Haithem Turki
turki.haithem@gmail.com ● https://haithemturki.com/ ● https://www.linkedin.com/in/haithemturki/

Education

Carnegie Mellon University, Pittsburgh, PA *08/24*
PhD in Computer Science.
Thesis: Towards City-Scale Neural Rendering
Committee: Deva Ramanan (chair), Jessica K. Hodgins, Martial Hebert, Jonathan T. Barron

Stanford University, Stanford, CA *06/18*
MS in Computer Science.

Stanford University, Stanford, CA *06/12*
BS in Computer Science, minor in Management Science and Engineering

Professional Experience

Senior Research Scientist, NVIDIA, Seattle, WA *09/24 - Present*
- Exploring neural rendering and data-driven simulation topics in the Spatial Intelligence Lab with Zan Gojcic

Research Scientist Intern, Meta, Pittsburgh, PA *05/23-12/23*
- Worked on real-time neural rendering in VR with Christian Richardt and Michael Zollhöfer

Research Intern, Argo AI, Pittsburgh, PA *05/22-08/22*
- Explored city-scale 3D reconstruction methods with Deva Ramanan and Francesco Ferroni

Engineering Manager and **Technical Lead,** Palantir Technologies, New York, NY *08/12-08/19*
- Led R&D efforts to improve Palantir's Foundry data platform:
  - Created interactive search and analysis infrastructure that evolved into one of the primary user interfaces of Palantir's Foundry data lake platform deployed across the company's customer fleet
  - Invented full-text search system that powers Palantir's internal log analysis infrastructure and a significant part of the Foundry data lake platform
- Led inception and rollout of one of Palantir's flagship data analytics products
- Managed the technical implementation of a variety of high-profile customer engagements in the commercial sector (leading teams of 10-100+ individuals)

Supervised and mentored a team of 7+ direct reports

Engineering Intern, Addepar, Mountain View, CA *04/11-09/11*
- Worked on developing a software platform to help private financial advisors and family offices perform quantitative analysis and interface with their clients

Research Assistant, Stanford University, Stanford, CA *06/10-12/10*
- Conducted semantic segmentation research under Professor Daphne Koller and Dr. M. Pawan Kumar

Haithem Turki
turki.haithem@gmail.com ● https://haithemturki.com/ ● https://www.linkedin.com/in/haithemturki/

Teaching Experience

Graduate Student Instructor, Dept. of Computer Science, Carnegie Mellon University *09/20-06-21*
- Teaching assistant for 15-640 (Distributed Systems) and inaugural 17-700 class (Data Science and Machine Learning at Scale)

Project Mentor, Dept. of Computer Science, Carnegie Mellon University *09/20-12/21*
- Designed a term project for 15-821 (Mobile and Pervasive Computing)
- Met and provided guidance to project members on a regular basis

Publications (Reverse Chronological Order)

[1] **Haithem Turki***, Qi Wu*, Xin Kang, Janick Martinez Esturo, Shengyu Huang, Ruilong Li, Zan Gojcic, Riccardo de Lutio, "SimULi: Real-Time LiDAR and Camera Simulation with Unscented Transforms." Preprint.

[2] Sherwin Bahmani, Tianchang Shen, Jiawei Ren, Jiahui Huang, Yifeng Jiang, **Haithem Turki**, Andrea Tagliasacchi, David B. Lindell, Zan Gojcic, Sanja Fidler, Huan Ling, Jun Gao*, Xuanchi Ren*, "Lyra: Generative 3D Scene Reconstruction via Video Diffusion Model Self-Distillation." Preprint.

[3] Zhangjie Wu*, Yuxuan Zhang*, **Haithem Turki**, Xuanchi Ren, Jun Gao, Mike Zheng Shou, Sanja Fidler, Zan Gojcic*, Huan Ling*, "DIFIX3D+: Improving 3D Reconstructions with Single-Step Diffusion Models." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2025 (best paper award candidate).

[4] **Haithem Turki**, Vasu Agrawal, Samuel Rota Bulò, Lorenzo Porzi, Peter Kontschieder, Deva Ramanan, Michael Zollhöfer, Christian Richardt, "HybridNeRF: Efficient Neural Rendering via Adaptive Volumetric Surfaces." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024 (highlight).

[5] Li Ma, Vasu Agrawal, **Haithem Turki**, Changil Kim, Chen Gao, Pedro Sander, Michael Zollhöfer, Christian Richardt, "SpecNeRF: Gaussian Directional Encoding for Specular Reflections." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024 (highlight).

