A PROBABILISTIC APPROACH TO ROOF EXTRACTION AND RECONSTRUCTION

S. Scholze[a,*], **T. Moons**[b,c], **L. Van Gool**[a,b]

[a]ETH Zurich, Computer Vision Laboratory (BIWI), Zurich, Switzerland
[b]Katholieke Universiteit Leuven, Center for Processing Speech and Images (ESAT / PSI), Leuven, Belgium
[c]Katholieke Universiteit Brussel, Group of Exact Sciences, Brussel, Belgium

Working Group III/7

KEY WORDS: Modelling, Reconstruction, Building, Three-dimensional, Aerial, Vision

ABSTRACT

This paper investigates into the model-based reconstruction of complex polyhedral building roofs. A set of 3D line segments, obtained from multiview correspondence analysis of high resolution colour imagery, is used as input data. The 3D line segments are grouped into planes by means of a Bayesian model selection procedure. In the resulting planes, models for polygonal patches are then instantiated. Driven by the Maximum Expected Utility principle, the algorithm chooses the optimal patch and plane configuration non-deterministically. Roof reconstruction is completed by further reasoning steps which are guided by the semantic interpretation of the intermediate patch configuration. Several successfully reconstructed complex roof structures corroborate the potential of the approach.

1 INTRODUCTION

For the automated reconstruction of building roofs from high resolution aerial images, a roof model consisting of planar polygonal patches has to be determined. As input data for this process, 3D line segments from a previous matching step are used. Due to the nature of the 3D scene under consideration, most of the 3D line segments represent boundaries of planar polygonal roof patches. However, one can not expect that all important image features are detected by the preprocessing steps and therefore not the complete patch outline is available in 3D. An aggravating fact is that the set of 3D line segments is contaminated by spurious matches, not corresponding to true image features. Consequently, to establish roof hypotheses, inaccurate and imprecise input data has to be handled. Nevertheless, the class of typical roof shapes is sufficiently constrained as to apply strong prior knowledge. This knowledge will be exploited implicitly by the proposed roof model and explicitly by deriving parameter distributions from a test dataset.

1.1 Previous work

Previous work on grouping straight line segments into planes in 3D space and on delineating patches lying in these planes mostly appears in the context of building reconstruction from aerial images. Some models are already quite advanced (Henricsson, 1998, Moons et al., 1998, Baillard and Zisserman, 2000), but many of these systems have no probabilistic underpinning for handling uncertainty.

(Kulschewski, 1997) studies the recognition of buildings from a single view using a dynamic Bayesian network. The Bayesian network approach allows the author to handle uncertainties in the input data, regarding accuracy and completeness. Extracted roof outlines are used to reliably classify the building type, yet modelling entire building types imposes a limitation to the system. Another approach (Heuel et al., 2000, Heuel and Förstner, 2001) is to guide the process of 3D grouping with topological and geometrical reasoning. Hypotheses for polyhedral surfaces are selected using topological relations and verified using geometry. The uncertainty of geometric relations (e.g. intersections) is modelled statistically. Since no patch models are imposed, gaps in the patch delineation have to be filled in a heuristic manner. (Baillard and Zisserman, 2000) propose a plane sweep strategy which is able to instantiate plane hypotheses from only one reconstructed 3D line segment. Although a probabilistic interpretation is mentioned, it is not exploited in the presented work. Missing segments in the outline of the piecewise planar models are mainly determined by back-projection in the input images and by heuristic completion rules. A Bayesian approach for generating plane hypotheses from dense point correspondences was proposed by (Cord et al., 1999). Since no patch model is imposed, the delineation of the patch is solved by an hybrid method using 3D segmentation results and grey-level intensity information of the input images.

Most recent developments in the field of automated building reconstruction are model based. In (Kim et al., 2001) a feature-based system is presented, which deals with flat, and to some extent, with complex roof types. Roof hypothesis verification is done in a probabilistic framework. Another model-based system is proposed in (Fuchs, 2001), which combines different types of features using graph theory. In this framework, roof reconstruction is formulated as graph matching problem.

