

Partial and Approximate Symmetry Detection for 3D Geometry

Niloy J. Mitra
Stanford University

Leonidas J. Guibas
Stanford University

Mark Pauly
ETH Zürich



Figure 1: Symmetry detection on a sculpted model. From left to right: Original model, detected partial and approximate symmetries, color-coded deviations from perfect symmetry as a fraction of the bounding box diagonal.

Abstract

“Symmetry is a complexity-reducing concept [...]; seek it everywhere.” - Alan J. Perlis

Many natural and man-made objects exhibit significant symmetries or contain repeated substructures. This paper presents a new algorithm that processes geometric models and efficiently discovers and extracts a compact representation of their Euclidean symmetries. These symmetries can be partial, approximate, or both. The method is based on matching simple local shape signatures in pairs and using these matches to accumulate evidence for symmetries in an appropriate transformation space. A clustering stage extracts potential significant symmetries of the object, followed by a verification step. Based on a statistical sampling analysis, we provide theoretical guarantees on the success rate of our algorithm. The extracted symmetry graph representation captures important high-level information about the structure of a geometric model which in turn enables a large set of further processing operations, including shape compression, segmentation, consistent editing, symmetrization, indexing for retrieval, etc.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling.

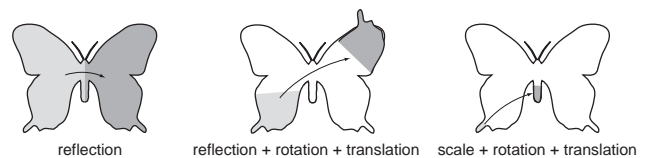
Keywords: geometric modeling, shape analysis, symmetry detection, shape descriptor, sampling guarantees.

1 Introduction

Symmetry is an essential and ubiquitous concept in nature, science, and art. For example, in geometry, the Erlanger program of Felix Klein [1893] has fueled for over a century mathematicians’ interest in invariance under certain group actions as a key principle for

understanding geometric spaces. Numerous biological, physical, or man-made structures exhibit symmetries as a fundamental design principle or as an essential aspect of their function. Whether by evolution or design, symmetry implies certain economies and efficiencies of structure that make it universally appealing. Symmetry also plays an important role in human visual perception and aesthetics. Arguably much of the understanding of the world around us is based on the perception and recognition of shared or repeated structures, and so is our sense of beauty [Thompson 1961].

In this paper we present a novel method for detecting meaningful symmetries in digital 3D shapes. We understand symmetry as the invariance under a set of transformations — in our case translation, rotation, reflection, and uniform scaling, the common generators of the Euclidean group. The figure below shows a 2D illustration. As can be seen in this example, symmetries or congruences that are quite apparent to us can be approximate and occur at different scales. Our goal is to define an algorithm that extracts (partial) symmetries at all scales, including approximate or imperfect symmetries of varying degree. This allows the user to select the subset of symmetries that are most meaningful for a specific application. Examples include scan registration and alignment, shape matching, segmentation and skeleton extraction, compression, advanced modeling and editing, and shape database retrieval.



To achieve this goal, we separate the symmetry computation into two phases: In the first step, we compute simple local shape descriptors at a selected set of points on the shape. These descriptors are chosen so that they are invariant under the group actions of interest. We use these local descriptors to pair up points that could be mapped to each other under a candidate symmetry action. We think of each such pair as depositing mass, or voting, for a specific symmetry in the transformation space of interest. In this space, pairs with similar transformations form clusters that provide evidence for the corresponding symmetry relation.

In the second step we use a stochastic clustering algorithm to extract the significant modes of this mass distribution. Since the mapping to transformation space does not preserve the spatial coherence or structure of samples on the input shape, we verify whether a meaningful symmetry has been found by checking the spatial consistency of the extracted subparts of the surface. Our clustering method provides the necessary surface correspondences, since every point mass in transformation space corresponds to a candidate pair of points in the spatial domain. Thus only a small set of candidates samples needs to be considered when detecting and extracting symmetric surface patches, avoiding a costly quadratic spatial search over the whole input data set.

This separation into two stages is crucial for the effectiveness of our algorithm. The underlying observation is the following: given a proposed symmetry relation, it is simple and efficient to verify whether this specific symmetry is present in the model; we just need to apply the symmetry transform and check whether the model is mapped onto itself, or a sub-part of the model is mapped to a corresponding sub-part. However, the number of all potential mappings is by far too large to do an exhaustive search. Therefore, we first accumulate statistical evidence for which symmetries are present via our clustering in transformation space. Only if this evidence is sufficient do we perform spatial verification to check whether a specific symmetry is actually valid. Thus the complexity of symmetry extraction depends primarily on the number and size of relevant symmetries present in the model and not on the complexity of the model itself or that of the underlying symmetry group. As part of our approach, we can provide a quantitative measure on the “exactness”, or saliency, of a symmetry relation, which allows the user to control the degree of perfection in the extracted symmetries. In addition, by specifying the size of the set of local shape descriptors, the user can trade accuracy for computational efficiency. While fewer samples are sufficient for detecting large global symmetries, small partial symmetries require a significantly denser sampling.

The final output of our algorithm is a “symmetry graph” of the object, which encodes the significant symmetries of the object, each described by a patch pair and the corresponding transformation between them. For objects that contain regular repeated structures, like windows or doors in architectural models, we can recover the symmetries of the repetition pattern through a basis reduction algorithm. This in effect leads to a sparser and more informative symmetry graph that contains only fundamental symmetry generators and avoids encoding separately symmetries that are just products of already recorded symmetries. This kind of repeated pattern discovery can be useful in consistent mesh editing applications.

1.1 Contributions

We propose a new algorithm for pairing sample points on 3D shapes with compatible local descriptors to generate a distribution in transformation space whose peaks capture relevant symmetries of the object. We show how a stochastic clustering algorithm over this distribution detects potential symmetry candidates, and provide a surface patching method that extracts a reduced symmetry graph from the extracted clusters. Our algorithms can be applied to 3D models of different shape characteristics and representations. Memory requirements are minimal and the computation is output-sensitive in the sense that its complexity depends mainly on the number and extent of symmetries actually present in the object. In addition, we provide theoretical bounds on the success rate of our algorithm as a function of the number of initial samples selected. These results indicate that the algorithm can be effective even for very large models that cannot fit in main memory.

1.2 Related work

The problem of symmetry detection has been extensively studied in numerous fields including visual perception, computer vision, robotics, and computational geometry. Early methods concentrated on finding perfect symmetries in 2D or 3D planar point sets [Atallah 1985], [Wolter et al. 1985]. Since the restriction to exact symmetries limits the use of these methods for real-world objects, Alt et al. [1988] introduced a method for computing approximate global symmetries in 3D point sets, but the complexity of the algorithm makes it impractical for large data sets. Zabrodsky et al. [1995] formalized the notion of approximate symmetry by expressing symmetry as a continuous feature. Sun et al. [1997] proposed to examine the correlation of the Gaussian image to recover global reflective and rotational symmetries. Kazhdan and co-workers [2002] introduced a shape descriptor that concisely encodes global reflective symmetries. Later they extended this work to rotational symmetries and used it for shape retrieval for database matching in [Kazhdan et al. 2004].

