

Efficient Volumetric Fusion of Airborne and Street-Side Data for Urban Reconstruction

András Bódis-Szomorú
Computer Vision Lab, ETH Zurich

Hayko Riemenschneider
Computer Vision Lab, ETH Zurich

Luc Van Gool
Computer Vision Lab, ETH Zurich

Abstract—Airborne acquisition and on-road mobile mapping provide complementary 3D information of an urban landscape: the former acquires roof structures, ground, and vegetation at a large scale, but lacks the facade and street-side details, while the latter is incomplete for higher floors and often totally misses out on pedestrian-only areas or undriven districts. In this work, we introduce an approach that efficiently unifies a detailed street-side Structure-from-Motion (SfM) or Multi-View Stereo (MVS) point cloud and a coarser but more complete point cloud from airborne acquisition in a joint surface mesh. We propose a point cloud blending and a volumetric fusion based on ray casting across a 3D tetrahedralization (3DT), extended with data reduction techniques to handle large datasets. To the best of our knowledge, we are the first to adopt a 3DT approach for airborne/street-side data fusion. Our pipeline exploits typical characteristics of airborne and ground data, and produces a seamless, watertight mesh that is both complete and detailed. Experiments on 3D urban data from multiple sources and different data densities show the effectiveness and benefits of our approach.

I. INTRODUCTION

Structure-from-Motion (SfM) techniques have recently been employed at unprecedented scales [10] to jointly reconstruct outdoor scenes and geo-locate images. However, the resulting point cloud is fragmented and inhomogeneous. Multi-View Stereo (MVS) methods are often used to compute a dense surface from known views. Advances in MVS have enabled the automated production of large-scale urban models from airborne imagery and it has become a solid alternative to airborne LiDAR. Unfortunately, such models still often lack facade and street-level details due to occlusions and shadows. In turn, MVS has also been applied to street-side mapping [21, 28], as an alternative to LiDAR mobile mapping [7]. Automated mobile mapping solutions are prepared to deliver facade details at city-scale, but have no coverage of roofs, higher floors (e.g. in narrow streets), or traffic-free areas such as courtyards (see Figure 1). Attempts to model and abstract buildings from such data typically result in facade models much like a floating Potemkin village, i.e. with roofs and ground missing [7, 24, 27]. We conclude that both airborne and mobile mapping data needs to be exploited in order to produce the next generation of large-scale city models that are to be both complete and detailed.

In this paper, we propose a solution to fuse an airborne point cloud of large coverage but possibly low detail and a detailed but incomplete street-side point cloud in a volumetric fashion. It is a two-step approach: a mutually exclusive point cloud blending based on graph-cut segmentation, followed by a volumetric optimization over a tetrahedral space-partitioning and based on lines of sight to the points, inspired by [18, 19].

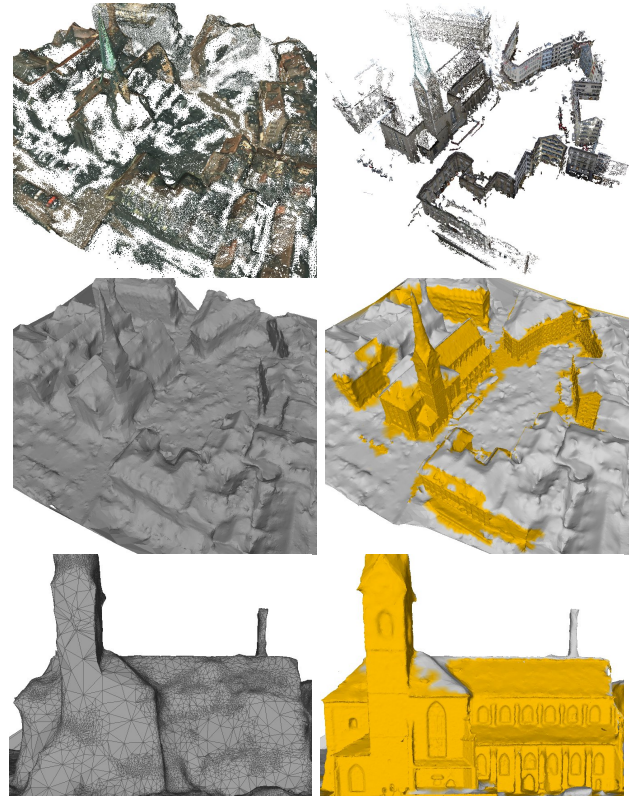


Fig. 1. Our method fuses a complete but inaccurate airborne point cloud (top left) and a detailed but incomplete street-side point cloud (top right) into a joint watertight mesh (right). In comparison, our mesh reconstructed solely from airborne nadir data is shown on the left.

We introduce several simplifications to the volumetric reasoning in order to cope with larger amounts of data, and we show that a drastic reduction of the number of visibility rays and limiting the range of ray-shooting have little effect on reconstruction quality, while the computational workload is reduced significantly. In the same vein, we carry out experiments on different input densities: sparse SfM and dense MVS data. The result is a watertight mesh that is complete at a large scale, but also detailed where street-side data is available.

Our contributions are the followings. (1) To the best of our knowledge, we are the first to apply volumetric fusion based on a 3D tetrahedralization (3DT) for joint airborne-streetside urban modeling. We propose (2) a mutually exclusive point cloud blending to cope with gross ray conflicts prior to 3DT-fusion (Sect. III-B) and (3) techniques to reduce the computational workload for larger scenes (Sect. III-C). Finally, we provide (4) detailed experimentation focusing on the trade-off between computational workload and surface detail (Sect. IV).

