Visual localization is the problem of estimating the 6 degrees of freedom camera pose from which a query image was taken relative to a known reference scene representation. It is the key for applications such as Augmented, Mixed, and Virtual Reality, as well as autonomous robotics such as drones or self-driving cars.

This thesis focuses on a visual localization pipeline, especially on its pose verification and reranking step. The pipeline uses 3D point clouds and 2D-3D correspondences between the query image and 3D scene points for candidate camera poses estimations. The thesis explores point cloud rendering approaches as they are utilized in the pipeline and the verification step—the render of the discretized scene from a given candidate position is compared to the actual query image to asses if the given couple depicts the same place.

One of the main challenges of such rendering is occlusion handling. Due to the sparsity of points employed for otherwise continuous real world representation, information about what lies in the front and what is hidden can be easily lost when projected to the 2D image. Rendering approaches explored in this thesis focus on the challenge directly or as a component of a novel view synthesis DNN-based renderer. Rendering influence on localization performance is investigated.