Stabilizing 3D Modelling with Geometric Constraints Propagation

Michela Farenzena, Andrea Fusiello

*Dipartimento di Informatica, Università di Verona,
Strada Le Grazie, 15 - 37134 Verona, Italy*

Abstract

This paper proposes a technique for estimating piecewise planar models of objects from their images and geometric constraints. First, assuming a bounded noise in the localization of 2D points, the position of the 3D point is estimated as a polyhedron containing all the possible solutions of the triangulation. Then, given the topological structure of the 3D points cloud, geometric relationships among facets, such as coplanarity, parallelism, orthogonality, and angle equality, are automatically detected. A subset of them that is sufficient to stabilize the 3D model estimation is selected with a flow-network based algorithm. Finally a feasible instance of the 3D model, i.e. one that satisfies the geometric constraints and whose 3D vertices lie within the associated polyhedral bounds, is computed by solving a Constraint Satisfaction Problem. The process accommodates uncertainty in a non-probabilistic fashion and thus provides rigorous results. Synthetic and real experiments illustrate the approach.

Key words:
PACS:

1 Introduction

Reconstruction of accurate and photorealistic 3D models is one of the challenging tasks in Computer Vision. It often shares problems and research fields with other communities such as Computer Graphics and Computer-Aided Design.

In this paper we address the problem of recovering 3D surface models from images and geometric clues. It is known in fact that using only the image in-

Email addresses: farenzena@sci.univr.it (Michela Farenzena), andrea.fusiello@univr.it (Andrea Fusiello).

formation is often an ill-conditioned problem; defining geometric constraints between scene primitives and incorporating them into the reconstruction system improves the quality of the model and limits the number of required images.

The methods proposed in literature can be mainly classified as *model-based* and *constraint-based*.

In *model-based methods* [1–4], the scene is defined as in CAD systems: objects are the assemblage of known primitive shapes. Reconstruction is carried out by fitting a 3D model to image data, thus determining its dimension, its position and orientation. The fact that the scene must be decomposable in primitive shapes is the main limitation of such methods.

Constraint-based methods [5–12] are more flexible, as they do not rely on a-priori models but use simple primitives like points and lines. Geometric information, such as orthogonality, parallelism, or planarity, is given in the form of constraints on 3D points and reconstruction is obtained as the solution of an optimization process.

In most of previous works, e.g. [4–6,9,10,12], the constraints detection phase requires the user to provide a geometrical description of the model, which can be very time-consuming. In other cases [7,8], geometric constraints are detected automatically thanks to prior knowledge about the model to reconstruct.

Besides, the *analysis* of constraints is usually overlooked. In fact, data sets with many points and geometric constraints do not necessarily define a consistent and unique 3D object. Parts of the scene may not be rigidly connected, so that there exist various shapes that verify the geometric constraints and project to identical image points. In addition, constraints may be redundant, making the optimization uselessly harder or even unfeasible. [10] proposes an algebraic method to check whether a configuration of points and constraints leads to a unique reconstruction, but it does not deal with redundancies. As far as we know, a principled analysis of constraints has not been proposed yet.

Geometric constraints may be directly embedded into the minimization of the reprojection error (or bundle adjustment) [5,7,8], but this causes a substantial increase of computational costs and both convergence and exact constraint satisfaction are not guaranteed. An alternative is to make the geometric constraints implicit in the parametrization of 3D points [6,10–12], so as they are satisfied exactly at every optimization step.

Our approach avoids altogether the non-linear least-squares problem arising in the methods above, for it casts the problem as a Constraint Satisfaction (CSP), where the 3D point positions are bounded by 3D boxes and a feasible solution is one that satisfies all selected geometric relationships and whose 3D points lie within the associated bounds.

These bounds are obtained as the result of a triangulation procedure that instead of producing *one* optimal solution, as customary, outputs the set of *all* the possible solutions, given a bounded error affecting the image points.

Accommodating uncertainty is crucial when the results are to be used as input for other processes. Albeit simple in concept, this triangulation is a principled, efficient and reliable approach for evaluating 3D point positions and the associated uncertainty.

Moreover, this work is enhanced by the automatic detection of constraints and by the subsequent analysis and pruning of these constraints that permits to verify if a unique solution can be obtained, and at the same time to prune redundancies. A preliminary version of this work is going to appear in [13].

2 Overview

The approach, summarized in Figure 1, consist of two stages.

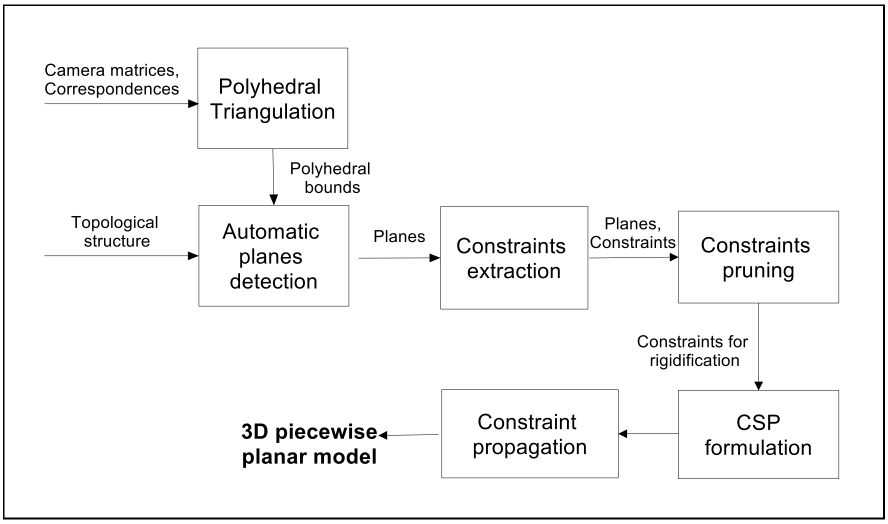

Fig. 1. Overview of the proposed method. External inputs are: camera matrices, 2D point correspondences, 3D points connectivity.

The first stage deals Section with triangulation, i.e., reconstructing 3D points from their corresponding image points (provided manually) and known camera matrices. Assuming that the error in the localization of image points is bounded by a rectangular region, our *polyhedral triangulation* (Section 3) computes the polyhedron that contains all the possible 3D point positions.

The second stage focuses on obtaining one single model by fixing the position of the 3D points in such a way that certain geometric relationships are satisfied. Given a set of reconstructed 3D points, represented by the polyhedra provided by the above 3D triangulation, and the connectivity of the points into triangular facets (provided manually), the method consists in a three-steps automatic process. First the geometric relationships (coplanarity, parallelism, orthogonality, and angle equality) are detected (Section 4); then a set of minimal relationships that allow a unique reconstruction is selected using the structural rigidity analysis (Section 5); finally, a feasible instance of the 3D model, i.e. one that satisfies all selected geometric relationships and whose 3D points lie within the associated polyhedral bounds, is computed using a constrained optimization technique (Section 6).

Experimental validation is reported in Section 7. Finally, conclusions are drawn in Section 8.