II. RELATED WORK

Many approaches have been proposed to convert a series of depth maps (range images) or a point cloud – obtained from SfM, MVS or LiDAR – into a consistent surface mesh.

Explicit methods directly construct mesh faces over the input points or depth data or join multiple meshes by zippering them, e.g. [7, 8, 21]. These are often interpolatory, sensitive to noise and can result in open meshes with holes and non-manifold areas. Implicit methods extract the surface as a level-set of a volumetric function evaluated over a spatial grid. The seminal work [5] integrates range images into a voxel grid via a weighted average over Truncated Signed Distance Functions (TSDF). It has inspired many indoor fusion techniques, e.g. KinectFusion [12]. Poisson reconstruction [16] is another popular approach due to its robustness to noise. In general, SDF-based methods tend to oscillate around noisy input. Thus, [11, 20] drop the sign of the distance function and find the surface as an s/t -cut in a graph over voxels, with regularization for a smooth result. Both only show results for small objects. Convex variational methods have also been applied for fusing noisy depth maps over a 3D voxel grid [31] or a 2.5D height map [22].

Volumetric reasoning over a voxel grid is prohibitively expensive for large outdoor scenes, motivating adaptive space partitioning methods. Octrees are used in Dual Contouring [14] and Poisson reconstruction [16] but 3D Delaunay-tetrahedralization (3DT) [17, 29], and a cell complex built from a plane arrangement [4] have also been applied to such scenes. Labatut et al. [17] build a 3DT on top of MVS points and do inside/outside classification of tetrahedra while enforcing line-of-sight and photoconsistency constraints. Later, the expensive photoconsistency term was dropped, and a remarkable robustness to outliers and to subsampling was shown [18]. Their street-side results appear superior in quality compared to TV-L¹ fusion [31]. As a drawback, the output mesh is interpolatory with the noisy input points as vertices. The 3DT approach has been extended by a photoconsistency-driven mesh optimization [29], and it has been adapted to point clouds with no visibility data by casting ray “tubes” [19]. Another extension addressing weakly-supported surfaces is implemented in CMP-MVS [13], which outputs a good-quality mesh in a matter of hours from hundreds of images of a street scene on a 3 GHz computer with GPU. It shall be noted that none of these techniques are specific to fusing aerial and street-side data.

Recently, we proposed an efficient view-driven meshing approach for street-side images and for large-scale height maps [2]. The benefits of combining different data sources for urban modeling have also been recognized by others [27, 23]. However, surface *fusion* of street-side and aerial data has received significantly less attention to date. Many works focus on geo-localization of either images or street-side reconstructions [27, 15, 26], a prerequisite for fusion. Their input is community photo collections, which are spatially fragmented, i.e. only represent popular landmarks. In turn, systematic industrial airborne/street-side mobile mapping deliver denser coverage and geo-registered data by means of GPS/IMU sensors and/or ground control points. Therefore, our approach assumes that the input is geo-registered (aligned) and focuses on the surface reconstruction step that tolerates minor misalignments.

Last but not least, there exist only a few works that combine street-side and aerial data for joint mesh reconstruction [8, 6, 25]. Früh and Zakhor [8] construct meshes over street-side LiDAR range maps and over a large-scale Digital Surface Model (DSM). Unlike in our approach, they reconstruct a facade and an airborne mesh separately without topological fusion. Shan et al. [25] solves the problem by directly applying Poisson surface reconstruction [16] over the joint dense point cloud computed by patch-based MVS [9] without a cross-consistency check between airborne and street-side data. Fiocco et al. [6] integrate over 200M points from a tripod-mounted ground LiDAR and an aerial DSM, by using a distance field over an octree and an out-of-core dual contouring approach. Despite the quite complex algorithm, results are noisy and contain many large holes. In contrast, our approach performs cross-consistency filtering and produces a good quality watertight surface.

III. THE PROPOSED METHOD

In what follows, we first detail the volumetric fusion method, then propose an *a priori* point cloud blending, and finally discuss how we reduce the computational effort.

A. Volumetric fusion

Inspired by the elegance of tetrahedral space partitioning methods [18, 19], we use a 3D Delaunay-tetrahedralization (3DT) over the joint point cloud from airborne and street-side acquisition (with the modifications discussed later). The underlying data structure is simpler than an octree or cell complex, yet it is adaptive to data density, a prerequisite for scalability. Denote the vertex set of the 3DT as $\mathcal{V} = \{v_i\}$, the set of tetrahedra by $\mathcal{T} = \{t_i\}$, and the triangular facet between any adjacent tetrahedra t_i and t_j as f_{ij} . Our goal is to assign a binary label $l_i \in \{in, out\}$ to every tetrahedron $t_i \in \mathcal{T}$, while minimizing the energy function

$$E(\mathcal{L}) = \sum_{i: t_i \in \mathcal{T}} E_i(l_i) + \sum_i \sum_{j: j < i} E_{ij} \cdot \mathbb{I}[l_i \neq l_j], \quad (1)$$

where the unary term $E_i(l)$ encodes the preference for tetrahedron t_i to obtain label l , the pairwise regularization term E_{ij} is the preference for two adjacent tetrahedra t_i and t_j to obtain the same label, $\mathbb{I}[l_i \neq l_j]$ is the indicator (1 if $l_i \neq l_j$, 0 otherwise), and $\mathcal{L} = (l_1, l_2, \dots, l_n)$ is a complete labelling.

