

Intrinsic Regularity Detection in 3D Geometry

Niloy J. Mitra¹, Alex Bronstein², and Michael Bronstein³

¹ Indian Institute of Technology, Delhi

² Tel Aviv University

³ Technion

Abstract. Automatic detection of symmetries, regularity, and repetitive structures in 3D geometry is a fundamental problem in shape analysis and pattern recognition with applications in computer vision and graphics. Especially challenging is to detect *intrinsic* regularity, where the repetitions are on an intrinsic grid, without any apparent Euclidean pattern to describe the shape, but rising out of (near) isometric deformation of the underlying surface. In this paper, we employ multidimensional scaling to reduce the problem of intrinsic structure detection to a simpler problem of 2D grid detection. Potential 2D grids are then identified using an autocorrelation analysis, refined using local fitting, validated, and finally projected back to the spatial domain. We test the detection algorithm on a variety of scanned plaster models in presence of imperfections like missing data, noise and outliers. We also present a range of applications including scan completion, shape editing, super-resolution, and structural correspondence.

1 Introduction

Symmetries and regular structures are ubiquitous in nature and in man-made objects, often being closely related to form, function, aesthetics, and manufacturing ease of geometrically complex but procedurally simple shapes. While humans are extremely skilled at perceiving and identifying such patterns, even under a cursory inspection [1], automatic detection of such regularity remains challenging. One of the main difficulties is the fact that neither the parts that are being repeated nor their repetition pattern is known *a priori*. Additionally, the surfaces are often warped in practice, making pattern detection challenging (see Figure 1). Such distortions have been widely studied in computer vision, specially work on shape-from-texture [2], and in the context of pose and articulation invariant shape representation and matching [3–7].

State-of-the-art methods [8–10] can detect structured repetitions in 3D geometry if the Euclidean transformations between repeated patches exhibit group-like behavior. In case of non-rigid and deformable shapes, however, the problem is challenging since no apparent structure is visible to simple Euclidean probes in the absence of repetitive Euclidean transformations to describe the shape (see Figure 4).

In this paper, we address the problem of identifying regularity and repeating structure on an *intrinsic* grid on the shape, i.e., regularity detection on (near) developable surfaces. The Euclidean structure of such a grid depends on the embedding of the shape in the ambient space. Yet, using an intrinsic notion of distance, the grid becomes deformation-invariant, and is simpler to identify and extract. To the best of our knowledge, this is the first attempt to detect grid-regularities invariant under isometric deformations. We demonstrate our algorithm on a variety of scanned plaster models of



Fig. 1. (Left) A cylindrical seal, (middle) impression left by a cylindrical seal on a (near) developable surface, and (right) near-intrinsic regular marking created by a car tyre on soft ground

stamped surfaces, with missing parts, and varying degree of noise. Such footprints are common among stenciled concrete, industrial patterns, impressions of cylindrical seals, and vehicle tire tracks just to mention a few.

Related work. Symmetry and structure detection in images and shapes is a well-researched topic in the computer vision and graphics communities (see e.g. [10–19]) with applications including segmentation [20], scan completion [9], pose invariant representation [21], image de-fencing [22], shape retrieval [12], and editing images with repeated elements [23].

Local structure and repetition detection for 3D geometry was addressed by Pauly et al. [8] where self-similarity is represented as a collection of local Euclidean transformations, parameterized in a suitably designed 7D transformation space. The paper observes that repeating self-similar structures correspond to regular grids in special slices of the transformation space, and presents an algorithm for detecting such grid and, subsequently, the repeating elements. However, this approach is limited to handle only regular Euclidean lattices. More recently, Park et al. [24] presented a computational framework using a fourth-degree Markov random fields and mean shift belief propagation, interleaved with thin plate spline warping, for detecting deformed lattices or 2D wallpaper patterns in images. To the best of our knowledge, there are no known extensions to handle intrinsic regularity in 3D geometry.

Raviv et al. [25] introduced *intrinsic symmetries* as a natural extension of the notion of symmetry to non-rigid objects, based on a model of shapes as metric spaces [3, 26,

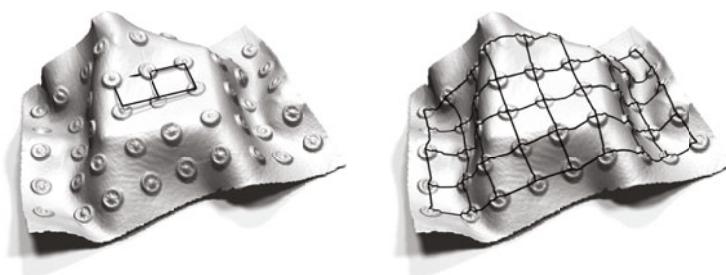


Fig. 2. Embedded intrinsic regularity on a shape as detected, shown in black, by Euclidean transformation analysis [8] (left) and by our algorithm (right)

