

Yinuo Xu

xuyinuo@cis.upenn.edu | yinuoxu.com | [Linkedin](#) | [Github](#) | [Google Scholar](#)

Research Interests

Computational Biology, Histopathology and Spatial Transcriptomics, Foundation Models for Digital Pathology

Education

University of Pennsylvania Aug 2024 – May 2026

M.S. in Computer and Information Science, Research Assistant at Zhi Huang Lab.

- **GPA:** 4.0 / 4.0
- **Relevant Coursework:** Machine Learning, Natural Language Processing, Advanced Topics in Deep Learning, Bayesian Optimization, Theory of Machine Learning
- **Relevant Skillsets:** Computational analysis of: spatial transcriptomics (Visium, Visium HD, Xenium, CosMX), spatial proteomics (multiplex immunofluorescence imaging), computational biology, histopathology (H&E, IHC).

New York University Feb 2021 – May 2024

B.A. in Computer Science and Mathematics

- **GPA:** 3.8 / 4.0
- **Honors:** University Honors Scholar, Dean's List (2020–2023), *Cum Laude*
- **Relevant Coursework:** Algorithms, Machine Learning, Computer Graphics, Computer Systems Organization, Computer Simulation, Data Management and Analysis, Parallel Computing, Numerical Analysis, Real Analysis, Ordinary Differential Equations

Publications

Adaptive Multi-Scale Integration Unlocks Robust Cell Annotation in Histopathology Images. Under Review, CVPR 2026

Preprint: [arXiv:2511.13586](https://arxiv.org/abs/2511.13586)

Yinuo Xu, Yan Cui, Mingyao Li, Zhi Huang

We propose an upgraded NuClass, a cell-wise multi-scale framework that integrates nuclear morphology with large-field tissue context using a two-path architecture and a learnable GT-free fusion gate. Trained on 2M+ marker-guided cells from **spatial transcriptomics** data across 8 organs, it achieves up to 96% F1 on fully held-out cohorts and consistently outperforms strong frozen and fine-tuned baselines in robust cross-organ phenotype prediction.

NuClass: An ontology-driven vision–language foundation model for zero-shot nuclei classification. CVPR 2025
MMFM-BIOMED Workshop

Poster

Yinuo Xu, Jina Kim, Zijie Zhao, Zhi Huang

We introduce **NuClass**, a scalable and interpretable vision–language foundation model that aligns hierarchical cell-type descriptions with ViT-based single-cell feature extraction to enable zero-shot nuclei classification. Across a wide range of benchmarks, NuClass delivers strong and consistent performance, with accuracy gains up to +85% over prior vision–language methods.

A co-evolving agentic AI system for medical imaging analysis. Under Revision, Nature

Preprint: [arXiv:2509.20279](https://arxiv.org/abs/2509.20279)

Songhao Li, Jonathan Xu, Tiancheng Bao, Yuxuan Liu, Yuchen Liu, Yihang Liu, Lilin Wang, Wenhui Lei, Sheng Wang, **Yinuo Xu**, Yan Cui, Jialu Yao, Shunsuke Koga, Zhi Huang

We present TissueLab, a co-evolving agentic AI system that integrates modular medical-imaging tool factories, guideline-grounded decision pathways, and human-in-the-loop active learning. TissueLab autonomously constructs explainable workflows, supports real-time model refinement, and achieves state-of-the-art performance across pathology, radiology, and spatial omics tasks—enabling rapid, transparent, and clinically aligned biomedical analysis.

Learning Fine-grained Rewards on Video Generation Fakeness via Multimodal Language Models.

Under Review, ICLR 2026

Preprint: [arXiv:2509.22646](https://arxiv.org/abs/2509.22646)

Xingyu Fu*, Siyi Liu*, **Yinuo Xu**, Pan Lu, Guangqiuse Hu, Tianbo Yang, Taran Anantasagar, Christopher Shen, Yikai Mao, Yuanzhe Liu, Keyush Shah, Chung Un Lee, Yejin Choi, James Zou, Dan Roth, Chris Callison-Burch

We introduce **DeepTraceReward**, the first fine-grained, expert-annotated benchmark for training reward models to identify fake clues in AI-generated videos. We further conduct extensive experiments to train video fake clue reward models leveraging multimodal language models (LMs), which surpasses GPT 4.1 by 20.5% in fake clue identification.