As input, we construct lines of sight (rays) from the visibility information of the SfM or MVS (or laser scanning) procedure. Every ray $r = (v, s)$ goes from a vertex $v \in \mathcal{V}$ to a sensor location s . The logic behind inside/outside reasoning is that tetrahedra crossed by rays should be labelled as *out*, while tetrahedra behind each ray’s vertex of origin should be labelled as *in* up to a certain distance δ_{max}^{in} . This distance should be conservative to avoid mistakes around corners and behind narrow objects. In practice, rays either penetrate or do not reach the true surface due to noise as shown in Fig. 2.

Each ray contributes to the unary terms $E_i(l)$ in Eq. (1). For $l = out$, we follow the path of each ray r by walking in the 3DT (Fig. 3). The walking procedure starts from the source vertex v of the ray, returns each traversed tetrahedron t , and the walking distance to the point where the ray exits t . The walk ends with the tetrahedron containing the sensor location

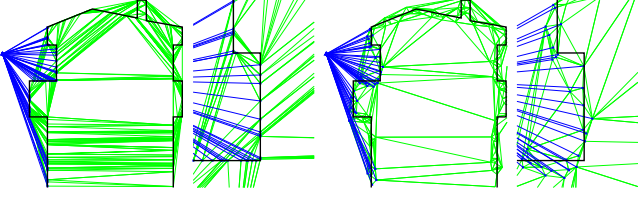


Fig. 2. 2D illustration of the effect of noise on tetrahedralization (green) and on the visibility rays (blue) to a single sensor. Point samples on the true surface (black) are Delaunay-triangulated without (left) and with noise (right). Close-ups are also shown. Noise increased the number of triangles by 8%.

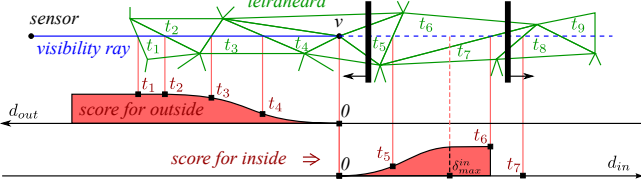


Fig. 3. Penalties arising from a single ray (solid blue) originated from vertex v . The true surface (black lines) is oriented towards outside (arrows). Vertices are off-surface due to noise. Crossed tetrahedra $\{t_1, t_2, t_3, t_4\}$ have a preference for $l = out$, while $\{t_5, t_6\}$ have preference for $l = in$.

s , or where the ray exits the 3DT, whichever occurs first. For $l = in$, the inverted ray is traced (dotted blue line in Fig. 3) until the tetrahedron at distance δ_{max}^{in} from v .

In order to take noise (Fig. 2) into account, we employ a soft-voting scheme, similar to [18]. Denote the exit distance of ray r from tetrahedron t by $d_l(r, t)$, where $l \in \{in, out\}$ distinguishes the two walking directions. $d_l(r, t) = 0$ for tetrahedra not affected by ray r . For a single ray r , the score of any tetrahedron t to be labelled as $l \in \{in, out\}$ is

$$S_l(r, t) = 1 - e^{-d_l^2(r, t)/(2\sigma_l^2)}, \quad (2)$$

where σ_{in} and σ_{out} are noise tolerance parameters. The truncation distance is fixed to $\delta_{max}^{in} = 3\sigma_{in}$, i.e. the last tetrahedron traversed backwards obtains the full score 1 for $l = in$ (see Fig. 3). The penalty $E_i(l)$ of tetrahedron t_i for label l is then defined as a monotonic function of the sum of the preferences of the opposite label \bar{l} over all rays:

$$E_i(l) = 1 - e^{-U_i(l)/\gamma_l} \quad \text{with} \quad U_i(l) = \sum_r S_{\bar{l}}(r, t_i), \quad (3)$$

where γ_{in} and γ_{out} control how many in and out votes collected from all rays per tetrahedron count as “many enough”, and can be used to trade-off robustness and accuracy.

The terms E_{ij} in Eq. (1) enforce spatial regularization for a smooth solution, and propagate labels to tetrahedra that have no unary preference. E_{ij} is a pairwise term that is a function of properties of two adjacent tetrahedra t_i and t_j . In this paper, we show experiments using a simple minimal area force

$$E_{ij} = \lambda A_{ij}, \quad (4)$$

where λ is a global regularization factor, and A_{ij} is the area of triangle f_{ij} between tetrahedra t_i and t_j . Other forces are also possible: an initial experimentation with forces against elongated triangles, long-edge triangles, or preference for faces f_{ij} between any two elongated tetrahedra [18] (see Fig. 2 for an intuition) shows that the area penalty works best with our datasets. Further comparison is out of the scope of this paper.

Note that the margins introduced by the soft-voting scheme (Fig. 3) for the unary term let the optimization select from a larger pool of faces.

The energy function (1) can be globally minimized using graph-cuts [3]. The sought optimal surface mesh can be extracted as the set of all faces between inside and outside tetrahedra. Although this mesh is watertight, it is interpolatory, since the 3DT is built on top of noisy points. This results in a jagged surface inherent in most 3DT-based methods. To obtain a smoother surface, we apply basic Laplacian smoothing by replacing every vertex by the mean of its neighbors. Finally, only the largest component is kept in order to eliminate (typically a few) minor components not attached to the ground.

B. Point cloud blending

Unfortunately, applying the proposed volumetric fusion procedure *directly* to the joint airborne/street-side set of points does not give satisfactory results in practice. The main reason is that airborne reconstructions are usually substantially lower quality and oversmoothed at facades, which results in duplicate surfaces and ray conflicts across the two data types. Rays from airborne data often intersect the true surface or show inside evidence far outside of the true surface. Setting our fusion parameters $\sigma_{in}, \sigma_{out}$ loose would result in loss of detail. To get the best out of both data types, we rather propose a mutually exclusive point cloud blending prior to volumetric fusion. Besides the considerations that (i) the airborne data has broad coverage, (ii) street-side input is incomplete or fragmented and (iii) airborne data is of substantially lower quality where street-side data is available, it is reasonable to eliminate airborne points where street side data is present. The removal of points comes with the additional computational benefit of lowering the number of tetrahedra in the partitioning.