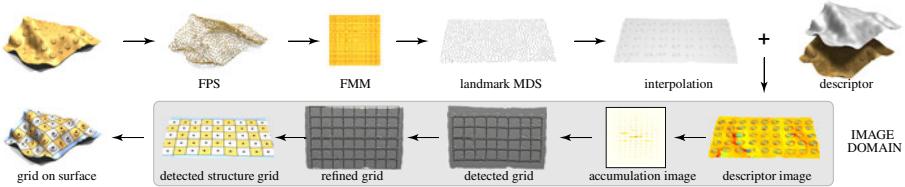


Fig. 3. Stages of the proposed algorithm for intrinsic structure detection: Flattening using MDS, computation of descriptor image, accumulation image, detection of grid generators by autocorrelation analysis, local grid refinement, validation and final detected structure (see Section 3)

27]. In a separate attempt, Ovsjanikov et al. [28] presented a method for symmetry detection based on the properties of eigenfunctions of the Laplace-Beltrami operator of the shape. Since the operator is invariant under isometric deformations, the resultant symmetry detection also detects intrinsic symmetry of objects. Recently, Xu et al. [29] introduced an algorithm to obtain intrinsic reflection symmetry axis (IRSA) transform of objects, followed by an iterative refinement to extract dominant IRSA curves. These efforts, however, are targeted towards detection of pairwise intrinsic symmetry, and not for extracting patterns among the detected symmetries.

Although, intrinsic distances has been employed in computer vision in various contexts like texture mapping [30], face animation and morphing [31], articulation invariant shape matching [3, 4], their use for repetition or regular structure detection has been largely unexplored.

Contribution. We extend the notion of regular structure detection in 3D geometry to handle isometric deformations. The detected structure grids are robust and invariant to bending and articulations of the shapes. State-of-the-art algorithm [8] in regularity detection under Euclidean transformations fails to identify such intrinsic structures in non-rigid shapes since their embedding distorts the Euclidean structures (see Figure 2, left). By using *multidimensional scaling* (MDS), the intrinsic geometry of the surface is mapped into a Euclidean one, thus reducing the problem to the case of Euclidean regularities. Such a planar embedding, however, reduces the problem to a simpler instance of regular grid detection in the plane, instead of intrinsic grid detection on the surface. The use of MDS removes the necessity to detect deformed lattices as proposed by Park et al. [24], leading to a simple, robust, and computationally-efficient algorithm.

2 Background

Let X be a surface modeled as a two-dimensional Riemannian manifold. A parametric curve $\gamma(t)$ on X is called a *geodesic* if parallel transport along the curve preserves the tangent vector $\dot{\gamma}$ to the curve, i.e., $\nabla_{\dot{\gamma}}\dot{\gamma} = 0$ for each point along the curve (∇ denotes the *covariant derivative* on the manifold, roughly equivalent to directional derivative in a vector space). For any point x and a tangent vector $v \in T_x X$, there exists a unique geodesic passing through x whose tangent vector is v . Also, given a pair of points on the manifold, the geodesic curve γ is the (locally) shortest path between the points. We denote by $d_X : X \times X \rightarrow \mathbb{R}$ the *geodesic metric* measuring the length of the shortest paths between points on X .



Fig. 4. Photographs of the plasticine imprints scanned in our experiments (top row), the detected grids in the parametrization domain obtained using MDS along with the locally refined grids (middle row), and the corresponding intrinsic structure depicted on the scanned models (last row). Parts of the shape that do not belong to any detected intrinsic regularities are in blue

Recent works [3, 26, 27] have considered non-rigid surfaces as metric spaces of the form (X, d_X) that are deformation-invariant. Raviv et al. [25] defined *intrinsic symmetry* as self-isometry with respect to the metric d_X : surface X is intrinsically symmetric if there exists a non-trivial bijection $\phi : X \rightarrow X'$ such that $d_X = d_{X'} \circ (\phi \times \phi)$.

Given a point x and a tangent vector $v \in T_x X$, we define translation t by a fixed length in the direction v using parallel transport on the manifold along the geodesic passing through point x in the direction v . Rotation R is defined as rotation of the tangent vector v in the tangent plane. With these two operations, we can define an *intrinsic grid* as a collection of points or nodes $G \subset X$ obtained by successive applications of translation and rotation operations. For example, an orthogonal intrinsic grid is defined by setting an origin x and a vector v , having R defined as a rotation by $\pi/2$ and t as a translation by fixed length (Figure 5).

