

Effective Shapes Generation for Bayesian CAD Model Reconstruction

A. Bey^{1,2,3}, R. Chaine^{1,2}, R. Marc³, G. Thibault³

¹Université de Lyon, CNRS

²Université Lyon 1, LIRIS, UMR5205, F-69622, France

³Électricité De France Recherche & Développement, F-92140, France

Abstract

We address the reconstruction of CAD models from 3D point clouds, assuming that an a priori CAD model, roughly similar to the scene to reconstruct, is given. The problem can be solved using a Bayesian approach which states the reconstruction task as the search of the most probable CAD model. This article presents a reliable cylinder generation algorithm that can be used to efficiently reach the configuration which maximizes the target probability, combined with an efficient specific optimization heuristic. Results illustrate how our approach can automatically reconstruct CAD models of industrial scenes.

Categories and Subject Descriptors (according to ACM CCS): J.6 [Computer Applications]: Computer-Aided Engineering—Computer-Aided Design (CAD) I.4.8 [Computing Methodologies]: Scene Analysis—Object Recognition, Surface fitting, Shape, Range data

1. Introduction

Context and motivations The "As Built" reconstruction task consists in building a 3D CAD model which describes the actual state of a facility. 3D LIDAR data are often used for that purpose, since they provide huge amount of 3D measures (point clouds) that can be used to infer an accurate as built CAD model of the facility. The CAD models that we consider here are assemblies of primitive shapes such as planes, cylinders, cones and tori which allow a concise representation of the scenes. However, the way shapes are assembled together must be driven so as to provide consistent representations of the scenes. Typically, the shapes connections must fulfill constraints which relate to the way the industrial facilities are built. Besides this expert knowledge, we also assume that a CAD model which roughly describes the scene to be reconstructed is available. For instance, we have a CAD model of a facility which is similar to the one where the point cloud has been acquired, although there may exist significant differences between these two scenes. This prior CAD model can be used to drive the reconstruction process, and to reinforce the reliability of the reconstructed 3D model. Hence, we aim at reconstructing a CAD model which fits to the point cloud and satisfies the expectations

embedded in the two sources of *a priori* knowledge mentioned above: the model must be consistent and it must also be similar to the prior CAD model which is known to be a fair estimate of the result.

Our contributions While usual approaches [SWK07, RV05] fit independent shapes to the point cloud, we state the reconstruction problem as a probability optimization task embedding *a priori* expectations about the model to be reconstructed [BCM*11, LKBH10]. This problem is solved using a greedy optimization algorithm based on a stochastic shapes generation tool which is meant to provide reliable shapes with a high probability, thus making the exploration of the solution space efficient. This framework enables the simultaneous recognition of several shape types by handling the conflicts that could arise, although we compute the tori (that mainly stand for elbows) upon the detected cylinders (that may stand for pipes). Unlike some other algorithms [LWC*11], our probabilistic approach may accept any configuration which seems to be relevant, even if it does not perfectly comply with the *a priori* constraints.

Data and notations We are given a point cloud $\mathcal{P} = \{(\mathbf{p}_1, \mathbf{n}_1), \dots, (\mathbf{p}_n, \mathbf{n}_n)\}$ coming from terrestrial LIDAR acquisitions in an industrial facility, where each \mathbf{p}_i de-