We formulate this as a segmentation over the airborne point cloud \mathcal{P} , which assigns a binary label $l \in \{0, 1\}$ to each point $p_i \in \mathcal{P}$, points to remove being marked by 0. Given the street-side point cloud \mathcal{Q} , an airborne point p_i has a (supposedly better quality) street-side substitute $q_i \in \mathcal{Q}$ if q_i is the nearest neighbor of p_i in \mathcal{Q} and if both the Euclidean distance d_i between them and the angle θ_i between their normals $\mathbf{n}(p_i)$ and $\mathbf{n}(q_i)$ are small. The likelihood for an airborne point to have a substitute can be formulated as

$$\phi_i = \phi(d_i, \theta_i) = e^{-d_i^2/(2\sigma_b^2)} \cdot \max\{0, \cos \theta_i\}, \quad (5)$$

where $\cos \theta_i = \mathbf{n}(p_i)^T \mathbf{n}(q_i)$, and ϕ ranges from 0 (no substitute) to 1 (perfect substitute). σ_b is a blending parameter to control our notion of vicinity, which should incorporate deviations of \mathcal{P} from \mathcal{Q} due to registration and reconstruction errors. Normals are computed in \mathcal{P} and \mathcal{Q} separately by least-squares plane fitting to the k -NN neighborhood of each point ($k = 10$), and by flipping normals according to visibility. For a smooth segmentation, we define the influence between adjacent nodes of the k -NN graph over \mathcal{P} as

$$\psi(p_i, p_j) = \exp(-d_{ij}/\text{med } d_{ij}), \quad (6)$$

where d_{ij} is the distance between any two adjacent points p_i and p_j , and $\text{med } d_{ij}$ is the median of all k -NN distances in \mathcal{P} . We seek the binary labelling over \mathcal{P} that minimizes

$$E^b(\mathcal{L}) = \sum_{i:p_i \in \mathcal{P}} E_i^b(l_i) + \lambda_b \sum_{i,j} \psi(p_i, p_j) \cdot \mathbb{I}[l_i \neq l_j], \quad (7)$$

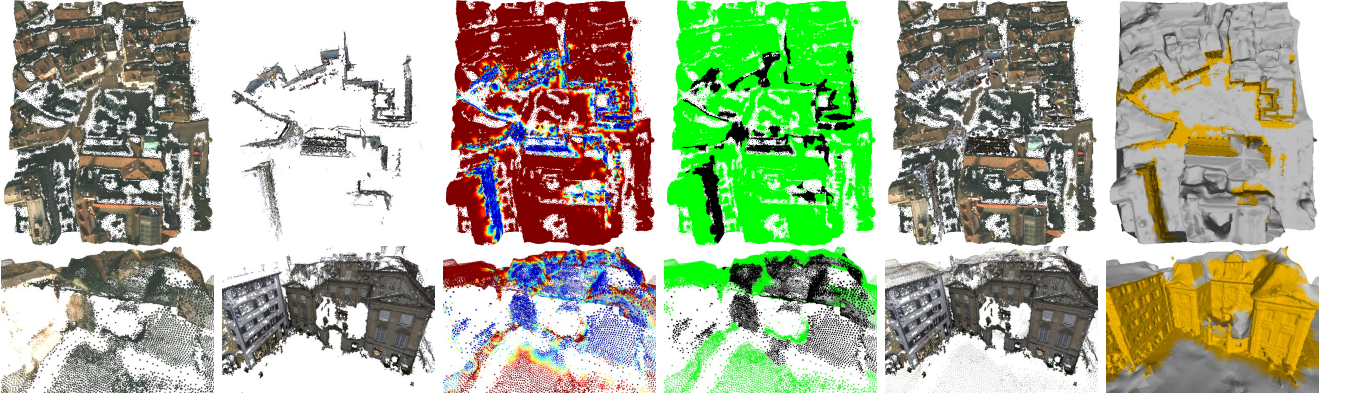


Fig. 4. Our pipeline applied to Münsterhof (PMVS). Overview (top) and a close-up (bottom). From left to right: input aerial point cloud, input street side point cloud, blending energies, airborne point cloud segmentation, blended aerial and street-side point clouds, output fused mesh with street-side parts in orange.

where $l_i \in \{0, 1\}$ is the label of point p_i , \mathbb{I} is the indicator function, and λ_b is a regularization parameter. The unary penalties $E_i^b(l)$ for point p_i to obtain label l are defined as

$$E_i^b(0) = 1 - \phi_i \text{ and } E_i^b(1) = \phi_i. \quad (8)$$

We minimize Eq. (7) by graph-cuts [3] and eliminate the airborne points labelled 0 prior to volumetric fusion. Although point quality measures could be incorporated into this scheme, we have found that the distance and normal cues prove sufficient to obtain a good mixture of the two point clouds.

C. Data reduction

Our aim is to apply the pipeline to large scenes. To reduce the computational complexity, we propose different simplifications, whose effect will be experimented in Section IV.

1) *Input point decimation*: We cluster the blended point cloud using a voxel grid, and replace the points in each voxel by their centroid. The visibility information is also merged accordingly. To keep the method scalable, only occupied voxels (a small fraction) are stored, using an associative array indexed with the global voxel index.