1.2 A generic model in a probabilistic setup

In this paper, a probabilistic formulation for the reconstruction of building roofs from sparse sets of incomplete and imprecise 3D line segments is presented. In a first step, all possible plane hypotheses – instantiated using at least one line segment – are created, using a Bayesian model selection procedure. This leads to a hierarchical interpretation of the scene. For each level in the hierarchy – corresponding to specific plane hypotheses – hypotheses for planar

polygonal faces (*patches*) are created. This is achieved by fitting a generic patch model to the data. After computing the Expected Utility for all patch and plane configurations, the system chooses its interpretation of the scene non-deterministically following the Maximum Expected Utility Principle. The semantic interpretation of the line segments in the resulting roof guides a search for further patches until a consistent roof model is obtained.

The organisation of this paper is as follows. Section 2 briefly summarizes the processing steps from the image data up to the reconstructed 3D line segments. In Section 3, the application of Bayes model selection to the problem of determining the number of planes in a scene is discussed. In Section 4 a parametric patch model is proposed together with a formalism to describe entire roofs composed out of patches. The instantiation and evaluation procedure for patches is presented in Section 5, and finally the complete roof reconstruction is given in Section 6. Results of the approach are presented in Section 7.

2 INPUT DATA

In this section we briefly summarize the steps leading to the 3D line segments which are used for roof reconstruction. A detailed description of the individual steps can be found in the references given below.

2.1 3D line segment reconstruction

The raw data consist of colour images of densely built up urban areas. After edge detection and straight line fitting, corresponding 2D line segments are sought in the stereo pairs. Besides the geometric constraints given by epipolar geometry and the specific image acquisition setup, chromatic constraints are also exploited to disambiguate the matching and thus reduce the number of wrong matches (Scholze et al., 2000). The chromatic constraints are obtained by comparing the image neighbourhood of the line segments under consideration in order to reject chromatically incompatible matches. The number of remaining matches is further reduced by considering the trifocal constraint and the chromatic constraints already mentioned. If a match passes all statistical and geometrical tests, its 3D position is recalculated using bundle adjustment to enhance the geometric precision of the reconstructed 3D line segment.

2.2 Initial semantic interpretation

Given a set of 3D line segments, these segments have to be grouped into planes before the actual patch outline can be determined. Only the most reliable line segments, obtained from three view matching in conjunction with the chromatic constraints are used at the moment. In the next section a method for plane instantiation with respect to a seed line segment is introduced. In order to invoke the appropriate instantiation procedure, the function of the seed segment in the roof should be known in advance. Hence, the segments have to be attributed with semantic labels (e.g. ridge, gutter, gable). Of course, at this stage it is impossible to unambiguously assign a label to an arbitrary 3D line segment. Nevertheless, class membership probabilities can be derived. The ground-truth of a test dataset (Institute of Geodesy and Photogrammetry, 2001) was used to learn the semantic labels. The dataset shows urban and sub-urban areas with four-fold overlap at an image scale of approximately 1:5000. Using some measured geometric attributes (e.g. height, slant), the probability distribution for the respective labels is learnt. For each class the algorithm creates a list of seed line segments sorted by probability.

3 PLANE HYPOTHESES

A key problem when fitting a model to data consists in the appropriate choice of the complexity of the model. On one side, the more complex model (that is, a model with more parameters) will usually fit the data more closely. On the other hand, the more complex model is likely to over-fit the data. In this section we are concerned with the question, how many planes are present in a subset of a 3D scene. For the sake of simplicity, the discussion will be tailored to the special case depicted in Figure 1, where a half-plane is rotated around a tentative ridge-line. However, it should be pointed out, that the framework is totally generic and is applied successfully to arbitrary reference lines.

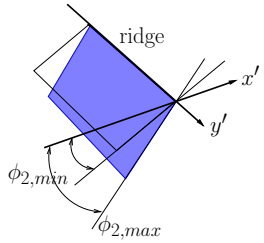

Figure 1: Illustration of rotating a half-plane around a reference line for plane instantiation. ϕ_2 is the slant angle of the plane, measured relative to a horizontal plane.

3.1 Probabilistic formulation

In the following, the term model refers to one or more plane hypotheses and will be denoted with m. A plane is given by a reference line and a slant angle (ϕ_2 in Figure 1). Models of different complexity correspond to a different number of planes. For a fixed reference line, the model is completely defined by the slant angle parameter(s) ϕ_2. Individual 3D line segments are denoted as x_i and the entire set of N segments with $\mathbf{x} = (x_1, \ldots, x_N)$.