The goal of intrinsic structure detection is to explain the surface using local self-similarity on an intrinsic grid, i.e., finding the largest grid G that for any pair of points $x_i, x_j \in G$ the surface is locally self-similar at x_i, x_j . The main idea of our approach is as follows: Let X (or its subset) has an isometric embedding into the plane, i.e., there exists a bijection $\psi : (X, d_X) \rightarrow (\mathbb{R}^2, d_{\mathbb{R}^2})$ satisfying $d_X = d_{\mathbb{R}^2} \circ (\psi \times \psi)$. Here ψ can be thought of as a flattening or parametrization of the surface that replaces

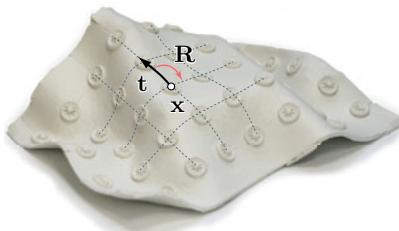


Fig. 5. Example of geodesic grid embedded on the surface

geodesics on X with straight lines in the plane. In particular, the intrinsic regularity manifests as a grid in the plane. Thus, by means of isometric embedding, the problem of grid detection on the surface is replaced by planar grid detection, a well-researched problem in computer vision and image processing. Since a warped surface rarely has an isometric parametrization in the plane, we find the best ψ that minimize distortion $\|d_X - d_{\mathbb{R}^2} \circ (\psi \times \psi)\|$ in the least squares sense, which, in the discrete case, is obtained using MDS [32].

3 Intrinsic Structure Detection

We now describe the different stages of our proposed pipeline for detecting intrinsic structures in 3D geometry (see Figure 3). The algorithm has three main stages: First, we use MDS for surface flattening, which transfers intrinsic grids defined on the surface onto the plane. Second, local shape structure is represented using intrinsic descriptors, and their repeating patterns detected in the plane. Finally, the detected planar grid is refined, validated, and mapped back to the surface.

In the following, we assume that the shape X is presented as a *triangular mesh* built upon a set of vertices $V \equiv \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. Let E denote the set of edges with $(i, j) \in E$ if vertices $\mathbf{x}_i, \mathbf{x}_j$ are connected by an edge. In our experiments, n is typically 50K–100K.

3.1 Shape Flattening

Distance computation. The geodesic metric on the triangular mesh is approximated using *fast marching* (FMM) [33], a numerical solver to the Eikonal equation that computes distance map from a point to the rest of the mesh vertices by simulating wavefront propagation. To reduce the computational cost of the MDS stage, we use a landmark-based approach [34]. The mesh is sampled at $m \ll n$ *landmark points* (denoted, without loss of generality, by $\mathbf{x}_1, \dots, \mathbf{x}_m$) using the *farthest point sampling* (FPS) procedure [35], performed as follows: Start with some vertex, say \mathbf{x}_1 , selected at random. The k -th point is selected from V to be the most distant point from the current selection of $k-1$ points, i.e.,

$$\mathbf{x}_k = \arg \max_{\mathbf{x} \in V} d_X(\{\mathbf{x}_1, \dots, \mathbf{x}_{k-1}\}, \mathbf{x}) = \arg \max_{\mathbf{x} \in V} \min_{i=1, \dots, k-1} d_X(\mathbf{x}_i, \mathbf{x}). \quad (1)$$

FPS produces a subsampling with m approximately equidistant points. Then FMM is employed to compute the $m \times n$ matrix of geodesic distances $d_X(\mathbf{x}_i, \mathbf{x}_j), i = 1, \dots, m; j = 1, \dots, n$ between the landmark points and all the vertices of the mesh. In our experiments, we set the sample size m to a default value of 500.

Flattening. The minimum-distortion parametrization ψ is computed using a variant of *landmark MDS* [34]. First, the landmark points are embedded into the plane. We denote by $\mathbf{u}_i = \psi(\mathbf{x}_i)$ for $i = 1, \dots, m$ their parametrization coordinates in the plane, i.e., $\mathbf{u}_i \in \mathbb{R}^2$, that are found by minimizing the *stress function*

$$\min_{\mathbf{u}_1, \dots, \mathbf{u}_m} \sum_{i=1}^m \sum_{j=i+1}^m (d_X(\mathbf{x}_i, \mathbf{x}_j) - \|\mathbf{u}_i - \mathbf{u}_j\|)^2. \quad (2)$$



Fig. 6. Robustness of our method to missing and corrupt data: occlusions due to imperfect acquisition and synthetic holes (left), simulated Gaussian noise (center), and shot noise (right). Regions not explained by detected intrinsic structure are in blue.