2) *Reducing the number of rays*: Lines-of-sight are distributed densely in free-space areas, leaving space for a sparsification. We reduce the number of rays per point to one. The ray that is closest to the normal direction of each point is kept, to minimize surface penetration due to noise in the points.

3) *Ray truncation*: When collecting inside votes in the fusion algorithm (Sect. III-A), the penetration depth has intentionally been limited to $\delta_{max}^{in} = 3\sigma_{in}$ to preserve narrow structures. We introduce a similar limit $\delta_{max}^{out} = 3\sigma_{out}$ when tracing the rays from each vertex v towards the sensor (solid blue line in Fig. 3) for a shorter walk in the 3DT.

IV. EXPERIMENTS

Our mixed C++/Matlab single-core implementation relies on Matlab/CGAL [1] for tetrahedralization, and on the GCOptimization library [3] for graph-cuts. We implemented 3DT adjacency computations and ray shooting in C++.

We are not aware of any publicly available dataset with both airborne and street-side data for the same geographic location. Thus, we show experiments on our datasets, Münsterhof

(140×160 m², Fig. 4) and Limmatquai (400×400 m², Fig. 5) captured in Zürich, Switzerland. 23 nadir images of 15 cm ground resolution, taken at a height of 3 km, were fed to a state-of-the-art MVS pipeline¹ to obtain an airborne point cloud. 629 (Münsterhof) and 847 (Limmatquai) street-side images were taken manually in the same areas, and we ran VisualSfM [30] to produce an SfM point cloud, as well as PMVS [9] for a denser cloud. The aerial data was geo-registered via ground control points, and the street-side models were registered manually to the aerial data (assuming that street-side geo-registration can be solved automatically in industrial mobile mapping). To quantify input misalignment (incl. noise), we measure distances between mutual nearest neighbors across the SfM/aerial clouds along the normal directions pointwise, and report the median, 90- and 99-percentiles. These are 0.2, 0.5, 1.0 m for Münsterhof and 0.2, 0.6, 1.2 m for Limmatquai.

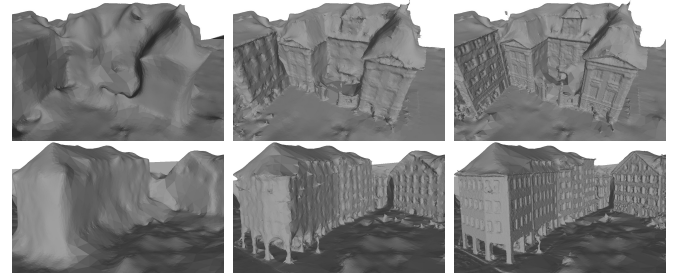


Fig. 6. Our reconstruction of airborne-only (left) airborne + street-side SfM (middle) and airborne + street-side PMVS input (right).

We applied our method to airborne-only data as a baseline and to airborne combined with street-side data from (i) SfM and (ii) PMVS and using different combinations of the reductions of Sect. III-C. Fig. 4 shows different stages of our pipeline. We focus on the trade-off between runtime and output quality. Fig. 6 compares models from different data types, while numerical results are summarized in Table I. In all experiments, we fix our parameters as $\sigma_b = 2$ m and $\lambda_b = 1$ for point cloud blending, and $\sigma_{in} = 0.1$ m, $\sigma_{out} = 0.5$ m and $\gamma_{in} = \gamma_{out} = 2$ for 3DT-fusion. The volumetric regularization parameter λ is shown in Table I. Denser input allows for better noise suppression without oversmoothing (higher λ). Input point decimation (Sect. III-C) is parametrized by the voxel size vox (0 for no decimation), ray truncation is marked by a \checkmark in column *tr* and reduction of the number of rays is marked

¹RealityCapture, <https://www.capturingreality.com/>

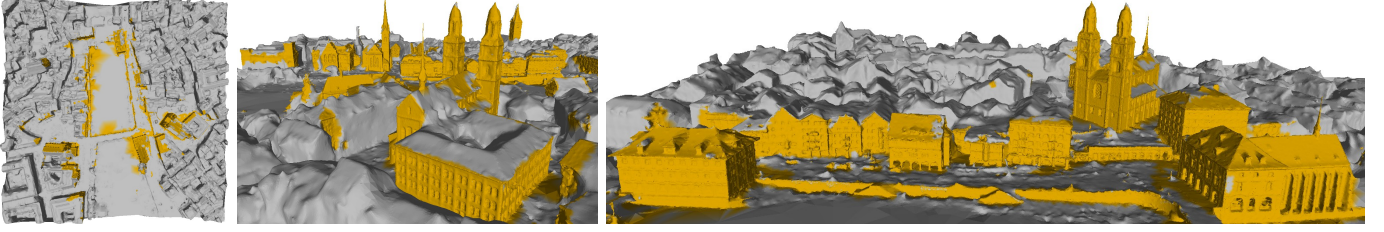


Fig. 5. Output of our fusion for the Limmatquai dataset using PMVS street-side input (see PMVS (*) in Table I). Vertices originated from street-side data are colored orange. The left image shows the full 400×400 m² result. Around 20% of surface area is covered by street-side data.

