Abstract

The increasing complexity of data in modern computational fields necessitates the development of specialized tools for data processing, analysis, and representation. This thesis addresses this need by introducing novel methods at the intersection of computer graphics, machine learning, and digital heritage. The work spans theoretical advances in aggregation functions, novel rendering techniques for neural representations, and practical applications in cultural heritage documentation.

First, we introduce k-restricted overlap and grouping functions, a generalized family of aggregation operators that extends traditional n-dimensional overlap and grouping functions by incorporating restrictions based on boundary element counts. These operators enable flexible control over conjunctive, disjunctive, and averaging behaviors through a parameter k, where the function output depends on at least k boundary elements rather than a single element. We establish existence and representation theorems for these functions and demonstrate their practical utility as pooling layer replacements in Convolutional Neural Networks (CNNs). Experimental validation across multiple architectures (LeNet-5, VGG16, ResNet variants) and datasets shows improved performance (e.g., up to .05 accuracy gain on already extremely optimized architectures), particularly on sparse training samples, with faster processing on both CPUs and GPUs compared to traditional pooling methods.

Second, we propose Neural Directional Distance Fields (NDDFs) as an efficient object representation for path-traced rendering. Unlike traditional signed distance functions, DDFs incorporate directional components, enabling constant-time ray-surface intersection queries while maintaining surface detail preservation. We develop a modified path-tracing algorithm specifically designed for NDDF rendering and introduce systematic sampling techniques for training neural networks to represent these fields. Our approach demonstrates significant rendering speedup compared to traditional bounding volume hierarchy methods while preserving geometric accuracy, as validated through chamfer distance metrics and visual quality assessments.

Third, we address practical challenges in Heritage Building Information Modeling (HBIM). We propose a comprehensive framework for digitizing, storing, and transmitting large-scale cultural heritage data. Using the 11th-century Rajarani Temple in Bhubaneswar, India, as a case study, we develop techniques for segmentation, decimation, mesh matching, and hierarchical storage of heritage models containing over 780 million points. Our methodology enables efficient transmission over limited bandwidth networks and exhibition on consumer-grade hardware while preserving culturally significant architectural details. The framework employs inverse density sampling, density-aware chamfer distance metrics, and kd-tree-based hierarchical storage to balance computational efficiency with preservation fidelity.

The three contributions demonstrate a structured approach to advancing computational methods, progressing from theoretical foundations to algorithmic innovations and real-world applications. The k-restricted aggregation functions provide mathematical tools with broad applicability beyond computer vision, the neural directional distance fields offer new possibilities for efficient 3D rendering, and the heritage digitization framework addresses urgent needs in cultural preservation. Together, these works contribute to the intersection of mathematics, computer graphics, and digital humanities, providing both theoretical insights and practical solutions for complex computational challenges.

Keywords: Aggregation functions, neural rendering, directional distance fields, heritage preservation, point cloud processing, convolutional neural networks, path tracing, HBIM