The plane model which best represents the data $\mathbf{x}$ is the one which maximizes the posterior model probability

$$p(m|\mathbf{x}, e) \rightarrow \max \tag{1}$$

The background evidence e includes all assumptions about the problem domain which are not explicitly present in the formulas, e.g. that roofs consist of planar patches. By means of Bayes' theorem, this probability can be expressed by the likelihood and the prior of the model

$$p(m|\mathbf{x}, e) = \frac{p(\mathbf{x}|m, e)p(m|e)}{p(\mathbf{x}|e)} \tag{2}$$

For the following discussion the prior model probability $p(m|e)$ can be neglected, hence the probability of a model is only affected by the likelihood of the data $\mathbf{x}$, given the plane model m

$$p(m|\mathbf{x}, e) \propto p(\mathbf{x}|m, e) \tag{3}$$

A model with more adjustable parameters can usually follow the data more precisely and thus will lead to a higher likelihood. But in cases where the data is subject to noise, the most appropriate description of the underlying distribution might not be obtained by modelling each data point as precisely as possible. This problem is addressed next.

3.2 Likelihood, Marginalization and Ockham's razor

An estimate of the model likelihood – without taking into account model complexity – can be obtained by a Maximum Likelihood calculation. If the model parameters would be known, the likelihood would be given by

$$p(x|m) \propto \exp(-\frac{1}{2}\chi^2) \tag{4}$$

there χ^2 is the sum of squared distances between the data points and the corresponding model predictions. For the case of the plane rotation algorithm, one has to deal with a one-dimensional problem. For fixed plane parameter ϕ_2, the distances between the 3D line segments and the plane can easily be determined by computing the orthogonal distance.

Assuming the measurements x_i are sampled from a Normal distribution with parameters $\theta = \{\mu(\phi_2), \sigma\}$

$$p(x_i|\theta, m, e) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{\mu(\phi_2) - x_i}{\sigma}\right)^2\right) \tag{5}$$

and assuming further, that the individual line segments are independent of each other, the *likelihood* of the entire dataset $\mathbf{x}$ is given by

$$p(\mathbf{x}|\theta, m, e) = \prod_{i=1}^{N} p(x_i|\theta, m, e) \tag{6}$$

$$= \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp\left(-\frac{1}{2}\sum_{i=1}^{N}\left(\frac{\mu(\phi_2) - x_i}{\sigma}\right)^2\right)$$

with N being the number of line segments. If needed, the normalization constant $p(\mathbf{x}|e)$ can be computed via the requirement

$$p(\mathbf{x}|e) = \sum_m p(\mathbf{x}, m|e) = \sum_m p(\mathbf{x}|m, e)p(m|e) \tag{7}$$

since the possible models form a set of mutually exclusive and exhaustive events. To deal with the unknown model parameters $p(\mathbf{x}|m, e)$ can be expressed as marginal density computed over θ (which corresponds to an integration over ϕ_2 in the given case)

$$\begin{aligned} p(\mathbf{x}|m, e) &= \int p(\mathbf{x}, \theta|m, e)d\theta \\ &= \int p(\mathbf{x}|\theta, m, e)p(\theta|m, e)d\theta \end{aligned} \tag{8}$$

To impose no prior knowledge about the probability distribution of parameter θ, the Maximum Entropy argument suggests a uniform density for $p(\theta|m, e)$

$$p(\theta|m, e) = \frac{1}{\mu_{max} - \mu_{min}} \tag{9}$$

As will turn out later, good estimates for μ_{max} and μ_{min} are available. Collecting the terms, following expression for the posterior probability for the one plane model is obtained

$$\begin{aligned} p(m|\mathbf{x}, e) \propto \frac{1}{\mu_{max} - \mu_{min}} \times \\ \int_{\theta_{min}}^{\theta_{max}} \exp\left(-\frac{1}{2}\sum_{i=1}^{N}\left(\frac{\mu - x_i}{\sigma}\right)^2\right)d\theta \end{aligned} \tag{10}$$