Our method bears some similarity to the Hough transform, a popular feature extraction method mainly used in image processing [Hough 1959]. Starting from a set of sample points obtained using edge detection, the method repeatedly selects small subsets of these samples to estimate the parameters of the feature curve. Analogous to our approach, votes cast by all of these estimates are accumulated and the final feature curve is extracted based on the majority of votes. Recently ideas based on the Hough transform have been used by [Loy and Eklundh 2006] to detect reflective and rotational symmetries in images.

The RANdom SAMple Consensus (RANSAC) method proposed by Fischler and Bolles [1981] is an algorithm for robust model fitting for data containing many outliers. In the context of shape matching the basic idea is to choose a random set of corresponding samples on the query and target shapes, apply the global transformation induced by these samples, and evaluate the matching error between the two shapes. If sufficiently many transformations are explored in this way, the relevant symmetries can eventually be determined. Since the evaluation of the matching error requires costly spatial proximity tests, geometric hashing [Lamdan and Wolfson 1988] pre-computes all possible alignments by densely sampling the space of transformations and storing the resulting shape distribution in a hash grid. Gal and Cohen-Or [2006] recently presented an effective method for shape matching based on this idea. Their algorithm computes local shape descriptors that are grouped to form salient shape features. Using an empirical saliency measure, shape features are then used to pre-compute a geometric hash table that allows efficient partial matching.

While sharing some similarities, our method is fundamentally different from both RANSAC and geometric hashing. We avoid the costly exhaustive search of the former by computing the matching error of a transformation only *after* we accumulate sufficient evidence for a symmetry. At the same time our method requires minimal storage, in contrast to geometric hashing, where hash tables of up to 3.5 GBytes have been reported for complex geometric shapes [Gal and Cohen-Or 2006].

1.3 Overview

We first give some intuition for our method by looking at the 2D example shown in Figure 3, where the goal is to detect reflective symmetries of the butterfly. Any pair of points (\mathbf{p}, \mathbf{q}) on the boundary of the model defines a unique reflection with respect to the bisector line through $(\mathbf{p} + \mathbf{q})/2$ with normal direction $\mathbf{p} - \mathbf{q}$. Hence

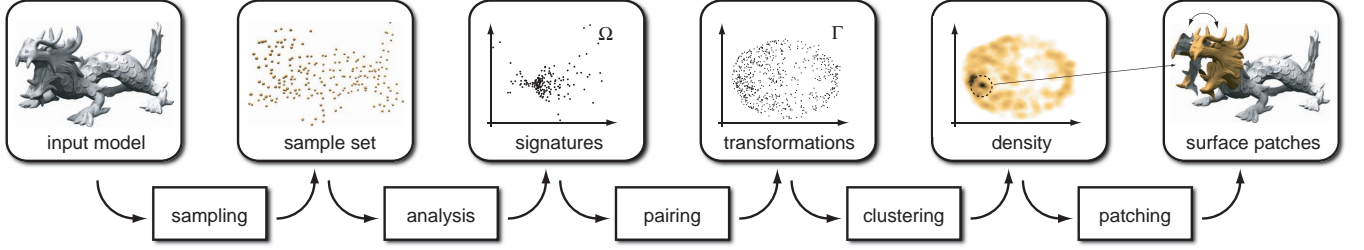


Figure 2: The symmetry extraction pipeline. Sampling yields a set P of surface points. For each $\mathbf{p}_i \in P$ a local signature is computed. Points $\mathbf{p}_i, \mathbf{p}_j$ with similar signatures are paired and a point in transformation space Γ is computed mapping the local frame of \mathbf{p}_i to the one at \mathbf{p}_j . Clustering in Γ yields subsets of P that remain invariant under a certain transformation, which can be extracted using spatial region growing.

such a pair can be understood as evidence for the existence of this specific reflective symmetry. By looking at all such pairs we can accumulate this evidence and extract the relevant symmetry relation(s). Only if many point pairs agree on (roughly) the same reflection line, do we have reason to believe that the corresponding symmetry is truly present in the model. Thus we can detect potential symmetries by looking at clusters of points in the space of transformations Γ , where each point corresponds to a specific reflection line. However, as shown in the illustration, the evidence of a single point pair is only reliable if the local geometry around the points is faithfully mirrored by the reflective transformation. This observation will allow us to significantly prune the set of all point pairs and avoid an exhaustive computation on a quadratic number of point pairs.

Since the mapping to Γ does not incorporate the spatial position of surface samples, pairs from unrelated parts of the object can be mapped to the same point in transformation space. Thus in a second phase we extract spatially coherent components of the model that are invariant under the extracted symmetry transformations. Using the point pair correspondences present in the cluster, we perform an incremental region growing algorithm to verify a specific symmetry. Figure 2 gives a high-level overview of our symmetry extraction pipeline. The following sections will elaborate on the individual stages and provide details of our approach.

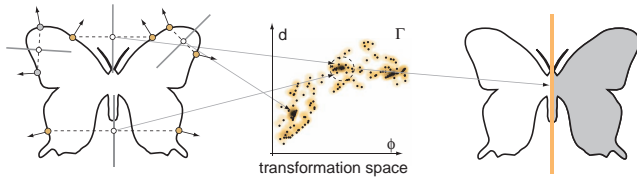
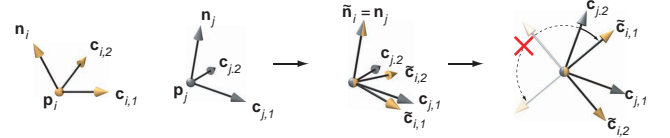


Figure 3: Illustration of symmetry detection for reflections. Every pair of points defines a symmetry line l that can be described by a distance d and an angle ϕ . Multiple points clustered in a small region in transformation space provide evidence of a symmetry. The pair on the top left is discarded due to normal inconsistency.

2 Signatures and Transformations

We consider the Euclidean transformation group generated by translations, rotations, reflections, and uniform scalings. Our goal is to find parts of a given 3D shape that are invariant under transformations in this symmetry group or some lower-dimensional subgroup.

