

A Symmetry Prior for Convex Variational 3D Reconstruction

Pablo Speciale¹(✉), Martin R. Oswald¹, Andrea Cohen¹, and Marc Pollefeys^{1,2}

¹ ETH Zürich, Zurich, Switzerland

`pablo.speciale@inf.ethz.ch`

² Microsoft, Redmond, USA

Abstract. We propose a novel prior for variational 3D reconstruction that favors symmetric solutions when dealing with noisy or incomplete data. We detect symmetries from incomplete data while explicitly handling unexplored areas to allow for plausible scene completions. The set of detected symmetries is then enforced on their respective support domain within a variational reconstruction framework. This formulation also handles multiple symmetries sharing the same support. The proposed approach is able to denoise and complete surface geometry and even hallucinate large scene parts. We demonstrate in several experiments the benefit of harnessing symmetries when regularizing a surface.

Keywords: Symmetry prior · 3D reconstruction · Variational methods · Convex optimization

1 Introduction

One of the long-time goals of computer vision algorithms is to imitate the numerous powerful abilities of the human visual system to achieve better scene understanding. Many methods have actually been inspired by the physiology of the visual cortex of mammalian brains. One of the strongest cues that humans use in order to infer the underlying geometry of a scene despite having access to only a partial view is symmetry, as shown in [20]. Moreover, symmetry is a very strong and useful concept because it applies to many natural and man-made environments. Following this inspiration, we propose a method which leverages symmetry information directly within a 3D reconstruction procedure in order to complete or denoise symmetric surface regions which have been partially occluded or where the input information has low quality. In contrast to the majority of 3D reconstruction methods which fit minimal surfaces in order to fill unobserved surface parts, our method favors solutions which align with symmetries and adhere to required smoothness properties at the same time. Similarly to how humans extrapolate occluded areas and 3D information from just a few view points, our method can hallucinate entire scene parts in unobserved areas, fill small holes, or denoise observed surface geometry once a symmetry has been detected. An example of our approach is shown in Fig. 1.

Equal contribution from P. Speciale and M.R. Oswald.

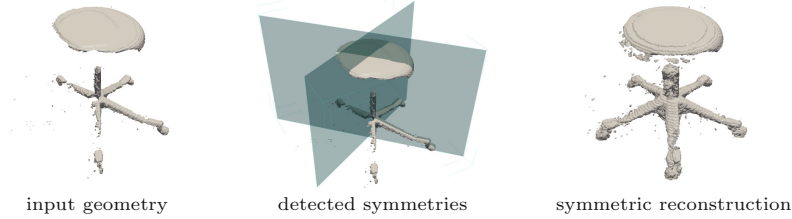


Fig. 1. Example application of our approach. A model of a stool was scanned by a depth camera and the result is incomplete due to occlusions. With only two detected symmetries we can complete the 5-way symmetry of the model.

1.1 Contributions

We propose to use symmetry information as a prior in 3D reconstruction in order to favor symmetric solutions when dealing with noisy and incomplete data. For this purpose, we extend standard symmetry detection algorithms to be able to exploit partially unexplored domains. Our framework naturally unifies the applications of symmetry-based surface denoising, completion and the hallucination of unexplored surface areas. To the best of our knowledge, we present the first method that handles multiple symmetries with a shared support region, since the proposed algorithm computes an approximation to a non-trivial projection to equally satisfy a set of symmetries. Finally, our method extends the toolbox of priors for many existing variational 3D reconstruction methods and we show that a symmetry prior can achieve quality improvements with a moderate runtime overhead.

1.2 Related Work

The works by Liu *et al.* [13] and Mitra *et al.* [16] give a broad overview of symmetry detection methods, although most of the work they discuss focuses on symmetry extraction for computer graphics applications. Therefore, their applicability is mostly shown on perfect synthetic data. In this section, we focus on works that detect and exploit symmetries from real data for vision applications.

Köser *et al.* [11] detect planar reflective symmetries in a single 2D image. They demonstrate that the arising stereo problem can be solved with standard stereo matching when the distance of the camera to the symmetry plane corresponds to a reasonable baseline.

Kazhdan *et al.* [7] define a reflective symmetry descriptor for 3D models by continuously measuring symmetry scores for all planes through the models' center of mass. We briefly repeat some of their theoretic results as we are going to use them in Sect. 2. Let $\gamma \in \Gamma$ be a symmetric transformation and u be a symmetric object, then u is symmetric with respect to γ if it is invariant under the symmetric transformation, that is, $\gamma(u) = u$. Using the group properties of the symmetry transform, Kazhdan *et al.* define the following symmetry distance

$$\text{SD}(u, \gamma) = \min_{v: \gamma(v)=v} \|u - v\| = \|u - \Pi_\gamma(u)\| = \left\| \frac{u - \gamma(u)}{2} \right\|, \quad (1)$$

in which $\Pi_\gamma(u) = \frac{1}{2}(u + \gamma(u))$ is the orthogonal projection of object u onto the set of objects which are symmetric with respect to the symmetry γ .

Based on these results, Podolak *et al.* [23] focused on planar reflective symmetries and generalized the descriptor to detect also non-object-centered symmetries. They further define geometric properties which can be used for model alignment or classification. Both [4, 22] detect symmetries in meshes and propose a symmetric remeshing of objects for quality improvements of the mesh and for more consistent mesh approximations during simplification operations. Similarly, Mitra *et al.* [15] use approximate symmetry information of objects to allow their transformation into perfectly symmetric objects.

Cohen *et al.* [2] detect symmetries in sparse point clouds by using appearance-based 3D-3D point matching in a RANSAC-based procedure. They then exploit these symmetries to remove drift from point-clouds. Although their work is done on sparse point clouds, we present a similar approach for voxel grids in Sect. 2.