3 Polyhedral Triangulation

The first and most important stage of model reconstruction consists in recovering the coordinates of points in 3D space given their images in two or more views. It is usually assumed that the camera matrices are known exactly, or at least with greater accuracy than point localization. In the absence of noise, i.e. when correspondences are perfectly detected, the problem is trivial, involving only finding the intersection of rays in the space. If data are perturbed, however, the rays corresponding to back-projections of image points do not intersect, and obtaining the 3D coordinates of the reconstructed points becomes far from trivial, as witnessed by the renewed interest aroused by this issue [14,15]. A statistical optimal solution, under the assumption of Gaussian noise, exists for two [16] and three views [15], but seems to be unfeasible beyond that.

This problem can be circumvented if one refrains from searching for *one* optimal solution and computes instead a *set* of possible solutions (defined in terms of errors affecting the image points) that contains the error-free solution. This permits to bound the exact solution in the 3D space for any number of views and to estimate, at the same time, the uncertainty of the result.

Let P^i, $i = 1, \ldots, n$ be a sequence of n known cameras and m^i be the image of some unknown point M in 3D space, both expressed in homogeneous coordinates. It is assumed that the localization error is bounded by a rectangular region $\mathcal{B}_i$ centered around each image point (one can imagine a uniform distribution inside $\mathcal{B}_i$). Each region $\mathcal{B}_i$ bounds the possible locus of the 3D point inside a semi-infinite pyramid $\mathcal{Q}_i$ with its apex in the camera center (see

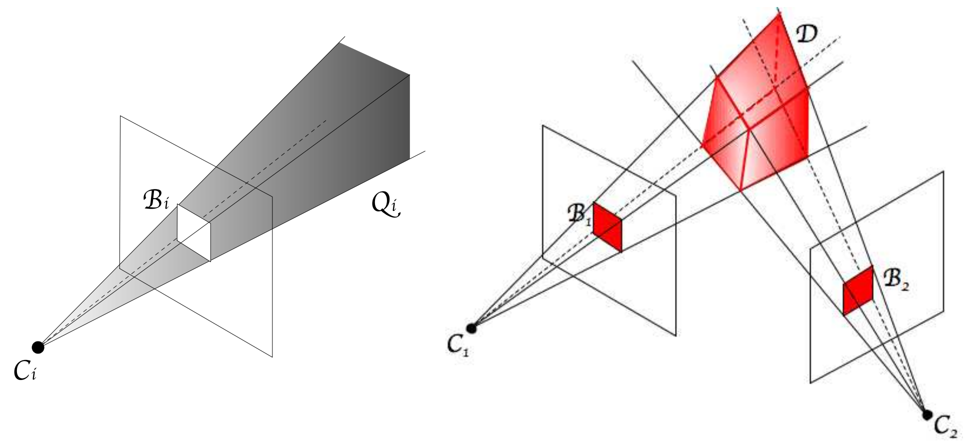

Fig. 2. On the left, the semi-infinite pyramid $\mathcal{Q}_i$ is defined from the camera centre $\mathcal{C}_i$ and the bound $\mathcal{B}_i$. On the right, the polyhedron $\mathcal{D}$ is the intersection of $\mathcal{Q}_1$ and $\mathcal{Q}_2$.

Figure 2). The solution set is defined as the polyhedron formed by the intersection of the n semi-infinite pyramids generated by the intervals $\mathcal{B}_1, \dots, \mathcal{B}_n$. Analytically, this region is defined as the following set:

$$\mathcal{D} = \mathcal{Q}_1 \cap \mathcal{Q}_2 \cdots \cap \mathcal{Q}_n = \{\mathrm{M}\colon \exists \mathrm{m}^i \in \mathcal{B}_i, i = 1 \dots n \text{ s.t. } \forall i\colon \mathrm{m}^i \simeq P^i\mathrm{M}\}. \quad (1)$$

This polyhedron is the best piece of information about the localization of the error-free 3D point one can deduce from the bounded error in image points. In other words, the solution set bounds the unknown probability distribution function over the possible 3D point positions. This approach represents the counterpoise of the max-likelihood approach using the Gaussian error model.

Instead of approximating $\mathcal{D}$ using Interval Analysis as in [9], it is computed precisely using Computational Geometry techniques. The semi-infinite pyramid $\mathcal{Q}_i$ can be written as the intersection of the four negative half-spaces $\mathcal{H}_1^i, \mathcal{H}_2^i, \mathcal{H}_3^i, \mathcal{H}_4^i$ defined by its supporting planes. Thus, the solution set D can be expressed as the intersection of $4n$ negative half-spaces:

$$\mathcal{D} = \bigcap_{\substack{i=1\dots n\\ \ell=1\dots 4}} \mathcal{H}_\ell^i \quad (2)$$

The vertices and the faces of $\mathcal{D}$ can be enumerated in $O(n \log n)$ time, being n the number of cameras [17].

As an example, Figure 3 shows the polyhedral triangulation result obtained from nine calibrated images of a toy object. Twenty-six points are manually matched in the sequence and a uniform error in the 2D point location bounded by a 10-pixel wide square is assumed. The mean volume of the polyhedra is $(5.8mm)^3$, with respect to a volume of $(2.6dm)^3$ of the object.

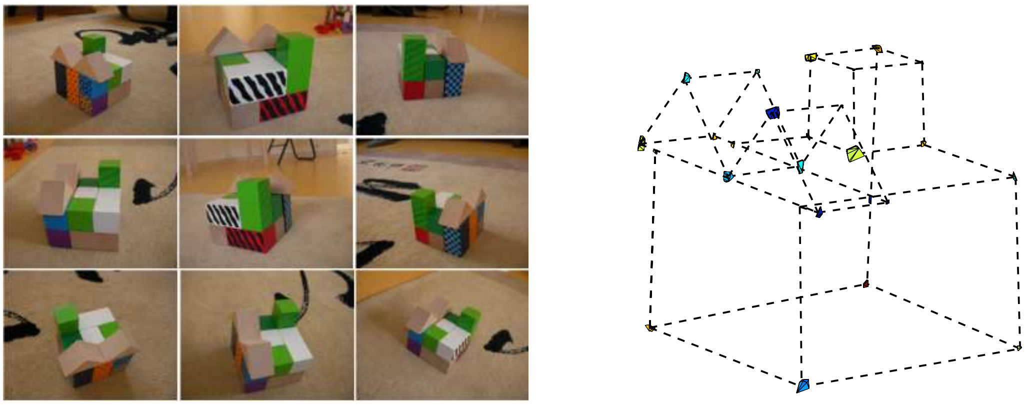

Fig. 3. Thumbnails of the nine images of the *Toy* sequence (left) and the polyhedral triangulation result (right). The small polyhedra bound the corners of the object. Dashed lines are added for visualization purpose.

4 Constraints Detection

From the polyhedral triangulation a bounded estimation of the position of reconstructed 3D points is obtained. Any random choice of 3D points inside the bound is an approximation of the exact 3D reconstruction. Considering one of these approximations, its points are connected (manually) into a triangular mesh, obtaining a piecewise planar surface model. This section describes how geometric constraints such as coplanarity, parallelism, orthogonality and angles equality are automatically detected on the approximate model.