In order to apply the ideas sketched above, we need to compute the transformation \mathbf{T}_{ij} that maps a point \mathbf{p}_i on the surface of the model to another point \mathbf{p}_j . While point positions are sufficient for defining a unique plane of reflection as in the example above, we cannot determine all degrees of freedom of a general Euclidean transform from the spatial positions alone. We therefore compute geometry signatures at each sample point \mathbf{p}_i based on the concept of normal cycles [Cohen-Steiner and Morvan 2003]. We apply the algorithm proposed in [Alliez et al. 2003] to approximate the curvature tensor at \mathbf{p}_i within a sphere of radius r and compute integrated principal curvatures $\kappa_{i,1} \leq \kappa_{i,2}$ and principal directions $\mathbf{c}_{i,1}$ and $\mathbf{c}_{i,2}$. The radius r should be on the order of the local sample spacing to achieve sufficient averaging when computing the curvature tensor and avoid a strong dependence on the specific location of the sample points.



The principal directions define a local frame $(\mathbf{c}_{i,1}, \mathbf{c}_{i,2}, \mathbf{n}_i)$, with normal vector $\mathbf{n}_i = \mathbf{c}_{i,1} \times \mathbf{c}_{i,2}$. We orient this frame as a right-handed coordinate frame that aligns with the outward pointing surface normal by flipping signs of the appropriate vectors if necessary. In order to obtain a canonical rotational component \mathbf{R}_{ij} of the transformation \mathbf{T}_{ij} we first align the two normals along their common plane and then pick the smaller of the two rotations around the normal that aligns to one of the two possible choices of orientation in tangent space. The uniform scale component of \mathbf{T}_{ij} is estimated from the ratio of principal curvatures as $s_{ij} = (\kappa_{i,1}/\kappa_{j,1} + \kappa_{i,2}/\kappa_{j,2})/2$, the translation is computed as $\mathbf{t}_{ij} = \mathbf{p}_j - s_{ij}\mathbf{R}_{ij}\mathbf{p}_i$. For a given pair $(\mathbf{p}_i, \mathbf{p}_j)$ we thus obtain a point in 7-dimensional transformation space Γ as $\mathbf{T}_{ij} = (s_{ij}, R_{ij}^x, R_{ij}^y, R_{ij}^z, t_{ij}^x, t_{ij}^y, t_{ij}^z)$, where $R_{ij}^x, R_{ij}^y, R_{ij}^z$ are the Euler angles derived from \mathbf{R}_{ij} and $\mathbf{t}_{ij} = [t_{ij}^x \ t_{ij}^y \ t_{ij}^z]^T$. In order to handle reflections, we also compute the transformation obtained when reflecting the model about an arbitrary but fixed plane.

2.1 Point Pruning

A differential surface patch at umbilic points, i.e., those for which $\kappa_{i,1} = \kappa_{i,2}$, is invariant under rotations around the surface normal. Pairs involving such points and their signatures do not define a unique transformation, but trace out curves in transformation space, which may quickly camouflage meaningful symmetry clusters. To avoid clutter in transformation space, we discard these points from the sample set, i.e., we only consider points on the surface with distinct principal curvatures (and hence stable principal directions),

which give rise to a unique transformation when paired with another compatible point. Apart from making the symmetry clustering more robust, point pruning has the additional advantage of reducing computation time. We obtain the adaptive sample set by applying a threshold $\gamma < 1$ on the ratio of curvatures: $\mathbf{p}_i \in P$, if $|\kappa_{i,1}/\kappa_{i,2}| < \gamma$. We use $\gamma = 0.75$ for all examples in this paper.

2.2 Pairing

Given the reduced set of surface samples P and their signatures, we can now compute transformations for pairs of points in P . We select a random subset $P' \subset P$ and find all pairs $(\mathbf{p}', \mathbf{p})$ with $\mathbf{p}' \in P'$ and $\mathbf{p} \in P$ that provide evidence for a symmetry relation. In the Appendix we give theoretical bounds on the size of P and P' required to successfully find symmetries of a certain size.

As indicated above, the evidence of a selected point pair for a specific symmetry relation is only reliable, if a local surface patch around each point is invariant under a transformation from the considered symmetry group G . In the 2D illustration of Figure 3, for example, we can reject a pair, if the curvature estimates at both points differ too much, since curvature is invariant under reflection. To obtain an efficient pairing algorithm we map all samples to a signature space Ω and use the metric of that space to estimate the deviation from perfect invariance. Only point pairs that are close in Ω are considered as suitable candidates for a local symmetry relation, which avoids an exhaustive computation of a quadratic number of point pairs.

For the full 7-dimensional Euclidean group in 3D, the mapping from P to $\Omega_7 = \mathbb{R}$ is given as $\sigma_7(\mathbf{p}_i) = \kappa_{i,1}/\kappa_{i,2}$, since uniform scaling, rotation, and translation leave the ratio of principal curvatures unchanged. The sub-index 7 indicates the dimension of the symmetry group. For purely rigid transforms, we define $\sigma_6(\mathbf{p}_i) = (\kappa_{i,1}, \kappa_{i,2})$ with $\Omega_6 = \mathbb{R}^2$. We can now for a given sample $\mathbf{p}_i \in P'$ determine all suitable partners in P by performing a range query in Ω . Using standard spatial proximity data structures, e.g., a kd -tree, we can perform pairing in $O(n' \log n)$ time, where $n = |P|$ and $n' = |P'|$. If only reflections and/or translations are considered, we can additionally reject pairs based on the orientation of the local frames, as illustrated in Figure 3.

Figure 4 shows that pruning not only reduces the complexity of the clustering algorithm, but, even more importantly, avoids clutter in transformation space. By focusing only on locally consistent symmetry pairs, meaningful clusters are stably detected in Γ .

3 Clustering

The pairing computed in the previous stage provides us with a set of transformations that map local surface patches onto each other. Each pair thus provides evidence for a symmetry relation at the level of the local sample spacing. To extract meaningful symmetries at larger scales we need to accumulate this local evidence, i.e., find groups of pairs with a similar transformation that correspond to symmetric subsets of the model surface. This requires the definition of a distance metric in Γ , which is non-trivial, since scaling, rotation, and translation need to be combined in a single metric. We follow the approach of [Amato et al. 2000] and define the norm of a transformation $\mathbf{T} = (s, R_x, R_y, R_z, t_x, t_y, t_z) \in \Gamma$ as the weighted sum $\|\mathbf{T}\|^2 = \beta_1 s^2 + \beta_2 (R_x^2 + R_y^2 + R_z^2) + \beta_3 (t_x^2 + t_y^2 + t_z^2)$. The weights β_i allow to adjust the relative influence of the individual components of the transformation. In all our examples we set these weights so that a rotation by 180 degrees corresponds to a displacement of half

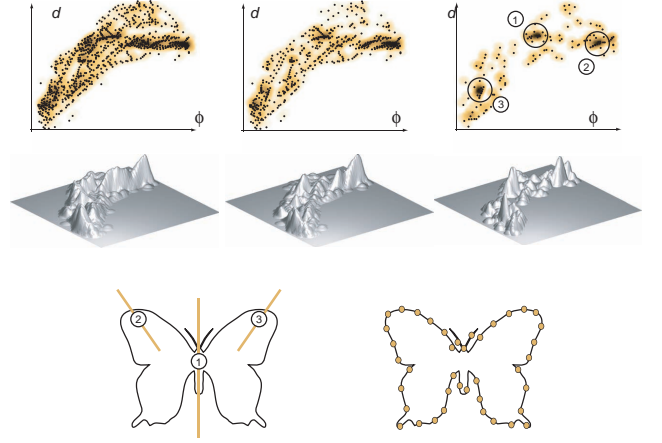


Figure 4: Pair pruning. 40 samples on the butterfly lead to $\binom{40}{2} = 780$ points in transformation space. Pruning based on curvature reduces the set to 503 points, while additionally normal-based pruning yields 138 points. The density plots show how the meaningful symmetry clusters become significantly more pronounced.

the bounding box diagonal and a scaling factor of 10. A metric for Γ can then be derived as $d(\mathbf{T}, \mathbf{T}') = \|\mathbf{T} - \mathbf{T}'\|$, where the subtraction is component-wise, see also [Hofer et al. 2004] for a detailed discussion.