Symmetries have also been used for many applications, e.g. shape matching and feature point matching [6], object part segmentation, and canonical coordinate frame selection [23]. Our goal is to apply these concepts into the domain of dense 3D reconstruction. Nevertheless, our work is not the first one leveraging symmetry information in this domain. One of the earliest attempts to incorporate symmetry priors into surface reconstruction methods was by Terzopoulos *et al.* [25] who use a deformable spine model to create a generalized cylinder shape from a single image.

Application-wise, the work by Thrun and Wegbreit [26] is closely related to ours. They detect an entire hierarchy of different symmetry types in point cloud data and subsequently demonstrate the completion of unexplored surface parts. As opposed to our approach, they do not simultaneously denoise the input data and they do not compute a water-tight surface.

In contrast, we propose to integrate knowledge about symmetries directly into the surface reconstruction process in order to better reason about noisy or incomplete input data which can come from image-based matching algorithms or 3D depth sensors. Furthermore, our approach can handle any number of symmetries and their support domains can arbitrarily overlap. To the best of our knowledge no other method deals with several symmetries that share the same domain in a way that the “symmetrized” result obeys the group structure of several symmetries at once.

We build our symmetry prior into the variational volumetric 3D reconstruction framework which has been used in various settings, e.g. depth map fusion [29], 3D reconstruction [9, 27], multi-label semantic reconstruction [5], and spatio-temporal reconstruction [17], with anisotropic regularization [5, 10, 18] or connectivity constraints [19]. The proposed method extends these lines of works by a symmetry prior which can be easily combined with any of these works.

Moreover, since the 3D reconstruction is formulated as a 3D segmentation problem, the proposed prior is also directly applicable to a large set of segmentation methods like [24,28].

2 Symmetry Detection

In order to exploit symmetry priors for reconstruction, we first need to detect the symmetries that best fit the data. In this paper, we focus on detecting planar symmetries. As input we use integrated depth information, which is represented by a truncated signed distance function (TSDF) on a volume $V \subset \mathbb{R}^3$ [3]. The TSDF assigns a positive value for voxels corresponding to free space and a negative value for occupied voxels (which are placed behind the observed depth values). A zero value in this function denotes an unobserved (or occluded) voxel. A surface is then implicitly defined as the transition between positive and negative values. Since we are only interested in voxels lying on the surface of the object, we only look at the voxels for which the gradient is very strong.

Furthermore, since planar surfaces are symmetric with respect to all planes perpendicular to them, which is not very informative for the global scene or even for a small object, we decide to look only at those voxels that exhibit a certain degree of curvature. Thus, the goal is to find the symmetry planes that best reflect these high-gradient and high-curvature voxels into positions that also have a large gradient and curvature or that are otherwise unknown or occluded. Note that [8] detect partial symmetries in volumetric data via sparse matching of extracted line features, which is a more sophisticated way of using the gradient on the data. While trying to find the symmetries of a scene, we define the symmetry support $V_\gamma \subseteq V \subset \mathbb{R}^3$ as a hole-free, connected subset of the reconstruction domain V that fits the detected symmetry, that is, we try to include all occupied and free space which complies with the symmetry γ . Unobserved regions will be included and treated in a way such that they perfectly fulfill the symmetry to allow for hole filling and hallucination.

2.1 RANSAC

We apply a RANSAC-based approach by taking as input the list of high-gradient and high-curvature voxels, which we will refer to as surface voxels, and the list of unknown voxels. First, we randomly sample two surface voxels, which define a unique symmetry plane that reflect one of these voxels into the other. Next, we look through all of the surface voxels and reflect them over this plane to look for inliers to this particular symmetry. If the reflection of a surface voxel falls into the position of another surface voxel, then they are both considered as inliers to this symmetry plane. However, if it falls into an unknown voxel position, then we also consider it as an inlier since this could be a potential occluded part of a symmetric object. We randomly sample planes from two surface voxels as many times as stated by the RANSAC termination formula.

The plane with the most inliers is chosen as the best global symmetry plane γ for this surface, and its inliers define the support V_γ . Since we are interested in potentially finding many symmetries for the same scene, RANSAC could be applied sequentially by removing the inliers for the best symmetry and then subsequently re-detecting the next best symmetry among the remaining surface voxels. However, since there could be many symmetries with the same support, we modify RANSAC and keep track of the set of best N solutions (where N has to be determined a-priori). This way, we extract the N symmetries that best fit the entire surface. We will refer to this set of symmetries as Γ .

Local Symmetry Detection. A scene can be composed of several objects with different sizes and symmetries. However, due to the size variability, applying RANSAC on the whole volume as described before would miss most of the symmetries for small objects and also cluster different objects with similar symmetry planes into one approximate, noisy, symmetry plane. Therefore, we apply the described RANSAC approach for sliding boxes of different sizes over the entire volume. This allows the segmentation of objects as the support for the different symmetries found. The symmetry planes with the bigger inlier ratios and their support are chosen as candidates for local object symmetries. For multiple detections of the same object (parts) at different scales we gave preference to larger support domains and rejected detections whose support is a subset of another.

2.2 Hough Transform

Alternatively to RANSAC, we also implement a method based on [23], which resembles the Hough transform approach, in order to have additional insights in the space of planar symmetries belonging to an object. As cost for the hough space voting, we use the Planar-Reflective Symmetry Transform (PRST), which is defined according to [23] as follows

$$\text{PRST}(u, \gamma) = 1 - \frac{\text{SD}^2(u, \gamma)}{\|u\|^2} = \frac{1 + u \cdot \gamma(u)}{2}. \quad (2)$$

We parametrize planes in 3D by the spherical coordinates of their normals $\theta \in [0, \pi]$, $\phi \in [0, \pi]$ and the distance to the origin $d \in [d_{min}, d_{max}]$. After finding the peaks in the Hough Space using a non-maximal suppression scheme, we can obtain planar symmetries with high PRST values, as illustrated in Fig. 2.

