



Project Type

- Master Thesis
- Bachelor Thesis
- Research Project

Supervisors

-  Enes Ulaş Dincer
-  enes.dincer@kit.edu

Difficulty

Algorithmic



Math



Application



Gaze-Adaptive Point Cloud Rendering in VR

Description

Point clouds provide a direct 3D representation of a remote robot workspace, but dense point clouds are expensive to render and stream in real time. This thesis investigates gaze-adaptive point-cloud rendering for VR teleoperation: regions around the user's gaze are rendered densely, while peripheral regions are downsampled. Beyond rendering performance and user experience, the project will study whether different visualization conditions affect the quality of collected teleoperation demonstrations and the performance of imitation-learning policies trained from them.




Figure 1: Real-robot teleoperation setup with VR headset, leader robot, follower robot, RGB-D camera, and point-cloud feedback.

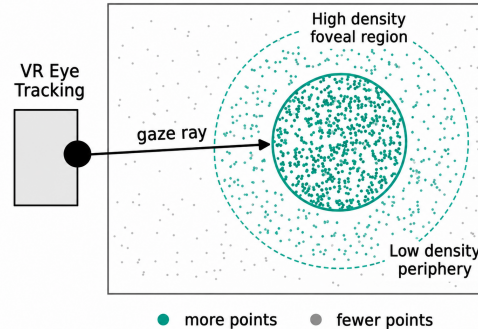


Figure 2: Concept of gaze-adaptive point-cloud rendering. The foveal region is dense; peripheral regions are downsampled.

Goals and Research Questions

The goal is to build and evaluate a real-time gaze-adaptive point-cloud renderer for VR teleoperation and study its effect on both human demonstrations and learned robot policies. Main questions:

- Can gaze-adaptive rendering reduce rendered points while preserving task performance and perceived quality?
- Do different rendering conditions affect the quality of human teleoperation demonstrations?
- How do policies trained from demonstrations collected under different visualization conditions compare in downstream task performance?

Tasks

- Eye-tracking integration: Use VR gaze data and map the gaze ray into the rendered 3D point-cloud scene.
- Baselines: Implement full-density, uniform low-density, fixed-center foveation, and gaze-adaptive foveation.
- Demonstration collection: Collect teleoperation demonstrations under different rendering conditions.
- Policy training: Train imitation-learning policies from the collected demonstrations and compare their performance.
- User study: Compare visualization conditions in terms of task performance, workload, comfort, and preference.

System pipeline



Study conditions




Figure 3: Pipeline and study conditions: full-density, uniform low-density, fixed-center foveation, and gaze-adaptive foveation.

Evaluation

Technical metrics: FPS, latency, GPU load, update rate, and rendered points. User-study metrics: completion time, success rate, accuracy, perceived quality, spatial understanding, NASA-TLX workload [3], SSQ cybersickness [4], and preference ranking. Learning metrics: demonstration quality, policy success rate, robustness, and generalization.

Expected Outcome

A working gaze-adaptive VR point-cloud renderer, a dataset of demonstrations collected under different visualization conditions, and an analysis of how rendering affects both human teleoperation and imitation-learning policy performance.

[1] Patney et al., ACM TOG, 2016. [2] Guenter et al., ACM TOG, 2012. [3] Hart, HFES, 2006. [4] Kennedy et al., IJAR, 1993.