[6] **Haithem Turki**, Michael Zollhoefer, Christian Richardt, Deva Ramanan, "PyNeRF: Pyramidal Neural Radiance Fields." In Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS), 2023.

[7] Shilpa George, **Haithem Turki**, Ziqiang Feng, Thomas Eiszler, Deva Ramanan, Padmanabhan Pillai, Mahadev Satyanarayanan, "Low-Bandwidth Self-Improving Transmission of Rare Training Data." In Proceedings of the 29th Annual International Conference on Mobile Computing And Networking (MobiCom), 2023.

[8] Shilpa George, **Haithem Turki**, Ziqiang Feng, Thomas Eiszler, Deva Ramanan, Padmanabhan Pillai, Mahadev Satyanarayanan, "Edge-based Privacy-Sensitive Live Learning for Discovery of Training Data." In Proceedings of the 1st International Workshop on Networked AI Systems (NetAISys '23).

[9] **Haithem Turki**, Jason Y. Zhang, Francesco Ferroni, Deva Ramanan. "SUDS: Scalable Urban Dynamic Scenes." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023.

[10] Mahadev Satyanarayanan, Ziqiang Feng, Shilpa George, Jan Harkes, Roger Iyengar, **Haithem Turki**, Padmanabhan Pillai, "Accelerating Silent Witness Storage," in IEEE Micro, Nov.-Dec. 2022.

[11] **Haithem Turki**, Deva Ramanan, Mahadev Satyanarayanan, "Mega-NeRF: Scalable Construction of Large-Scale NeRFs for Virtual Fly-Throughs." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022.

[12]Shilpa George, Thomas Eiszler, Roger Iyengar, **Haithem Turki**, Ziqiang Feng, Junjue Wang, Padmanabhan Pillai, Mahadev Satyanarayanan, "OpenRTiST: End-to-End Benchmarking for Edge Computing," in IEEE Pervasive Computing, vol. 19, no. 4, pp. 10-18, 1 Oct.-Dec. 2020
[13] Mahadev Satyanarayanan, Thomas Eiszler, Jan Harkes, **Haithem Turki** and Ziqiang Feng, "Edge Computing for Legacy Applications," in IEEE Pervasive Computing, 1 Oct.-Dec. 2020
[14] M. Pawan Kumar, **Haithem Turki**, Dan Preston and Daphne Koller, "Parameter Estimation and Energy Minimization for Region-Based Semantic Segmentation," in IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 37, no. 7, pp. 1373-1386, 1 July 2015.
[15] M. Pawan Kumar, **Haithem Turki**, Dan Preston and Daphne Koller, "Learning specific-class segmentation from diverse data," 2011 International Conference on Computer Vision.

Patents

[1] **Haithem Turki,** Robert Fink, Amr Al Mallah. "Systems and Methods for Indexing and Searching." US Patent 2019/0347343 A1, filed June 8, 2018, and issued November 14, 2019.
[2] **Haithem Turki,** Sander Kromwijk, Stephen Cohen, Yixun Xu, Feridun Arda Kara. "Systems and Methods for Constraint Driven Database Searching." US Patent 2018/293239 A1, filed April 11, 2017, and issued October 11, 2018.

Technical Reports

[1] Ziqiang Feng, Shilpa George, **Haithem Turki,** Roger Iyengar, Padmanabhan Pillai, Jan Harkes, Mahadev Satyanarayanan. "Improving Edge Elasticity via Decode Offload." CMU-CS-21-139, September 2021.

Invited Talks

Towards City-Scale Neural Rendering *01/24-02/24*
- Stanford Vision Lab
- Waymo Research (virtual)
- Google (virtual)
- NVIDIA Toronto AI Lab (virtual)

Mega-NeRF: Scalable Construction of Large-Scale NeRFs for Virtual Fly-Throughs *04/22*
- Argo AI

Additional Information

- **Reviewer:** CVPR, ACCV, ECCV, ICCV, NeurIPS, ICLR, SIGGRAPH, SIGGRAPH Asia, AAAI, Transactions on Image Processing
- **Volunteer Activities:** Alchemist Accelerator (mentor and investor)
- **Languages:** Fluent in French and Arabic. Currently learning Japanese and German.
- **Other Qualifications:** Top Secret security clearance (expired). FINRA Series 65. Graduate of the Stanford StartX startup accelerator.