If we consider the special case when u values are binaries, representing an occupancy grid, the Eq. (2) becomes the number of inliers of the γ symmetric plane, a metric also used in the RANSAC method described previously. Therefore, the methods are essentially very similar and the decision between one or the other depends only on technical considerations. For example, the runtime of Hough Transform is fixed, which is the number of iterations used in the plane sampling; on the other hand, the number of iterations ran by RANSAC depends on the current inlier threshold and, therefore, could possibly finish sooner. Another advantage of RANSAC is its low memory footprint

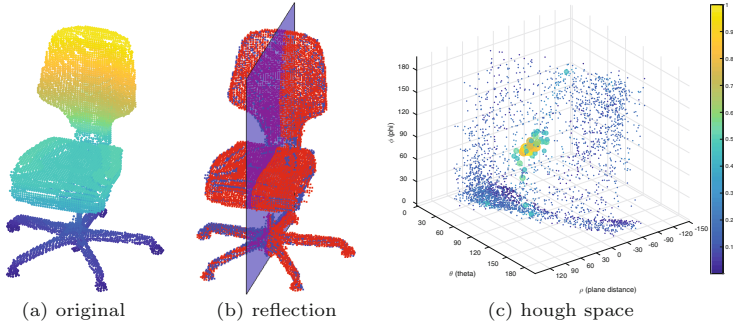


Fig. 2. Hough transform example: (a) Shows a 3D scan of a real chair. (b) Illustrates the planar symmetry reflection (red) of the high-gradient voxels (blue). Finally, (c) is the Hough space obtained by *importance sampling* in a Monte Carlo framework, as described in [23]. (Color figure online)

and the fact that it doesn’t require a non-maximal suppression step. However, knowing the Hough space can be a handy tool to visually understand the type of symmetries obtained in the process.

One should notice that there are more sophisticated methods in the literature to better extract symmetries, for instance [15,26]. However, the main contribution of the paper is the use of these detected symmetries in a variational optimization framework, described in the following section.

3 Surface Reconstruction with a Symmetry Prior

Given the set of detected symmetries Γ as described in Sect. 2, the goal is to find the 3D reconstruction of a surface which fulfills three simultaneous conditions: it should interpolate the given depth data, align with the given symmetries, and adhere to the defined degree of smoothness. We represent the surface by the implicit binary labeling function $u : V \subset \mathbb{R}^3 \rightarrow \{0, 1\}$. The depth information is fused into the data cost $f : V \rightarrow \mathbb{R}$ which encodes the depth measurements volumetrically as a truncated signed distance function similar as in [29]. Similarly to this work, we integrate the data cost on the entire ray from the measured depth towards the camera. This has the advantage of being able to directly identify unexplored areas within the data cost, since unseen voxels will keep their initial value $f = 0$. In this way, a zero value in the computed TSFD shows no preference towards occupied or free-space, thus implying an unobserved voxel. On the other hand, this approach has the disadvantage that outliers in the depth map carve incorrect holes into the aggregated cost. Using the symmetry distance from Eq. (1), the surface can then be found as the minimizer of the following energy

$$E(u) = \int_V (|Du| + \lambda fu) dx + \mu \sum_{\gamma \in \Gamma} \omega_\gamma \cdot \text{SD}_{V_\gamma}^2(u, \gamma), \tag{3}$$

in which $\lambda, \mu \in \mathbb{R}_{\geq 0}$ respectively weigh the contributions of each term to control the amount of surface smoothness and symmetry. The first term is the Total Variation regularizer in which D is the derivative in a distributional sense. The last term of Eq. (3) minimizes the distance of all symmetries given by Γ for which we use a slightly modified distance measure (Eq. (1)) in such a way that the distance is only evaluated for points within the corresponding support domain V_γ . Furthermore, the weights ω_γ can be used to change the impact among individual symmetries.

Symmetry Projection. For a given set of m symmetries $\Gamma = \{\gamma_1, \gamma_2, \dots, \gamma_m\}$, minimizing energy (3) approximates the joint projection onto a set of symmetries

$$\Pi_\Gamma(u) = \Pi_{\gamma_1, \gamma_2, \dots, \gamma_m}(u) = \min_{v: \forall \gamma_i \in \Gamma: \gamma_i(v) = v} \|u - v\|, \quad (4)$$

and if the regularization is turned off, i.e. for $\lambda, \mu \rightarrow \infty$ and uniform weights ω_γ this projection corresponds to the minimizer of (3), i.e. $\Pi_\Gamma(u) = \arg \min_u E(u)$. This relation can be seen by replacing the and-condition over the symmetries in Eq. (4) by the sum of the costs which leads to a similar expression as in Eq. (3).

Remark. As shown in [7] and stated in Eq. (1), the projection $\Pi_\gamma(u)$ of function u onto a single symmetry has a simple analytic solution: being the average of u and $\gamma(u)$. This is not the case for the projection onto a set of symmetries. Figure 3 illustrates that even the projection $\Pi_{\gamma_1, \gamma_2}(u)$ onto only two symmetries is not a simple combination of the individual projections $\Pi_{\gamma_1}(u), \Pi_{\gamma_2}(u)$. Since our method minimizes the symmetry distance to an arbitrary set of symmetries, the minimizer of energy (3) for an infinite symmetry term weight ($\mu \rightarrow \infty$) approaches the joint projection Π_Γ onto all symmetries in Γ which inherently generates a complex symmetry group as a combination of the input symmetries.

