation stack that processes raw visual frames into occupancy grids using **YOLOv5** for real-time game content (play, obstacle, coin) detection.
- **Deep Reinforcement Learning:** Trained a **Deep Q-Network (DQN)** to solve dynamic path planning in stochastic environments, achieving a high score of 177 and significantly outperforming Supervised Learning baselines (average score 35) which failed to generalize.
- **Custom State Representation:** Created a discrete grid-based state space from continuous visual inputs, enabling the agent to handle dynamic occlusion and fast-moving traffic with reactive control policies.

Cross Modal Shape Generation Introspection (COMPSCI 674) | *PyTorch, Blender, SAM, SDF*

- **Deep Generative Reconstruction:** Implemented a pipeline using **Variational Auto-Decoders (VAD)** and **Signed Distance Functions (SDF)** to recover dense 3D geometry from sparse 2D sketches, addressing the inverse problem of shape generation.
- **Regularization for ill-posed problems:** Designed a custom loss function with regularization terms to stabilize the latent space, forcing the model to ignore random input noise and ensuring 3D consistency during editing tasks.
- **Vision-Language Model Integration:** Integrated the **Segment Anything Model (SAM)** and **3DHighlighter** to distinguish and pinpoint object parts via text prompts, allowing for precise, language-guided modifications of 3D structures.

Book Recommender System Using Review History (COMPSCI 646) | *PyTorch, UltraGCN*

- **Large-Scale Data Processing:** Processed the Goodreads dataset (15M+ interactions), implementing rigorous data cleaning and feature engineering pipelines to handle data sparsity and class imbalance.
- **Graph Representation Modeling:** Used **UltraGCN** to model complex user-item interactions, capturing high-dimensional structural connectivity in sparse data to improve recommendation diversity.
- **Baseline MLP Model:** Designed and trained a 4-layer Deep Neural Network (DNN) to predict user ratings based on genre-specific feature vectors, achieving a 92% hit rate for top-3 recommendations.

Scalable Image Rating System (COMPSCI 532) | *Python, PyTorch, MongoDB, CNN*

- **Complete Data Pipeline:** Built a scalable automated pipeline to scrape, index, and store over 70,000 images using **MongoDB**, built the system for continuous online data ingestion.

- **Data Cleaning:** Developed a statistical standardization method to decouple intrinsic image quality from popularity bias (follower count), ensuring the model learned robust visual features rather than metadata trends.
- **Image to Score Model:** Trained a custom **Convolutional Neural Network (CNN)** to predict aesthetic scores, outperforming linear baselines by effectively capturing non-linear visual patterns in noisy real-world data.

Technical Skills

Tools: NVIDIA Isaac Sim, ROS 2, Unity, Docker, Git
Languages: Python, Matlab, C/C++, Java, Swift, SQL, Assembly
Frameworks: Hugging Face Transformers, PyTorch, NumPy, SciPy, OpenCV, Flask, Pandas, Seaborn

SELECTED Coursework

Probabilistic Graphical Models | *COMPSCI 688* Spring 2022

- Implemented deep generative models, including Variational Autoencoders (VAE) and Normalizing Flows, to approximate complex data distributions.
- Built sampling algorithms from scratch, such as Hamiltonian Monte Carlo (HMC) and Metropolis-Hastings, to estimate posterior distributions and quantify uncertainty.
- Applied Variational Inference (VI) and Expectation-Maximization (EM) to solve parameter estimation problems in latent variable models.

Reinforcement Learning | *COMPSCI 687* Fall 2021

- Modeled sequential decision-making problems using Markov Decision Processes (MDPs) to optimize long-term rewards in stochastic and uncertain environments.
- Studied model-free control algorithms, including Monte Carlo and Policy Gradients, to solve problems where the environment dynamics are complex or unknown.
- Applied function approximation techniques, such as neural networks, to estimate value functions and scale solutions to large, continuous state spaces.
- Implemented and validated Deep RL algorithms (including Actor-Critic and n-step SARSA) for a stochastic Blackjack environment, performing extensive hyperparameter tuning to ensure stable convergence.

Advanced Natural Language Processing | *COMPSCI 685* Spring 2022

- Studied deep learning architectures with a focus on Transformers and self-attention mechanisms, specifically how transfer learning adapts pre-trained models to new tasks.
- Covered advanced modeling techniques including Vision-Language models, model distillation, and efficient Transformers to understand how to optimize large models.
- Implemented and fine-tuned neural networks (like BERT) using PyTorch, gaining practical experience in training models and evaluating experimental results.

Neural Networks & Machine Learning | *CS 682, CS 689* Fall 2021

- Studied the mathematical foundations of machine learning, focusing on probabilistic modeling and statistical estimation (MAP, Expectation-Maximization) to handle data uncertainty.
- Covered modern computer vision techniques, including Convolutional Neural Networks (CNNs) and standard architectures (ResNet) for image analysis.
- Learned about numerical optimization and generative models (GANs), which are fundamental for solving reconstruction and generation problems.
- Implemented learning algorithms from scratch in Python to understand the core math, while using PyTorch for building practical deep learning applications.

Signal Processing & Linear Systems | *ENEE 322, ENEE 222* Undergrad

- Learned the theoretical foundations of continuous and discrete-time signals, using Fourier, Laplace, and Z-transforms to analyze system behavior in both time and frequency domains.
- Studied sampling theory and signal reconstruction, understanding critical concepts like the Nyquist-Shannon theorem and aliasing, which are fundamental for processing raw sensor data.
- Analyzed Linear Time-Invariant (LTI) systems and designed digital filters (FIR/IIR) to process noisy signals, implementing these algorithms mathematically and in MATLAB.