4.1 Planes Detection

Planes in the model are extracted using a Mean Shift clustering procedure [18] on the triangular facets. The proposed technique is composed by a two-step, hierarchical strategy.

(1) First facets are clustered according to their normal, thereby grouping together (approximately) coplanar and (approximately) parallel facets.
(2) Then, within each group, the clustering is refined by taking into account also the distance to the origin of the plane containing the facet. In this way facets belonging to parallel planes are separated.

We adopted in both cases the uniform kernel, i.e., a multidimensional unit sphere, with bandwidth automatically selected as described in the following.

4.1.1 *Bandwidth Automatic Estimation*

The bandwidth parameter h defines the level of detail of the clustering analysis. Large bandwidth values lead to global but coarse separation, whereas small bandwidth values better identify local modes, but at the risk of over-partitioning the data space. Good segmentation results could be obtained after an accurate parameters tuning. In line with the concept of stable segmentation [19] we use the strategy developed in [20]. We single out extreme values for h, we uniformly sample the range $[h_{min}, h_{max}]$ and perform a Mean Shift clustering for each value of h. Finally, we consider the plot of the number of clusters obtained versus h and we choose the centre of the largest plateau as the optimal bandwidth (see Figure 4).

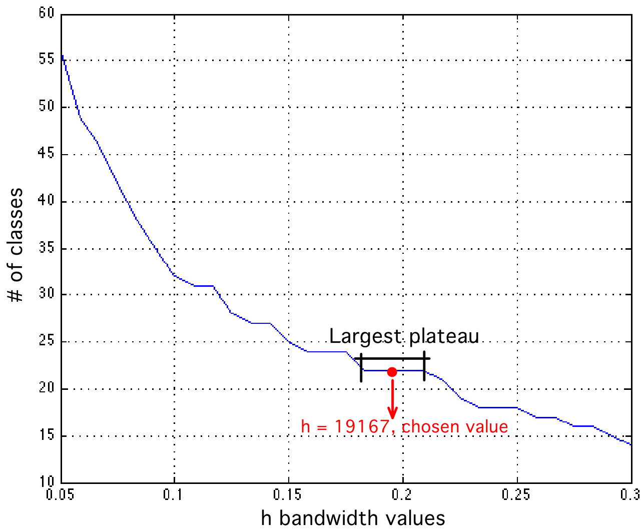

Fig. 4. Example of automatic bandwidth selection: on the graph of the number of clusters obtained from each trial with a h value, from the interval $[0.05, 0.3]$, the centre of the largest plateau is selected.

Please note that the process clusters together facets belonging to the same plane, regardless of their distance.

4.2 *Constraints Extraction*

Geometric constraints involving planes can now be automatically inferred.

Facets belonging to the same group after the second clustering step are related by coplanarity constraints. Each plane is identified by one reference facet. As to parallelism constraints, if two different reference facets belong to the same group after the first clustering step, their respective planes are parallel. Finally, angular constraints are deduced from *grouping heuristics*: if two or more planes nearly satisfy a constraint then they are forced to satisfy it. Orthogonality is checked for every pair of reference facets: whenever two of them are found to be approximately orthogonal (within 5 degrees), then they are linked by an orthogonality constraint. Likewise, equality of angles is checked for every quartet of reference facets.

The constraints form a hierarchy (Figure 5): at the bottom level there are facets, grouped into planes by coplanarity constraints, then planes, grouped into equivalence classes modulo parallelism, and, finally, these equivalence classes related by angular constraints.

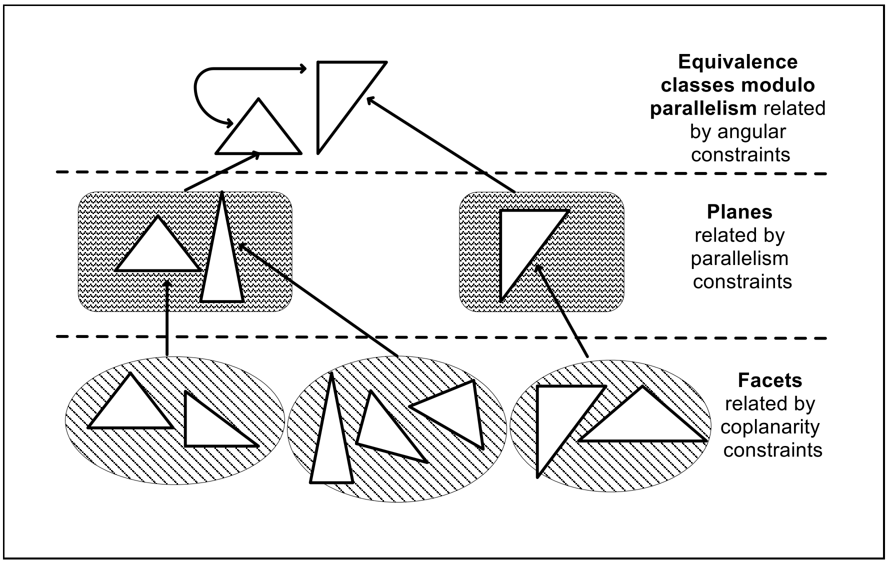

Fig. 5. The hierarchy induced by constraints detection.

At the highest level the position of the planes does not matter, as only the orientation is considered. This is consistent with the fact that 3D point's positions are not determined by the polyhedral triangulation.

Carrying on with the *Toy* model, the 26 polyhedra are manually connected into triangular facets. Then, 13 planes are automatically correctly extracted by the algorithm, as depicted in Figure 6. This clustering process implies 8 parallelism constraints after the first clustering step (middle level of the hierarchy) and 29 coplanarity constraints on the triangular facets after the second step (bottom level of the hierarchy). At the higher level of the constraints hierarchy five equivalence classes modulo parallelism are found. Automatic constraints detection identifies 6 orthogonality and 3 angle equality constraints; labelling the reference planes as in Figure 6, these constraints are the following:

Perpendicularity: 1-2, 1-3, 1-4, 1-5, 2-5, 3-4;
Angle Equality: (5 4)-(5 3), (5 4)-(2 3), (5 3)-(2 3).

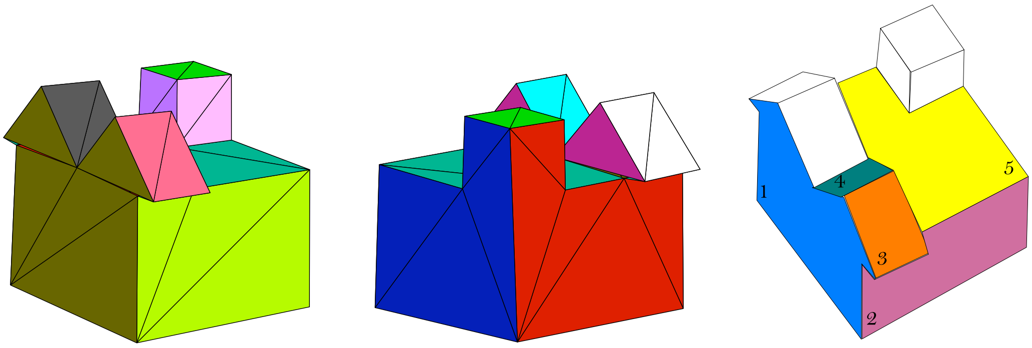

Fig. 6. Automatic extraction of planes for the *Toy* model (left and centre image). Each plane is identified by a different colour. The right image shows the five planes representative of the equivalence classes modulo parallellism.

