

Microtiles: Extracting Building Blocks from Correspondences

Javor Kalojanov¹, Martin Bokeloh², Michael Wand^{1,3}, Leonidas Guibas², Hans-Peter Seidel³, Philipp Slusallek¹

¹Saarland University ²Stanford University ³MPI Informatik

Abstract

In this paper, we develop a theoretical framework for characterizing shapes by building blocks. We address two questions: First, how do shape correspondences induce building blocks? For this, we introduce a new representation for structuring partial symmetries (partial self-correspondences), which we call “microtiles”. Starting from input correspondences that form point-wise equivalence relations, microtiles are obtained by grouping connected components of points that share the same set of symmetry transformations. The decomposition is unique, requires no parameters beyond the input correspondences, and encodes the partial symmetries of all subsets of the input. The second question is: What is the class of shapes that can be assembled from these building blocks? Here, we specifically consider r -similarity as correspondence model, i.e., matching of local r -neighborhoods. Our main result is that the microtiles of the partial r -symmetries of an object \mathcal{S} can build all objects that are $(r + \epsilon)$ -similar to \mathcal{S} for any $\epsilon > 0$. Again, the construction is unique. Furthermore, we give necessary conditions for a set of assembly rules for the pairwise connection of tiles. We describe a practical algorithm for computing microtile decompositions under rigid motions, a corresponding prototype implementation, and conduct a number of experiments to visualize the structural properties in practice.

1. Introduction

In contemporary computer graphics, the creation of high-fidelity 3D models still remains a difficult and time consuming task. Correspondingly, approaches that facilitate the reuse of existing 3D models have recently gained quite some attention. One promising approach is *inverse procedural modeling* [HF97, ARB07, ŠBM*10, BWS10]: The goal is to decompose a model into *building blocks* and associated rules (i.e., a *shape grammar*) describing how these pieces can be put together to create variations of the original. At the heart of this problem, there are two questions: First, given a notion of similarity, what is the right choice of building blocks and rules? And second, what is the class of objects that can be constructed in this way?

Representing partial symmetries with “microtiles”:

We consider a shape \mathcal{S} and set of mappings \mathcal{F} that maps subsets of \mathcal{S} back to itself, in a way that it defines an equivalence relation among the points of \mathcal{S} . In other words, we consider *pairwise correspondences* within \mathcal{S} , a structure also often referred to as partial symmetries [MGP06] (a stricter notion is to use the term “symmetry” only when algebraic regularity is involved [MPWC12]; we stick to the common convention of using the terms interchangeably). Usually, the

regions matched by the pairwise relations overlap arbitrarily. Hence, this provides a notion of redundancy, but not yet a decomposition into disjoint building blocks.

To form building blocks, we split the original surface into maximal fragments for which all points are mapped by the same set of transformations (see Fig. 1a-d). Equivalently, we can cut the surface at all boundaries of all partial matches (Fig. 1c). Because the resulting tiles are obtained by cutting as much as possible, we call them “microtiles”. Yet another characterization is to form a graph with all points of \mathcal{S} as nodes, and transformations from \mathcal{F} annotating edges that connect equivalent points (shown schematically in Fig. 1e). Microtiles are maximal connected subsets of \mathcal{S} that have cliques with the same edge annotations. This means, all microtiles in the same clique are always exchanged with each other as a whole (c.f. Fig. 3).

Microtiles have some interesting further properties: First, the construction is unique and canonical: it does not require any choices or parameters in addition to the input correspondences. Furthermore, microtiles are disjoint and different types of tiles do not have partial correspondences among each other or themselves, only global symmetries are possible. Finally, unions of microtiles have the intersection of

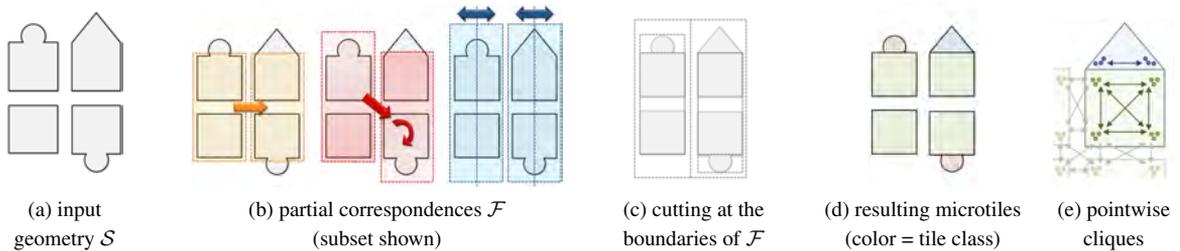


Figure 1: Construction of microtiles: (a) Input geometry \mathcal{S} . (b) We consider partial correspondences \mathcal{F} within \mathcal{S} . (c) The geometry is cut at the boundaries of the partial matches. (d) This yields the microtile decomposition. (e) Each microtile is characterized by cliques of equivalent points; cliques with the same set of transformations form the same class of microtiles.

their associated cliques of transformations as permissible operations. Thus, we can derive all partial symmetry properties of a shape from our decomposition by simple set operations.