3.1 Mean-Shift

If the symmetries in the model are perfect (and the sampling includes point pairs that are perfectly symmetric), then all pairs of the same (discrete) symmetry relation map to a single point in Γ . Many real-world objects exhibit approximate symmetries, however, and the sampling will not be precisely symmetric in general. We thus need a method to find clusters in transformation space. When looking at the distribution of points in Γ , we immediately see that standard parametric clustering methods, such as k -means clustering, are not suitable for our purposes. In general we have no a priori knowledge on the number of (partial) symmetries of the input model, i.e., selecting k would be difficult. Furthermore, clusters are not necessarily isotropic, especially for approximate symmetries like the ones shown in Figure 1. A more suitable clustering method is *mean shift clustering*, a non-parametric method based on gradient ascent on a density function ρ [Comaniciu and Meer 2002]¹. This density function is defined as a sum of kernel functions K centered at each point \mathbf{T}_i in Γ as

$$\rho(\mathbf{T}) = \sum_i K(\|\mathbf{T} - \mathbf{T}_i\|/h).$$

We use the radially symmetric Epanechnikov kernel with bandwidth h as suggested in [Comaniciu and Meer 2002]. The significant modes of ρ are determined using gradient ascent. All points that flow into a local maximum of sufficient height are considered samples of a significant cluster C_k . The corresponding symmetry transformation \mathbf{T}_k is then defined by the cluster's maximum. Essentially, the algorithm can be understood as a voting scheme: Every point pair votes for the symmetry relation that has been extracted

¹Mean shift clustering has also been used in [James and Twigg 2005] for skinning mesh animations and in [Tuzel et al. 2005] for 3D motion estimation.

from its local frames. If many votes are cast for the same symmetry, a local peak is created in the accumulated density function. For more details on mean-shift clustering we refer to [Comaniciu and Meer 2002].

4 Verification

A significant mode detected by the mean-shift clustering algorithm does not necessarily correspond to a meaningful symmetry. Since the spatial relation of sample points is lost during the mapping to transformation space, sample pairs from uncorrelated parts of the object can accumulate to form discernible clusters. The effectiveness of our method is based on the observation that statistically such spurious modes are rare (see also the analysis in the Appendix): It is highly unlikely that many uncorrelated point pairs agree on the same transformation, i.e., are mapped to the same point in 7D transformation space. We can thus afford to perform a spatial verification for each cluster C_k by extracting the connected components of the model that are invariant under the corresponding transformation \mathbf{T}_k . We compute these surface patches using an incremental patch growing process, starting with a random point of C_k , which corresponds to a pair $(\mathbf{p}_i, \mathbf{p}_j)$ of points on the model surface. Now we look at the one-ring neighbors of \mathbf{p}_i , apply \mathbf{T}_k , and check whether the distance of the transformed points to the surface around \mathbf{p}_j is below a given error threshold. If so, we add them to the current patch. We keep extending this patch along its boundary until no more points can be added. During the growth process, we mark all visited samples on the surface and remove points in C_k that correspond to these samples. This process is then repeated using the next point in C_k until all points have been considered.

Since the transformation \mathbf{T}_k at the cluster's maximum does not necessarily provide the best possible transformation for matching the surface patches, we incrementally refine \mathbf{T}_k during the patch growing using the iterated closest points (ICP) algorithm [Rusinkiewicz and Levoy 2001]. The normalized residual of the ICP matching then provides a quantitative measure for the exactness of the symmetry [Mitra et al. 2004]. Other measures, such as the Hausdorff distance can also be used. We end up with a collection of pairs of patches on the model surface that are mapped onto each other by the cluster's transformation \mathbf{T}_k . This information can be encoded in a weighted graph, where each node corresponds to a patch and each edge denotes the transformation that maps two patches onto each other, weighted by the matching error.

4.1 Compound Transforms

Many geometric objects exhibit symmetries in a structured or repetitive fashion resulting in a large number of clusters in transformation space [Liu et al. 2004]. Encoding all pair-wise symmetry relations for such models leads to a complex and highly redundant symmetry graph and thus a costly verification stage. In this section we describe a simple basis reduction algorithm that computes a compact set of generators for all detected symmetries in transformation space [Magnus et al. 2004]. This significantly reduces the number of spatial consistency checks required for verification and yields a more informative symmetry graph that supports advanced editing operations and high level shape comparisons.

The algorithm shown below takes as input all extracted symmetry transformations T sorted in descending order of cluster height and iteratively processes each transformation $\mathbf{T}_i \in T$. During execution we maintain an alphabet A of generators and the language L that encodes T in terms of the alphabet A . A user parameter η controls

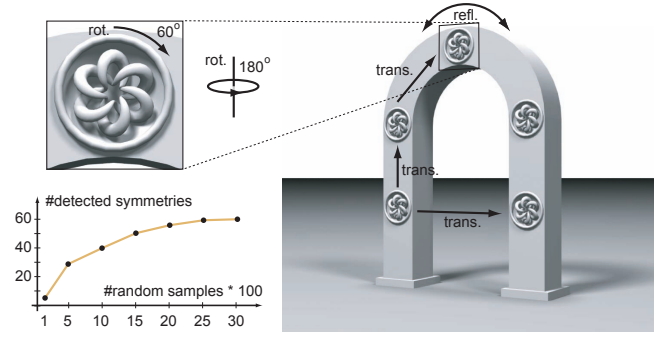


Figure 5: Symmetry graph reduction for a model with structured symmetries at two different scales. 60 significant modes have been extracted in the clustering stage. The reduced basis contains 6 transformations, as indicated by the arrows. The graph on the left shows the number of detected symmetries as a function of random samples in the subset $P' \subset P$.

the complexity of the algorithm by limiting the search to loops of length $\eta + 1$. The threshold δ measures the allowed deviation from the exact transformation.