5 Constraints Analysis

In this section we will discuss how the angular constraints, which – in general – are redundant, can be pruned while maintaining their capacity of stabilizing the estimation of the 3D model.

The concept of rigidity (or *constriction*) for geometric systems, has been studied in several scientific fields like Computational Geometry and Structural Topology, with application mainly in Computer-Aided Design (CAD). We are applying here the notion of *structural rigidity* to systems of planes (modulo parallelism) in order to remove redundant constraints while keeping the "rigidity" of the system. Some definitions, taken from [21], are in order here to introduce notation and concepts.

Definition 1 (Geometric Constraint System) *A Geometric Constraint System (GCS) is a pair $S = (O, C)$, where O is a set of geometric objects (represented by some variables), and C is a set of constraints.*

Our geometric objects are equivalence classes of planes modulo parallelism. They are identified by their direction (the normal vector). The constraints are orthogonality and angle equality.

Definition 2 *Let $S = (O, C)$ be a GCS. A solution to S is an evaluation θ_O of the variables in O such that every predicate in C is true. The set of solutions to S is denoted by $Sol(S)$.*

Definition 3 (Constriction) [1] *A GCS S is well-constrained if $Sol(S)$ is finite, over-constrained if $Sol(S) = \emptyset$ and under-constrained if $Sol(S)$ is infinite.*

[1] In fact, this is the definition of *global* [22] or *generic* [23] constriction.

In practice, a GCS can be under-constrained, but its solutions be identical modulo a geometric transformation (e.g., translation, rotation). Constriction modulo direct isometries (also called *rigidity*) is the type of constriction usually sought in CAD. In our case, translations are factorized out by the parallelism equivalence, hence only rotations are left. As a consequence, we consider constriction modulo orthogonal transformations.

Constriction depends on the number of solutions, but computing all of them is intractable. Hence, approximate characterizations that can be checked in polynomial time are frequently used. A characterization known as *structural constriction* is based on the *degrees of freedom* abstraction of the geometric constraints and objects.

Definition 4 *The number of degrees of freedom (DOFs) of a geometric object is the number of independent parameters used to represent it. The number of DOFs of a geometric constraint is the number of independent equations needed to represent it.*

In the following, we denote by $\mathrm{dof}(\cdot)$ the number of DOFs of an object or a constraint.

In our case, geometric objects have 2 DOF, because normals are unit vectors, and angle constraints have 1 DOF.

Definition 5 (Structural G-constriction) *Let $S = (O, C)$ be a GCS. Let G be an invariance group of dimension $\mathcal{D}$. The system S is structurally G-over-constrained if there exists a subsystem $S' = (O', C')$ of S such that $\sum_{x \in O'} \mathrm{dof}(x) - \sum_{c \in C'} \mathrm{dof}(c) < \mathcal{D}$.*

The system S is structurally G-well-constrained if it is not structurally G-over-constrained and $\sum_{x \in O} \mathrm{dof}(x) - \sum_{c \in C} \mathrm{dof}(c) = \mathcal{D}$.

The system S is structurally G-under-constrained if it is not structurally G-over-constrained and $\sum_{x \in O} \mathrm{dof}(x) - \sum_{c \in C'} \mathrm{dof}(c) > \mathcal{D}$.

In our case, structural constriction modulo orthogonal transformations can be checked using $\mathcal{D} = 3$.

Definition 6 (Constraint graph) *Let $S = (O, C)$ be a GCS. Its constraint graph, denoted by $G_S = (V, E)$, is a bipartite undirected graph where $V = O \cup C$ (every object in S and every constraint in C is a vertex in G_S) and an edge connects each constraint to each entity it constrains.*

Hoffmann et al. in [23] introduced the Dense algorithm that checks structural constriction in polynomial time, considering a flow-network derived from the bipartite constraint graph. The source is linked to each constraint, and each object is linked to the sink. The capacities correspond to the DOFs of the constraints (edges from the source to constraints) and to the DOFs of the objects (edges from objects to the sink). Edges from constraints to objects have infinite capacity. A maximum flow in this network represents an optimal distribution of the constraints DOFs onto the objects DOFs. To identify over-rigid subsystems, the method adds an additional $\mathcal{D}$ capacity to one constraint at a time.

If a maximum flow distribution cannot saturate all the edges from the sink to the constraints, this means that some constraints DOFs cannot be absorbed by the objects. Thus there exists a subsystem with less DOFs than $\mathcal{D}$, and the GCS is over-constrained (Figure 7). Please note that, being structural constriction an abstraction, a GCS is deemed over-constrained as soon a redundant constraints are detected, regardless of the fact that they are consistent or not.

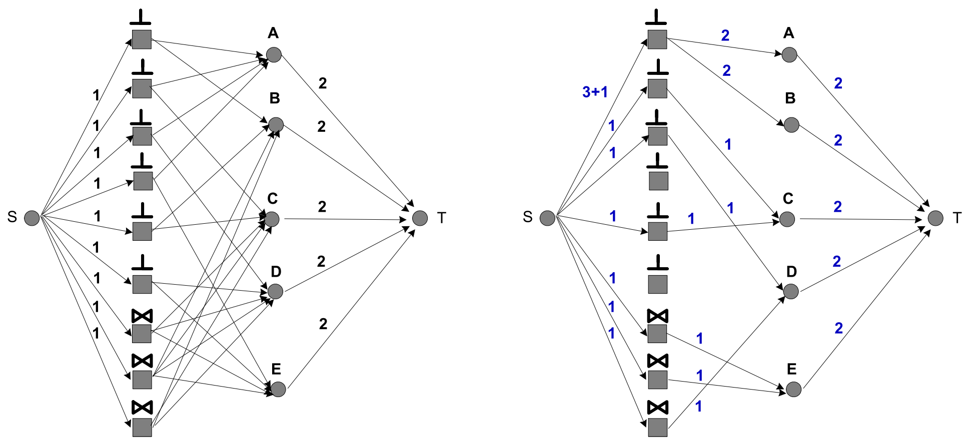

Fig. 7. Constraint network for the *Toy* example (left). The capacity of the arcs connecting the constraint and the object vertices is infinite. On the right a case of flow distribution on the network, when the first constraint edge is overloaded (right). The GCS is over-constrained, because not all the arcs from the sink to the constraints are saturated. $\perp$ indicates the orthogonality constraint and $\bowtie$ the equality of angles.