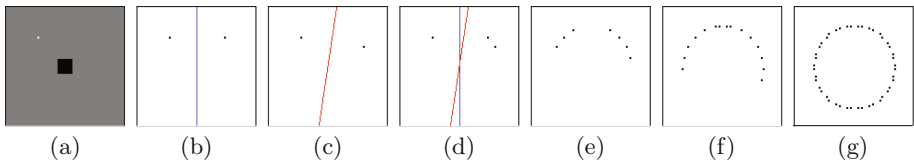


Fig. 3. A toy example illustrating the projection onto multiple symmetries. (a) Shows data cost f which enforces a single point (white corresponds to $f < 0$) as being occupied, and a small region to be free-space (black $\hat{=}$ $f > 0$). The rest of the image has no data cost (gray $\hat{=}$ $f = 0$). (b) A planar-reflective symmetry γ_1 and corresponding projection $\Pi_{\gamma_1}(u)$. (c) $\gamma_2, \Pi_{\gamma_2}(u)$. (d) Overlaid solution of $\Pi_{\gamma_1}(u), \Pi_{\gamma_2}(u)$ - this image corresponds to a solution of Eq. (3) with a low weight μ for the symmetry term. (d)–(g) In our setup we can continuously steer the amount of enforced symmetry by increasing μ until the image eventually fully adheres to the group structure of both symmetries.

4 Numerical Optimization

In order to minimize Eq. (3) we discretize the volume domain V on a regular voxel grid. In a discrete setting u is a stacked vector of all voxels in the volume domain. The symmetry dependencies between different points can be represented as a linear transformation, that is, the symmetry distance in Eq. (3) can be rewritten as $\text{SD}^2(u, \gamma) = \|A_\gamma u - b\|^2$. As a result, all three terms of the functional are convex and we can efficiently minimize energy (3) with the preconditioned first-order primal-dual algorithm [21]. In our setting the algorithm alternately iterates the following projected gradient ascent, gradient decent and linear extrapolation steps.

$$\begin{aligned} p^{n+1} &= \Pi_{\|p\| \leq 1} [p^n + \sigma D \bar{u}^n] \\ u^{n+1} &= \Pi_{[0,1]} \left[u^n + \tau \left(\text{div } p^{n+1} - \lambda f - \mu \sum_{\gamma \in \Gamma} \omega_\gamma A_\gamma^T (A_\gamma u^n - b) \right) \right] \\ \bar{u}^{n+1} &= 2u^{n+1} - u^n \end{aligned} \quad (5)$$

with $\Pi_{\|p\| \leq 1}[p] = p / \max(1, \|p\|)$ being the projection onto the unit ball and $\Pi_{[0,1]}(x) = \max(0, \min(1, x))$ being a simple clamping. The derivative and divergence operators are discretized by forward and backward differences, respectively.

The primal-dual surface optimization lends itself to a parallel implementation for which we used the CUDA framework. Without further processing, we finally extract a mesh as the 0.5 iso-surface from the implicit surface representation u using the Marching cubes [14] algorithm.

Although the most common types of symmetries, like planar reflective symmetries, rotational, and translational symmetries are represented by linear transformations, the convexity of the energy is also not violated if the symmetry transformation is non-linear. This is because the transformation affects only the argument x of u , but not u itself. Hence, our formulation can handle any type of symmetry. In this work, we focus, however, on planar reflective symmetries since these are the most common ones in many environments.

Planar Reflective Symmetries. In this case the symmetry γ is parametrized by a 3D plane given in Hessian normal form by $nx - d = 0$ with unit normal $n \in \mathbb{R}^3$ and the distance from the origin $d \in \mathbb{R}$. The linear transformation inside the symmetry term in Eq. (5) is then given by $b = 0$, $A_\gamma = (\mathbb{I} - M_\gamma)$, with \mathbb{I} being the identity matrix. The binary matrix $M_\gamma \in \{0, 1\}^{|V| \times |V|}$ (with $|V|$ being the number of voxels in the discretized volume domain) encodes all pairwise dependencies between voxels being linked by the symmetry γ and is defined as

$$(M_\gamma)_{ij} = \begin{cases} 1 & \text{if } m(j) = m(i) - 2n(m(i)^T n - d) \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

Here function $m : \mathbb{Z} \rightarrow \mathbb{Z}^3$ converts between the stacked 1D voxel index and the corresponding 3D voxel index. Further, pairwise dependencies (non-zero entries

in matrix M_γ) are only added for voxels inside the corresponding symmetry support domain V_γ . In sum, for the case of planar reflective symmetries the relations in Eq. (6) essentially add pairwise interactions according to the symmetry for all points in the symmetry support domain V_γ .

5 Experiments

Although our approach inherently combines surface denoising, completion and hallucination, we try to isolate and evaluate these properties separately in the following experiments.

5.1 Surface Denoising and Completion

In order to evaluate the denoising properties of our approach, we artificially degraded the 3D reconstruction of a building by dropping every second depth map from the original reconstruction. These depth maps were created by a plane sweep stereo matching approach from a sequential set of 118 images taken from all sides of the building. To improve the degraded model, we used the symmetry prior with the best scoring vertical symmetries in the scene, see Fig. 4.