Figure 5 shows an example of a reduced symmetry basis. Verification can now be applied more efficiently on the set L of compound transformations. For more details on basis reduction we refer to [Magnus et al. 2004].

Algorithm 1 Symmetry basis reduction.

```

Input:  $T = \{\mathbf{T}_1, \mathbf{T}_2, \dots, \mathbf{T}_n\}$ 
 $A \leftarrow \{\mathbf{I}\}$ 
 $L \leftarrow \emptyset$ 
for  $i = 1$  to  $n$  do
  if  $\exists (\mathbf{A}_1, \dots, \mathbf{A}_\eta)$  with  $\mathbf{A}_j \in A$  s.t.  $|\mathbf{T}_i - \prod_{j=1}^\eta \mathbf{A}_j| \leq \delta$  then
     $L \leftarrow L \cup \{(\mathbf{A}_1, \dots, \mathbf{A}_\eta)\}$ 
  else
     $A \leftarrow A \cup \{\mathbf{T}_i, \mathbf{T}_i^{-1}\}$ 
     $L \leftarrow L \cup \{(\mathbf{T}_i)\}$ 
  end if
end for

```

5 Results and Applications

We have implemented the pipeline sketched in Figure 2. An initial sample set is created by uniformly sampling the input model. After computing signatures using the method of [Alliez et al. 2003], point pruning yields the reduced sample set P . We then select a random subset $P' \subset P$, find all suitable pairs $(\mathbf{p}' \in P', \mathbf{p} \in P)$ based on the proximity in signature space, and compute the corresponding transformations. We perform mean-shift clustering using the method proposed in [Arya et al. 1998] to efficiently compute neighborhoods in 7D transformation space. Basis reduction and verification finally yield the symmetric patches. We show in the Appendix that our method is guaranteed to find existing symmetry relations provided the sampling is dense enough with respect to the size of the symmetric patches. The following examples verify this claim and demonstrate that practical results can be obtained even if the theoretical sampling requirements are not met.

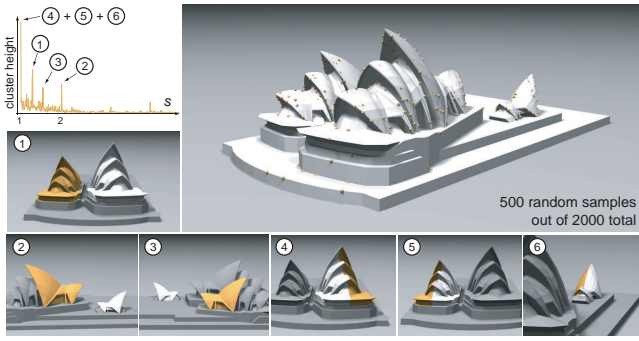


Figure 6: The six most significant modes of the Sydney Opera with the full 7-dimensional symmetries (1, 2, 3) and pure reflections (4, 5, 6). The graph shows the distribution of scaling factors.

Figure 1 shows partial and approximate symmetry detection on a laser scan of a hand-sculpted model. The dragon has been sampled with 2470 points out of which 800 have been randomly selected to extract the five most significant modes. The deviations from perfect symmetry are visualized as the signed distances to the closest point on the perfectly symmetric patch. Since these displacements can be compactly encoded, a compressed representation of the surface can be computed based on the extracted symmetry graph.

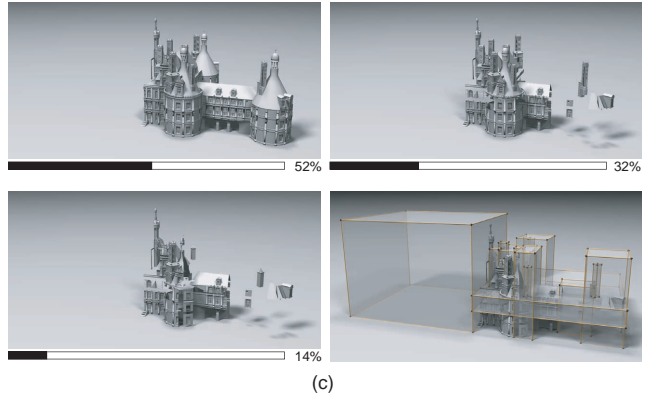
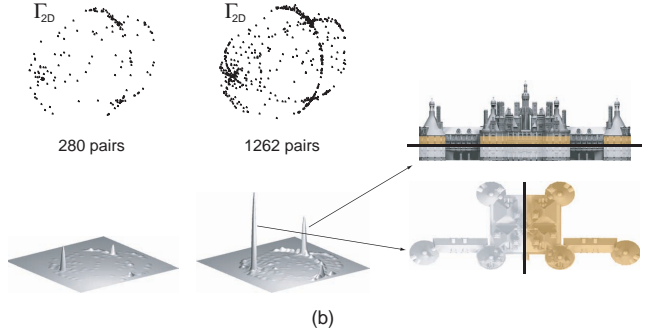
Figure 6 shows an example using the full 7-dimensional symmetry group composed of uniform scaling, rotation, reflection, and translation. All major symmetries are faithfully recovered from only 500 random samples, drawn from an initial sample set of 2000 points.

Figure 7 shows a complex architectural model with symmetries at many different scales. The model has been sampled with 2254 points out of which 100 points (black spheres) and 500 points (yellow spheres) were randomly chosen leading to 280 and 1262 points in Γ , respectively. For visualization purposes we project the samples in transformation space to 2D using metric multi-dimensional scaling [Cox and Cox 1994] as shown in (b). Note that the elliptical structures are due to errors caused by this projection. The two biggest modes map to the symmetries shown on the right, where the perfect global symmetry is faithfully recovered from only 100 random samples. Automatic model reduction and instantiation is shown in (c). Using the first eight significant modes, a reduction to only 14% of the original model size is achieved by taking out the corresponding symmetric patches. The resulting bounding box hierarchy shown in the lower right corner supports efficient spatial queries for applications such as ray-tracing or collision detection directly from the reduced geometry and the corresponding symmetry relations. In (d) we utilized the extracted symmetries to perform advanced editing operations. The user can select a specific symmetry relation and modify certain parts of the model. The system will then automatically apply these modifications to all corresponding patches to maintain the original symmetry.

Figure 8 illustrates an application of our method for segmentation. Two poses of the horse have been sampled with 1000 points. We then selected 500 random points on pose A and paired these with samples from pose B, as shown in (d). The mapping to transformation space is thus restricted to only include pairs that contain one sample from either pose. We can then extract the rigid segments of the model as those parts that are invariant under a rigid transformation between the two poses (e). The projected density function is shown in (b). The biggest modes on the left correspond to the torso and head of the horse. The plot on the right is obtained after removing these parts from the model and adding 192 additional samples to the set of random points. Note that in the 6D transformation space



(a) ● 100 random samples
● 500 random samples



(d)

Figure 7: Chambord castle. (a) input model with random surface samples drawn from a total of 2254 samples, (b) points in transformation space projected to 2D and associated density plots; the symmetries corresponding to the biggest two modes are shown on the right, (c) successive reduction by taking out symmetric patches and resulting bounding box hierarchy, (d) advanced editing using the extracted symmetry relations.

