

Structural Symmetries from Motion for Scene Reconstruction and Understanding

Natesh Srinivasan

<http://www.cc.gatech.edu/grads/n/nsriniva/>

Luca Carlone

<http://www.lucacarlone.com/>

Frank Dellaert

http://www.cc.gatech.edu/~dellaert/FrankDellaert/Frank_Dellaert/Frank_Dellaert.html/

Robotics and Intelligent Machines Lab,

College of Computing,

Georgia Institute of Technology




Figure 1: The proposed approach uses (a) a set of 2D images, and (b) a sparse 3D reconstruction to detect repetitive patterns (regular structure) and exploits them for SfM. (c) Generative model of the regular structure obtained from the proposed approach, for the Leuven dataset.

The identification and description of partial symmetries in man-made structures is a powerful tool to improve the quality of 3D reconstruction from unordered images and to enable high-level understanding of scene geometry. In this work we propose an approach to identify symmetries and exploit them in Structure from Motion (SfM). Our first contribution is a symmetry detection approach that uses the 3D geometry of the scene as well as 2D appearance cues. We show that a particular parametrization of the *transformation space* (space in which each point represents a candidate symmetry relation) exposes the dominant symmetries in the scene. Then, we use appearance information to prune incorrect symmetry hypotheses. The second contribution is a *constrained bundle adjustment* (CBA) scheme that jointly optimizes for the best 3D reconstruction and the symmetry generators. In contrast to related work on CBA, our approach models n -fold (rotational and translational) repetitions of architectural elements, and allows estimating a *generative model* of the 3D geometry. Experimental results confirm that our method can correctly identify and exploit partial symmetries in noisy and sparse SfM datasets.

We propose an approach for 2D-3D symmetry detection and we show how to leverage the presence of repetitive structure to improve SfM reconstruction. Our approach includes three building blocks. The first is a multi-hypotheses estimator for 3D symmetry generators. We borrow key insights from [4], which shows how to map putative symmetry transformations into 1D or 2D lattices. However, we skip grid fitting (which is unreliable on SfM data), and we show that a polar parametrization of the transformation space clearly exposes dominant symmetries. The second block prunes the multiple hypotheses on the generators and returns the generators that are most consistent with 2D appearance. This is similar in spirit to [3], while we avoid 3D surface fitting. Finally, the last block takes the estimate for the generators and jointly refines this estimate and the 3D reconstruction. We present an on-manifold optimization scheme that respects the regular structure and the manifold nature of the variables involved in the problem (rotations, poses). Contrarily to standard CBA we explicitly model n -fold repetitions, we avoid priors [2] and user intervention [1], and we also output a *generative model* of the scene.

We first detect 3D features on the sparse point cloud. While related work obtains these features by uniform sampling [4], we use 3D SIFT keypoint. For each feature, we associate a set of neighboring points which are within a ball of predetermined radius. Then, for each pair of features, we compute a relative transformation by applying ICP to local patches around the features. This gives a set of putative, highly noisy set of transformations similar to [4]. Our goal is to obtain a good guess for the parameters of the underlying symmetry (a.k.a generators) from these noisy transformations.

The key insight for obtaining a good guess for the generators is that the presence of the repetitive structure in the 3D model implies that many of the transformations will be (approximately) in the form $x_{ij} = n \cdot g_k$ ($k \in \{1, 2\}$) i.e., will be produced by an n -fold repetition along the generator $g_k \in \mathbb{R}^2$ (this is also one of the motivations for the grid fitting of [4]).

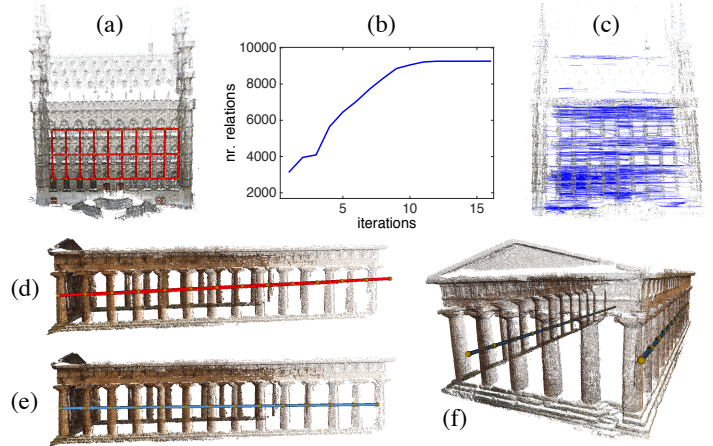


Figure 2: Leuven dataset: (a) initial estimate for the generators (before CBA); (b) number of symmetry relations discovered during the CBA iterations; (c) symmetry relations shown as lines connecting points considered n -fold repetitions of each other. Neptune dataset: (d) initial estimate of the generators, (e, f): refined generator estimate from CBA.

Moreover, the 2-vector g_k can be written as $g_k = \delta_k \cdot u_k$, where $\delta_k \in \mathbb{R}$ is the *repetition period* and the unit vector $u_k \in \mathbb{R}^2$ ($\|u_k\| = 1$) is the *repetition direction*. Therefore, all the putatives generated by g_k become $x_{ij} = (n\delta_k) \cdot u_k$, i.e., they share the same direction, while having different norms depending on the unknown number of repetitions n . Polar coordinates naturally adapt well to recover the underlying symmetries for the 2 parameters family of generators. We further refine the guess of the generator by using a appearance based verification step. The intuition is that a good generator is one that maps a patch to another one having similar appearance (this is essentially the concept of regular structure). Since each point in the 3D model is obtained via SfM, it has a corresponding feature descriptor. We use this insight to devise a generator ranking scheme.

We refine the generators that is obtained in the previous step using a bundle adjustment scheme that implicitly encode n -fold repetitions in them. Contrary to existing approaches [2], that consider only translational symmetries or 1-fold repetitions, we model n -fold rotational and translational repetitions. We jointly optimize for the value of the generators and also refine the sparse point cloud and the associated camera poses similar to the method described in [1]. We distinguish ourselves from previous work here by dealing with all kinds of 2-parameter transforms described in [4] but applied to noisy point clouds reconstructed from SfM.

- [1] D. Ceylan, N.J. Mitra, Y. Zheng, and M. Pauly. Coupled structure-from-motion and 3D symmetry detection for urban facades. *ACM Transactions on Graphics*, 33(2), January 2014.
- [2] A. Cohen, C. Zach, S. Sinha, and M. Pollefeys. Discovering and exploiting 3d symmetries in structure from motion. In *CVPR. Computer Vision and Pattern Recognition*, June 2012. URL <http://research.microsoft.com/apps/pubs/default.aspx?id=163609>.
- [3] N. Jiang, P. Tan, and L.-F. Cheong. Multi-view repetitive structure detection. In *ICCV*, pages 535–542, 2011.
- [4] M. Pauly, N. J. Mitra, J. Wallner, H. Pottmann, and L. Guibas. Discovering structural regularity in 3D geometry. *ACM Transactions on Graphics*, 27(3):#43, 1–11, 2008.