Thomas Walker

(+44) 7387172915 | thomaswalker0519@gmail.com | LinkedIn | Website

Summary

Final-year Ph.D. candidate at the University of Edinburgh and researcher at the Generative AI Laboratory (GAIL). My current research focus is improving the physical realism of interactive world models. In particular, persistent memory, viewpoint consistency, causality, and interactivity. I have a formal background in mathematics and first-author publications at top-tier CV/ML conferences.

Research Interests

World Modeling, Diffusion Models, Differentiable Rendering, 4D Reconstruction, Geometry Processing

Education

University of Edinburgh — Edinburgh, Scotland
PhD, 3D Vision and Computer Graphics (Expected 2026) — *2022 – Present*

- Supervised by Dr. Hakan Bilen and Dr. Amir Vaxman

University of Edinburgh — Edinburgh, Scotland
MSc, Artificial Intelligence (Distinction) — *2021 – 2022*

- Thesis: "Disentangling Neural Implicit Representations For 3D Assets"
- Specialization: Computer Graphics and Computer Vision
- **The IBM Prize:** awarded for the best project ($\approx 1/100$) on the machine learning course

University of Durham — Durham, England
MSci, Mathematics and Physics (First-Class Honours) — *2016 – 2020*

- Thesis: "Geometrical Aspects Of Gauge Theories and Gravity" (87%)
- Specialization: Geometry and Quantum Field Theory
- **The Peter Rowe Memorial Prize:** $\approx 1/200$ in Natural Sciences with Mathematics
- **Outstanding Performance Award:** for averaging 85% in final year

Selected Publications

PERSIST: World Models with Persistent 3D State — **ICML 2026**
Thomas Walker*, *Samuel Garcin**, *Steven McDonagh, Tim Pearce, Hakan Bilen, Kaixin Wang, Tianyu He, Jiang Bian*

- A new paradigm of world model with a rendered, dynamic 3D representation. We demonstrate improved long-horizon stability, 3D consistency, spatial memory, editing capabilities and out-of-view physics modeling. [Website]

Semantic Foam: Unifying Spatial and Semantic Scene Decomposition — ***Highlight*** **CVPR 2026**
Amr Sharafeldin, Shrisudhan Govindarajan, ***Thomas Walker****, Aryan Mikaeili, Daniel Rebain, Kwang Moo Yi, Andrea Tagliasacchi*

- **Highlight, top 10%.** Extends "Radiant Foam" for scene understanding. We achieve state-of-the-art 3D segmentation from images, improving over 3D Gaussian splatting methods while supporting real-time ray-tracing capabilities. [Website]

DynaSlice: Multi-Class 4D Reconstruction from Cross Sections
Thomas Walker*, João Paulo Silva do Monte Lima, Diana Aldana Moreno, Oisín Nolan, Salvatore Esposito, Tiago Novello*

- **Under Review.** We reconstruct multi-organ sequences from real CT-scans, achieving significantly improved performance in a fraction of the training time. [Paper]

OmniLiDAR: Multi-Sensor Controllable 4D LiDAR Generation
Khanh Nguyen, ***Thomas Walker****, Hichem Telli, Carlos De la Torre-Ortiz, Stefano Fiorini, Salah Eddine Bekhouche, Mohammed-En-Nadhir Zighem, Wassim Swaileh, Tinghuai Wang*

- **Under Review.** Work completed during Huawei internship. OmniLiDAR enables calibration-consistent multi-sensor LiDAR rollouts over long horizons for controllable, scenario-based autonomous-driving simulation.

CrossSDF: Reconstruction Of Thin Structures From Cross Sections **CVPR 2025**
Thomas Walker, Salvatore Esposito*, Daniel Rebain, Amir Vaxman, Arno Onken, Changjian Li, Oisin Mac Aodha*

- State-of-the-art 3D reconstruction from cross-sections, a modality commonly found in CT and MRI data. Significantly improved performance (4×) for challenging thin vessel structures. [Paper]

Spatially Adaptive Hash Encodings For Neural Surface Reconstruction **WACV 2025**
***Thomas Walker**, Octave Mariotti, Amir Vaxman, Hakan Bilen*

- Image-based surface reconstruction addressing limitations of hash-grid methods, allowing recovery of both fine and coarse details while mitigating noise. [Paper]

Explicit Neural Surfaces **(Workshop) NeurIPS 2023**
***Thomas Walker**, Octave Mariotti, Amir Vaxman, Hakan Bilen*

- Novel image-based surface reconstruction approach which learns an explicit neural surface, enabling 1000× improved sampling and rendering efficiency. [Paper]

Experience

Multi-modal Generative AI Intern Helsinki, Finland
Huawei *Oct. 2025 – Mar. 2026*

- Built generative LiDAR world model "OmniLiDAR", enabling stable long-horizon generations with state-of-the-art techniques for auto-regressive diffusion.

Research Internship Edinburgh, Scotland
GAIL, Generative AI Laboratory *Jan. 2025 – Jun. 2025*

- Developed robust and undetectable frequency-space watermarking algorithms for diffusion models.
- Hosted workshop "Beyond Deepfakes: advancing watermarking in Generative AI".

Teaching Assistant: Computational Geometry and Simulation 2024 – 2025

Teaching Assistant: Machine Learning 2022 – 2023

Additional Projects

Spherical Feature Pyramid Networks
***Thomas Walker**, Varun Anand, Pavlos Andreadis*

- Adapted feature pyramid networks to the spherical domain, achieving state-of-the-art semantic segmentation performance on spherical images. [Paper]

Calabi-Yau Metrics with U(1)-Invariant Neural Networks **LOGML Summer School 2024**
*Yidi Qi, **Thomas Walker**, Albert Wood, Marek Cerny, Philipp Dahlinger, Heath Pearson*

- Developed a new U(1)-invariant physics-informed network (PINN) to learn Riemannian metrics of Calabi-Yau manifolds, geometries underpinning string theory.

Technical Skills

Languages: Python, C++, MATLAB, LaTeX
Frameworks: PyTorch, TensorFlow, OpenCV, NumPy, SciPy, Scikit-learn, Scikit-image
Tools: Git, Blender, PolyScope