The space of r -similar shapes: We now use microtiles to study inverse procedural modeling problems. We look specifically at r -similarity as a notion of correspondence: two points on a shape are matched if their r -neighborhoods match under a rigid mapping. This yields a point-wise equivalence relation, r -symmetry, for which we can build r -microtiles. Furthermore, a shape \mathcal{S}_2 is considered r -similar to \mathcal{S}_1 if all points $\mathbf{y} \in \mathcal{S}_2$ are r -similar to some point $\mathbf{x} \in \mathcal{S}_1$. We can now characterize the space of r -similar shapes: We prove that all shapes \mathcal{S}_2 that are $(r + \epsilon)$ -similar to \mathcal{S}_1 can be constructed by rigidly assembling r -microtiles of \mathcal{S}_1 for any $\epsilon > 0$ and that the assembly is disjoint and unique. We further make a first step towards a microtile-based shape grammar, i.e., giving explicit rules for building all $(r + \epsilon)$ -similar shapes: We show, as a necessary condition, that a valid assembly is restricted to pairwise adjacencies that have been observed in the exemplar.

While our paper focuses on the theory, we also perform a brief practical study to examine and illustrate the decomposition in practice. We describe an algorithm for computing microtiles and propose a prototype implementation based on a simple volumetric discretization [BWS10]. We compute the decomposition for a number of example scenes and visualize building blocks and redundancy.

In summary, we make two main contributions: First, we introduce a canonical decomposition of partial symmetry, converting pairwise correspondences to microtiles as building blocks. Second, we look at local rigid matching as notion of correspondence and show that the space of shapes locally similar to an exemplar \mathcal{S} can be build from microtiles of \mathcal{S} .

2. Related Work

Symmetry detection has become an indispensable tool for shape analysis [GCO06, MGP06, PSG*06, SKS06, Mar07,

OSG08, PMW*08, BBW*09, LCDF10, RBBK10], leading to numerous applications such as symmetry preserving editing [GSMCO09, BWKS11, WXL*11], non-local reconstruction [GSH*07, BBW*09], and shape understanding [MY*10].

The classic mathematical notion of symmetry is global in nature [MPWC12]: Let \mathcal{T} be a group of bijections of \mathbb{R}^3 acting on $\mathcal{S} \subset \mathbb{R}^3$. The functions $f \in \mathcal{T}$ that map \mathcal{S} back to itself form a group, the symmetry group of \mathcal{S} under \mathcal{T} . While this structure is well understood, structuring the space of partial symmetries is more complicated because they usually do not form closed groups in the transformation domain.

Early approaches to structuring partial symmetries by building blocks used simple region growing [MGP06, BBW*09], which does not lead to canonical results because they depend on the initialization. It is also not clear how to reassemble such tilings to form new shapes. Another way of structuring pairwise correspondences is to look for algebraic regularity in the domain of the transformations, for example by detecting commutative grids [PMW*08, MBB10]. This describes the structure of the correspondences only partially, as the input contains only excerpts of the symmetry groups (which are usually infinite), and irregularly placed instances are not captured.

An alternative is to just enumerate the overlapping regions of pairwise matches [BWS10]. This implicitly encodes all symmetry information, but does not expose its structure. Based on this, Bokeloh et al. [BWS10] use the boundaries of partial r -symmetry in order to cut out *dockers* that form a context-free shape grammar. It encodes a set of objects that are r -similar to the input exemplar. Due to the restriction to context-free grammars, their method must avoid intersections of cuts. Our approach is the opposite, performing all cuts, even for continuous symmetries (which their method explicitly excludes). Thereby, we can span the complete space of r -similar objects with assemblies of microtiles, while context-free tiles only cover a restricted sub-

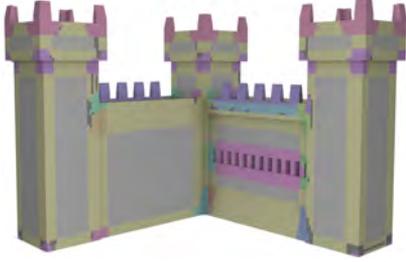


Figure 2: A castle wing decomposed into microtiles w.r.t. r -symmetry. The colors indicate the class of the tiles.

set. Lifting this restriction and understanding the set of r -similar models was the main motivation of our work.

We should also note that local similarity is frequently used in texture synthesis [EL99, WL00] and their analogs in geometry synthesis [BIT04, ZHW*06, Mer07]. However, these methods use variational formulations that do not provide insight into the structure of the shape space.

In terms of representing partial symmetry, our method is strongly related to, and has been motivated by, the work of Lipman et al. [LCDF10]: Their method also starts from point-wise equivalence relations and computes cliques in correspondence graphs. Unlike our approach, they use a spectral clustering technique that is robust even under noisy and ambiguous data. The result is a symmetry factored embedding that maps points with a similar symmetry structure to nearby points in a Euclidean space. A subsequent clustering in this space produces a partitioning that resembles our microtiles. However, our model for dealing with partiality is different. Rather than weighting the percentage of mapped points, we model the functions and domains of partial matches explicitly. This allows us to formally study the resulting tile decompositions and gives us a formal framework for how general mapping functions should be treated. Overall, our paper provides complementary insights: We focus on the study of the resulting tiles and their properties, and study how the class of shapes that can be assembled out of those tiles.