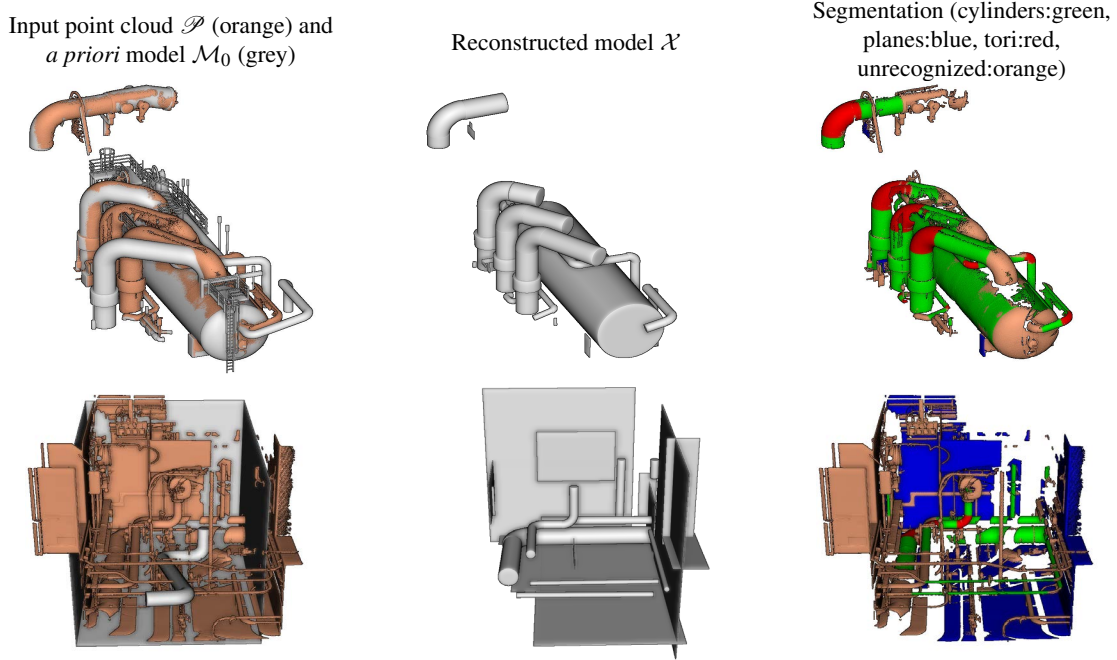


Figure 1: Results on industrial scenes with complete or partial a priori models.

notes a 3D point and \mathbf{n}_i denotes the normal vector to the scanned surface at point \mathbf{p}_i (either provided with LIDAR data, or computed using local planes fitting [MN04]). Moreover, we also have $\mathcal{M}_0 = \{\mathcal{S}^1, \dots, \mathcal{S}^m\}$ a prior CAD model standing for the theoretical state of the facility. Each element $\mathcal{S}^i \in \mathcal{M}_0$ is a parametrized primitive shape (the parameters define the orientation, the position and the geometry of the shape). In this paper, we only consider the \mathcal{S}^i 's that are either planes, cylinders or tori. $\mathcal{X} = \{\mathcal{S}_0, \dots, \mathcal{S}_x\}$ denotes the solution which is computed. Again, each \mathcal{S}_i is either a plane, a cylinder or a torus.

2. Bayesian reconstruction

Probabilistic approach The reconstruction problem can be stated as the search of the most probable CAD model. Thus we have to find the configuration \mathcal{X}^* that maximizes a given *a posteriori* probability. According to the Bayes rule, this probability quantifies the relevancy of any configuration \mathcal{X} with respect to both the observed data (point cloud) and the *a priori* knowledge. Using decreasing exponential probabilities, the *a posteriori* probability maximization problem is turned into the minimization of the negative log-likelihood. Thus we have to find the configuration \mathcal{X}^* that minimizes the energy:

$$E(\mathcal{X}) = \lambda_D E_D(\mathcal{X}, \mathcal{P}) + \lambda_G E_G(\mathcal{X}, \mathcal{M}_0) + \lambda_T E_T(\mathcal{X}) \quad (1)$$

The data fitting energy E_D quantifies how well the data \mathcal{P}

fits to the model \mathcal{X} . It is modelled as the sum of the contributions $e_D(\mathbf{p}_i)$ over all the $\mathbf{p}_i \in \mathcal{P}$. Each point \mathbf{p}_i is associated with its nearest shape $\mathcal{S}_j \in \mathcal{X}$, and brings the following contribution:

$$e_D(\mathbf{p}_i) = \begin{cases} \left(\frac{d(\mathbf{p}_i, \mathcal{S}_j)}{\epsilon} \right)^2 - 1 & \text{if } d(\mathbf{p}_i, \mathcal{S}_j) \leq 3\epsilon \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where d is the euclidean distance function and the parameter ϵ stands for the noise in the point cloud. The points close to their nearest shape decrease the energy while those that are a bit too far increase it. The points that are obviously too far from all the shapes (beyond 3ϵ) do not contribute, since we do not want to penalize the partial reconstruction of the point cloud. However, we consider that any shape in \mathcal{X} must have at least one neighbor (closer than 3ϵ).

