2. Publishable summary

Increasing interest in the modeling and reconstruction of digital architecture and urban scenes may largely be attributed to recent advances in laser scanning technology and the proliferation of GIS services such as those offered by Google Earth and Google StreetView. We are thus witnessing a strong trend towards the application of 3D modeling, reconstruction and visualization of urban scenes based on satellite and aerial photography combined with geometry capture enabled by street-level and aerial laser scanners. We expect this trend to continue as more cities and a large variety of architectural structures are scanned.

In the proposed research I have developed a set of novel algorithms and tools for analysis, modeling, reconstruction and manipulation of scanned 3D urban scenes and structures. This research makes a pioneering effort towards obtaining full 3D digital cities using various acquisition modalities. It will contribute to significant progress in the domain of 3D digital cities, not only in their acquisition and modeling but also in their visualization and manipulation.

The main problem inherent in 3D geometry of large-scale urban environments is data quality (or lack thereof). For example, in the case of street-level acquisition, many parts of buildings are often occluded and therefore cannot be captured accurately. Furthermore, large distances between the scanner and scanned objects lead to a reduction in precision, and the presence of reflective objects, such as windows and other glossy structures, yields outliers. In contrast to the scanning of relatively small artifacts where the scanner can be optimally positioned to achieve adequate point density and coverage, such control is limited in large-scale urban scanning. As a result, the obtained point clouds typically suffer from significant missing data (due to occlusion), as well as uneven point density, noise and outliers.

In this report I will describe my development of techniques for efficient point-set processing, understanding and manipulation. Specifically, my research objectives are:

Objective 1. Interactive techniques

The aim of this objective is to develop an interactive solution for reconstruction of complex structures and fine details in scanned urban scenes.

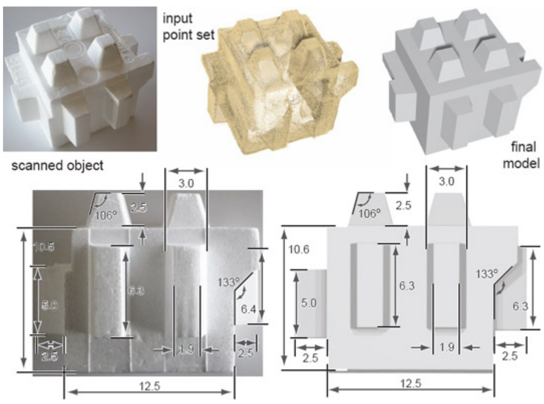

Figure 1: Starting from a noisy scan, our algorithm recovers the primitive faces along with their global mutual relations, when are then used to produce a final model (all lengths in mm).

In the realm of reconstruction of complex structures, we introduce a method that simultaneously recovers a set of locally fitted primitives along with their global mutual relations. We operate under the assumption that the data corresponds to a man-made engineering object consisting of basic primitives, possibly repeated and globally aligned under common relations. We introduce an algorithm to directly couple the local and global aspects of the problem. The local fit of the model is determined by how well the inferred model agrees to the observed data, while the global relations are iteratively learned and enforced through a constrained optimization (Figure 1). **This work has been published in: ACM Transactions on Graphics, (30) 52:1-12, 2011 (Proceedings of SIGGRAPH)**