Now consider the case of a M plane model. Here one has $\theta_j = \{\mu_j(\phi_{2,j}), \sigma_j\}$, $j = 1, \ldots, M$. The probability of measurement $\mathbf{x}$ can be modelled as mixture of M Gaussians

$$p(\mathbf{x}|\mathbf{m}, e) = \sum_{j=1}^{M} p(\mathbf{x}|m_j, e)P(m_j|e) \tag{11}$$

Each component of the mixture corresponds to a plane. The *mixing parameters* $P(m_j|e)$ denote the probability of a segment, being generated from component j. The conditional densities $p(\mathbf{x}|m_j, e)$ are called *component densities* or *class–conditional densities*. In order to obtain the posterior probability of the mixture model $\mathbf{m} = \{m_1, \ldots, m_M\}$ one follows the steps described above to finally obtain

$$\begin{aligned} p(\mathbf{m}|\mathbf{x}, e) \propto \frac{1}{M(\mu_{max} - \mu_{min})^M} \\ \int_{\theta_{min}}^{\theta_{max}} \exp\left(-\frac{1}{2}\sum_{j=1}^{M}\sum_{i=1}^{N}\left(\frac{\mu_j - x_i}{\sigma_j}\right)^2 z_{ij}\right)d\theta \end{aligned} \tag{12}$$

with z_{ij} being a *class indicator variable* defined as

$$z_{ij} = \begin{cases} 1 & : \text{ if } x_i \text{ was generated by component } j \\ 0 & : \text{ else} \end{cases} \tag{13}$$

The details of the calculation can be found in (Scholze, 2002). The fraction in front of the integral is commonly referred to as *Ockham* factor. It penalizes more complex models by decreasing the posterior model probability.

3.3 Hierarchical interpretation

In this section, the probabilistic plane model selection framework is turned into an algorithm. The problem will be formulated as divide-and-conquer algorithm. In a first step, all 3D segments in the neighbourhood of a seed line segment are collected. Now the goal is to determine how many planes are needed to describe these segments. Using Equation 12, the posterior model probabilities for a one-plane model $p(\mathbf{m}^{(1)}|\mathbf{x}, e)$ and for a two-plane model $p(\mathbf{m}^{(2)}|\mathbf{x}, e)$ are determined. If the probability for the one-plane model is higher than for the two-plane model the algorithm stops. In the opposite case (there more than one

plane is to be expected) the set of line segments is split into two subsets $\mathbf{x}'_1, \mathbf{x}'_2$, according to the indicator variables (Equation 13) of the two-plane model. The above procedure is repeated recursively for the individual subsets until the condition

$$p(m^{(2)}|\mathbf{x}') < p(m^{(1)}|\mathbf{x}') \tag{14}$$

is reached for all subsets. Formally, the system constructs a hierarchical scene interpretation which can be represented as a tree as depicted in Figure 2. Due to a lack of space, the publication if a significant empirical validation of this findings has to be postponed, however it is currently in preparation.

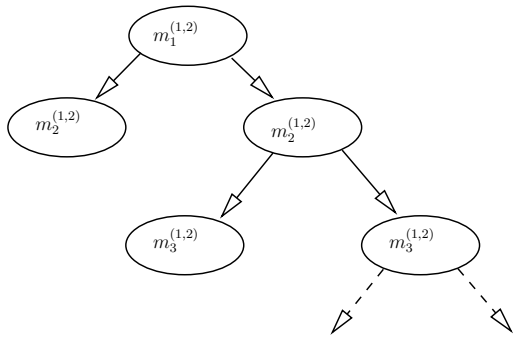

Figure 2: Illustration of Divide and Conquer algorithm for plane hypothesis creation. Each node represents a subset of line segments and the two model hypotheses $m^{(1)}$ and $m^{(2)}$. (Only the root node corresponds to the entire set of neighbouring segments.) The lower indices denote the level of recursion. For a leave node at level i the relation $p(m_i^{(2)}|\mathbf{x}'_i) < p(m_i^{(1)}|\mathbf{x}'_i)$ holds.