3. Correspondences and Microtiles

Input: In the following, we use $S \subset \mathbb{R}^3$ to denote the exemplar piece of geometry for which we want to compute a microtile decomposition. Let

$$\mathcal{F} \subseteq \{(\mathcal{P}, f) \mid \mathcal{P} \subseteq S, f: \mathcal{P} \rightarrow S, f \text{ is a homeomorphism}\} \quad (1)$$

be a set of functions f that map subsets $\mathcal{P} \subset S$ of the exemplar back to the itself, in a topology preserving way. In other words, \mathcal{F} is a set of partial correspondences on S . The sets \mathcal{P} identify the part of S the functions f act upon.

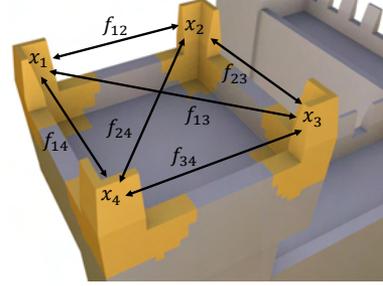


Figure 3: A graph representation of an example set of partial mappings. The nodes are (colored) points on the surface and the edges are mappings. As illustrated, all connected components of the graph have to be cliques by definition.

Equivalence of points: Given a set of mappings \mathcal{F} , we can now say that two points $\mathbf{x}, \mathbf{y} \in S$ are equivalent or similar, if there is a mapping $f \in \mathcal{F}$, which maps \mathbf{x} to \mathbf{y} :

$$\mathbf{x} \equiv_f \mathbf{y} \quad :\Leftrightarrow \quad \exists (\mathcal{P}, f) \in \mathcal{F} : \mathbf{x} \in \mathcal{P} \text{ and } f(\mathbf{x}) = \mathbf{y} \quad (2)$$

We require that this relation be an equivalence relation. This means, we must choose \mathcal{F} in a way this induced point-wise relation is:

- Reflexive: $(S, id) \in \mathcal{F}$
- Symmetric: $\mathbf{x} \equiv_f \mathbf{y} \Rightarrow \mathbf{y} \equiv_g \mathbf{x}$ for some $g \in \mathcal{F}$
- Transitive: $\mathbf{x} \equiv_f \mathbf{y}$ and $\mathbf{y} \equiv_g \mathbf{z} \Rightarrow \mathbf{x} \equiv_h \mathbf{z}$ for some $h \in \mathcal{F}$

This is illustrated in Figure 3 where points on a surface are mapped based on partial symmetries. Having an equivalence relation of such mapping means that if we construct a graph with the surface points as nodes and the mappings between them as edges, then each connected component of this graphs will be a clique. The vertices in each clique form an equivalence class with respect to \mathcal{F} .

The equivalence classes given by the set of mappings already provide a decomposition of the input into sets of points. However, there are usually infinitely many such classes (or cliques), which makes this kind of decomposition impractical. In the following, we therefore cluster cliques with the same set of mappings.

Pointwise correspondence sets: For any $\mathbf{x} \in S$, let $F_o(\mathbf{x})$ denote the set of all outgoing pointwise correspondences:

$$F_o(\mathbf{x}) = \{f \mid (\mathcal{P}, f) \in \mathcal{F}, \mathbf{x} \in \mathcal{P}\} \quad (3)$$

Instances of microtiles: Comparing F_o introduces a new equivalence relation on S : A maximal connected component of points \mathbf{x} that have the same set $F_o(\mathbf{x})$ forms the same *instance* of a *microtile*. We require connected components at this point because, as we show later, these form the building blocks required for constructing new shapes.

Classes of Microtiles: Two instances of microtiles τ_1, τ_2 are from the same *class* if and only if they are matched by a correspondences $f \in \mathcal{F}$. This is because if two instances share a single pointwise correspondence, the whole instances must already be mapped by a correspondence from \mathcal{F} , by the following arguments: First, by definition, the same set of transformation applies to all points in a tile, implying that all points are mapped simultaneously. Second, any transformation is required to be a homeomorphism. Thus, connected components are preserved. In the following, we will denote the microtile decomposition of a shape \mathcal{S} by $\mu(\mathcal{S})$. Examples of microtiles are visualized in Figures 2 and 3. Microtiles are colored by their classes.

Properties of Microtiles

Elementary properties: There are several important properties of microtiles: First of all, for any \mathcal{S} there **exists a unique** decomposition $\mu(\mathcal{S})$. This follows directly from the definition. In addition, if a point \mathbf{x} is not equivalent to any other points, it belongs to a microtile characterized by the set $F_o(\mathbf{x})$ containing only the identity mapping.

For any point $\mathbf{x} \in \mathcal{S}$ there is **exactly one microtile** $\tau \in \mu(\mathcal{S})$ that covers \mathbf{x} in \mathcal{S} . This is also straight-forward: for each point $\mathbf{x} \in \mathcal{S}$ there is exactly one set of mappings to equivalent points. Thus there is exactly one microtile that contains \mathbf{x} . Thus, microtiles form a **disjoint** partition of \mathcal{S} .