Research Experience

Research Assistant, supervised by: [Prof. Zhi Huang](#) and [Prof. Lyle Ungar @ UPenn](#)

Feb 2025 - Present

- Developed **NuClass**, a cell-wise multi-scale framework integrating nucleus-centered morphology with wide-field tissue context via a dual-path architecture (Path_local + Path_global) and a GT-free gated fusion module.
- Built a **marker-guided Xenium dataset** with **2M+ single-cell labels across 8 organs and 16 cell types**, enabling scalable, reproducible, and fine-grained cell annotations.
- Designed uncertainty-aware training for contextual specialization, implemented tissue-conditioned FiLM adaptation, and optimized multi-scale fusion to achieve **up to 96% F1** on fully held-out cohorts.
- Performed comprehensive evaluation under cross-organ distribution shifts, demonstrating strong generalization, calibrated predictions, and interpretable per-cell reasoning.
- **Ongoing:** Developing a **cell-level world model** to simulate cell states and microenvironmental perturbations, enabling in-silico experiments for gene/protein expression prediction, cell-state transitions, and environment-conditioned generative modeling.

Research Assistant, supervised by: [Prof. Dan Roth @ Upenn](#)

Aug 2024 - Aug 2025

- Co-led the development of DeepTraceReward, a benchmark for evaluating video fakeness using temporally and spatially localized human annotations.
- Designed and standardized annotation schema for identifying multimodal artifacts (e.g., deformation, inconsistency, trajectory shifts) in **3.3K+** AI-generated videos.
- Trained reward models with multimodal LLMs, achieving a 20.5% improvement over GPT-4.1 in fine-grained fake clue identification.
- Investigated trade-offs between global video-level classification and fine-grained temporal signal modeling.

Research Assistant, supervised by: [Dr. Jake Hofman @ Microsoft NYC](#)

Jun 2024 - Aug 2024

- Completed intensive data science program, acquiring, cleaning, and analyzing real-world data.
- Collaborated on research projects with Microsoft Research scientists, enhancing teamwork skills.
- Applied machine learning to assess CitiBike's impact in NYC, providing measurable insights.
- Utilized Python and R for data analysis, delivering actionable insights on CitiBike deployment.

Research Assistant, supervised by: [Prof. Nan Xu @ NYU Shanghai](#)

April 2023 - May 2024

- Collected data in diverse formats, including pre-processing of CSV/JSON files by the pandas library.
- Performed web scraping and analysis using BeautifulSoup for efficient data extraction.
- Inspected and cleansed CUSMA data from 96 chapters, each containing over 50,000 data points, using pandas DataFrame and numpy for streamlined data manipulation and processing.
- Analyzed essential data metrics using pandas for improved data comprehension.

Selected Engineering Project

Classification of Model-Generated Text: Comparing GPT-4, Grok, and LLaMA

Mar 2023 - May 2023

- Investigated the classification of AI-generated text by developing a system to distinguish outputs from three leading Large Language Models (LLMs): GPT-4, Grok, and LLaMA.
- Achieved near-perfect classification accuracy (e.g., 97.9% in a three-way task) by implementing a Gradient

Boosting Classifier that leveraged linguistic features such as POS tags, n-grams, and text length.

- Demonstrated the scalability and robustness of the classification methodology by successfully extending it from binary to three-way classification and validating its performance across different content domains (technology and culture/arts), highlighting the existence of model-specific "fingerprints."

Digital Acoustic Simulation of the Piano

Mar 2023 - May 2023

- Utilized MATLAB, finite difference methods, and second order wave equations to simulate piano strings.
- Applied finite difference methods and a time-stepping scheme to discretize the wave equation
- Designed functions to handle the construction of complex musical compositions, demonstrating the practical applicability of the simulation model.
- Conducted validation by comparing the synthesized sounds against actual piano recordings, optimizing parameters like the number of points per string to balance simulation accuracy and computational efficiency

Decider

Sep 2022 - Dec 2022

- Designed a **full-stack** web application to help people make decisions, including features such as user registration, login, password change, adding/selecting tasks on the whiteboard, and selecting a random number within a specified range.
- Used **Sass, React, and Bootstrap** to optimize the main interface and customer input interface. Implemented real-time multi-user interaction and data retention on the website using socket.io.
- Used Selenium as the Automated functional testing to test password reset, login, add new items on the form, random number generator.

Languages and Technologies

Programming Languages: Python, Java, Javascript, C++, C, HTML, CSS, MATLAB, Stata

Technologies & Tools: LaTeX, Adobe Illustrator, Selenium, SQLAlchemy, Sass, Mocha, Chai, Figma, Git, PostgreSQL, MongoDB, Neo4j, Tailwind CSS

Frameworks & Libraries: PyTorch, NumPy, ScanPy, Pandas, Matplotlib, TensorFlow

Language Abilities: Chinese (Native), English(Proficient, GRE:332/3.5, TOEFL:105, R27 L26 S26 W26)

Outreach

STEM Outreach and Mentorship Coordinator, NYU Women in Science (WINS)

JUN 2022 - MAY 2024

- Guide NYC high school girls in STEM learning and competitions, fostering academic growth and confidence.
- Plan events connecting students with top scientists, enhancing learning and networking opportunities.
- Contribute to programs featuring women leaders and inspiring peers on women's contributions in STEM.
- Mentor WINS Scholars academically and in research, supporting their STEM career paths.