| | | Params | | | Airborne data | | | Street-side data | | | 3DT & Ray shooting | | | Mesh | Mesh distance from ref (cm) | | | | Timings | | | |
|------------|------------|--------|----|-----------|---------------|------------------|------|------------------|------------------|------|--------------------|-------|-------|--------|-----------------------------|-------------------|-------------------|-------------------|---------|-------|------|-------|
| | Experiment | vox | tr | λ | #pts | #/m ² | #vis | #pts | #/m ² | #vis | #verts | #tets | #rays | #verts | mean ^a | mean ^s | 10cm ^s | 50cm ^s | tet | ray | gco | total |
| Münsterhof | PMVS (*) | 0 | | 3 | 272k | 8.0 | 11.9 | 1.54M | 108.0 | 6.8 | 1.75M | 11.0M | 13.2M | 1.44M | (ref) | (ref) | (ref) | (ref) | 32s | 182s | 27s | 285s |
| | PMVS (a) | 0 | ✓ | 3 | 272k | 7.9 | 11.9 | 1.54M | 106.8 | 6.8 | 1.75M | 11.0M | 13.2M | 1.43M | 2.9cm | 1.3cm | 0.7% | 0.0% | 32s | 80s | 28s | 182s |
| | PMVS (b) | 0 | | 3 | 272k | 8.4 | 1.0 | 1.54M | 108.9 | 1.0 | 1.75M | 11.0M | 1.75M | 1.24M | 6.7cm | 1.5cm | 0.8% | 0.0% | 32s | 28s | 34s | 137s |
| | PMVS (c) | 0 | ✓ | 3 | 272k | 8.5 | 1.0 | 1.54M | 108.5 | 1.0 | 1.75M | 11.0M | 1.75M | 1.23M | 8.6cm | 1.6cm | 0.9% | 0.0% | 32s | 15s | 34s | 125s |
| | PMVS (d) | 0.10 | | 3 | 272k | 8.4 | 1.0 | 780k | 56.3 | 1.0 | 996k | 6.27M | 996k | 756k | 6.8cm | 1.7cm | 1.0% | 0.0% | 18s | 16s | 24s | 82s |
| | PMVS (e) | 0.20 | | 3 | 266k | 8.2 | 1.0 | 310k | 23.5 | 1.0 | 522k | 3.33M | 522k | 416k | 7.1cm | 2.2cm | 2.5% | 0.1% | 9.4s | 7.9s | 16s | 46s |
| | PMVS (f) | 0.35 | | 1 | 213k | 6.3 | 1.0 | 127k | 6.3 | 1.0 | 293k | 1.89M | 293k | 262k | 3.0cm | 3.1cm | 5.1% | 0.3% | 5.2s | 4.2s | 6.0s | 23s |
| | SfM (*) | 0 | | 1 | 272k | 7.8 | 11.9 | 233k | 18.6 | 4.6 | 452k | 2.84M | 3.91M | 410k | 3.6cm | 7.7cm | 19.8% | 1.4% | 8.0s | 45s | 7.6s | 70s |
| | SfM (a) | 0 | ✓ | 1 | 272k | 7.6 | 11.9 | 233k | 18.7 | 4.6 | 452k | 2.84M | 3.91M | 412k | 3.6cm | 7.1cm | 18.1% | 1.1% | 9.8s | 19s | 6.5s | 47s |
| | SfM (b) | 0 | | 1 | 272k | 7.8 | 1.0 | 233k | 18.7 | 1.0 | 452k | 2.84M | 452k | 382k | 3.6cm | 7.5cm | 18.9% | 1.3% | 8.1s | 6.4s | 8.4s | 33s |
| | SfM (c) | 0 | ✓ | 1 | 272k | 7.7 | 1.0 | 233k | 19.0 | 1.0 | 452k | 2.84M | 452k | 383k | 3.5cm | 7.5cm | 18.7% | 1.3% | 8.5s | 3.2s | 8.4s | 29s |
| | aerial (*) | 0 | | 1 | 272k | 6.2 | 11.9 | not used | | | 272k | 1.77M | 3.24M | 269k | 3.1cm | 69cm | 86.4% | 49.1% | 5.0s | 41s | 3.3s | 52s |
| | aerial (a) | 0 | ✓ | 1 | 272k | 6.3 | 1.0 | not used | | | 272k | 1.77M | 272k | 253k | 3.1cm | 70cm | 86.5% | 49.2% | 7.3s | 2.0s | 4.0s | 16s |
| Limmatquai | PMVS (*) | 0.20 | | 3 | 1.65M | 6.8 | 11.5 | 1.14M | 26.0 | 43.6 | 2.60M | 16.7M | 67.8M | 2.41M | (ref) | (ref) | (ref) | (ref) | 49s | 1059s | 53s | 1261s |
| | PMVS (a) | 0.20 | ✓ | 3 | 1.65M | 6.7 | 11.5 | 1.14M | 26.3 | 43.6 | 2.60M | 16.7M | 67.8M | 2.45M | 0.2cm | 0.3cm | 0.9% | 0.1% | 51s | 314s | 53s | 517s |
| | PMVS (b) | 0.20 | | 3 | 1.65M | 6.9 | 1.0 | 1.14M | 26.5 | 1.0 | 2.60M | 16.7M | 2.60M | 2.11M | 3.7cm | 2.6cm | 8.1% | 0.3% | 49s | 50s | 79s | 277s |
| | PMVS (c) | 0.20 | ✓ | 3 | 1.65M | 6.8 | 1.0 | 1.14M | 26.9 | 1.0 | 2.60M | 16.7M | 2.60M | 2.13M | 3.8cm | 3.0cm | 8.9% | 0.4% | 50s | 19s | 78s | 248s |
| | SfM (*) | 0 | | 1 | 1.68M | 6.5 | 11.5 | 259k | 9.1 | 4.6 | 1.79M | 11.5M | 19.6M | 1.72M | 0.3cm | 20cm | 44.4% | 8.9% | 35s | 262s | 35s | 380s |
| | SfM (a) | 0 | ✓ | 1 | 1.68M | 6.5 | 11.5 | 259k | 9.2 | 4.6 | 1.79M | 11.5M | 19.6M | 1.73M | 0.3cm | 19cm | 42.4% | 7.7% | 35s | 95s | 35s | 213s |
| | SfM (b) | 0 | | 1 | 1.68M | 6.6 | 1.0 | 259k | 9.0 | 1.0 | 1.79M | 11.5M | 1.79M | 1.62M | 0.6cm | 19cm | 44.0% | 8.4% | 34s | 30s | 39s | 153s |
| | SfM (c) | 0 | ✓ | 1 | 1.68M | 6.5 | 1.0 | 259k | 9.5 | 1.0 | 1.79M | 11.5M | 1.79M | 1.62M | 0.6cm | 21cm | 44.0% | 8.8% | 34s | 12s | 39s | 134s |
| | aerial (*) | 0 | | 1 | 1.68M | 6.0 | 11.5 | not used | | | 1.68M | 10.9M | 19.3M | 1.66M | 0.0cm | 81cm | 88% | 47.3% | 33s | 269s | 33s | 340s |
| | aerial (a) | 0 | ✓ | 1 | 1.68M | 6.0 | 1.0 | not used | | | 1.68M | 10.9M | 1.68M | 1.57M | 0.3cm | 81cm | 88% | 47.4% | 33s | 12s | 26s | 103s |