The minimizer of the stress function is the minimum-distortion parametrization of the surface in the least squares sense. We use SMACOF iterations [32]

$$U^{(k+1)} = \frac{1}{m} B(U^{(k)}) U^{(k)}, \quad (3)$$

repeated until convergence, to solve the LS-MDS problem (2) iteratively¹. Here, U is a $m \times 2$ matrix of landmark point parametrization coordinates with

$$b_{ij}(U) = \begin{cases} \frac{d_X(\mathbf{x}_i, \mathbf{x}_j)}{\|\mathbf{u}_i - \mathbf{u}_j\|} & i \neq j \text{ and } \|\mathbf{u}_i - \mathbf{u}_j\| \neq 0, \\ 0 & i \neq j \text{ and } \|\mathbf{u}_i - \mathbf{u}_j\| = 0, \\ -\sum_{k \neq i} b_{ik} & i = j. \end{cases}$$

We initialize the LS-MDS solver using classical scaling based on a globally-convergent algebraic MDS method minimizing the Frobenius norm of the distance distortion. SMACOF iterations are guaranteed to produce a monotonically decreasing sequence of stress values [37].

Interpolation. The obtained landmark parametrization coordinates are employed to interpolate the parametrization coordinates for the remaining mesh vertices using a distance-based interpolation proposed in [34]. The interpolated coordinates are inferred from the landmark coordinates using

$$\mathbf{u}_j = -\frac{1}{2} U^\dagger (\delta_j - \bar{\delta}),$$

¹ SMACOF iteration is equivalent to a weighted gradient descent and generally does not guarantee global convergence. However, using a sufficiently good initialization or multiscale optimization, reasonable convergence is obtained [36].

for $j = m + 1, \dots, n$, where U^\dagger denotes the pseudoinverse of the matrix of landmark point parametrization coordinates, $\delta_j = (d_X^2(\mathbf{x}_j, \mathbf{x}_1), \dots, d_X^2(\mathbf{x}_j, \mathbf{x}_m))^T$ is the $m \times 1$ vector of squared distances from \mathbf{x}_j to the landmark points, and

$$\bar{\delta} = \frac{1}{m} \sum_{i=1}^m (d_X^2(\mathbf{x}_i, \mathbf{x}_1), \dots, d_X^2(\mathbf{x}_i, \mathbf{x}_m))^T$$

is the average squared distance between the landmark points. Thus, for every vertex $\mathbf{x} \in V$ on the mesh, we get a mapping to a corresponding point $\mathbf{u} = \psi(\mathbf{x})$ in the plane.

3.2 Grid Detection

Descriptor. We now compute a simple scalar descriptor at each point to facilitate repetition detection in the next stage. The input mesh is smoothed using a discrete Laplacian,

$$(\Delta \mathbf{x})_i = \frac{1}{v(i)} \sum_{j:(i,j) \in E} \mathbf{x}_j, \quad (4)$$

producing a smoothed version $X_p = \Delta X$. Here, $v(i)$ denotes the valence of the vertex \mathbf{x}_i , i.e., the number of vertices adjacent to it. We define a scalar descriptor at each vertex of the mesh as $c \equiv \langle \mathbf{x} - \mathbf{x}_p, \mathbf{n} \rangle$, where \mathbf{x} and \mathbf{x}_p are the original and smoothed shape coordinates, respectively, and \mathbf{n} is the normal vector to the smoothed mesh at that point. Such a descriptor captures the high-frequency geometric details or the *coating* of the surface, and is insensitive to low-frequency bending and non-rigid deformations. Mapping the descriptor to the plane using the parametrization ψ results in a *descriptor image* $c \circ \psi^{-1}$, which contains regular Euclidean 2D patterns. In our experiments, we sampled the descriptor image on a regular planar Cartesian grid with the largest dimension of 128.

Accumulation. The 3D surface of the height-field descriptor image contains regular structures in the Euclidean sense. While one can detect the repetitions on this derived surface using the method proposed by Pauly et al. [8], given the nature of the repetitions in the MDS domain, a much simpler approach is to directly detect the grids on the descriptor image. An *accumulation image* representing the repeating patterns in the descriptor image is constructed as

$$A(\mathbf{w}) = \sum_{\mathbf{z}} \exp(-\langle P(\mathbf{z}), P(\mathbf{z} + \mathbf{w}) \rangle / 2\sigma^2) \quad (5)$$

where $P(\mathbf{z})$ and $P(\mathbf{z} + \mathbf{w})$ are normalized descriptor image patches, and the inner product between them is weighted by a Gaussian window with kernel width σ . Thus, if the descriptor image contains many patches that are similar up to a displacement by a vector \mathbf{w} , the accumulation image will exhibit a peak at \mathbf{w} . In our experiments, by default, the accumulation image was of size 257×257 , the patch size was 21×21 , and $\sigma = 0.075$ was used.