The geometrical energy E_G quantifies the similarity between \mathcal{X} and the *a priori* model \mathcal{M}_0 . It is modelled as the sum of the contributions $e_G(\mathcal{S}_i, \mathcal{M}_0)$ over all the shapes $\mathcal{S}_i \in \mathcal{X}$. Each contribution is computed by associating the shape $\mathcal{S}_i \in \mathcal{X}$ with its most similar match $\mathcal{S}^j \in \mathcal{M}_0$, *i.e.* the shape \mathcal{S}^j minimizing the similarity measurement $e_G(\mathcal{S}_i, \mathcal{S}^j)$.

The topological energy E_T quantifies how well \mathcal{X} fulfills the *a priori* expectation on the way shapes should be assembled. It is modelled as the sum of the contributions $e_T(\mathcal{S}_i, \mathcal{S}_j)$ for each pair of connected shapes $(\mathcal{S}_i, \mathcal{S}_j) \in \mathcal{X}^2$ (shapes are connected if they are in contact).

The expression of the geometrical and topological con-

tributions depends on the shapes that are involved. We consider cylinders and planes because they stand for a significant amount of components in industrial scenes. Moreover, the reconstructed cylinders can be used subsequently to complete the reconstruction of pipe lines by using tori as elbows.

Energies for cylinders The similarity measurement $e_G(C_i, C^j)$ between two cylinders $C_i \in \mathcal{X}$ and $C^j \in \mathcal{M}_0$ combines the angle α between the shapes axes, the distance δ from C_i 's center to C^j and the difference ρ between their radii, given user specified tolerances σ_α , σ_δ and σ_ρ to each of these three changes:

$$e_G(C_i, C_j) = \left(\frac{\alpha}{\sigma_\alpha}\right)^2 + \left(\frac{\delta}{\sigma_\delta}\right)^2 + \left(\frac{\rho}{\sigma_\rho}\right)^2 \quad (3)$$

The contribution $e_T(C_i, C_j)$ of two connected cylinders C_i and C_j combines three distinct constraints: cylinders axes should intersect, the angle between cylinders axes should be either flat or right and connected cylinders having the same radius should be favored. Each of these three constraints is modeled by an inverted Gaussian mixture $g(x) = \sum_i b - a \exp(-c(x - d_i)^2)$ (with $a > 0$ and $c < 0$). These functions respectively depend on the cylinders radii difference ζ , the distance χ and the angle η between their axes. Their minima stand for the ideal cases $\chi = 0$, $\eta \in \{0^\circ, 90^\circ\}$ and $\zeta = 0$. The whole contribution is the sum of these three Gaussian mixtures, whose parameters are computed so that each function becomes positive (increases the energy) as soon as the error gets greater than user specified tolerances. However, the Gaussian mixture $g(\zeta)$ is negative whatever ζ , because we favor radii equality but do not penalize their difference.

Extending energies to planes The similarity measurement $e_G(\mathcal{P}_i, \mathcal{P}^j)$ between two planes $\mathcal{P}_i \in \mathcal{X}$ and $\mathcal{P}^j \in \mathcal{M}_0$ can be computed using equation 3 by considering that α is the angle between the normal vectors of these two planes, and δ is the distance from \mathcal{P}_i 's center to \mathcal{P}^j (the leftmost term of equation 3 vanishes for the planes).

Concerning e_T , we expect connected planes to be orthogonal, and that any plane which is connected to a cylinder is orthogonal to the cylinder axis. These two constraints are modeled by two Gaussian mixtures that respectively depend on the angle η_P between the normals vectors of two planes and the angle η_C between the normal vector of a plane and a cylinder axis, with minima $\eta_P = 0^\circ$ and $\eta_C = 90^\circ$.

3. Strategy for the optimization

We propose to adapt the stochastic birth and death optimization approach [DMZ08] to the point cloud reconstruction, by merging the birth (random shape generation) and death (probabilistic shape removal) phases into one single greedy transition process. This greedy optimization does not rely on a simulated annealing scheme, contrary to the complete stochastic approach [DMZ08].

A single random candidate \mathcal{C} is generated at each step (c.f. section 4). Once a candidate has been built, we try to find the elements in \mathcal{X} that are no longer relevant, i.e. the set of shapes \mathcal{K}^* that lead to the best configuration (the lowest energy) when removed from \mathcal{X} . The new configuration $\mathcal{X}' = \mathcal{X} \setminus \mathcal{K}^* \cup \{\mathcal{C}\}$ is accepted if and only if it decreases the energy ($E(\mathcal{X}') < E(\mathcal{X})$).