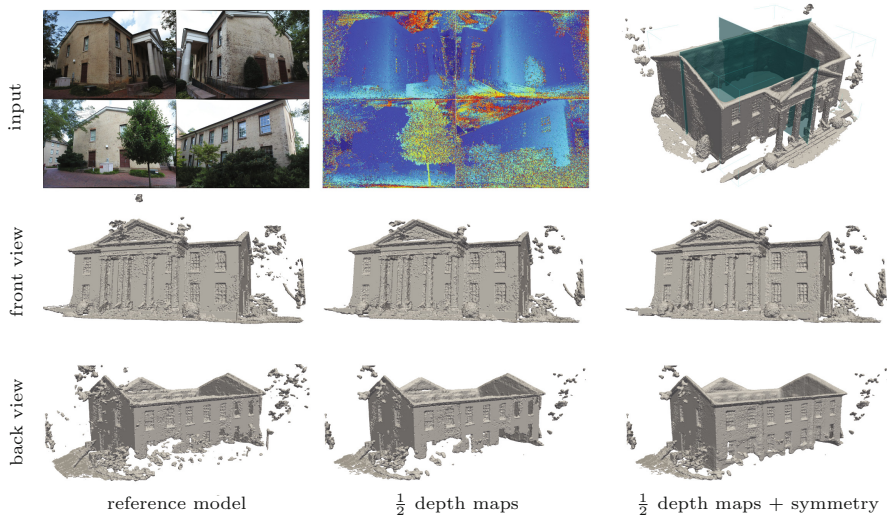


Fig. 4. Reconstruction of the building “g-hall” with large occlusions in the input data due to vegetation. The top row shows example input images, corresponding depth maps and the two best scoring planar symmetries. For comparison purposes, we degraded the original reconstruction (left column) by taking out half of the depth maps (center column) in order to introduce noise and missing data. Applying a symmetry prior to the degraded model completes many occluded areas on the backside, reduces the noise and enhances details like the window frames.

5.2 Surface Hallucination

In Fig. 5 we demonstrate the ability of our method to hallucinate large parts of a scene in unexplored areas. To this end, we took the Capitol data set from [1] consisting of a large building captured with 359 images. For the experiment, we selected one of the connected components with 129 images of the data set, representing half of the building as shown in the center column of Fig. 5. Using the two best scoring planar reflective symmetries, we were able to hallucinate the other half of the building and compared the result to the full stitched model presented as a result in [1]. An overlay of the two reconstructions reveals that the center part of this building is actually not exactly symmetric. Naturally, surrounding objects like the stairs do not fit to the reflected surface, but it is also visible that large parts of the walls and details like the windows align well.

5.3 Local Symmetries

Figure 6 shows an experiment with multiple local symmetries of objects in a desk scene. The local symmetries were detected using the RANSAC approach with sliding search boxes as described in Sect. 2.1. The first row in Fig. 6 depicts the input data which was scanned with a structured light RGB-D sensor and is missing surface information for several unobserved scene parts. The baseline reconstruction without the symmetry term ($\mu = 0$) is shown in the second row of Fig. 6. Some of its areas are filled with minimal surfaces as visible in the backside of the monitor and the lower part of the office chair. In contrast, the results with our approach (third row of Fig. 6) demonstrate that these areas are filled in a more meaningful way with the help of the symmetry information.

Another advantage of our approach is the combination of the symmetry prior and the classical minimal surface smoothness prior. While the lower part of the chair and the backside of the monitor clearly show the benefit of the symmetry term while regularizing the surface, the impact of the total variation term is not clear in these cases. Nevertheless, its impact is still important as it helps to denoise the surface and close smaller holes. For instance, the input surface information of the monitor stand is not sufficient to fully reconstruct it without holes by solely using the symmetry prior. As highlighted in Fig. 8, the combination of the two priors unifies their desired properties and yields superior reconstruction results in comparison to using only one of the priors alone.

5.4 Discussion and Limitations

While the proposed symmetry detection is rather robust to noise, the subsequent surface optimization with the proposed symmetry prior is sensitive even to very small changes of the symmetry parameters: (1) The accuracy of the symmetry support (up to surface noise) is essential, because otherwise inconsistent scene parts are forced to take consistent occupancy labels and the resulting label depends on the data support, the smoothness parameter and the number of pairwise inconsistencies within the group structure of the favored symmetry.