4 ROOF MODEL

A roof is modelled as a collection of planar patches. It suffices to distinguish between triangular and quadrangular patches, since more complex patch shapes (e.g. L-shapes) are obtained by patch composition. A patch can be accessed through two different types of representations which are kept in parallel. On the one hand, a parametric representation is provided, allowing direct inference of quantities such as angles or lengths. On the other hand, the coordinates of patch corners in the 3D world coordinate system are made accessible. The importance of this dual representation transpires when modelling an entire roof. A roof is defined by its patches and the relations between them. Possible relations are given by coincidence or equality constraints and can hold either between coordinates or parameters of different patches. Consequently, the dual representation allows to integrate diverse constraints such as geometric symmetry or topological connectivity simultaneously and at the same level of complexity. Once a relation between patches has been established, changes to one patch are propagated to all other affected patches in the roof.

4.1 Patch model

In Figure 3 the parametrization for a quadrangular patch is shown. The advantages of the parametric representation are that the quantities involved have an obvious meaning, and that probability distributions for the parameters can be obtained (e.g. from a test dataset). Additionally, this representation allows us to incorporate symmetries into the resulting roof model.

In parallel to the parametric representation, a representation based on the 3D world coordinates of the corner points $P_0, \ldots, P_3$ of the patch is kept. Although this dual representation is somewhat redundant, the coordinate based representation will be especially useful when introducing coincidence constraints as discussed in the next section.

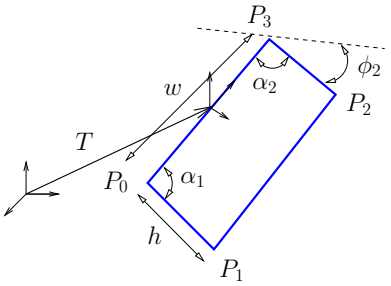

Figure 3: Quadrangular patch model. T is the transformation from world coordinates to the local patch coordinates system including a translation R and a rotation around the z-axis by ϕ_1. The slant of the patch plane is given by ϕ_2. The inclination of the bordering segments is denoted with α_1 and α_2 respectively. The parametrization of the patch is completed by specifying the width w and the height h.

4.2 Roof $= \sum$ patches + constraints

A roof consists of its constituting patches *plus* constraints which describe the relations between the patches. As an illustrative example, consider the L-shaped patch in Figure 4, constructed out of two quadrangular patches. In order to compose the two patches into one, two constraints are imposed. By making use of the parametric representation, a unique slant angle can be achieved by requiring $\phi_2^{(1)} = \phi_2^{(2)}$. Additionally, the coordinate based representation allows to glue corner points together by setting $P_0^{(1)} = P_1^{(2)}$.

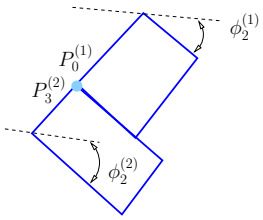

Figure 4: Illustration of the two different types of constraints. See description in text.

5 PATCH HYPOTHESES

This section focuses on the instantiation of the roof models. During the Bayesian model selection procedure (Section 3) 3D line segments are associated with different plane hypotheses. Now, patch hypotheses containing the segments in the different planes are sought. The outlines of the

patches are determined by computing the extremal points of the convex hull of the plane segments. Then, the patch parameters are adjusted to contain the extremal points.

5.1 Utility Theory

In order to decide which of the possible patch and plane configurations should be chosen for the final model, Utility Theory is used. Consider a non-deterministic action a which changes a system from state S to a set of n possible outcome states $\{S'_1, \ldots, S'_n\}$. The *expected utility* EU of action a is given by

$$EU(a) = \sum_i p(S'_i|a)U(S'_i) \tag{15}$$

where $p(S'_i|a)$ is the probability of state S'_i given action a is preformed, and $U(S'_i)$ is the utility of outcome state S'_i. If several different actions are possible, the *Maximum Expected Utility* (MEU) principle assures a rational behaviour of the system (Pearl, 1988). It says, that a rational system should choose the action which maximizes the expected utility. For roof reconstruction, the state of the system is described by one or more nodes in the interpretation tree from Figure 2. Possible actions correspond to choosing a specific state as final result.