Finally, a microtile can be globally mapped to itself, but **not partially**. If a point on the tile can be mapped to another one by f , then f maps all other points on the same tile as well, because otherwise, we would have introduced varying cliques of mappings within the same tile, contradicting the definition.

Cutting: We can define instances of microtiles in an alternative way: An instance τ of a microtile is a maximal, connected set such that:

$$\forall \mathbf{x}_1, \mathbf{x}_2 \in \tau : \forall (\mathcal{P}, f) \in \mathcal{F} : (\mathbf{x}_1 \in \mathcal{P} \Leftrightarrow \mathbf{x}_2 \in \mathcal{P})$$

The equivalence to the definition above is straightforward. This formulation shows that we can characterize the tile instances by the intersection of the domains of all partial mappings. Intuitively, we can think of *cutting* at the boundaries of all partial matches.

Combinations of tiles: Let $Q \subseteq \mu(\mathcal{S})$ be a set of tile instances. We then consider a union of tiles $\bigcup_{\tau \in Q} \tau$. We now want to find the set of symmetry transformations of Q , i.e., the set of $f \in \mathcal{F}$ that map all tile instances in Q simultaneously back to \mathcal{S} . From the definition, this is obviously $\bigcap_{\tau \in Q} F_o(\tau)$, i.e., the intersection of all mapping functions associated with the individual tiles. This also holds for if Q contains arbitrary fragments of tiles: Otherwise, we would

have partial symmetries of tiles. Hence, the microtile decomposition encodes all partial symmetries of \mathcal{S} and they can be computed easily and efficiently by set operations.

Permutation groups: Microtiles also unveil an algebraic regularity model: If we restrict any pair of mappings $\{f, f^{-1}\} \in \mathcal{F}$ that acts on a tile τ to the domain of this tile instance, we obtain an operation that swaps two tiles $\tau, f(\tau)$, not affecting the shape of \mathcal{S} . The set of all such permutations for all microtiles generates a group of permutations that characterize the symmetries of the shape. Please note that neither the transformations involved nor the input set \mathcal{F} has a group structure in general. By definition, the permutation group of each tile class is maximal with respect to sets of possible permutations, and the union of tiles of the same class are the maximal subset of \mathcal{S} with that property.

4. Microtiles for r -Symmetry

The above definition of microtiles is still abstract; whether the concept of microtiles is useful or not depends on the input set of mappings that determine the building blocks. In order to demonstrate a concrete application, we derive some of properties for a microtile decomposition based on partial r -symmetry as defined by Bokeloh et al. [BWS10].

r -Similarity and r -Symmetry

From now on, let $\mathcal{T} = E(3)$ be the group of rigid transformations of \mathbb{R}^3 . Let $N_r(\mathbf{x})$ be the spherical r -neighborhood of \mathbf{x} in \mathcal{S} with respect to Euclidean distance. We first define a local notion of equivalence by matching neighborhoods of points: The *points* \mathbf{x} and \mathbf{y} are r -similar under a transformation $\mathbf{T} \in \mathcal{T}$ if and only if their local neighborhoods match under a rigid motion. We denote this by:

$$\mathbf{x} \overset{r, \mathbf{T}}{\leftrightarrow} \mathbf{y} \quad :\Leftrightarrow \quad \mathbf{T}(N_r(\mathbf{x})) = N_r(\mathbf{y})$$

r -symmetry: This relation, which we call *r -symmetry*, is an equivalence relation because $E(3)$ forms a group and its actions are isometries with respect to Euclidean distance, i.e., they do not change r . In the following, let \mathcal{F} denote the partial mappings induced by the set of all pairwise r -symmetry relations.

r -similarity: We can also use the same matching model to compare shapes: Let $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{S}$ be two shapes in \mathcal{S} , with distance larger than r , and $\mathbf{x} \in \mathcal{S}_1, \mathbf{y} \in \mathcal{S}_2$. We call \mathcal{S}_2 r -similar to \mathcal{S}_1 if every point in \mathcal{S}_2 is equivalent to at least one (arbitrary) point in \mathcal{S}_1 . Formally:

$$\mathcal{S}_2 \text{ is } r\text{-similar to } \mathcal{S}_1 \quad :\Leftrightarrow \quad \forall \mathbf{y} \in \mathcal{S}_2 : \exists \mathbf{x} \in \mathcal{S}_1 : \mathbf{x} \overset{r, \mathbf{T}}{\leftrightarrow} \mathbf{y}.$$

Obviously, r -similarity is reflexive and transitive but in general not symmetric.

In the following, we will consider decomposition of surfaces (e.g. triangle meshes) into microtiles with respect to r -symmetry, and we will show a number of interesting properties that these particular microtiles have. The main result of this paper is the proof that these microtiles are a canonical description of the space of r -similar shapes to the input shape. Before we continue, we will make further restrictions to the input to make the further model well defined:

- S is bounded.
- S is a 2-manifold (with or without boundary).
- S is piecewise smooth; we assume a union of a finite number of facets, each of which is an algebraic surface bounded by a finite set of algebraic curves.