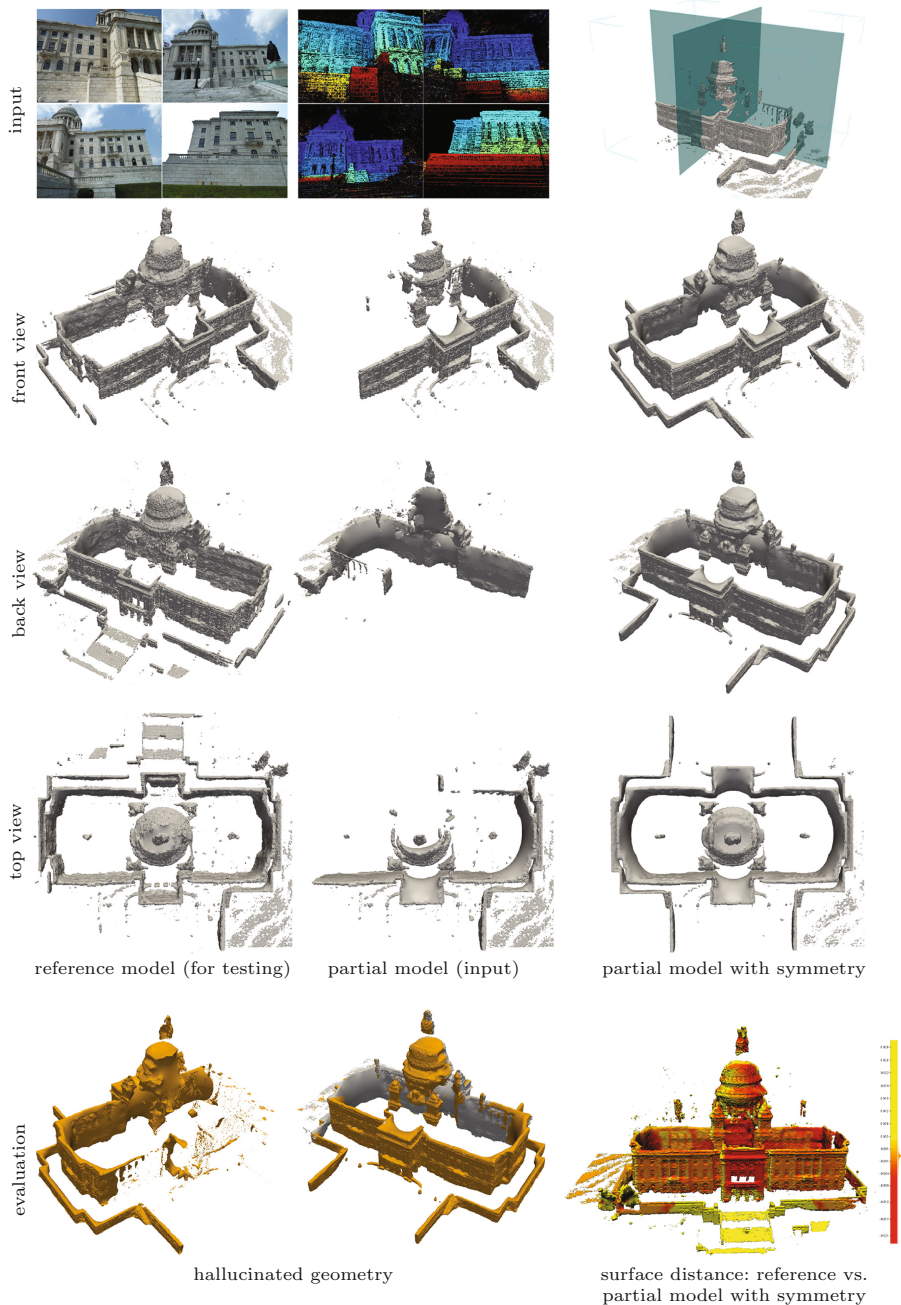


Fig. 5. This experiment shows a symmetric reconstruction of the capitol building in Providence (Rhode Island, USA) with a large unexplored area. We created a partial model by leaving out depth maps and then detected symmetries and hallucinated the rest of the building into the unexplored area.

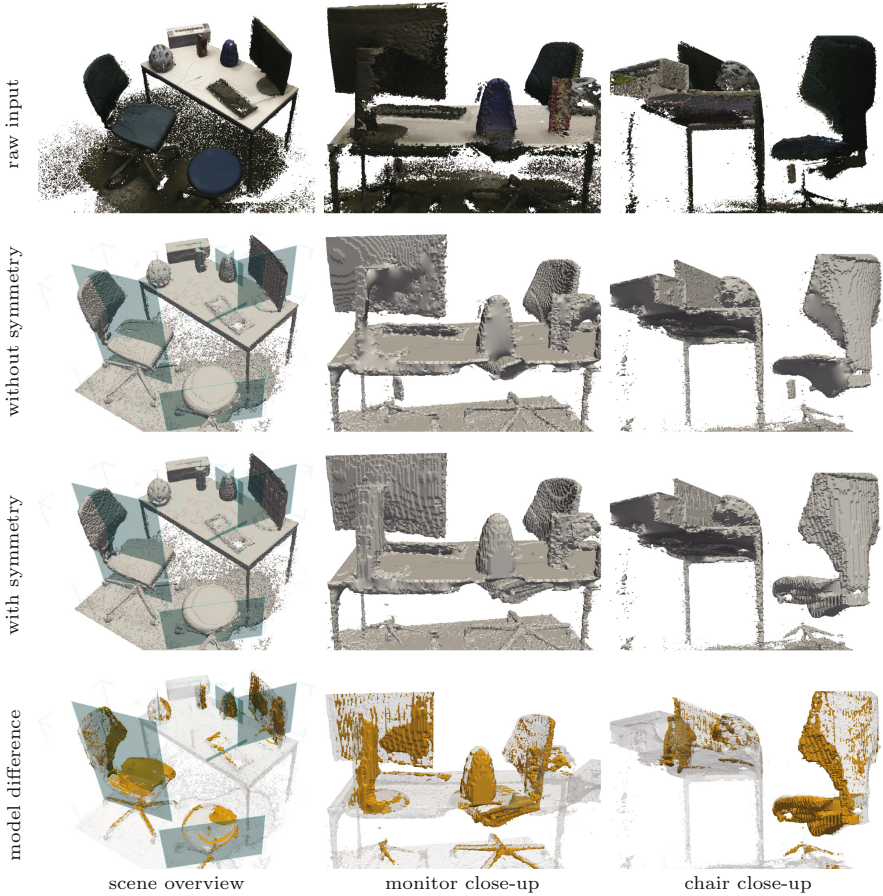


Fig. 6. Experiment with a desk scene containing multiple local planar symmetries shown as teal-colored planes in the first column. Several scene parts such as the backside of the monitor and the lower part of the chair were occluded during data acquisition. The figure shows the reconstruction results of the baseline approach (second row) which fills-in a minimal surface into the unknown regions in comparison to reconstruction with the proposed symmetry prior (third row) which completes these regions while obeying to the previously detected symmetries.

(2) The accuracy of the symmetry plane normal has strong influence on the result, because even small angular errors lead to large reconstruction errors for points that are far away from the symmetry plane as explained in Fig. 7.

Parameter Settings. We experimented with several settings for the individual symmetry weights ω_γ , like e.g. the support size, inlier ratios, or number of inliers, following the idea that symmetries with a stronger data support should also be enforced in a stronger way. However, we found that uniform weights ($\forall \gamma : \omega_\gamma = 1$) gave the best results and used them in all of our experiments.