We exploit over-rigidity in order to detect redundancies. Dense returns the over-constrained subsystem S' if the system is over-constrained, or an empty set otherwise. This S' is the subsystem induced by the objects traversed during the last search for an augmenting path, in the max flow computation. Constraints binding the objects in S' can be removed until the system itself becomes structurally well-constrained. This is implemented in the Prune procedure:

Algorithm 1 Prune

Input $S = (O, C)$: GCS
Output $S_o = (O, C_o)$: GCS such that $C_o \subseteq C$ and S_o is structurally well-constrained

$S' = (O', C') \leftarrow$ Dense(S)
if isEmpty(S')
 Return S
else
 select $c \in C'$
 $S \leftarrow S(O, C \backslash \{c\})$
 Prune(S)
end

The selection of c is random, provided that its removal do not leave any object node with less inbounding arcs than the object's DOFs in the constraint graph.

The correctness of Prune is proven by the following:

Theorem 1 (Object Condition) *Necessary condition for a constraint graph $G_S = (V, E)$ to be structurally rigid is that the sum of DOFs of the constraints bounding each object vertex is greater or equal to the object's DOFs.*

Proof. Let $o_j \in V$ be an object vertex and $c_1, \ldots c_k \in V$ be the constraints vertices bounding o_j. Let $\mathrm{dof}(o_j) = n$, and $\mathrm{dof}(c_i) = m_i$, for $i = 1, \ldots k$. The condition $\sum_{i=1}^{k} \mathrm{dof}(c_i) < \mathrm{dof}(o_j)$ means that not all the parameters in o_j are constrained. Thus, the system can not be rigid.

In our case, since all constraints have one DOF it is sufficient to check the number of inbounding arcs.

Then, if one starts from an over-rigid constraint graph that satisfies the *Object Condition* of Theorem 1 and executes Prune on that graph such that after each constraint removal the *Object Condition* is verified, it is guaranteed to obtain as output a structurally rigid constraint system.

As shown in Figure 7 (left), the constraints analysis on the *Toy* model correctly reveals that the constraint system is over-rigid. The Prune algorithm on the constraint graph eliminates two constraints, one perpendicularity and one angle equality. The resulting constrained system has then 7 angular constraints and $\sum_{x \in O} \mathrm{dof}(o) - \sum_{c \in C} \mathrm{dof}(c) = 10 - 7 = 3 = \mathcal{D}$.

6 Finding a Feasible Solution

Finally, a feasible instance of the 3D model, i.e. one that satisfies all selected geometric constraints and whose 3D points lie within the associated polyhedral bounds, is computed. This is formalized in the following Constraint Satisfaction Problem (CSP):

$$\begin{aligned} \text{find} \quad & \mathrm{X} \\ \text{subject to} \quad & \mathrm{X}_L \leq \mathrm{X} \leq \mathrm{X}_U \\ & \mathrm{c}_L \leq c(X) \leq \mathrm{c}_U \end{aligned} \tag{3}$$

where X is the variables vector, i.e. the 3D points of the model; X_L and X_U delimit the domain of each variable, and they derive from polyhedral triangulation; $c(\mathrm{X})$ are algebraic equations containing the non linear constraints on X, with bounds c_L and c_U.

The geometric constraints must then be translated into constraints among points and formalized as algebraic equations.

For each equivalence class modulo parallelism a plane is chosen as the reference one. Angular constraints are applied among the reference planes. These are then linked to each other plane in the same equivalence class by a parallelism constraint. Each constraint among planes (i.e. parallelism and angular constraints) is translated into a constraint on the normals of the reference facets, as shown in Table 1. The normal vector, in turn, is a function of the three vertices of the facet, in Cartesian coordinates. In order to simplify the complexity of the algebraic equations, in the orthogonality and parallelism constraints the normal vector is not normalized.

Table 1
Translation of constraints among facets $\{f_i\}_{i=1,\dots,n}$ of the model into algebraic equations among points.

Constraint	*Algebraic equation*
Orthogonal(f_1, f_2)	$\mathrm{n}_1 \cdot \mathrm{n}_2 = 0$
SameAngle(f_1, f_2, f_3, f_4)	$(\mathrm{n}_1 \cdot \mathrm{n}_2) - (\mathrm{n}_3 \cdot \mathrm{n}_4) = 0$
Parallel(f_1, f_2)	$(\mathrm{n}_1 \times \mathrm{n}_2) = 0$

The reference facets are linked to all the other facets belonging to the same plane by coplanarity constraints. Let $C = \{\mathrm{M}_1, \mathrm{M}_2, \dots, \mathrm{M}_n\}$ be the vertices of a group of coplanar facets, then all these points must lie on the same plane. This can be translated into a set of overlapping coplanarity constraints among four points at a time:

$$\begin{aligned}&\text{Coplanar}(\mathrm{M}_1, \mathrm{M}_2, \mathrm{M}_3, \mathrm{M}_4) \wedge \\ &\text{Coplanar}(\mathrm{M}_2, \mathrm{M}_3, \mathrm{M}_4, \mathrm{M}_5) \wedge \ldots \\ &\text{Coplanar}(\mathrm{M}_{n-3}, \mathrm{M}_{n-2}, \mathrm{M}_{n-1}, \mathrm{M}_n)\end{aligned} \tag{4}$$

where $\text{Coplanar}(\mathrm{M}_1, \mathrm{M}_2, \mathrm{M}_3, \mathrm{M}_4)$ is equivalent to:

$$[(\mathrm{M}_1 - \mathrm{M}_2) \times (\mathrm{M}_1 - \mathrm{M}_3)]^{\mathsf{T}} \cdot (\mathrm{M}_1 - \mathrm{M}_4) = 0. \tag{5}$$

Once all constraints are translated into algebraic equations, a solution can be found using a constraint solver. In our case we use SNOPT [24], a general-purpose system for solving optimization problems involving many variables and constraints. It is suitable for large-scale linear and quadratic programming and for linearly constrained optimization, as well as for general nonlinear programs.

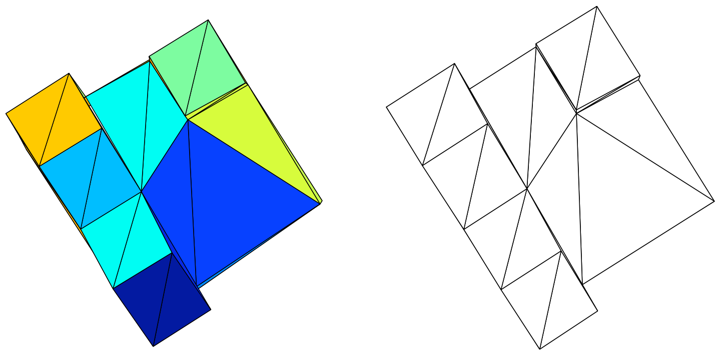

Fig. 8. A top view of the *Toy* model before (left) and after (right) constraints propagation.

Table 2
[Min,Max] deviation from the constraints in the *Toy* model before and after propagation (in degrees). Legend: $\perp$ is orthogonality, $\bowtie$ is angle equality, $\square$ is coplanarity and $\|$ is parallelism.