The last point allows representations like meshes of trimmed NURBS, which can represent (exact) rotational symmetries. Our current practical implementation is actually based on triangle meshes, where exact matching rules these out.

Types of Microtiles

When constructing the microtiles according to r -symmetry, one can distinguish between two types of microtiles – those with finitely and those with infinitely many equivalent counterparts. The restrictions on the input surface to a finite mesh of polygons or patches allow us to simplify the cases where infinitely many transformations map a point to an r -symmetric one. If a tile is characterized by infinitely many transformations, then these can be represented by a continuous function parameterizable in one or two dimensions. This kind of symmetries (here r -symmetries) are analyzed by Gelfand and Guibas [GG04] and the symmetric points are called *slippable*. Thus we have two types of tiles with respect to the size of their defining transformation sets:

- **Discrete pieces:** A single piece of geometry that is instantiated once or more by a (finite) discrete set of symmetry transformations.
- **Slippable pieces:** These pieces (and their r -neighborhood in their instantiation) are slippable. They can be planar, spherical, cylindrical, helical, surfaces of revolution, or extrusions, e.g. edges (see Gelfand and Guibas [GG04]).
 - Planar, spherical, cylindrical microtiles: These form area elements, represented by a single point and its slippage properties. They can be expanded to arbitrary area covering the underlying primitive, but need to be bounded by other tiles (or self-closed, for spheres)
 - Helical tiles, extrusions, surfaces of revolution: These can be expanded in one dimension, forming curve primitives. They are either self-closed (surfaces of revolution) or bounded by other tiles (helical tiles, edges).

In Figure 2 the gray-colored points are all instances of a single class of microtiles because their r -neighborhoods

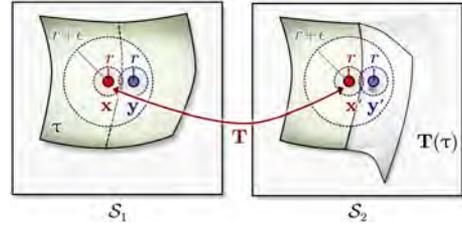


Figure 4: $(r + \epsilon)$ -similar surfaces must be constructed out of r -microtiles. If not, a contradiction would arise at the boundary, between matching (x, x') and non-matching (y, y') points.

are planar. We call them *2-slippable* because their set of r -symmetry transformations can be parametrized in 2 dimensions. Analogously the yellow-colored points correspond to multiple *1-slippable* classes of microtiles. For triangle meshes, only planar 2-slippable elements and 1-slippable extrusions are possible.

Complexity: The set of microtiles can be encoded with a finite number of real parameters. For a triangle mesh, this is easy to see. A single triangle is always described by a finite set of different microtile classes. Inductively, adding a triangle to a set of finite classes of tiles can only create a finite number of new such classes, bounded by a finite graph of straight edges. We conjecture that a similar property holds for a mesh of general algebraic surfaces as well, but a formal proof is beyond the scope of this paper.

5. The Space of r -Similar Shapes

In this section, we discuss how the microtiles of an exemplar S characterize the space of all shapes that are r -similar to S . We discuss two different aspects: First, we show that all such shapes can be assembled from a disjoint union of microtiles, and that this decomposition is unique (up to global symmetries of the tiles, which has no effect). Second, we give some necessary conditions for a shape grammar that constrains how the tiles can possibly be fit together.

Theorem: The main result that we want to show is the following: Let S_1, S_2 be valid input surfaces in the sense of our definition. Furthermore, let S_2 be $(r + \epsilon)$ -similar to S_1 for $\epsilon > 0$. Then S_2 can be constructed (completely covered) by a disjoint union of r -microtiles from $\mu(S_1)$, using only transformations from \mathcal{T} to arrange the pieces.

Proof: Any point in S_2 must correspond to a point in S_1 , because of $(r + \epsilon)$ -similarity. This implies we must take the whole tile at a time (Fig. 4). Assume that this was not the case. Let $x' \in S_2$ be a point within such a tile fragment. It

must be $(r + \epsilon)$ -similar to a point $\mathbf{x} \in S_1$ under a transformation $\mathbf{T} \in \mathcal{T}$, i.e., $\mathbf{T}(\mathbf{x}) = \mathbf{x}'$. Assume $\mathbf{x} \in \tau$, where τ is the microtile \mathbf{x} is contained in. It always exists because the tiles completely and disjointly partition S_1 . We have assumed that $\mathbf{T}(\tau) \not\subseteq S_2$. Thus, there is a boundary within τ that separates the part that maps to S_2 from the part that does not. Let \mathbf{y} be a point that is not r -similar to $\mathbf{y}' = \mathbf{T}(\mathbf{y}) \notin S_2$. Because τ is a connected component and \mathbf{T} preserves connectivity within τ , we can choose \mathbf{x}' (and thereby \mathbf{x}) and \mathbf{y} such that the distance is smaller than $\epsilon/2$. This yields a contradiction: Because of $(r + \epsilon)$ -similarity of \mathbf{x} and \mathbf{x}' , both points \mathbf{x}, \mathbf{y} match under \mathbf{T} along with their r -neighborhood. Therefore, \mathbf{y} and \mathbf{y}' must be r -similar, contradicting the assumptions of the contrary.