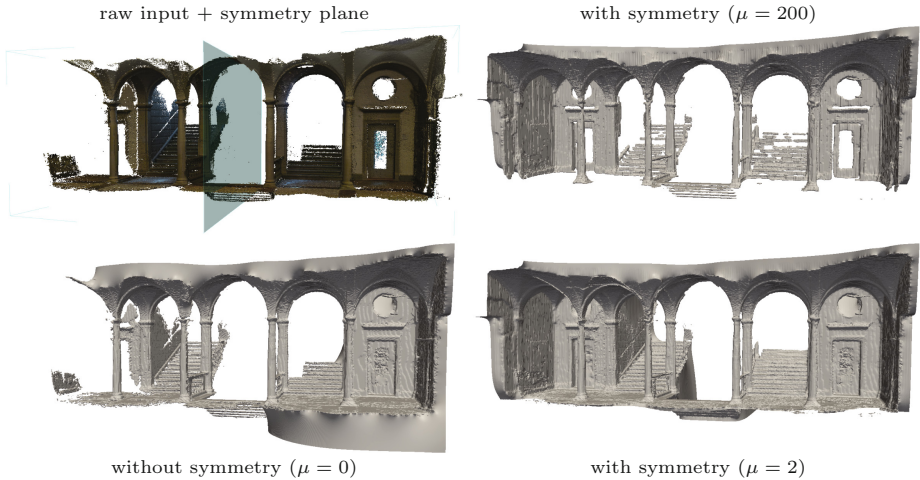


Fig. 7. Partial scan of a symmetric hallway. The figure shows different values of symmetry term weights μ in comparison. Even the very small angular error in the symmetry plane leads to inconsistencies in surface areas like the floor and walls and makes them disappear if the symmetry is enforced strongly ($\mu = 200$). Our approach allows for a compromise by only putting a low symmetry weight ($\mu = 2$), which leads to a not fully symmetric scene, but allows to recover the floor and all the walls in this dataset.

The other model parameters (volume resolution, data term weight λ , and symmetry term weight μ) are summarized in Table 1. The choice of these parameters is intuitive and careful tuning was not necessary. In order to weigh the amount of smoothness, data fidelity and symmetry against each other, we found as a rule of thumb that the data term weight and the symmetry term weight should be changed in a similar manner in order to keep a similar amount of smoothness, e.g. when enforcing more symmetry, the data term weight should be raised as well to maintain the amount of data fidelity. Conversely, for less smoothness the data term weight is raised and the symmetry term weight needs to be raised as well in order to maintain a comparable impact of the symmetry term. For more sophisticated symmetries like the stool or the toy example in Fig. 3, large symmetry term weights are required to enforce the full group structure.

For the detection of the global symmetries, we achieved better results and shorter runtimes on the larger data sets with the Hough transform, while for the local symmetry detection, the RANSAC approach gave better symmetry proposals. As mentioned in Sect. 2, we use a sliding box approach (3D convolution) with different box sizes to better detect all the objects in the scene. We chose sliding boxes of sizes 20%, 30%, 40%, and 60% of the total volume size and moved them by quarter box length quantities. For the symmetry support regions V_γ , we simply took the tight bounding box of the inlier points of the symmetry, but a better symmetry-based segmentation could also be used [26].

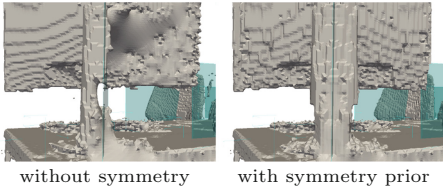


Fig. 8. Benefit of combining the minimal surface and the symmetry prior: less than half of the monitor stand was observed, the symmetry prior alone would leave a hole in the stand, but the combination of the priors yields the desired result.

Table 1. Overview of parameters for all experiments. The first two columns show the voxel resolution next to the corresponding dataset and the last two columns show the data term weight λ and the symmetry term weight μ (see Eq. (3)).

Dataset	Voxel resolution $ V $	λ	μ
Stool	$191 \times 158 \times 135$ (4M)	10^8	10^7
Capitol	$700 \times 560 \times 560$ (220M)	1.5	6
g-hall	$512 \times 512 \times 213$ (56M)	0.75	6
Desk	$311 \times 289 \times 239$ (21M)	1500	300
Hallway	$730 \times 436 \times 301$ (96M)	200	200

Computation Time. All results have been computed on a Linux-based i7 CPU with a GeForce GTX 780 graphics card. While the symmetry detection is performed by unoptimized CPU code, the surface optimization with a symmetry prior has been implemented on the GPU. The computation times for the symmetry detection vary depending on the dataset and are within seconds for simple scenes like the stool and up to 2 h for the sliding box approach on the desk dataset. The computation time for the surface optimization with and without symmetry constraints with 80M voxels were 2:20 min and 1:03 min, respectively. This highly depends on the grid resolution, e.g., the reconstruction of the stool in Fig. 1 with 4M voxels took only 5 s with- and 2 s without symmetry constraints.

6 Conclusion

In this paper we proposed a novel symmetry prior for variational 3D reconstruction. Our method is able to enforce several symmetries with the same support area, allowing for symmetry interactions that proved to be very useful for several applications such as surface denoising and completion, as well as surface hallucination in the case of highly incomplete data. We also discussed two planar symmetry detection approaches and how to handle and exploit unobserved areas in order to more robustly detect such symmetries. We showed the results of our method in several datasets, ranging from noisy and slightly incomplete reconstructions, to models with almost half of its surface missing. We also showed results on scenes with several local symmetries with different support sizes. In future work, we would like to explore better symmetry detection methods, as well as experiment with different kinds of symmetries, such as rotational, translational or curved symmetries (e.g. as in [12]).

Acknowledgments. This work was supported by the Horizon 2020 research and innovation programme under grant agreement No. 637221, and by the CTI Switzerland grant No. 17136.1 Geometric and Semantic Structuring of 3D point clouds.