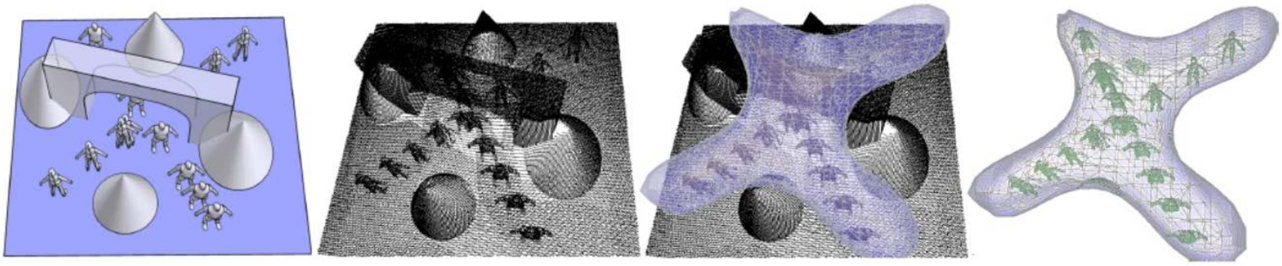

Figure 2: Interactive segmentation of dynamic point-cloud scenes

In the realm of interactive reconstruction, we introduce a novel interactive learning technique for point-cloud motion segmentation, we present a novel approach for motion segmentation in dynamic point-cloud scenes designed to cater to the unique properties of such data. Our key idea is to augment fg/bg classification with an active learning framework by refining the segmentation process in an adaptive manner. Our method initially classifies the scene points as either fg or bg in an un-supervised manner by training discriminative RBF-SVM classifiers on automatically labeled, high-certainty fg/bg points. We present a unique interactive paradigm for enhancing this learning process, by using a manual editing tool. The user explicitly edits the RBF-SVM decision borders in unreliable regions in order to refine and correct the classification.
(Figure 2). **This work has been published in: Computer Graphics Forum (CGF), 32:51–60, 2013 (Proceedings of Pacific Graphics)**

Objective 2. Photo-scan registration and enhancement of urban data

The aim of this objective is to explore the mutual relations between two modalities: 2D photographs and 3D scans. The individual characteristics of these modalities are different: point cloud scans are coherent and inherently 3D, but are often sparse, noisy, and incomplete; photographs, on the other hand, are of high resolution, easy to acquire and dense, but view-dependent and inherently 2D, lacking in depth information.

Figure 3: Multi-modal data fusion produces textured polygonal model, with missing parts in LiDAR data being completed using extracted

In the realm of Objective 2: photo-scan registration, we developed a method for fusing two acquisition modes, 2D photographs and 3D LiDAR scans, for depth-layer decomposition of urban facades. Our key observation is that with an initial registration of the 2D and 3D datasets we can decompose the input photographs into rectified depth layers. The layer decomposition enables accurate repetition detection in each planar layer, and using the detected repetitions to propagate geometry, remove outliers and enhance the 3D scan. Finally, the algorithm produces an enhanced, layered, textured model (Figure 3).

This work has been published in: Proceedings of the International Conference on Computer Vision (ICCV), 882-889, 2011

Figure 4: Given a partial 3D scan and a 2D photograph (left), we lift 2D shape structures into 3D yielding a faithful textured mesh reconstruction (right).

We also present an algorithm for shape reconstruction from incomplete 3D scans by fusing together two acquisition modes: 2D photographs and 3D scans. We compute geometrical and topological shape properties in 2D photographs and use them to reconstruct a shape from an incomplete 3D scan in a principled manner. Our key observation is that shape properties such as boundaries, smooth patches and local connectivity, can be inferred with high confidence from 2D photographs. Thus, we register the 3D scan with the 2D photograph and use scanned points as 3D depth cues for lifting 2D shape structures into 3D. Our contribution is an algorithm which significantly regularizes and enhances the problem of 3D reconstruction from partial scans by lifting 2D shape structures into 3D (Figure 4).

This work has been published in: Computer Graphics Forum (CGF), 33: 249–258, 2014 (Proceedings of Pacific Graphics)

Objective 3. Visualization of 3D urban structures

This third objective lies in the realm of cartography, where 3D digital cities will be utilized for novel map visualization and navigation. Our goal is to explore ways of visualization of complex 3D environments and their integration with 2D maps for a multi-layered visualization. Thus, we generate effective visualizations that take advantage of the full 3D scene while generating symbolic representations, simplifications and abstractions of urban structures.

Figure 5: Our method models the conjoining gestalts and correctly groups horizontal and vertical window structures and formations from a complex building façade

In the realm of Objective 3: visualization in 3D urban scenes, we developed a method for structural summarization and abstraction of complex spatial arrangements found in architectural drawings. The method is based on the well-known Gestalt rules, which summarize how forms, patterns, and semantics are perceived by humans from bits and pieces of geometric information. We develop a computational framework which models Gestalt rules and more importantly, their complex interactions. We show an application of our method to line drawings of architectural models of various styles (Figure 5).

This work has been published in: ACM Transactions on Graphics, (30) 185:1-10, 2011 (Proceedings of SIGGRAPH)

Figure 6: Our method creates abstract stylized objects from a given input model

We also introduce the novel paradigm of non-realistic 3D stylization, where the expressiveness of a given 3D model is manifested in the 3D shape itself, rather than only in its rendering. We analyze the input model using abstraction, simplification, and symmetrization operators to determine important features that are later represented by new geometry. Doing so, we create a stylized and expressive representation of the input that can be rendered or might be printed using a 3D printer. (Figure 6).

This work has been published in: Symposium on Computational Aesthetics in Graphics, Visualization and Imaging