	$\perp$	$\bowtie$	$\square$	$\|$
Before	$[0.26, 2.47]^\circ$	$[0.002, 0.09]^\circ$	$[0.0, 0.02]^\circ$	$[0.83, 6.55]^\circ$
After	$[0.0, 0.48]^\circ$	$[0.0, 0.01]^\circ$	$[0, 0]^\circ$	$[0.12, 3.87]^\circ$

As to the *Toy* model, the formalization of the problem into algebraic equations yields 62 non linear constraints. The results are summarized in Table 2 and Figure 8. As the reader can notice, the whole pipeline described throughout the paper leads to an accurate 3D model. The fact that the constraints are not exactly satisfied is due to the optimizer, that stops when it deems the solution cannot improved further.

7 Experimental Results

Synthetic experiments. Constraints detection and analysis was tested on the synthetic models shown in Figure 9. The 3D points were replaced by boxes to simulate the output of polyhedral triangulation. The size of the box varied from 0.5% to 3.0% of a "size gauge" computed as the median over the model's points of the farthest point distance. For each value, 20 perturbed models were generated using a uniform random distribution inside the boxes. Planes and constraints were automatically detected by our algorithm, the constraints were cut down using Prune, and a feasible solution of the resulting CSP was found using SNOPT.

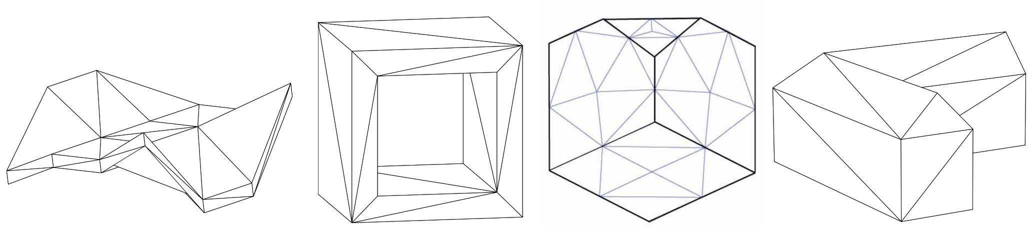

Fig. 9. The four objects used for the synthetic experiments, here referred as (from the left) *Test*, *Boxhole*, *Cutcube* and *House*.

The number of planes correctly extracted is 22 for *Test*, 10 for *Boxhole*, 7 for *Cutcube*, and 11 for *House*. Table 3 shows the constraints extracted, before and after the structural rigidity analysis. Please note that at each step Prune chooses randomly which constraint to eliminate, so the final set of constraints varies from time to time; here, the most common case is reported. The solution of the CSP took about 20 s for *Cutcube* and *Boxhole*, 50 s for *House* and 90 s for *Test*.

Real experiments. The whole method was tested on real images.

The *Palmanova* set is composed by 15 calibrated images of a monument (Figure 10). Polyhedral triangulation was carried out, considering a uniform error in the 2D point location bounded by a 10-pixel wide square (Figure 11). The mean volume of the polyhedra is $(13cm)^3$, with respect to a volume of $(5.2m)^3$. Then, we selected the centre of each polyhedra as a point-wise approximate solution and we connected them manually. Plane detection algorithm extracted 34 planes. Constraints detection and analysis were performed, and the results are outlined in Table 3. As the reader can notice, in this case constraints pruning is essential to simplify the problem. The CSP solver (SNOPT) produced, in a few minutes, the result shown in Figure 12, with the errors reported in Table 4.

Table 3
Number of constraints automatically detected and number of remaining constraints after the structural rigidity analysis for the 3D models. The rightmost column reports the total number final of constraints.

	Automatic constraints				Pruning		
	$\perp$	$\bowtie$	$\square$	$\parallel$	$\perp$	$\bowtie$	Total
Test	22	61	60	8	9	16	93
Boxhole	12	1	24	4	9	0	37
Cutcube	0	3	40	5	0	3	48
House	9	6	25	4	7	4	40
Palmanova	21	1460	134	4	4	53	195
Pozzoveggiani	59	2043	100	5	1	58	164
Tribuna	29	799	478	34	2	51	565

Fig. 10. Three of the 16 images of the *Palmanova* set.

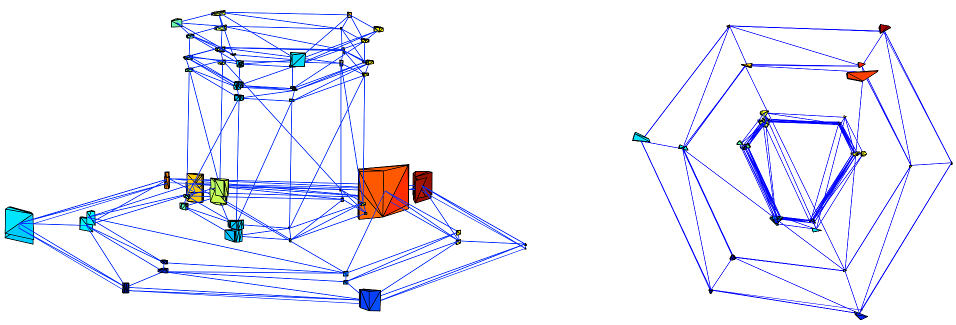

Fig. 11. Polyhedral triangulation for *Palmanova*.

The *Pozzoveggiani* set is composed by 16 calibrated images of a church (Figure 13). Polyhedral triangulation was performed, assuming a uniform error in the 2D point location bounded by a 7-pixel wide square (Figure 14). The mean volume of the resulting polyhedra is $(13cm)^3$, with respect to a volume of $(16.88m)^3$. Then, starting from the approximate solution obtained by randomly choosing one point inside each polyhedron, and connecting them

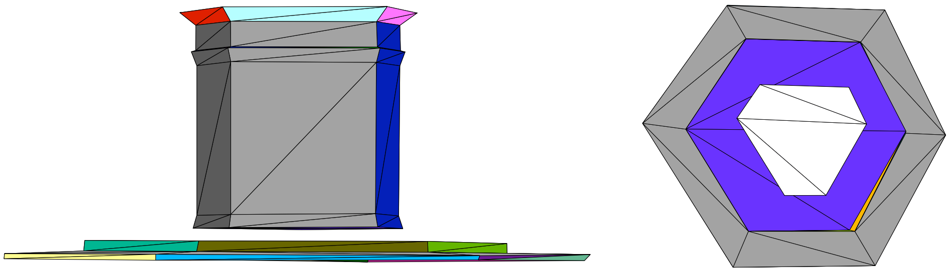

Fig. 12. Final geometric reconstruction of *Palmanova* after the constraints propagation. Each plane is identified by a different colour.

Table 4
[Min,Max] deviation from the constraints in the initial and final model (in degrees).