References

1. Cohen, A., Sattler, T., Pollefeys, M.: Merging the unmatched: stitching visually disconnected SFM models. In: ICCV, December 2015
2. Cohen, A., Zach, C., Sinha, S., Pollefeys, M.: Discovering and exploiting 3D symmetries in structure from motion. In: CVPR, June 2012
3. Curless, B., Levoy, M.: A volumetric method for building complex models from range images. In: SIGGRAPH (1996)
4. Golovinskiy, A., Podolak, J., Funkhouser, T.: Symmetry-aware mesh processing. In: Hancock, E.R., Martin, R.R., Sabin, M.A. (eds.) *Mathematics of Surfaces XIII*. LNCS, vol. 5654, pp. 170–188. Springer, Heidelberg (2009)
5. Häne, C., Zach, C., Cohen, A., Angst, R., Pollefeys, M.: Joint 3D scene reconstruction and class segmentation. In: CVPR, pp. 97–104 (2013)
6. Hauage, D.C., Snavely, N.: Image matching using local symmetry features. In: CVPR, pp. 206–213 (2012)
7. Kazhdan, M.M., Chazelle, B., Dobkin, D.P., Funkhouser, T.A., Rusinkiewicz, S.: A reflective symmetry descriptor for 3D models. *Algorithmica* **38**(1), 201–225 (2003)
8. Kerber, J., Wand, M., Krüger, J.H., Seidel, H.: Partial symmetry detection in volume data. In: *Proceedings of the Vision, Modeling, and Visualization Workshop*, Berlin, Germany, 4–6 October 2011, pp. 41–48 (2011)
9. Kolev, K., Klodt, M., Brox, T., Cremers, D.: Continuous global optimization in multiview 3D reconstruction. *IJCV* **84**(1), 80–96 (2009)
10. Kolev, K., Pock, T., Cremers, D.: Anisotropic minimal surfaces integrating photoconsistency and normal information for multiview stereo. In: *ECCV*, Heraklion, Greece, September 2010
11. Köser, K., Zach, C., Pollefeys, M.: Dense 3D reconstruction of symmetric scenes from a single image. In: Mester, R., Felsberg, M. (eds.) *DAGM 2011*. LNCS, vol. 6835, pp. 266–275. Springer, Heidelberg (2011)
12. Liu, J., Liu, Y.: Curved reflection symmetry detection with self-validation. In: Kimmel, R., Klette, R., Sugimoto, A. (eds.) *ACCV 2010, Part IV*. LNCS, vol. 6495, pp. 102–114. Springer, Heidelberg (2011)
13. Liu, Y., Hel-Or, H., Kaplan, C.S., Gool, L.V.: Computational symmetry in computer vision and computer graphics. *Found. Trends Comput. Graph. Vis.* **5**(12), 1–195 (2010)
14. Lorensen, W.E., Cline, H.E.: Marching cubes: a high resolution 3D surface construction algorithm. *SIGGRAPH Comput. Graph.* **21**, 163–169 (1987)
15. Mitra, N.J., Guibas, L.J., Pauly, M.: Symmetrization. *ACM Trans. Graph.* **26**(3), 63 (2007)
16. Mitra, N.J., Pauly, M., Wand, M., Ceylan, D.: Symmetry in 3D geometry: extraction and applications. *Comput. Graph. Forum* **32**(6), 1–23 (2013)
17. Oswald, M.R., Cremers, D.: A convex relaxation approach to space time multi-view 3D reconstruction. In: *ICCV - Workshop on Dynamic Shape Capture and Analysis (4DMOD)* (2013)
18. Oswald, M.R., Cremers, D.: Surface normal integration for convex space-time multi-view reconstruction. In: *Proceedings of the British Machine and Vision Conference (BMVC)* (2014)
19. Oswald, M.R., Stühmer, J., Cremers, D.: Generalized connectivity constraints for spatio-temporal 3D reconstruction. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014, Part IV*. LNCS, vol. 8692, pp. 32–46. Springer, Heidelberg (2014)

20. Pizlo, Z., Sawada, T., Li, Y., Kropatsch, W.G., Steinman, R.M.: New approach to the perception of 3D shape based on veridicality, complexity, symmetry and volume. *Vision. Res.* **50**, 1–11 (2010)
21. Pock, T., Chambolle, A.: Diagonal preconditioning for first order primal-dual algorithms in convex optimization. In: *ICCV*, Washington, DC, USA, pp. 1762–1769 (2011)
22. Podolak, J., Golovinskiy, A., Rusinkiewicz, S.: Symmetry-enhanced remeshing of surfaces. In: *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*, Barcelona, Spain, 4–6 July 2007, pp. 235–242 (2007)
23. Podolak, J., Shilane, P., Golovinskiy, A., Rusinkiewicz, S., Funkhouser, T.A.: A planar-reflective symmetry transform for 3D shapes. *ACM Trans. Graph.* **25**(3), 549–559 (2006)
24. Reinbacher, C., Pock, T., Bauer, C., Bischof, H.: Variational segmentation of elongated volumetric structures. In: *CVPR* (2010)
25. Terzopoulos, D., Witkin, A., Kass, M.: Symmetry-seeking models and 3D object reconstruction. *IJCV* **1**, 211–221 (1987)
26. Thrun, S., Wegbreit, B.: Shape from symmetry. In: *ICCV*, Beijing, China. *IEEE* (2005)
27. Ummenhofer, B., Brox, T.: Dense 3D reconstruction with a hand-held camera. In: Pinz, A., Pock, T., Bischof, H., Leberl, F. (eds.) *DAGM and OAGM 2012*. LNCS, vol. 7476, pp. 103–112. Springer, Heidelberg (2012)
28. Unger, M., Pock, T., Cremers, D., Bischof, H.: TVSeg - interactive total variation based image segmentation. In: *Proceedings of the British Machine and Vision Conference (BMVC)*, Leeds, UK, September 2008
29. Zach, C., Pock, T., Bischof, H.: A globally optimal algorithm for robust TV-l1 range image integration. In: *ICCV*, pp. 1–8 (2007)