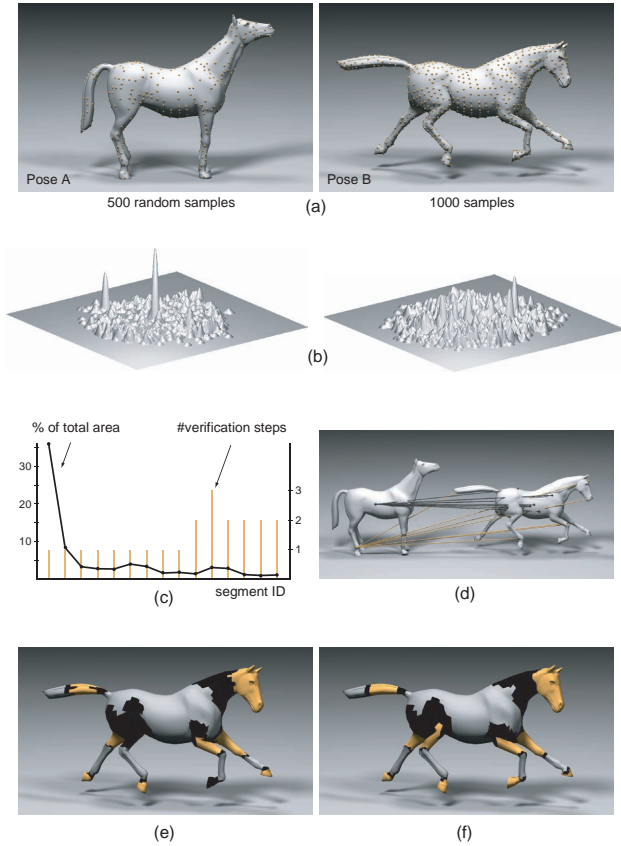


Figure 8: Segmentation and correspondence for the horse model in two different poses. (a) sampling distribution on both poses, (b) projected density plots, (c) extracted segments and verification effort, (d) sample pairing, (e) segmentation using our method, (f) segmentation obtained with explicit, exact point correspondences.

clusters are significantly better separated than in the projected 2D space. As shown in (c), only few clusters have to be discarded in the verification stage. The image in (f) shows the result obtained when using the explicit and perfect one-to-one correspondence available for this specific example, i.e., each vertex is paired with exactly one vertex on the other pose (c.f. [James and Twigg 2005]). This example illustrates that our method can be used to compute a global correspondence map for articulated models *without* requiring any user-specified marker points. Potential applications include partial scan registration and skeleton extraction.

The performance data of Table 1 indicates how the computation depends on the symmetries of the model. The castle has a sig-

Model	# Vertices	Sign.	Pairing	Cluster.	Verif.
Dragon	160,947	3.44	49.24	13.63	7.45
Opera	9,376	0.96	0.02	0.03	0.86
Castle	172,606	5.61	117.81	159.73	5.63
Horse	8,431	0.92	0.01	0.01	1.63
Arch	16,921	0.08	5.86	26.89	2.42

Table 1: Timings in seconds for the different stages of the pipeline on a 2.8 GHz Pentium IV with 2GBytes main memory.




Figure 9: Reflective symmetries of a 2D lizard. Approximate reflective symmetry across the spline gets mapped to a curve in the 2D transformation space.

nificantly more complex symmetry structure as compared to the dragon, hence computation times are substantially higher even though the models have roughly the same size. Note that unlike geometric hashing our method requires minimal additional storage, in all our example less than 500 KBytes of additional main memory were needed.

The set of symmetries extracted by our algorithm is limited to *discrete* symmetries, i.e., ones that can be described by a discrete set of points in transformation space. Continuous symmetries, such as those found in rotational or helical surfaces, lead to smooth curves in transformation space and are currently not detected by our method. This is primarily a limitation of our clustering algorithm, which is not well suited for extracting these types of continuous structures. Figure 9 illustrates this limitation in a slightly different context. The mean-shift algorithm finds two dominant modes that correspond to the reflections of the feet as shown in the illustration. The global reflective symmetry across the spine of the lizard, however, is obscured by the deformation of the specific pose of the model and hence not extracted. As the density plot reveals, this deformation is sufficiently small for the reflection planes to vary smoothly from the head to the tail. We can extract the corresponding curve in 2D transformation space using the method proposed in [Arias-Castro et al. 2006], but have not yet extended the implementation to higher dimensions. Another limitation of our method is inherent in the sample-based approach that we take. Although our algorithm successfully detects the major symmetries of an object, it may fail to identify small partial symmetries during the clustering stage due to the presence of sampling noise. In such a situation, pre-smoothing of the model is required, or the user has to manually tune the patch radii used for curvature estimation (see Section 2), until noise is sufficiently blurred out. Naturally, this leads to less distinct curvature estimates, which might cause the subsequent clustering algorithm to miss less pronounced symmetries.

6 Conclusion

We have demonstrated how matching local shape signatures followed by clustering in transformation space leads to a provably efficient method for discovering and extracting partial and approximate symmetries of 3D geometric models. Our algorithms are easy to implement, require minimal space and computation time, and are applicable to a wide range of 3D shapes. Since our method captures important global and local aspects of the shape it can be exploited in a variety of geometry processing applications. An interesting question for future work is how to automatically make a given model more symmetric, e.g., undo a deformation that obscures a symmetry as in the example of Figure 9. This would allow pose-independent shape matching, which could lead to more effective database retrieval methods. We also want to explore the potential of our algorithm for unsupervised registration in the context of partial scan alignment and protein docking.

Acknowledgements

The work was supported by DARPA grant 32905, NSF grants FRG 0354543, CARGO 0310661, ITR 0205671 and NIH grant GM-072970. Niloy Mitra was partially supported by the Stanford Graduate Fellowship. The authors are grateful to Pierre Alliez, Mario Botsch, Doo Young Kwon, Marc Levoy, Ren Ng, Bob Sumner, and Dilys Thomas for the helpful discussions, suggestions and their help at various stages of the project. Special thanks to the anonymous reviewers for their valuable comments.