Unique, disjoint construction: Let us assume that we have two different decompositions of S_2 into microtiles from $\mu(S_1)$. This means that the tiles overlap. Let $\mathbf{x}' \in S_2$ be a point from such an overlap region, with tiles $\mathbf{T}_1(\tau_1), \mathbf{T}_2(\tau_2)$. Because of $(r + \epsilon)$ -similarity, $\mathbf{T}_1^{-1}(\mathbf{x}')$ and $\mathbf{T}_2^{-1}(\mathbf{x}')$ must be r -similar points in S_1 . According to Section 3, this implies that they are from the same tile. Thus, the overlapping tiles $\mathbf{T}_1(\tau_1), \mathbf{T}_2(\tau_2)$ either match up completely, under a global symmetry of $\tau_1 = \tau_2$, which is possible, or we have found a partial symmetry within a tile, which is impossible. The same arguments also show that the tiles must be disjoint. \square

Discussion: The theorem shows that $(r + \epsilon)$ -similar shapes can be build out of r -microtiles, so there is a small gap, which is required to guarantee continuity across boundaries. Can we close this gap? While we so far cannot show a similar result for $\epsilon = 0$, we can at least give an informal argument that the gap almost always does not matter: First, let us assume that for the chosen value of r , the shapes of microtiles changes continuously with r . In particular, no tiles are newly introduced or disappear for infinitesimal changes. Then, the construction out of r - and $(r + \epsilon)$ -microtiles differs only by a set of points with zero measure because we can choose $\epsilon > 0$ as small as we like. The continuity assumption holds for all r except from a finite set of numbers, i.e., the likelihood of choosing a bad value of r at random from a given, non-trivial interval is zero. The reason is that for discontinuous changes, the regions around the discrete boundaries of our (triangle) mesh need to connect or disconnect, and there is only a finite number of such events. In practice, precision limits of a floating point implementation will probably override these considerations anyway.

6. Towards a Microtile Shape Grammar

In addition to having a unique decomposition into microtiles, we can constrain the construction of r -similar shapes further. We will now show that we cannot only learn the tiles from the exemplar, but also some rules of how these can possibly be assembled. We first have to define boundaries of a tile:

Let τ_1, τ_2 be two instances of microtiles of different type.

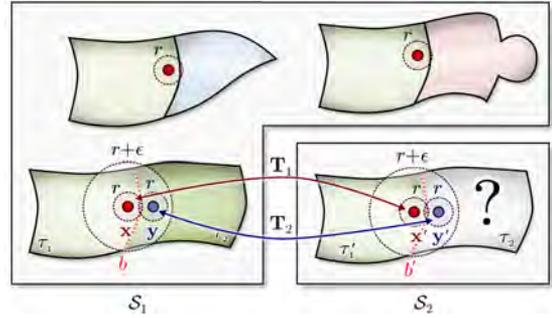


Figure 5: Microtiles can only be assembled with pairwise adjacency relations as in the exemplar (left).

The boundary b between τ_1, τ_2 is the set of points that have distance zero to both τ_1, τ_2 , i.e., $b = \overline{\tau_1} \cap \overline{\tau_2}$, where the bar denotes set closure. Because we assume manifold input, and tiles are disjoint, the boundaries are (multi-)curves (possibly consisting of multiple fragments). Only two tiles can share a common curve except from isolated points that can be a boundary to more than two tiles.

We now show that tiles can only be assembled along boundaries that have been observed in the exemplar. We have the following theorem:

Theorem: Let S_2 be $(r + \epsilon)$ -similar to S_1 . Let τ_1, τ_2 be two different instances of r -microtiles of S_1 . Let $\tau'_1 = \mathbf{T}_1(\tau_1), \tau'_2 = \mathbf{T}_2(\tau_2), \mathbf{T}_1, \mathbf{T}_2 \in \mathcal{T}$ be instances of tiles of these two types in S_2 , i.e., $\tau'_1, \tau'_2 \subset S_2$. Let τ'_1, τ'_2 share a common boundary b' . Then, there must exist instances of τ_1, τ_2 in S_1 that share a boundary b . Furthermore, we can always find a pair of instances $\tau_1, \tau_2 \in S_1$ such that $\mathbf{T}_1 = \mathbf{T}_2$ and $b' = \mathbf{T}_1(b)$.