The computation of the actual set \mathcal{K}^* is a combinatorial optimization problem whose resolution may take a prohibitive time. Instead, we use heuristics to compute a set \mathcal{K}' approximating \mathcal{K}^* . First, we reduce the solution space by considering only the set \mathcal{L} of shapes in \mathcal{X} that are connected to the new candidate \mathcal{C} . Indeed, the elements that are not connected to \mathcal{C} are unlikely to get involved in the decrease of the energy. Then we use a greedy approach in order to efficiently estimate the subset of \mathcal{L} reaching the lowest energy. Starting with $\mathcal{K}_0 = \emptyset$, we incrementally build the set $\mathcal{K}_i = \mathcal{K}_{i-1} \cup \{\mathcal{S}_i^*\}$ at each step i , where $\mathcal{S}_i^* \in \mathcal{L} \setminus \mathcal{K}_{i-1}$ is the shape whose removal brings the best contribution (the greatest decrease of the energy). The set \mathcal{K}' approximating \mathcal{K}^* is the set \mathcal{K}_i whose removal leads to the optimal configuration (the lowest energy). and can be computed in a time which quadratically depends on the size of \mathcal{L} .

4. Effective Cylinders Generation

Overview To make the optimization tractable, we have to propose a random method which mostly generates shapes that are relevant with respect to the energy E . We assume we are given a cylinder $\mathcal{C}^i \in \mathcal{M}_0$ (e.g. randomly picked in \mathcal{M}_0), and we would like to build a random \mathcal{C} which looks like \mathcal{C}^i and fits well to the point cloud. For that purpose, we propose an approach inspired by the RANSAC method [SWK07], in which the candidate shapes are build from points randomly chosen in the point cloud. But the points selection process is biased in order to favor the generation of relevant shapes.

Localisation First, we have to find a part of \mathcal{P} that is likely to represent a cylinder matching \mathcal{C}^i . To do so, we define a probability P_S over \mathcal{P} favoring the points that are close to \mathcal{C}^i and whose normal is compliant with \mathcal{C}^i . We use a Gaussian function of the distance d_{C^i} from the point to \mathcal{C}^i and the angle a_{C^i} between the normal and \mathcal{C}^i 's axis, whose maximum lies at ($d_{C^i} = 0$, $a_{C^i} = 90^\circ$), using deviances σ_δ and σ_α introduced in equation 3. A first point \mathbf{p}_0 is picked by simulating a sampling with probability P_S over \mathcal{P} [CG95].

Candidate generation Then we would like to find a set of points $\mathcal{Q}_{\mathbf{p}_0}$ sampling the same surface as \mathbf{p}_0 . The points of \mathcal{P} that are close to \mathbf{p}_0 have great chances to represent the same surface as \mathbf{p}_0 . Therefore we pick the remaining point in a spherical neighborhood $\mathcal{N}_{\mathbf{p}_0}$ of \mathbf{p}_0 , whose size depends on \mathcal{C}^i 's radius and whose center is shifted along the normal vector associated with \mathbf{p}_0 in order to encompass the surface sampled by \mathcal{P} . We then compute a cylinder \mathcal{C}'' grossly approximating the points in $\mathcal{N}_{\mathbf{p}_0}$. The axis of \mathcal{C}'' is estimated

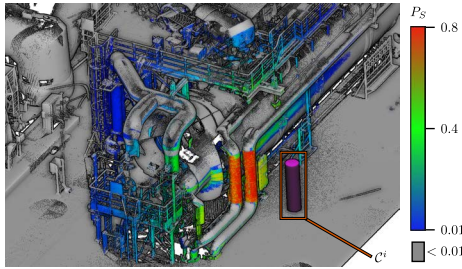


Figure 2: Probability P_S for a point to be picked with respect to some a priori cylinder C^i (violet). We can see that P_S is high on surfaces matching C^i .

using a plane fitting in the Gaussian sphere of \mathcal{N}_{p_0} , while its position and radius are computed using a circle detection after the points have been projected in the plane orthogonal to the detected axis [RV05]. The points in \mathcal{N}_{p_0} whose distance to C'' is greater than ϵ are then discarded, and we build the candidate \mathcal{C} using a least square fitting on the remaining points \mathcal{P}_{p_0} (Levenberg-Marquardt algorithm initialized with C'').