		$\perp$	$\bowtie$	$\square$	$\parallel$
Palmanova	Before	$[0.15, 3.78]^\circ$	$[0.0, 0.02]^\circ$	$[0.0, 0.0]^\circ$	$[1.45, 3.68]^\circ$
	After	$[0.0, 2.61]^\circ$	$[0.0, 0.02]^\circ$	$[0.0, 0.0]^\circ$	$[0.05, 1.25]^\circ$
Pozzoveggiani	Before	$[2.05, 2.05]^\circ$	$[0, 0.04]^\circ$	$[0, 0.19]^\circ$	$[1.2, 6.8]^\circ$
	After	$[1.73, 1.73]^\circ$	$[0, 0.03]^\circ$	$[0, 0.13]^\circ$	$[1.2, 4.7]^\circ$
Tribuna	Before	$[0.03, 3.3]^\circ$	$[0, 0.03]^\circ$	$[0, 0.15]^\circ$	$[0.9, 10.4]^\circ$
	After	$[0.03, 0.3]^\circ$	$[0, 0.17]^\circ$	$[0, 0]^\circ$	$[0.9, 7.9]^\circ$

manually, 36 planes were automatically extracted. Results from constraints detection and analysis are outlined in Table 3. The CSP solver (SNOPT) produced, after less than one minute, the result shown in Figure 15, with the errors reported in Table 4.

Fig. 13. Three of the 16 images of the *Pozzoveggiani* set.

The *Tribuna* set consists of 10 calibrated images of an apse (Figure 16). Polyhedral triangulation was performed, assuming a uniform error in the 2D point location bounded by a 2-pixel wide square (Figure 17). The mean volume of the resulting polyhedra is $(2cm)^3$, with respect to a volume of $(4.05m)^3$. Then, starting from the approximate solution obtained by randomly selecting one point inside each polyhedron, and connecting them manually, 62 planes were automatically extracted. Constraints detection and analysis results are

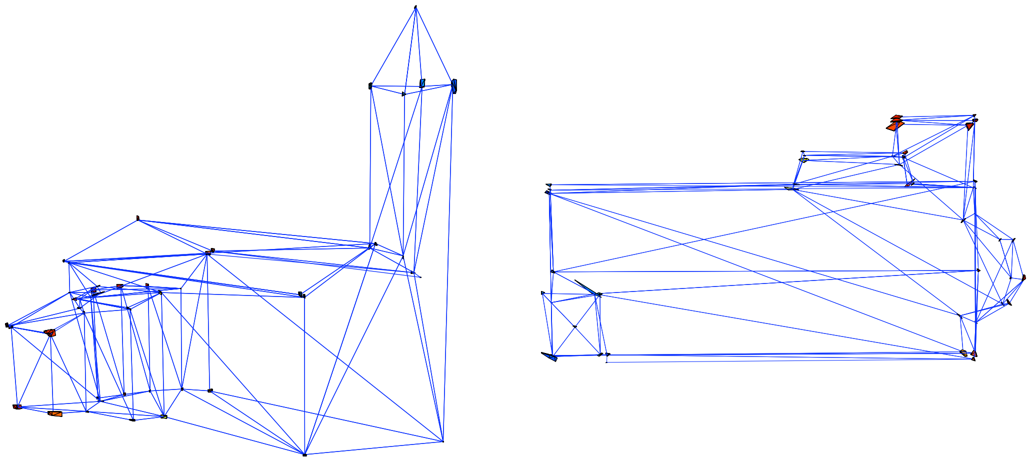

Fig. 14. Polyhedral triangulation for *Pozzoveggiani*.

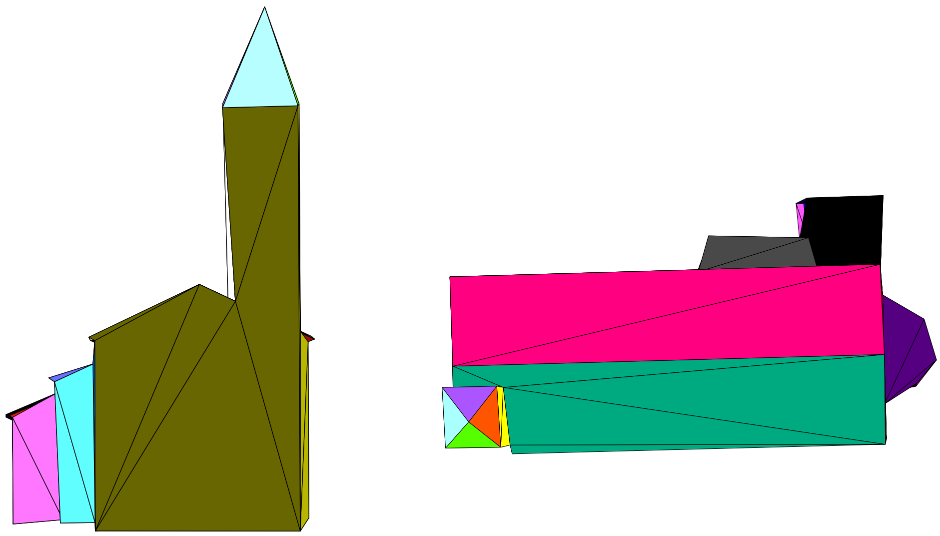

Fig. 15. Final geometric reconstruction of *Pozzoveggiani* after the constraints propagation. Each plane is identified by a different colour.

summarized in Table 3. The CSP solver (SNOPT) produced, after a few minutes, the result shown in Figure 18, with the errors reported in Table 4.

8 Conclusions

In this paper we presented a new approach to constrained modeling from many calibrated views. We demonstrated how polyhedral triangulation and a suitable constraint analysis and propagation can be used to obtain an accurate geometric model of a scene.

Fig. 16. Three of the 10 images of the *Tribuna* set.

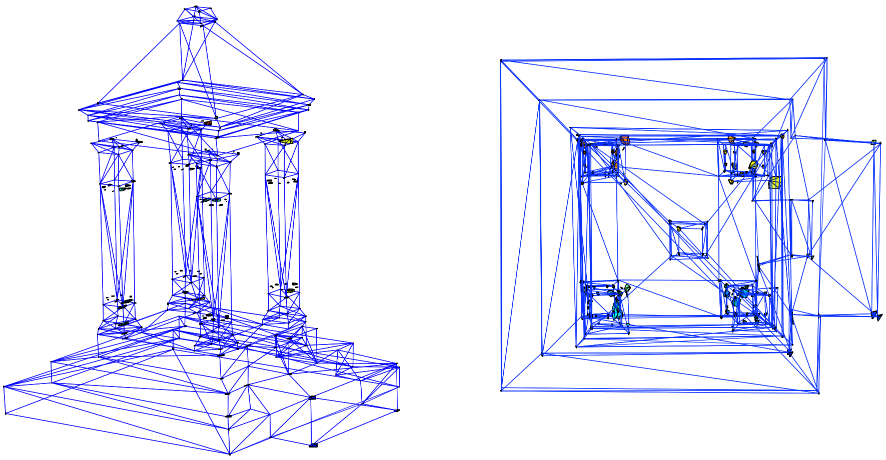

Fig. 17. Polyhedral triangulation for *Trubuna*.