Proof: We show that $b \subset S_1$ must exist (Fig. 5) with arguments similar to the proof of the previous theorem. We know that $\tau'_1 \subset S_2$ corresponds to one or more tiles in S_1 . Let $\mathbf{x}' \in \tau'_1$ and $\mathbf{y}' \in \tau'_2$ be two points, both with distance less than $\epsilon/2$ to the boundary b' . Because S_2 is $(r + \epsilon)$ -similar to S_1 , \mathbf{x}' and its $(r + \epsilon)$ -neighborhood map to a point $\mathbf{x} \in S_1$ which is contained in one of these tiles. Let τ_1 denote this tile, and $\mathbf{T}_1^{-1} \in \mathcal{T}$ denote the mapping. Because of the $(r + \epsilon)$ -similarity, $N_r(\mathbf{x}')$ and $N_r(\mathbf{y}')$ are both mapped back to S_1 under \mathbf{T}_1^{-1} . As before, this implies that $\mathbf{T}_1^{-1}(\tau'_2) \subset S_1$ is also an instance of the microtile τ_2 (otherwise, we would have found a partial symmetry of τ_2). Therefore, we have constructed a pair of corresponding tiles with $\mathbf{T}_1 = \mathbf{T}_2$, which implies that the boundary $b = \mathbf{T}_1^{-1}(b')$ exists in S_1 . \square

This shows that we can learn restrictions how to connect microtiles from the exemplar: We can collect all boundaries along which microtiles are neighboring in the exemplar S_1 and allow only these combinations for building new shapes. For *discrete* microtiles, this is straightforward (we

just need to enumerate the observed combinations). For *continuous* microtiles, the situation is more complicated: continuous tiles can be neighboring to themselves. For example, the tile of a plane consists of a single point only; finite pieces are formed by collections of points. This implies that continuous microtiles that have themselves as neighbors (which actually can be expected to be almost always the case) can be extended to form arbitrary kinematic surfaces of the type the tile corresponds to (planes, spheres, helical surfaces, etc). The only constraint is the boundary: The continuous tile must be bounded by observed boundaries to tiles of other types, which again can be continuous or discrete (the sphere type is an exception, because we can form a closed sphere, which is a bounded set that respects the neighborhood rule). Any combinations of these boundaries are possible, as long as the adjacency of the boundary elements themselves is compatible as well. As an example, consider a flat wall with a single window in it that forms a single tile. This means, that in a newly constructed shape, we can insert an arbitrary number of windows into the plane, rotated and translated arbitrarily, as long as their distance is larger than r .

Discussion: The rules are necessary for assembling $(r + \epsilon)$ -similar objects, but not sufficient. Following the (informal) arguments of Bokeloh et al. [BWS10], we conjecture that our rules of assembling microtiles according to previously observed “docking sites” are sufficient for creating r -similar objects, leaving only a gap of ϵ in the class of shapes described. However, a formal proof is subject to future work, as well as a constructive characterization of a shape grammar that can directly create new objects.

7. Computation of the Decomposition

We now describe an algorithm that computes the microtile decomposition of a manifold triangle mesh in polynomial time. We start with the abstract algorithm; its correctness and a concrete prototype implementation are discussed in Appendix A and B. Following Mitra et al. [MPWC12], we proceed in three stages: **feature extraction**, **aggregation** and **extraction**.

Feature Selection: We cannot test infinitely many transformations as they appear in slippable regions. Therefore, we first perform slippage analysis for all r -neighborhoods [GG04]. Afterwards, we ignore slippable regions in the remaining computation. Then, we compute *line features* [BBW*09]. For a triangle mesh, this is the subset of the edges with adjacent non-coplanar faces. We then generate *candidate transformations* by matching line features. We discuss how the transformations should be computed in Appendix A.

Aggregation: For each candidate transformation \mathbf{T} and its inverse we match the whole scene \mathcal{S} against $\mathbf{T}(\mathcal{S})$. An exact algorithm would compute an intersection of the two meshes (in practice, we will resort to an approximation, as discussed in Section B, using voxels rather than triangle fragments as representation [BWS10]). For each matching fragment, we record the matching element and the transformation indices in a table. After all transformations are processed the table encodes all detected partial r -symmetries for the shape.

Extraction: We extract a segmentation of the input scene into microtiles by region growing starting at an arbitrary (non-processed) element and expanding the current tile with elements that have the same set of symmetry transformations. We use the table we computed in the previous step to look up the transformations that map the geometry to r -symmetric parts of the surface. After the initial segmentation, we compute the equivalence classes of points (the cliques discussed in Section 3 and Figure 3). We transform the voxels that belong to a given microtile, and search in the overlapping voxels for the equivalent microtile instance.

8. Practical Experiments

We have implemented a simple prototype of the algorithm outlined above. We follow the method of [BWS10] and use a volumetric grid to discretize the symmetry information: cubes of side length h are annotated with transformations.

We have applied our prototype implementation to a few scenes to visualize the structure of the decomposition. For the tests we set the radius of symmetry to 0.008 (Figure 6) or 0.016 (all other tests) of the diagonal of the bounding box of the scene. The voxel size was set to $1/512$ of the diagonal ($1/256$ for Fig. 8). To prevent errors due to coarse discretization, classes of microtiles were computed only for microtiles larger than 32 voxels. Very small tiles usually indicate places where a finer discretization is required and we could not reliably compute the equivalence classes of such microtiles. Computing of the candidate transformations and the table that stores the set of transformations for each voxel are implemented in parallel. All test were performed on a single Intel Core 2 Quad Q9400 CPU with 4 cores running at 2.66GHz.

