A light field capturing system using multiple omnidirectional cameras

Li Qiaoge

School of Integrative and Global Majors, Ph.D. Program in Empowerment Informatics, University of Tsukuba

1. Project Purpose

We are developing a method that uses implicit neural representations to reconstruct omnidirectional scenes solely from RGB data captured from various spherical panoramas. Commercially available omnidirectional cameras, which can capture comprehensive views of an environment in real-time at a low cost, facilitate this. Our approach, based on the foundational principles of NeRF, seeks to learn and utilize ray information from these panoramas to create a continuous, high-quality omnidirectional radiance field. This technology has shown promising results on synthetic datasets featuring random camera positions, and further testing with real-world data suggests potential for consistent, real-time reconstruction of full environmental perspectives. We aim to extend this to dynamic scenes and integrate it seamlessly with major platforms like Unity and Unreal, without the need for calibration. This endeavor could significantly broaden the practical applications of our research in fields such as virtual tours, product demonstrations, and remote education, albeit at a high computational cost.

2. Results

Our preliminary results have validated the feasibility of using 3D Gaussian splatting for large-scale urban representations. By utilizing supercomputing resources to pre-train large models, we have successfully quantized these models to perform efficiently on consumer-level hardware, achieving impressive outcomes. Additionally, we have trained models on several other scenarios aimed at generating city layout distributions, as well as providing interfaces suitable for autonomous driving applications. These findings highlight the robust versatility and potential of our techniques in handling extensive urban environments effectively.

3. Roles of the MCRP and its significance

The MCRP plays a crucial role in our ambitious project, particularly due to the intensive IO speeds and GPU memory demands required for training models like

VAE and Gaussian splatting on ultra-large-scale scenes. The capabilities of MCRP's H100 GPUs, coupled with 2TB of non-volatile memory, are essential for handling the vast data throughput and complex computations needed. Without these advanced components, our project objectives would be unachievable. Thus, we deeply value this opportunity to utilize such cutting-edge technology, which is pivotal for the success of our plans.

4. Future plan

Looking ahead, our plans are ambitious and focused on advancing urban scene rendering. We aim to replace traditional Level of Detail (LOD) rendering techniques with our approach, which promises higher fidelity and detailed scene reconstructions that rival or even surpass the current state-of-the-art NeRF methods in terms of realism. Our ongoing work has been submitted for review at ACM Multimedia (ACMMM), and we plan to present our findings at this year's CVPR. By 2024, we are committed to achieving our targets, further refining our techniques, and showcasing the effectiveness of our models in producing highly realistic urban environments.

5. Publications and conference presentations

(1) Journal papers

(2) Presentations

None

(One ACMMM under review)

(3) Others

Supercomputer	Use	Allocated resources*		
		Initial resources	Transferred resources**	Additional resources
Cygnus	Yes/No			
Pegasus	Yes/No	10000		
Wisteria/BDEC-01	Yes/No			

*in units of node-hour product

** If the budget transfer was performed, fill in here, such as "+2000" and "−1000".