Autocorrelation. The autocorrelation of the accumulation image allows to find the grid generators, i.e., the vectors that define the grid in the parametrization domain. Autocorrelation is computed as the similarity of the accumulation image to its version shifted

by a vector \mathbf{u} , represented in polar coordinates. Two peaks are detected in the polar autocorrelation image, representing the grid generators. Often, there might exist more than one pair of generators explaining the same grid, as visualized in Figure 10. In such cases, we give preference to shorter ones resulting in a denser grid.

Phase selection. The detected grid is inherently ambiguous to *phase*, i.e., shift along the grid generator vectors (Figure 10). While there is no theoretical preference to a specific phase of the detected grid, some phases produce semantically and visually more meaningful results. We perform phase selection by shifting the grid to maximize the local variation of the descriptor at the grid nodes. This way, the grid locks onto *interesting* geometric features.

3.3 Refinement and Validation

Refinement. Peaks of the autocorrelation function of the accumulation image provide good generators for grids in the MDS domain. However, because of them being only approximate, we locally refine the grid point locations as well as remove grid points and connections that do not correspond to any structural element (see Figure 3). This local correlation and refinement is performed using descriptor images with twice the resolution used in the previous stage. We used images with maximum dimension of 256 pixels as default.

Validation. The input mesh and the corresponding MDS descriptor image do not solely constitute of regular structures. Hence we explicitly identify and extract the structural elements around each grid point using a greedy growth with validation (see Figure 10). Finally, we project back the detected structural elements to the surface. In, Figure 4 the unstructured parts are indicated in blue.

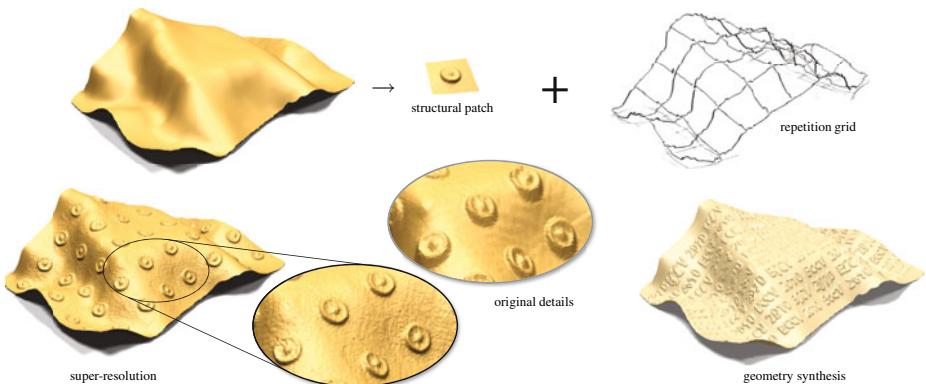


Fig. 7. Square-shape (see Figure 4) is decomposed into a smooth base and a detail layer consisting of a structural patch and a repetition grid (top). This enables geometry processing like super-resolution and synthesis of a shape with the same structural layout but different with geometric details (bottom).

Table 1. Mesh sizes and run times (in seconds) for different stages of the algorithm. Performance measured on a 2 GHz Core Duo Pentium CPU with 3GB of RAM.

| | square | oblique | Colosseum | ball |
|---------------|--------|---------|-----------|--------|
| # vertices | 54,880 | 54,817 | 57,563 | 40,601 |
| FMM + MDS | 5 | 5 | 6 | 4 |
| interpolation | 11 | 10 | 12 | 5 |
| descriptor | 3 | 4 | 4 | 4 |
| accumulation | 23 | 34 | 17 | 54 |
| correlation | 31 | 41 | 28 | 59 |
| refinement | 13 | 16 | 12 | 18 |

4 Results and Applications

For the experiments in this paper, we used scans of objects sculpted and stamped using plasticine depicted in Figure 4 (first row). The objects were designed to contain intrinsically repeating structure, which is hard to perceive if considered in a Euclidean way. Four objects were used: deformed surfaces with square and oblique grids, a detail of a curved architectural shape (non-rigid Colosseum), and part of a ball with square grid structure. Each of the objects presents a different challenge in structure recovery. Thus, in the Colosseum shape, the structural elements are holes (windows), and the ball has a non-developable geometry. The objects were scanned using a coded light range camera, producing triangular meshes that were cleaned up and resampled to about 50K vertices. Data and code are available for academic use from the project webpage. Our algorithm is robust to a range of parameter settings, and all the reported results are with a default set of values.