References

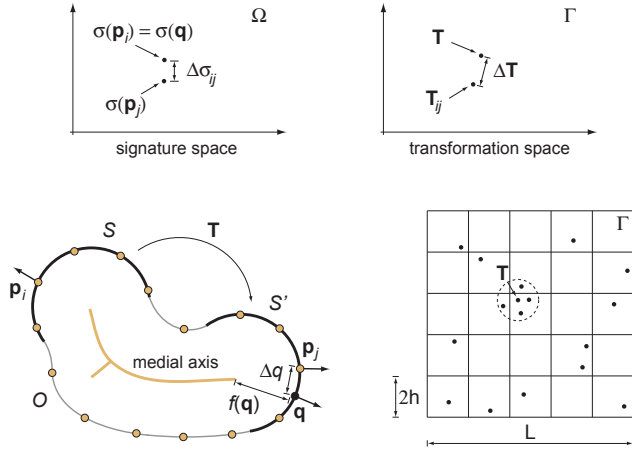
- ALLIEZ, P., COHEN-STEINER, D., DEVILLERS, O., LEVY, B., AND DESBRUN, M. 2003. Anisotropic polygonal remeshing. *ACM TOG* 22, 3, 485–493.
- ALT, H., MEHLHORN, K., WAGENER, H., AND WELZL, E. 1988. Congruence, similarity and symmetries of geometric objects. *Discrete Comput. Geom.* 3, 237–256.
- AMATO, N. M., BAYAZIT, O. B., DALE, L. K., JONES, C., AND VALLEJO, D. 2000. Choosing good distance metrics and local planners for probabilistic roadmap methods. In *IEEE Trans. on Robotics and Automation*, 442–447.
- AMENTA, N., AND BERN, M. 1998. Surface reconstruction by voronoi filtering. In *Symposium on Computational Geometry*, 39–48.
- ARIAS-CASTRO, E., DONOHO, D. L., AND HUO, X. 2006. Adaptive multiscale detection of filamentary structures in a background of uniform random points. *Annals of Statistics* 34, 1.
- ARYA, S., MOUNT, D. M., NETANYAHU, N. S., SILVERMAN, R., AND WU, A. Y. 1998. An optimal algorithm for approximate nearest neighbor searching. In *Journal of the ACM*, vol. 45, 891–923.
- ATALLAH, M. 1985. On symmetry detection. *IEEE Trans. on Computers* 34, 7, 663–666.
- BOISSONNAT, J. D., AND OUDOT, S. 2003. Provably good surface sampling and approximation. In *Symposium on Geometry Processing*, 9–18.
- COHEN-STEINER, D., AND MORVAN, J.-M. 2003. Restricted delaunay triangulations and normal cycle. In *Symposium on Computational Geometry*, 312–321.
- COMANICIU, D., AND MEER, P. 2002. Mean shift: A robust approach toward feature space analysis. *IEEE PAMI* 24, 603–609.
- COX, T., AND COX, M. 1994. *Multidimensional Scaling*. Chapman and Hall, London.
- FISCHLER, M. A., AND BOLLES, R. C. 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. In *Comm. of the ACM*, vol. 24, 381–395.
- GAL, R., AND COHEN-OR, D. 2006. Salient geometric features for partial shape matching and similarity. vol. 25(1), to appear.
- GONNET, G. 1981. Expected length of the longest probe sequence in hash code searching. In *Journal of the Association for Computing Machinery*, 289–304.
- HOFER, M., POTTMANN, H., AND RAVANI, B. 2004. From curve design algorithms to the design of rigid body motions. In *The Visual Computer*, 279–297.
- HOUGH, P. 1959. Machine analysis of bubble chamber pictures. In *International Conference on High Energy Accelerators and Instrumentation*.
- JAMES, D. L., AND TWIGG, C. D. 2005. Skinning mesh animations. *ACM TOG* 24, 3, 399–407.
- KAZHDAN, M. M., CHAZELLE, B., DOBKIN, D. P., FINKELSTEIN, A., AND FUNKHOUSER, T. A. 2002. A reflective symmetry descriptor. In *Proceedings of ECCV*, 642–656.
- KAZHDAN, M., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2004. Symmetry descriptors and 3d shape matching. In *Symposium on Geometry Processing*, 116–125.
- KLEIN, F. 1893. Vergleichende betrachtungen ber neuere geometrische forschungen. *Mathematische Annalen* 43.
- LAMDAN, Y., AND WOLFSON, H. J. 1988. Geometric hashing: A general and efficient model-based recognition scheme. In *Proceedings of ICCV*, 238–249.
- LIU, Y., COLLINS, R., AND TSIN, Y. 2004. A computational model for periodic pattern perception based on frieze and wallpaper groups. In *IEEE PAMI*, 354–371.
- LOY, G., AND EKLUNDH, J. 2006. Detecting symmetry and symmetric constellations of features. In *Proceedings of ECCV*, to appear.
- MAGNUS, W., KARRASS, A., AND SOLITAR, D. 2004. *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*. Dover.
- MANAY, S., HONG, B.-W., YEZZI, A. J., AND SOATTO, S. 2004. Integral invariant signatures. In *Proceedings of ECCV*, 87–99.
- MITRA, N. J., GELFAND, N., POTTMANN, H., AND GUIBAS, L. 2004. Registration of point cloud data from a geometric optimization perspective. In *Symposium on Geometry Processing*, 23–32.
- MOTWANI, R., AND RAGHAVAN, P. 1995. *Randomized Algorithms*. Cambridge University Press.
- RAAB, M., AND STEGER, A. 1998. “balls into bins” — A simple and tight analysis. *Lecture Notes in Computer Science*.
- RUSINKIEWICZ, S., AND LEVOY, M. 2001. Efficient variants of the icp algorithm. In *3DIM*, 145–152.
- SUN, C., AND SHERRAH, J. 1997. 3d symmetry detection using the extended gaussian image. *IEEE PAMI* 19.
- THOMPSON, D. W. 1961. *On Growth and Form*. Cambridge.
- TUZEL, O., SUBBARAO, R., AND MEER, P. 2005. Simultaneous multiple 3d motion estimation via mode finding on lie groups. In *Proceedings of ICCV*, vol. 1, 18–25.
- WOLTER, J., WOO, T., AND VOLZ, R. 1985. Optimal algorithms for symmetry detection in two and three dimensions. *The Visual Computer*.
- ZABRODSKY, H., PELEG, S., AND AVNIR, D. 1995. Symmetry as a continuous feature. *IEEE PAMI* 17.

A Theoretical Analysis

In this section we provide probabilistic bounds on the sampling requirements of our algorithm. More precisely, we define conditions on the sample set P and the number $n' = |P'|$ with $P' \subset P$ of random samples required to find a symmetry of a certain size with high probability.

Suppose we are given a smooth manifold surface O with a symmetric patch $S \subseteq O$ and a partial symmetry transformation $\mathbf{T} \in \Gamma$, such that $S' = \mathbf{T}(S) \subseteq O$. For conciseness of the exposition, we restrict the derivations to the group of rigid transformation, i.e., ignore uniform scaling. The analysis extends in a natural way, however. Assume $P = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ is an ε -sampling of the surface O , i.e., for every point $\mathbf{x} \in O$ there exists a sample $\mathbf{p} \in P$ such that $|\mathbf{p}_i - \mathbf{x}| < \varepsilon f(\mathbf{x})$, where $f(\mathbf{x})$ denotes the local feature size at \mathbf{x} , i.e., the smallest distance of \mathbf{x} to the medial axis of O (see [Amenta and Bern 1998]).

