PRIORITIZED IRLS SURFACE NORMAL ESTIMATION AND POINT CLASSIFICATION

By

Cody James Phillips

An Abstract of a Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic Institute

in Partial Fulfillment of the

Requirements for the Degree of

MASTER OF SCIENCE

Major Subject: COMPUTER SCIENCE

The original of the complete thesis is on file
in the Rensselaer Polytechnic Institute Library

Approved:

Charles V. Stewart, Thesis Adviser

Rensselaer Polytechnic Institute
Troy, New York

May 2009

This thesis presents a statistical method called PINE for robust surface normal estimation and point classification of range images. This is an important preprocessing step required in many three dimensional modeling algorithms. This thesis is concerned with range data of outdoor scenes acquired by a ground-based multi-view laser range scanner. New statistically robust methods are required to process such data, as outdoor environments pose significant challenges that limit the success of previous techniques. Poor viewpoints for outdoor scans due to limited selection may cause undersampling along oblique surfaces and areas of missing data due to occluding geometry. Outdoor objects that are non-reflective or highly specular may also result in missing or incorrect data points. Factors such as these cause scans of uncontrolled outdoor scenes to be of lesser quality than those of indoor scenes or objects.

The algorithm presented segments the range data into surface and non-surface points of several types and generates surface normal estimates for each surface point. The hallmark of our algorithm is the use of a local coherency measure to prioritize the order of surface normal estimation to better exploit the cohesive relationship among neighboring samples. Surface normals are assigned to each point by a robust fitting of local tangent planes using iteratively reweighted least-squares (IRLS). The algorithm begins by assigning normals to points in regions that are locally flat and should therefore have relatively stable normals, and then proceeds outwards towards areas of high curvature and discontinuity. The data is smoothed by re-estimating point locations, and ultimately each point's normal is robustly averaged with its neighbors' normals.

The results are directly compared with a state of the art multi-scale formulation of tensor voting (MSTV) and the major differences between the two algorithms are identified. MSTV was chosen for comparison as it shows great promise for surface reconstruction with its theoretical foundation based in perceptual organization. PINE shows strong performance across many physical scales with varying sample densities, producing consistent normal estimates on flat surfaces with sample densities down to 2500 $points/m^2$ and surface areas to $0.05m^2$. PINE tends to produce smoother results than MSTV, with less surface noise accompanied by a loss of sur-

face detail. Undersampling along oblique surfaces is handled quite well by PINE, whereas MSTV is unable to cope with extremely anisotropic sampling.