Figure 4 shows the intrinsic grids detected using the proposed algorithm. The algorithm was implemented in Matlab without optimization. Overall run time in these examples is about a minute (see Table 1 for detailed timing of each stage of the algorithm). In the 2D processing stages of the algorithm (accumulation image creation, correlation, and refinement) the complexity is dictated mainly by the descriptor image

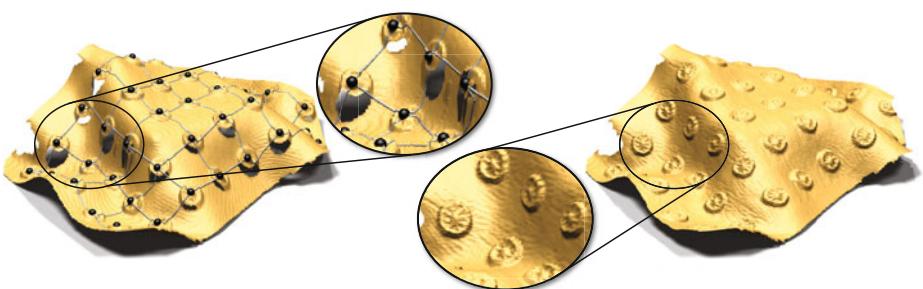


Fig. 8. Intrinsic structures can be detected on scans with missing data (left). The detected structure is used to propagate structural elements from *healthy* regions to conceal the damaged or missing areas (right).



Fig. 9. Structural correspondence, up to phase ambiguity, between two geometrically different shapes having similar repeating structures, established using the detected correspondence between the respective intrinsic grids. Combinatorially similar respective grids are shown at the top right and top left, respectively. High-frequency texture mapping was used to visualize the accuracy of local correspondence.

size, which, in turn, depends on the way MDS embedding maps the surface into the plane. For comparison, Figure 2 (left) shows the results produced by Euclidean structure detection of Pauly et al. [8]. Only a small part of the grid (four points located in the flat part of the shape) is recovered. The behavior is similar for the other scans.

Robustness. The proposed algorithm is insensitive to noise and can work even when large parts of the grid are missing. Figure 6 (left) shows examples of grid detection in shapes suffering from missing details (resulting from real occlusion artifacts in 3D acquisition and synthetic removal of parts of the shape). Despite large portions of the shape missing (up to 35% in Figure 6, middle), most of the grid structure is correctly detected. Figure 6 (middle) shows a shape contaminated by random Gaussian noise with standard deviation of about 70% of the average feature elevation, while Figure 6 (right) shows corruption by shot noise appearing as spikes. Besides a few missing grid nodes, the repeating structure is detected correctly in these cases as well.

Geometry substitution. A shape can be decomposed into a low-frequency base governing the embedding in \mathbb{R}^3 , and the high-frequency detail admitting the intrinsic repeating pattern. Replacing the detail allows to synthesize new shapes sharing with the original one the 3D layout, while retaining the repeating structure. An example of such a substitution is presented in Figure 7. The low-frequency surface was obtained by solving the Laplace equation for the surface coordinates with the grid lines serving as the boundary conditions. The new detail was mapped onto the low-frequency base using a normal displacement map. Detail substitution can be used to conceal irregularities of a regular shape due to manufacturing or acquisition imperfections. In this case, the structural element from a “healthy” region of the shape is transferred to a damaged one as shown in Figure 8. While MDS mapping by itself can be influenced by topological errors, the grid detection and refinement phases make the pipeline robust to small holes and perforations (see also the Colosseum example). Use of topologically-consistent weighted MDS [38] can increase the stability of the system. In this example, we first closed the holes using smooth interpolation; the interpolated regions were healed by detail transfer from healthy regions resulting in a nearly perfectly regular shape. Substituting the detail with its higher resolution version (obtained, for example, from a

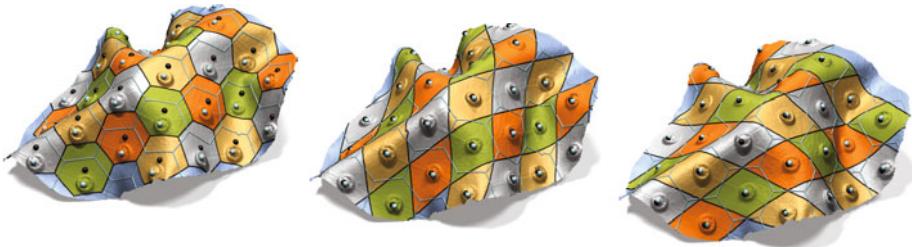


Fig. 10. Ambiguities inherent to intrinsic grid detection: phase (left), shape of the structural element (middle), and different generators explaining the same grid (right). All such results provide plausible explanations to the intrinsic structure of the oblique-shape. For reference, results from Figure 4 are overlaid as blue curves/spheres.

close-up scan or a CAD model) produces *super-resolution* of the original shape, allowing to overcome the classical field-of-view versus resolution tradeoff or combine different acquisition modalities (see Figure 7 for an example).

