Angela Dai – 3D in a Large-Data World

March 18th, 2025 – 11:00 am

Recent advances in machine learning have shown remarkable progress in the 2D and video domain, fueled by very large-scale data and compute. 3D, however, which is critical for applications spanning content creation, mixed reality, and robotics, remains constrained by significantly more limited data representing higher-dimensional information. In this talk, we first address real-world 3D data availability by introducing a large-scale, high-fidelity 3D dataset of real scans. We then explore methods for generating and editing large-scale 3D scenes, to enable new possibilities in content creation. Finally, as real-world 3D scenes are rarely fully static, we address the unique challenges of data scarcity in the 4D domain, showcasing how structured priors and distillation from large-scale models can enable more robust 4D understanding. These 3D learning strategies promise to usher in a new paradigm of generalized 3D perception, beyond the limits of existing 3D datasets, to enable in-the-wild 3D analysis of environments.

Bio: Angela Dai is an Associate Professor at the Technical University of Munich where she leads the 3D AI Lab. Angela's research focuses on understanding how real-world 3D scenes around us can be modeled and semantically understood. Previously, she received her PhD in computer science from Stanford in 2018, advised by Pat Hanrahan, and her BSE in computer science from Princeton in 2013. Her research has been recognized through an ECVA Young Researcher Award, ERC Starting Grant, Eurographics Young Researcher Award, German Pattern Recognition Award, Google Research Scholar Award, and an ACM SIGGRAPH Outstanding Doctoral Dissertation Honorable Mention.