5.2 Definition of the utility function

The outline of a patch consists of reconstructed 3D line segments and line segments stemming from the convex-hull completion of the outline. The most reliable description of a patch is certainly given, when all of its borders are reconstructed 3D line segments. Gaps between the reconstructed 3D line segments are closed using a convex-hull calculation, hence these border segments are more doubtful. This line of thought leads to the definition of the utility function for the task at hand. The utility function is defined as the ratio between the length of the border given by the reconstructed 3D line segments and the total length of the outline

$$U = \frac{\sum ||\mathbf{l}_{reconstructed}||}{\sum ||\mathbf{l}_{reconstructed}|| + \sum ||\mathbf{l}_{computed}||} \tag{16}$$

Besides the completeness of the boundary description of the individual patches, the compatibility within patches is taken into account. The *mutual overlap* turned out to be an appropriate measure

$$U = 1 - \text{normalized mutual overlap} \tag{17}$$

there the area of overlap is normalized to lie within the interval $[0 \ldots 1]$.

Since the entities for both Utility Functions are independent, a combined Utility Function is found by multiplication. The patch (and consequently plane) configuration with Maximum Expected Utility is entered into the final roof model.

6 ROOF RECONSTRUCTION

The patch hypothesis generation procedure described in Section 5 is sufficient to reconstruct roofs whose patches lie on planes which contain the seed line. To complete the reconstruction of more complicated roofs, a semantic interpretation of already reconstructed patches is performed. The key idea is to identify patch segments which might correspond to an internal boundary of the roof – that is a concave or convex joint of roof patches. Given such segments could be identified (if present at all), these in turn form a set of seed segments which are fed into the reconstruction algorithm again.

6.1 Semantic labels

We distinguish five different semantic labels for patch segments. The names of the labels are chosen to be coherent, although they should not be taken literally. Hence, a *gutter* segment just corresponds to the lower boundary of a patch, no matter if there is a gutter in the scene or not. Figure 5 gives an overview. The labels for the reconstructed patch borders are obtained using a probabilistic relaxation labelling algorithm (Christmas et al., 1995), which will be described in more detail elsewhere. Basically, geometric features of the individual line segment (unary relations) and geometric relations between pairs of line segments (binary relations) are used to obtain a semantic classification.

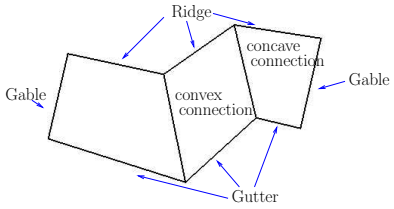

Figure 5: Functional parts of a roof and their semantic labels. The label *ridge* is generally used for the upper boundary of a patch, *gutter* for the lower one. The boundaries of a patch are either labelled *gable*, *convex connection* or *concave connection* depending on the neighbourhood.

6.2 Final reconstruction

It is desirable to have a topologically correct, that is, a point-wise connected reconstruction of the roof. Therefore, corner points of patches with compatible semantic labels are forced to coincide in a final *gluing step*. Thus, in this step, topological correctness is preferred over geometric precision.

7 RESULTS AND CONCLUSION

The presented results are obtained using a state-of-the-art dataset, produced by Eurosense Belfotop n.v. The image characteristics are: 1:4000 image scale and geometrically accurate film scanning with 20×20 $microns^2$ pixel size corresponding to 8×8 cm^2 on the ground. The 3D line segments are obtained using three overlapping views. The precise sensor orientation is known. To emphasise the quality

of the reconstructed roof geometry no texture mapping is applied.

Figure 6 shows a reconstruction result for a building roof, which was completely reconstructed from its seed (here: ridge) line by one pass of the reconstruction algorithm. The reconstruction results are detailed and topologically correct. For instance the small difference in the slope of the two patches on the right side of the roof has been correctly detected. Figure 7 shows the reconstruction of another building roof. The triangular patches on the front side do not lie in a plane given by the seed line. However, driven by the semantic labels attributed to the partially reconstructed roof, the missing patches could be successfully found.

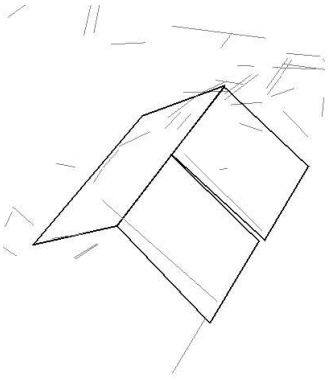

Figure 6: Three-dimensional view of a reconstructed building roof. The structure is modelled correctly, especially the two different slopes on the right hand side of the roof.