TABLE I. SUMMARY OF OUR RESULTS FOR TWO DATASETS. CHANGES W.R.T. THE LINE (*) IN EACH GROUP ARE MARKED BY BOLD TEXT.

by value 1.0 in the columns reporting the number of visibility rays per point (#vis) in Table I. The input point density (#/m²) is estimated as the number of input points divided by output surface area. The table also reports the number of points (#pts), tetrahedra (#tets), and mesh vertices (#verts) in each case. Note that the number of triangles in the 3DT (resp. output mesh) is nearly twice the vertex count. Since no detailed ground-truth 3D model exists for the fusion task to date, we evaluate the effect of the proposed simplifications w.r.t. the best-quality 3D model obtained by our method (marked by *ref*) in Table I). The distances from each vertex of the (dense) reference mesh to the output mesh is measured using AABB-trees [1]. As the aerial input is not altered in our experiments, and to report results independently from street-side coverage, we partition the reference mesh into street-side and aerial regions based on the origin of its vertices, and plot the error distributions for the street-side part in Fig. 7. Similarly, Table I reports the mean error over the aerial (mean^a) and street-side part (mean^s) separately, and the percentage of distances larger than 10 and 50 cm over the street-side part.

Runtimes are measured in Linux on a single i7-2600K 3.4GHz CPU core, 16 GB RAM. Table I also lists runtimes for computing the 3DT and its adjacencies (tet), ray shooting (ray), energy computations and volumetric optimization (gco) and total time, which also includes point cloud *k*-NN and normal computation, point cloud blending and post-processing. For Limmatquai, we limited the input density to around 26 points/m² (via *vox* = 0.2 m) having a memory peak of around 11.6 GB, the most memory-intensive steps being the

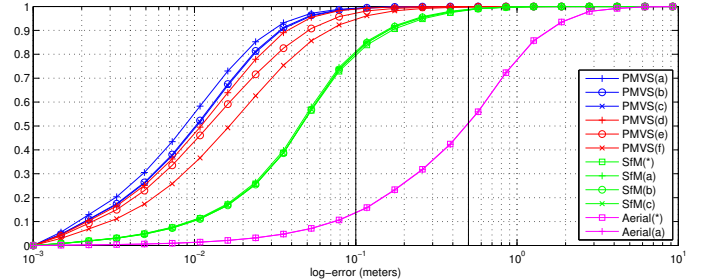


Fig. 7. Cumulated distribution of mesh errors for the cases in Table I over the street-side data regions of Münsterhof, the reference being PMVS(*). Vertical lines mark the thresholds 10cm and 50cm used for error reporting in the Table.

volumetric decomposition and graph-cut optimization over it.

Our results in Table I and Fig. 7 consistently show that the quality of the model improves substantially by fusing street-side details into the airborne model, as expected. Improvements of using PMVS data instead of SfM are in the 1-10 cm range. Depending on the use-case, the quality provided by the street-side SfM may suffice. As an example, our method could be used for improving satellite or airborne Digital Surface Models (DSMs). It is also clear from the results that the proposed simplifications cause only little harm in quality (curves of the same color are close in Fig. 7) while a significant gain in runtime (and memory requirements) is achieved. We also found that ray truncation improves robustness to gross outliers. In particular, PMVS results in larger groups of outlier rays crossing the interior of certain buildings, while their effect remains local due to truncation. Although ray truncation makes

the optimization local (unary terms are only non-zero in a crest around the data) [20], there are no serious consequences, once δ_{max}^{out} is set high enough so that street-side rays sweep out protrusions of the airborne input at the bottom of walls (see left of Fig. 6). As a failure case, point cloud blending tends not to disambiguate narrow structures (e.g. towers) if these are duplicated distinctively due to misregistration in the input.

V. CONCLUSION

This paper proposes an efficient method for reconstructing a surface mesh from a detailed but incomplete street-side and a complete but low-detail airborne point cloud, which are pre-aligned. Although the next generation of 3D urban models should be both complete and contain street-side details, the problem received only little attention to date. Our method joins strengths of the two data types by fusing them via volumetric reasoning over a tetrahedralization, preceded by a point cloud blending to avoid gross line-of-sight conflicts due to inaccurate airborne measurements where more accurate street-side data is available. Our detailed experimentation with both SfM and dense MVS data over large urban scenes shows good surface quality and proves that additional simplifications can be used with little harm in output quality but substantial reduction in runtime. Applying the method at city scale via mobile mapping data is part of our future work. Ways to auto-adapt the parameters to input noise, ray density or tetrahedra volumes would also be interesting to study in the future.