Structural correspondence. The knowledge of intrinsic structure allows us to establish correspondence between objects significantly different geometrically and topologically, yet resembling in their self-similarity structure. This concept has been recently explored in image analysis applications for comparison of images depicting similar concept in visually different ways [5]. Given two shapes with similar intrinsic structures, we first extract the intrinsic grids and then find the correspondence between these grids. The structural elements and the extrinsic geometry of the shapes can be wildly different, as exemplified in Figure 9. Such a great difference in geometry and topology is an obstacle that most state-of-the-art non-rigid correspondence methods find very challenging to overcome.

5 Conclusions

We presented an approach for intrinsic local self-similarity detection in 3D shapes. Unlike previous approaches limited to Euclidean self-similarity, our approach is able to detect warped and curved grids. By using MDS, we reduce the problem of intrinsic grid detection on the surface to regular grid detection in the plane. We demonstrated the efficiency and robustness of the method on various scanned (stamped) models with different geometries, topologies, and structures, as well as real and simulated artifacts. Examples of applications to scan completion, detail substitution, super-resolution, and correspondence between structurally similar yet geometrically and topologically different shapes were presented.

Limitations and extensions. The current limitation of our approach is that the use of planar parametrization implies a tacit assumption that the topology of the surface is coarsely similar to that of the plane. While this is true in many cases of shapes acquired by means of a range scanner and represented as geometry images, a generic shape may have more complicated topology, e.g. of a sphere. Trying to embed such shapes into the plane would result in large distortions such that intrinsic grids would be no more

mapped into planar regular grids. A possible way to handle complex topologies is by applying MDS in a local manner, to disk-like regions on the shape, and then stitch together the detected grids, which we leave to future work.

Other limitations, inherent to regular structure detection in general, are the ambiguities in phase and non-uniqueness of the grid generators and the structural elements (see Figure 10). Optimization over these parameters with the goal to achieve optimal packing of full structural elements over the shape can be a way to resolve such ambiguities.

Acknowledgement

We thank Helmut Pottmann and the anonymous reviewers for their comments and helpful suggestions. Niloy Mitra was partially supported by a Microsoft outstanding young faculty fellowship.

References

1. Leyton, M.: *Shape as memory*. Springer, Heidelberg (2006)
2. White, R., Forsyth, D.A.: Combining cues: Shape from shading and texture. In: IEEE CVPR, vol. II, pp. 1809–1816 (2006)
3. Elad, A., Kimmel, R.: Bending invariant representations for surfaces. In: CVPR, pp. 168–174 (2001)
4. Ling, H., Jacobs, D.W.: Deformation invariant image matching. In: IEEE ICCV, Washington, DC, USA, pp. 1466–1473. IEEE Computer Society, Los Alamitos (2005)
5. Shechtman, E., Irani, M.: Matching local self-similarities across images and videos. In: CVPR, pp. 511–518 (2007)
6. Zaharescu, A., Boyer, E., Varanasi, K., Horaud, R.: Surface feature detection and description with applications to mesh matching. In: IEEE CVPR, pp. 373–380 (2009)
7. Mitra, N.J., Flory, S., Ovsjanikov, M., Gelfand, N., Guibas, L., Pottmann, H.: Dynamic geometry registration. In: Symposium on Geometry Processing, pp. 173–182 (2007)
8. Pauly, M., Mitra, N.J., Wallner, J., Pottmann, H., Guibas, L.: Discovering structural regularity in 3D geometry. ACM ToG (Proc. SIGGRAPH) 27, #43, 1–11 (2008)
9. Thrun, S., Wegbreit, B.: Shape from symmetry. In: IEEE ICCV, pp. 1824–1831 (2005)
10. Mitra, N.J., Guibas, L., Pauly, M.: Partial and approximate symmetry detection for 3d geometry. ACM ToG (Proc. SIGGRAPH) 25, 560–568 (2006)
11. Zabrodsky, H., Peleg, S., Avnir, D.: Symmetry as a continuous feature. IEEE Trans. Pattern Anal. Mach. Intell. 17, 1154–1166 (1995)
12. Kazhdan, M., Funkhouser, T., Rusinkiewicz, S.: Symmetry descriptors and 3D shape matching. In: Proc. of Symp. of Geometry Processing (2004)
13. Liu, Y., Collins, R.T., Tsin, Y.: A computational model for periodic pattern perception based on frieze and wallpaper groups. IEEE Trans. Pattern Anal. Mach. Intell. 26, 354–371 (2004)
14. Martinet, A., Soler, C., Holzschuch, N., Sillion, F.: Accurate detection of symmetries in 3d shapes. ACM ToG 25, 439–464 (2006)
15. Podolak, J., Shilane, P., Golovinskiy, A., Rusinkiewicz, S., Funkhouser, T.: A planar-reflective symmetry transform for 3D shapes. ACM ToG (Proc. SIGGRAPH) 25 (2006)
16. Loy, G., Eklundh, J.O.: Detecting symmetry and symmetric constellations of features. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3952, pp. 508–521. Springer, Heidelberg (2006)