The final model satisfy the geometric constraints and its reprojection onto the images is guaranteed to fall within the bounds set on the points localization error.

Experiments show the effectiveness and the accuracy of the approach.

Future work will aim at removing the need for manually entering points and connectivity, thereby making the system fully automatic. Preliminary results in this direction are reported in [25].

Acknowledgments

We are grateful to P. Sturm for valuable comments and to L. Corazza for contributing to the implementation of the software. Some of the 3D models

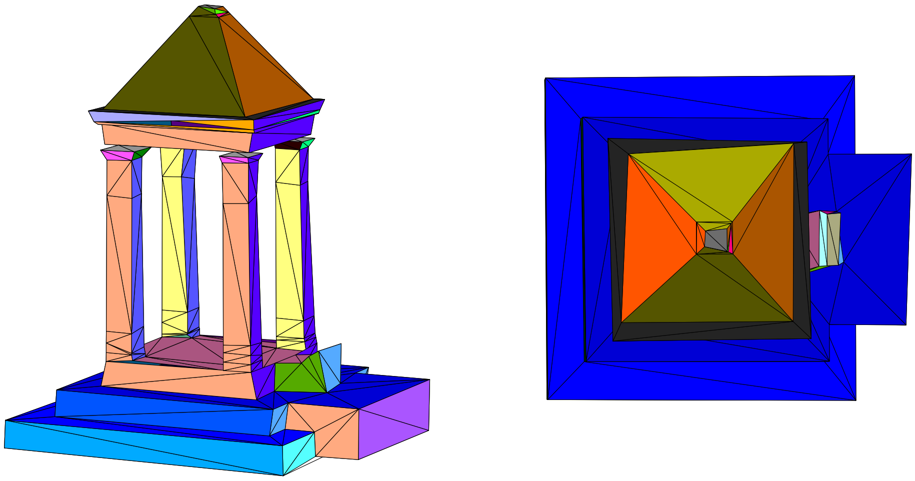

Fig. 18. Final geometric reconstruction of *Tribuna* after the constraints propagation. Each plane is identified by a different colour.

shown in Figure 4 are courtesy of Mountaz Hascoët[2]. The SNOPT solver is available inside TOMLAB.

References

[1] P. E. Debevec, C. J. Taylor, J. Malik, Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach, in: H. Rushmeier (Ed.), SIGGRAPH: International Conference on Computer Graphics and Interactive Techniques, New Orleans, Louisiana, 1996, pp. 11–20.

[2] D. Jelinek, C. J. Taylor, Reconstruction of linearly parametrized models from single images with camera of known focal length, IEEE Transactions on Pattern Analysis and Machine Intelligence 23 (7) (2001) 767–774.

[3] D. Lowe, Fitting parameterized three-dimensional models to images, IEEE Transactions on Pattern Analysis and Machine Intelligence 13 (5) (1991) 441–450.

[4] M. Wilczkowiak, P. Sturm, E. Boyer, Using geometric constraints through parallelepipeds for calibration and 3D modeling, IEEE Transactions on Pattern Analysis and Machine Intelligence 27 (2) (2005) 194–207.

[5] M. I. A. Lourakis, A. A. Argiros, Enforcing scene constraints in single view reconstruction, in: Proceedings of EuroGraphics 2007, Prague, Czech Republic, 2007.

[6] A. Bartoli, P. Sturm, Constrained structure and motion from multiple uncalibrated views of a piecewise planar scene, International Journal of Computer Vision 52 (1) (2003) 45–64.

[2] http://www.lirmm.fr/ mountaz/Ens/DessTni/OpenGL/Exemples/tutors/data/

[7] H. Cantzler, R. B. Fisher, M. Devy, Improving architectural 3D reconstruction by plane and edge constraining, in: British Machine Vision Conference, Cardiff (UK), 2002, pp. 43–52.

[8] A. R. Dick, P. H. S. Torr, S. J. Ruffle, R. Cipolla, Combining single view recognition and multiple view stereo for architectural scenes, in: Proceedings of the International Conference on Computer Vision, Vol. 1, 2001, p. 268.

[9] M. Farenzena, A. Fusiello, A. Dovier, Reconstruction with interval constraints propagation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2006, pp. 1185–1190.

[10] E. Grossman, J. Santos-Victor, Least-square 3D reconstruction from one or more views and geometric clues, Computer Vision and Image Understanding 99 (2005) 151–174.

[11] M. Wilczkowiak, G. Trombettoni, C. Jermann, P. Sturm, E. Boyer, Scene modelling based on constraint system decomposition techniques, in: Proceedings of the International Conference on Computer Vision, Vol. II, IEEE, IEEE, 2003, pp. 1004–1010.

[12] S. Cornou, M. Dhome, P. Sayd, Architectural reconstruction with multiple views and geometric constraints.

[13] M. Farenzena, A. Fusiello, 3d surface models by geometric constraints propagation, to appear in CVPR 2008.

[14] F. Kahl, Multiple view geometry and the l_∞-norm, in: Proceedings of the IEEE International Conference on Computer Vision, Beijing, China, 2005, pp. 510–517.

[15] H. Stewenius, F. Schaffalitzky, D. Nister, How hard is 3-view triagulation really?, in: Proceedings of the International Conference on Computer Vision, Beijing, China, 2005, pp. 510–517.

[16] R. I. Hartley, P. Sturm, Triangulation, Computer Vision and Image Understanding 68 (2) (1997) 146–157.

[17] F. P. Preparata, M. I. Shamos, Computational Geometry. An Introduction, 1st Edition, Springer-Verlag, 1985, Ch. 2, pp. 72–77.

[18] D. Comaniciu, P. Meer, Mean shift: A robust approach toward feature space analysis, IEEE Transactions on Pattern Analysis and Machine Intelligence 24 (5) (2002) 603–619.

[19] K. Fukunaga, Statistical Pattern Recognition, Academic Press, second edition, 1990.

[20] M. Cristani, U. Castellani, V. Murino, Adaptive feature integration for segmentation of 3D data by unsupervised density estimation, in: Proceedings of the International Conference on Pattern Recognition, Vol. 4, 2006, pp. 21–24.

[21] C. Jermann, G. Trombettoni, B. Neveu, P. Mathis, Decomposition of geometric constraint systems: a survey, Internation Journal of Computational Geometry and Applications 16 (5-6) (2006) 379–414.

[22] B. N. C. Jermann, G. Trombettoni, A new structural rigidity for geometric constraint systems, in: Fifth International Workshop on Automated Deduction in Geometry, Linz (Hagenberg), 2002.

[23] C. M. Hoffmann, A. Lomonosov, M. Sitharam, Finding solvable subsets of constraint graphs., in: Constraint Programming, 1997, pp. 463–477.

[24] W. M. P. E. Gill, M. A. Sauders, SNOPT: An SQP algorithm for large-scale constrained optimization, SIAM J. Optimization 12 (2002) 979–1006.

[25] M. Farenzena, A. Fusiello, R. Gherardi, R. Toldo, Towards unsupervised reconstruction of architectural models, submitted to SGP 2008.