ACKNOWLEDGMENT

This work was supported by the European Research Council project VarCity (#273940). We thank Dr. Michal Havlena, Wilfried Hartmann and Prof. Konrad Schindler for the Zürich aerial dataset, Michal Havlena for his comments on the paper.

REFERENCES

- [1] CGAL, The Computational Geometry Algorithms Library. <http://www.cgal.org>, version 4.6.3.
- [2] A. Bódis-Szomorú, H. Riemenschneider, and L. Van Gool. Efficient edge-aware surface mesh reconstruction for urban scenes. *CVIU*, 2016.
- [3] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *PAMI*, 23(11):1222–1239, Nov 2001.
- [4] A. Chauve, P. Labatut, and J. Pons. Robust piecewise-planar 3D reconstruction and completion from large-scale unstructured point data. In *CVPR*, 2010.
- [5] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH*, 1996.
- [6] M. Fiocco, G. Bostrom, J. G. M. Goncalves, and V. Sequeira. Multisensor fusion for volumetric reconstruction of large outdoor areas. In *3DIM*, 2005.
- [7] C. Früh, S. Jain, and A. Zakhor. Data processing algorithms for generating textured 3D building facade meshes from laser scans and camera images. *IJCV*, 61(2):159–184, 2005.
- [8] C. Früh and A. Zakhor. Constructing 3D city models by merging aerial and ground views. *IEEE Computer Graphics and Applications*, 23(6):52–61, 2003.
- [9] Y. Furukawa and J. Ponce. Accurate, Dense, and Robust Multi-View Stereopsis. *PAMI*, 32(8):1362–1376, 2010.
- [10] J. Heinly, J. L. Schönberger, E. Dunn, and J.-M. Frahm. Reconstructing the world in six days. In *CVPR*, 2015.
- [11] A. Hornung and L. Kobbelt. Robust reconstruction of watertight 3D models from non-uniformly sampled point clouds without normal information. In *Eurographics*, 2006.
- [12] S. Izadi, R. Newcombe, D. Kim, O. Hilliges, D. Molyneaux, S. Hodges, P. Kohli, J. Shotton, A. J. Davison, and A. Fitzgibbon. KinectFusion: Real-time dynamic 3D surface reconstruction and interaction. In *SIGGRAPH*, 2011.
- [13] M. Jancosek and T. Pajdla. Multi-view reconstruction preserving weakly-supported surfaces. In *CVPR*, 2011.
- [14] T. Ju, F. Losasso, S. Schaefer, and J. Warren. Dual contouring of hermite data. In *SIGGRAPH*, 2002.
- [15] R. S. Kaminsky, N. Snavely, S. M. Seitz, and R. Szeliski. Alignment of 3D point clouds to overhead images. In *Second IEEE Workshop on Internet Vision*, 2009.
- [16] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Eurographics*, 2006.
- [17] P. Labatut, J. Pons, and R. Keriven. Efficient multi-view reconstruction of large-scale scenes using interest points, Delaunay triangulation and graph cuts. In *ICCV*, 2007.
- [18] P. Labatut, J.-P. Pons, and R. Keriven. Robust and efficient surface reconstruction from range data. *Computer Graphics Forum*, 28(8):2275–2290, 2009.
- [19] F. Lafarge and P. Alliez. Surface reconstruction through point set structuring. In *Eurographics*, 2013.
- [20] V. Lempitsky and Y. Boykov. Global optimization for shape fitting. In *CVPR*, 2007.
- [21] L. Wang P. Mordohai J.-M. Frahm R. Yang D. Nistr P. Merrell, A. Akbarzadeh and M. Pollefeys. Real-time visibility-based fusion of depth maps. In *ICCV*, 2007.
- [22] T. Pock, L. Zebedin, and H. Bischof. *TGV-fusion*. Springer, 2011.
- [23] Cabezas R., O. Freifeld, Rosman G., and J Fisher III. Aerial reconstructions via probabilistic data fusion. In *CVPR*, 2014.
- [24] H. Riemenschneider, A. Bódis-Szomorú, J. Weissenberg, and L. Van Gool. Learning where to classify in multi-view semantic segmentation. In *ECCV*, 2014.
- [25] Q. Shan, R. Adams, B. Curless, Y. Furukawa, and S. M. Seitz. The Visual Turing Test for Scene Reconstruction. In *3DPVT*, 2013.
- [26] Q. Shan, C. Wu, B. Curless, Y. Furukawa, C. Hernandez, and S. Seitz. Accurate geo-registration by ground-to-aerial image matching. In *3DPVT*, 2014.
- [27] C. Strecha, T. Pylvänäinen, and P. Fua. Dynamic and scalable large scale image reconstruction. In *CVPR*, 2010.
- [28] E. Tola, C. Strecha, and P. Fua. Efficient large scale multi-view stereo for ultra high resolution image sets. *Machine Vision and Applications*, 23(5):903–920, 2011.
- [29] H. Vu, P. Labatut, J. Pons, and R. Keriven. High accuracy and visibility-consistent dense multi-view stereo. *PAMI*, 34(5):889–901, 2012.
- [30] C. Wu. VisualSFM: A Visual Structure from Motion System. <http://ccwu.me/vsfm/>, version 0.5.22.
- [31] C. Zach, T. Pock, and H. Bischof. A globally optimal algorithm for robust TV-L¹ range image integration. In *ICCV*, 2007.