Dr. Héctor Carrión

World Model Team Lead · Multimodal & Spatio-Temporal AI
Function · Ph.D. in Computer Science, UC Santa Cruz / UC Berkeley

hi@hectorcarrion.com
hectorcarrion.com
github.com/hectorcarrion
linkedin.com/in/hectorcarrion

AI researcher (6+ YoE) building **multimodal world models** across longitudinal health, 3D vision, and autonomy. Leads the World Model team at **Function** (JEPA and autoregressive models over longitudinal labs and imaging); previously Founding Lead Research Scientist at **Tera AI** (1B-parameter self-driving worldmodel) and Ph.D. researcher at **Google Research**, **Google X** (granted patents), and **UC Santa Cruz / UC Berkeley**. First-author papers at **MICCAI** (x4) and **PLOS Comp. Biol.**

Focus: World Models • JEPA • 3D Vision • Diffusion • Medical AI • Physical AI

Research & Professional Experience

Sep 2026 – now — Function

World Model Team Lead, Function

- Directing the World Model team and research roadmap on JEPA and autoregressive multimodal architectures over longitudinal lab tests, imaging, and member health history.
- Demonstrated novel SSL methods outperform and outscale FST (>2x) on clinical trajectory predictions.

2024 – 2026 — Tera AI

Founding Lead Research Scientist, Tera AI

- Set and executed the research agenda across spatial reasoning, generative models, and embodied autonomy, managing multi-node compute budgets and iteration schedules.
- Led TARDIS and STRIDE (with P. Perona and J. Malik): a spatio-temporal transformer unifying street-view synthesis, action prediction, and sequential georeferencing (60% within 10 m), trained >7B-param model.

2020 – 2026 — UC Santa Cruz / UC Berkeley

Graduate Student Researcher, UC Santa Cruz & UC Berkeley

- Built label-efficient medical vision models (<100 labeled samples) using controllable diffusion (cgDDI: 90.9% DDI accuracy, FEDD: +18 IoU points) and monocular metric 3D reconstruction (DermDepth: 16x better depth scale ratio vs SOTA tuning 0.6% of parameters).
- Developed temporal SSL stage discovery (HealNet) and automated biomarker imaging pipelines for multi-institutional NIH bioelectronic studies. Funded by DARPA.

Summer 2023 — Google

Student Researcher (Ph.D.), Google Research

- Trained unified multimodal video-image foundation encoders, improving downstream transfer by >25% over deployed single-modality baselines.

2021 – 2022 — Google

AI Resident (Ph.D.), X, The Moonshot Factory (Google X)

- Built simulation, edge-perception, and aerial/ground computer vision systems for autonomous platforms, exceeding deployment benchmarks and earning Google patents.

Education

2020 – 2026 — **Ph.D. in Computer Science**, University of California, Santa Cruz
Advisor: Dr. Narges Norouzi (UC Berkeley). GPA: 3.92 / 4.00.

2017 – 2021 — **B.Sc. in Computer Science**, University of Puerto Rico, Río Piedras
Research PI: Dr. Rémi Mégret. GPA: 3.75 / 4.00.

2018 – 2019 — **Software Engineering & Design**, Apple Developer Academy, Univ. of Naples Federico II
Full-stack iOS engineering, interface design, and on-device CoreML applications.

Selected Publications

MICCAI 2026

DermDepth: Toward Monocular Metric Scale 3D Reconstruction Models for Dermatology
H. Carrión, N. Norouzi
16.1×→1.11× SKINL2 depth scale with 0.6% tuned params paper | arXiv | code | data | model | demo

MICCAI 2026

cgDDI: Controllable Generation of Diverse Dermatological Imagery for Fair and Efficient Malignancy Classification
H. Carrión, N. Norouzi
266K+ synthetic images; 90.9% DDI accuracy & 86.6% EOM paper | arXiv | code | data