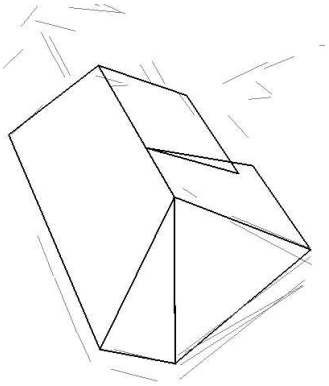

Figure 7: Three-dimensional view of a reconstructed building roof. The non trivial roof structure is captured to its full extent.

Concluding, in this paper we have presented a novel approach for the probabilistic modelling of building roofs. The proposed theoretical foundation leads to stable and reliable results. Future work will focus on exploiting the knowledge available in form of test datasets to a broader extent, with the goal to initialize roof models with even less evidence in 3D space. Another research direction which is pursued at the moment is oriented toward the optimization of the geometric accuracy of the models using back-projection into the input images.

Acknowledgement: The authors gracefully acknowledge support of the European IST project *CogViSys* and the Flemish Institute of Research in Industry and Agriculture IWT.

REFERENCES

Baltsavias, E., Grün, A., Van Gool, L.: Automatic Extraction of Man-Made Objects from Aerial and Space Images. Balkema Publishers, Rotterdam (2001)

Baillard, C. and Zisserman, A., 2000. A Plane-Sweep Strategy for the 3D Reconstruction of Buildings from Multiple Images. In: IAPRS, Vol. 33 part B2, pp. 56–62.

Christmas, W., Kittler, J. and Petrou, M., 1995. Structural Matching in Computer Vision Using Probabilistic Relaxation. PAMI 17, pp. 749–764.

Cord, C., Jordan, M., Cocquerez, J.-P. and Paparoditis, N., 1999. Automatic Extraction and Modelling of Urban Buildings from High Resolution Aerial Images. In: IAPRS, Vol. 32 part 3-2W5, pp. 187–192.

Fuchs, F., 2001. Building Reconstruction in Urban Environment: a Graph-Based Approach. In: [Baltsavias, E., Grün, A., Van Gool, L.]

Henricsson, O., 1998. The Role of Color Attributes and Similarity Grouping in 3-D Building Reconstruction. CVIU 72(2), pp. 163–184.

Heuel, S. and Förstner, W., 2001. Topological and Geometrical Models for Building Reconstruction from Multiple Images. In: [Baltsavias, E., Grün, A., Van Gool, L.]

Heuel, S., Lang, F. and Förstner, W., 2000. Topological and Geometrical Reasoning in 3D Grouping for Reconstructing Polyhedral Surfaces. In: IAPRS, Vol. 33

Institute of Geodesy and Photogrammetry, 2001. Dataset of Zürich Hoengg. Swiss Federal Institute of Technology (ETH) Zürich, CH-8093 ETH Hoenggerberg, Switzerland.

Kim, Z., Huertas, A. and Nevatia, R., 2001. A Model-Based Approach for Multi-View Complex Building Description. In: [Baltsavias, E., Grün, A., Van Gool, L.]

Kulschewski, K., 1997. Building Recognition with Bayesian Networks. In: W. Förstner and L. Plümer (eds), SMATI.

Moons, T., Frere, D., Vandekerckhove, J. and Van Gool, L., 1998. Automatic Modelling and 3D Reconstruction of Urban House Roofs from High Resolution Aerial Imagery. In: ECCV, Vol. 1, pp. 410–425.

Pearl, J., 1988. Probabillistic Reasoning in Intelligent Systems. Morgan Kaufmann Publishers.

Scholze, S., 2002. Bayesian Model Selection for Plane Reconstruction. Technical Report BIWI–TR–259, Communication Technology Laboratory, Computer Vision Group, ETH, Zürich, Switzerland.

Scholze, S., Moons, T., Ade, F. and Van Gool, L., 2000. Exploiting Color for Edge Extraction and Line Segment Stereo Matching. In: IAPRS, Vol. 33, pp. 815–822.