17. Bokeloh, M., Berner, A., Wand, M., Seidel, H.P., Schilling, A.: Symmetry detection using line features. *Computer Graphics Forum (Proc. EUROGRAPHICS)* 28, 697–706 (2009)
18. Park, M., Lee, S., Chen, P.C., Kashyap, S., Butt, A.A., Liu, Y.: Performance evaluation of state-of-the-art discrete symmetry detection algorithms. In: *IEEE CVPR*, pp. 1–8 (2008)
19. Gong, Y., Wang, Q., Yang, C., Gao, Y., Li, C.: Symmetry detection for multi-object using local polar coordinate. In: Jiang, X., Petkov, N. (eds.) *Computer Analysis of Images and Patterns. LNCS*, vol. 5702, pp. 277–284. Springer, Heidelberg (2009)
20. Simari, P., Kalogerakis, E., Singh, K.: Folding meshes: hierarchical mesh segmentation based on planar symmetry. In: *Proc. of Symp. of Geometry Processing*, pp. 111–119 (2006)
21. Mitra, N.J., Guibas, L.J., Pauly, M.: Symmetrization. *ACM ToG (Proc. SIGGRAPH)* 26, #63, 1–8 (2007)
22. Liu, Y., Belkina, T., Hays, J.H., Lublinerman, R.: Image de-fencing. In: *IEEE CVPR* (2008)
23. Cheng, M.M., Zhang, F.L., Mitra, N.J., Huang, X., Hu, S.M.: Repfinder: Finding approximately repeated scene elements for image editing. *ACM ToG (Proc. SIGGRAPH)* 29 (2010) (to appear)
24. Park, M., Collins, R., Liu, Y.: Deformed lattice detection via mean-shift belief propagation. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008, Part II. LNCS*, vol. 5303, pp. 474–485. Springer, Heidelberg (2008)
25. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Full and partial symmetries of non-rigid shapes. *IJCV* (preprint)
26. Mémoli, F., Sapiro, G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Comp. Mathematics* 5, 313–346 (2005)
27. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* 103, 1168–1172 (2006)
28. Ovsjanikov, M., Sun, J., Guibas, L.J.: Global intrinsic symmetries of shapes. *Comput. Graph. Forum (Proc. SGP)* 27, 1341–1348 (2008)
29. Xu, K., Zhang, H., Tagliasacchi, A., Liu, L., Li, G., Meng, M., Xiong, Y.: Partial intrinsic reflectional symmetry of 3D shapes. *ACM ToG (Proc. SIGGRAPH Asia)* 28 (2009)
30. Zigelman, G., Kimmel, R., Kiryati, N.: Texture mapping using surface flattening via multidimensional scaling. *IEEE Trans. on Vis. and Comp. Graphics* 9, 198–207 (2002)
31. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Calculus of non-rigid surfaces for geometry and texture manipulation. *IEEE Trans. on Vis. and Comp. Graphics* (2008)
32. Borg, I., Groenen, P.: *Modern multidimensional scaling - theory and apps.* Springer, Heidelberg (1997)
33. Kimmel, R., Sethian, J.A.: Computing geodesic paths on manifolds. *PNAS*, 8431–8435 (1998)
34. De Silva, V., Tenenbaum, J.: Global versus local methods in nonlinear dimensionality reduction. In: *NIPS*, pp. 721–728 (2003)
35. Hochbaum, D., Shmoys, D.: A best possible heuristic for the k-center problem. *Mathematics of Operations Research* 10, 180–184 (1985)
36. Bronstein, M.M., Bronstein, A.M., Kimmel, R., Yavneh, I.: Multigrid multidimensional scaling. *Numerical Linear Algebra with Applications* 13, 149–171 (2006)
37. Borg, I., Groenen, P.: *Modern multidimensional scaling: Theory and applications* (2005)
38. Rosman, G., Bronstein, A., Bronstein, M., Kimmel, R.: Topologically constrained isometric embedding. *Human Motion: Understanding, Modelling, Capture, and Animation* (2008)