Comm. Biol. 2025
A Modular Fluorescent Camera Unit for Wound Imaging
M. Tebyani, G. Keller, W. S. Hee, P. Baniya, ..., **H. Carrión,** ..., R. R. Isseroff, M. Rolandi, M. Teodorescu
Nature Communications Biology paper

arXiv 2025
TARDIS STRIDE: A Spatio-Temporal Road Image Dataset and World Model for Autonomy
H. Carrión, Y. Bai, V. A. Hernández Castro, K. Panaganti, A. Zenith, M. Trang, T. Zhang, P. Perona, J. Malik
1B-param world model; 60% within 10 m georeferencing arXiv | project | code | data

Adv. Mater. Tech.
Remote-Controlled Wireless Bioelectronics for Fluoxetine Therapy to Promote Wound Healing in a Porcine Model
H. Li, N. Asefifeyzabadi, ..., **H. Carrión,** N. Norouzi, M. Gomez, M. Zhao, R. R. Isseroff, M. Teodorescu, M. Rolandi
Advanced Materials Technologies, 2025 paper

Adv. Mater. Tech.
Programmable Delivery of Fluoxetine via Wearable Bioelectronics for Wound Healing In Vivo
H. Li, H.-Y. Yang, N. Asefifeyzabadi, ..., **H. Carrión,** ..., N. Norouzi, M. Gomez, M. Zhao, M. Teodorescu, R. R. Isseroff, M. Rolandi
Advanced Materials Technologies, 2024 paper

MICCAI 2023
FEDD: Fair, Efficient, and Diverse Diffusion-Based Lesion Segmentation and Malignancy Classification
H. Carrión, N. Norouzi
81% DDI malignancy accuracy from 60 labeled images paper | arXiv | code

MICCAI-W 2022
HealNet: Self-Supervised Acute Wound Heal-Stage Classification
H. Carrión, M. Jafari, H.-Y. Yang, R. R. Isseroff, M. Rolandi, M. Gomez, N. Norouzi
Machine Learning in Medical Imaging (MLMI); 90.6% stage accuracy paper | arXiv

PLOS CB 2022
Automatic Wound Detection and Size Estimation Using Deep Learning Algorithms
H. Carrión, M. Jafari, M. D. Bagood, H.-Y. Yang, R. R. Isseroff, M. Gomez
PLOS Computational Biology paper | code

VISAPP 2022
Honeybee Re-Identification in Video: New Datasets and Impact of Self-Supervision
J. Chan, **H. Carrión,** R. Mégret, J. Rivera, T. Giray
109K-image self-supervised track dataset paper

Open-Source Code, Models & Datasets

DermDepth	Metric 3D weights, D-Synth dataset, and live demo → github.com/hectorcarrion/dermdepth
cgDDI	Diffusion code and 266K+ synthetic dermatology dataset → github.com/hectorcarrion/ControllableGenDDI
TARDIS	1B-param world model (JAX/EasyLM) and STRIDE dataset → github.com/tera-ai/tardis
FEDD	Diffusion feature probes and DDI lesion segmentation masks → github.com/hectorcarrion/fedd

Honors, Patents & Advisory

Patents	**Google Patent Award**: Granted perception & autonomous systems patents at X (Google X)
Fellowships	**Google–CAHSI Doctoral Fellowship** | **NIH BD2K Biomedical Data Science Fellowship**
Awards	**MICCAI STAR Award** | **SIDIM Outstanding Research Award**
Advisory	**Technical Advisor**, Visia AI (2021–present): computer vision for recyclability & hazard monitoring
Mentorship	**Machine Learning Tutor**, UC Santa Cruz (2020–2022): deep learning & computer vision

Skills & Languages

Programming	Python, C++, Swift, JAX, Bash, LaTeX
Deep learning	PyTorch, JAX/EasyLM, TensorFlow, Keras, JEPA & autoregressive world models, diffusion, multimodal LLMs
Domains	Longitudinal health modeling, metric 3D vision, medical imaging, spatial autonomy, fairness
Leadership	Research roadmap & team management, compute budgeting, cross-functional execution, product design
Languages	English (native), Spanish (native)

Glyph key: World Model | 3D Vision | Diffusion | Med AI | Self-Supervised / JEPA | Autonomy | Bioelectronics