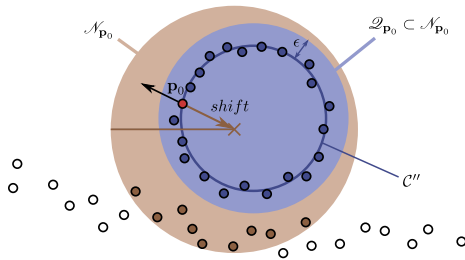


Figure 3: Selection of the points \mathcal{P}_{p_0} (blue) given a first point p_0 (red) for the cylinder generation.

5. Generate planes

This approach can be extended to the planes generation assuming we are given a prior plane $\mathcal{P}^i \in \mathcal{M}_0$ instead of a cylinder C^i . To do so, we propose a probability for a point to match \mathcal{P}^i which increases as the point gets closer to \mathcal{P}^i , and as its associated normals aligns with \mathcal{P}^i 's normal. Moreover, instead of using a spherical neighborhood, we use cylindrical neighborhood whose radius depends on the area of \mathcal{P}^i and whose height depends on the noise estimate ϵ .

6. Finalize the reconstruction

In the proposed approach, we do not handle shapes such as tori and cones because these shapes usually stand for connection pieces in pipe lines. Instead, the computed model is completed afterwards: for each pair of connected cylinders, a candidate torus linking the cylinders is built and fitted to the

points lying at the cylinders junction. This candidate elbow is accepted if it decreases the data fitting energy E_D .

7. Results and conclusion

The results presented in figure 1 show that the proposed method can handle the recognition of a priori models, despite changes in orientation, position and geometry of the shapes to be detected. This algorithm can handle partial a priori models, and only recognizes the parts that actually match \mathcal{M}_0 up to the specified tolerances. The reconstructed models fulfill the expectations on the shapes connections, and fit well to the point cloud. The proposed method is meant to favor the reliability of the results over the speed of convergence. Each of the tested point clouds has several millions of points (from 1 to 20 millions), and the algorithm takes hours to provide such results (from 20 minutes to 20 hours in the experiments that we have made). However, the user can decide when to stop the computation: a great part of the computation time in the whole process is spent on slight improvements of the model, and some good results can be obtained at the early stage of the reconstruction process. However, the shape generation effectiveness could be significantly improved by considering high order geometry knowledge (e.g. curvatures and normals uncertainties) in the points selection process, and thus increase the probability to generate relevant shapes.

References

- [BCM*11] BEY A., CHAINE R., MARC R., THIBAUT G., AKKOCHE S.: Reconstruction of consistent 3d CAD models from point cloud data using a priori cad models. In *ISPRS Workshop on Laser Scanning* (2011). 1
- [CG95] CHIB S., GREENBERG E.: Understanding the metropolis-hastings algorithm. *The American Statistician* 49, 4 (1995), 327–335. 3
- [DMZ08] DESCOMBES X., MINLOS R., ZHIZHINA E.: Object Extraction Using a Stochastic Birth-and-Death Dynamics in Continuum. *Journal of Mathematical Imaging and Vision* 33, 3 (Oct. 2008), 347–359. 3
- [LKBH10] LAFARGE F., KERIVEN R., BRÉDIF M., HIEP V.: Hybrid multi-view reconstruction by jump-diffusion. In *Computer Vision and Pattern Recognition, 2010* (2010), IEEE, pp. 350–357. 1
- [LWC*11] LI Y., WU X., CHRYSATHOU Y., SHARF A., COHEN-OR D., MITRA N.: GlobFit: consistently fitting primitives by discovering global relations. *ACM Transactions on Graphics* 30, 4 (2011), 52. 1
- [MN04] MITRA N. J., NGUYEN A.: Estimating Surface Normals in Noisy Point Cloud Data. *Special Issue of International Journal of Computational Geometry and Applications* 14 (2004), 261–276. 2
- [RV05] RABBANI T., VAN DEN HEUVEL F.: Efficient hough transform for automatic detection of cylinders in point clouds. *ISPRS WG III/3, III/4 3* (2005), 60–65. 1, 4
- [SWK07] SCHNABEL R., WAHL R., KLEIN R.: Efficient RANSAC for Point-Cloud Shape Detection. *Computer Graphics Forum* 26, 2 (2007), 214–226. 1, 3