For a given sample $\mathbf{p}_i \in P \cap S$, let $\mathbf{q} \in S'$ be the symmetric point of \mathbf{p}_i on the surface, i.e., $\mathbf{q} = \mathbf{T}(\mathbf{p}_i)$. In general, this point will not be part of the sample set, i.e., $\mathbf{q} \notin P$. However, we can show that there exists a point $\mathbf{p}_j \in P$ such that $\mathbf{p}_j = \mathbf{T}_{ij}(\mathbf{p}_i)$ and $\|\Delta \mathbf{T}\| = \|\mathbf{T} - \mathbf{T}_{ij}\|$ is small.



Let $F_{\mathbf{q}} = [\mathbf{n}_{\mathbf{q}} \ \mathbf{c}_{\mathbf{q},1} \ \mathbf{c}_{\mathbf{q},2}]$ denote the local frame at \mathbf{q} spanned by the normal and the principal curvature directions (see Section 2). Then the transform mapping $\mathbf{p}_i \mapsto \mathbf{q}$ can then be expressed as a rotation $\mathbf{R} = F_{\mathbf{q}} F_{\mathbf{p}_i}^{-1}$ followed by a translation $\mathbf{t} = \mathbf{q} - \mathbf{R}\mathbf{p}_i$. Since P is an ε -sampling of O , $\exists \mathbf{p}_j \in P$ such that $|\mathbf{p}_j - \mathbf{q}| < \varepsilon f(\mathbf{q})$. Using results from [Amenta and Bern 1998; Cohen-Steiner and Morvan 2003] it follows that if $\varepsilon < 0.08$, $\|\mathbf{F}_{\mathbf{q}} - \mathbf{F}_{\mathbf{p}_j}\| \leq c_1 \varepsilon$ where c_1 is a constant depending on the radius of the ball used for estimating the curvature tensor. Let $\mathbf{T}_{ij} = (\mathbf{R}', \mathbf{t}')$ denote the transform mapping $\mathbf{p}_i \mapsto \mathbf{p}_j$. Using the previous relations and the triangle inequality one can show that

$$\|\Delta \mathbf{T}\|^2 = \|\mathbf{R} - \mathbf{R}'\|^2 + \beta \|\mathbf{t} - \mathbf{t}'\|^2 \leq c_2^2 \varepsilon^2,$$

where c_2 is a constant depending on c_1, β and the diameter of O .

Due to the stability of local signatures on a smooth surface [Manay et al. 2004], we can choose a small $\Delta \sigma$ so that $\Delta \sigma_{ij} = \|\sigma_{\mathbf{p}_i} - \sigma_{\mathbf{p}_j}\| < \Delta \sigma$ when $|\mathbf{p}_j - \mathbf{q}| < \varepsilon f(\mathbf{q})$. In other words, the signatures of \mathbf{p}_i and \mathbf{p}_j are sufficiently similar for our pairing algorithm of Section 2.2 to group these points and compute their transformation \mathbf{T}_{ij} as a sample point in transformation space Γ . At the same time this transformation is close to the unknown transformation \mathbf{T} , hence \mathbf{T}_{ij} provides reliable evidence for the symmetry relation that we want to find.

Thus if we choose a clustering radius h larger than $c_2 \varepsilon$ it follows that for any point $\mathbf{p}_i \in S$, our algorithm will deposit at least one point in Γ within distance h of \mathbf{T} . If m is the number of points in $P \cap S$, then for any random sample from P we get a vote within h of \mathbf{T} with probability $p = m/n$.

Using the Chernoff bound [Motwani and Raghavan 1995] we can show that if n' points are independently and randomly chosen from P , then with probability greater than $1 - \alpha$ there will be at least k points within h of \mathbf{T} in Γ , where

$$k = \left(1 - \sqrt{-2 \log \alpha / n' p}\right) n' p$$

and $\alpha \in (0, 1)$. Until now we have shown that with high probability a cluster of height at least k that includes the transformation \mathbf{T} will appear in Γ . To complete the analysis, we now need to ensure that this cluster is in fact a pronounced local maximum of the transformation density function and will thus be successfully retrieved by the mean-shift clustering algorithm. We prove this claim using a counting argument on Γ partitioned into a set of bins.

Suppose the average number of neighbors in Ω for a query radius of $\Delta \sigma$ is μ . Then n' random samples results in roughly $M = \mu n'$ points in Γ . Let the maximum extent along any dimension in Γ be L . So partitioning Γ using a grid of size $2h$ results in $N = (L/2h)^d$ bins, where d represents the dimension of Γ . It is easy to see that if there are more than $k/2^d$ samples. Assuming that point pairs that are not related by any meaningful symmetry relation map to a random bin in Γ , we observe that our scenario is identical to M balls being independently and uniformly thrown into N urns. It is known that the maximum number of balls in any urn with high probability is given by [Gonnet 1981; Raab and Steger 1998]

$$E(n', \mu, \Delta \sigma) = \log N / \log(N \log N / M).$$

In order for the bin corresponding to \mathbf{T} to stand above the noise level, we select n' such that the following inequality holds:

$$E(n', \mu, \Delta \sigma) < \left(1 - \sqrt{-2 \log \alpha / n' p}\right) n' p / 2^d. \quad (1)$$

Thus if we mark a bin in Γ as interesting only if its height is more than $k/2^d$ then: (1) with probability $1 - \alpha$ a bin corresponding to the desired transformation \mathbf{T} is correctly marked, and (2) only a few bins corresponding to spurious transforms are falsely marked. These outlier bins are easily pruned away in the verification step.

Requiring an ε -sampling with $\varepsilon < 0.08$ for P is a fairly severe restriction. While algorithms exist for computing such a point set for a given smooth surface [Boissonnat and Oudot 2003], the sampling density would be prohibitively high. As commented in [Amenta and Bern 1998], empirical evidence suggests that these bounds are quite conservative. In practice, we successfully found the existing symmetries using a much less restricted sampling. Consider the example of Figure 7, where $d = 6$, $L/h \approx 20$, and $\mu = 10$. If we hypothetically assume our initial point set P satisfies the sampling requirements stated above (which is clearly not the case since the surface is not even smooth), then Equation 1 prescribes $n' \approx 300$ to detect a global symmetries where $p = 0.5$ with probability more than 95%. We found that even with only 100 samples we can reliably detect the global reflective symmetry of the castle (see Figure 7). Observe that as the size of the symmetric patch becomes smaller, i.e., p decreases, Equation 1 suggests higher values for n' . Also note that while our analysis has been restricted to perfect symmetries, it can easily be extended to approximate symmetries by increasing h and thereby decreasing the number of bins N .