Fig. 6 shows a very simple test scene composed out of boxes. The left hand side shows a scene of three different shapes, decomposed simultaneously. On the right, a simpler scene of independent boxes is decomposed. For these simple scenes, we obtain accurate results up to the resolution of the discretization. In Fig. 7, we apply the algorithm to more complex meshes of architectural objects. We depict 2-slippable tiles in gray, 1-slippable in yellow (irrespective of the class), and only show the large tiles, as explained above. The corresponding unassigned area is shown in dark gray.

The computed decompositions are in most regions qualitatively correct, however, the grid-discretization leads to certain variations at the boundary. We observe some unassigned area, but its diameter is below r in all of the examples. Because the voxel-discretization does not permit a 1 : 1 mapping, boundaries show some variability within voxel resolution (particularly visible at the sides of the courthouse). Furthermore, rotational patterns are numerically problematic (e.g., oversegmentation of the steps of the staircase). Similar results are obtained for the models in Figure 8. We compare our results to the previous method by Bokeloh et al. [BBW*09], which is computationally mostly similar but uses (as most others) simple region growing for segmentation. The method is similarly susceptible to discretization and boundary artifacts. It does not capture all symmetries, but samples prominent representatives due to the area/instance ratio heuristic employed. Global symmetries of the steps are detected, which do not affect the microtiles but are obtained implicitly with our new approach.

Our implementation is only intended as a proof of concept, but there are some direct applications: We can determine whether two shapes are r -similar, by matching their respective microtiles. The three box-sculptures in Fig. 6 are made of the same tiles, except from the leftmost, which contains one extra, unique tile, colored violet. Similarly, the isolated tower at the left of the castle in Fig. 7 is r -similar to the castle, which contains additional tiles. A further example is demonstrated in Fig. 8 (right). A sequence of models with increasing complexity is decomposed into microtiles, revealing the redundancy in the model collection.

The runtime of the decomposition is still rather large, as the algorithm performs all pairwise comparisons explicitly. Small test scenes compute in a few minutes, medium complexity scenes such as the castle require 1 hour (see Fig. 2,7). Both the number of features and the required resolution for representing the symmetries are limiting factors, and both act quadratically on the run-time.

9. Conclusions and Future Work

We have presented a formal model for decomposing a model into building blocks when we have correspondences that identify equivalent surface points. The basic idea is very simple: we group all connected points that are mapped by the same set of mapping functions. Nevertheless, the decomposition has a number of remarkable properties. As an abstract decomposition, it is a unique, disjoint partition that encodes all consistent matches of subsets of the input surface. Applying it to the concrete correspondence model of rigid matches of local neighborhoods, we get the even stronger result that the tiles of the decomposition are sufficient to uniquely construct any model that is similar to the decomposed one in this sense.

Discussion: We see two main application areas where our approach is of utility. First, it provides a canonical representation of redundancy in shapes. With respect to given correspondences, we can consider the microtiles as elementary units. This has obvious applications in shape compression and could be used to analyze collections of shapes for similarity. In addition to relating the tiles themselves, the relative arrangement of the tiles could probably, in future work, be used to characterize geometry independent of the actual geometry. The second application area is inverse procedural modeling, where we can characterize the space of shapes that is locally similar to an exemplar.

Limitations and future work: Our current practical implementation is limited to exact input. Even for synthetic triangle meshes, the method is not very robust; in particular a rotational matching of local features can lead to artifacts. Unlike the approach of Bokeloh [BWS10], which uses a similar computational framework, we are very dependent on not missing any pairwise matches, as this directly impacts the results. We think that an approximate formulation of fitting a consistent decomposition to presumably noisy or ambiguous data would be the appropriate approach to overcome these problems (the obvious alternative, exact mesh intersections, would be complex and costly, and almost any real-world 3D model would still break the assumptions). Incorporating some of the ideas of Lipman et al. [LCDF10] might be one avenue towards a more robust computation. Finally, unlike Bokeloh et al.'s representation [BWS10], our assembly rules do not yet lend themselves to automatic modeling. So far, we have necessary but not yet sufficient conditions for $(r + \epsilon)$ -similarity. Furthermore, the resulting rules form a pairwise jigsaw-puzzle so that the automated assembly is a hard combinatorial problem. In future work, we would like to improve the characterization of the grammar and study algorithms for solving the resulting assembly problems. In particular, there might different choices of rich enough subsets of all r -similar shapes that permit an efficient assembly.

Acknowledgments

The authors wish to thank the anonymous reviewers for their valuable comments. This work was supported by the International Max-Planck Research School for Computer Science, the Cluster of Excellence on Multi-Modal Computing and Interaction, the Max-Planck-Center for Visual Computing and Communication, by NSF grants FODAVA 808515 and CCF 1011228, and by a Google Research Award.

References

- [ARB07] ALIAGA D., ROSEN P., BEKINS D.: Style grammars for interactive visualization of architecture. *IEEE Trans. Vis. Comp. Graph.* 13, 4 (2007), 786–797. doi:10.1109/TVCG.2007.1024.1

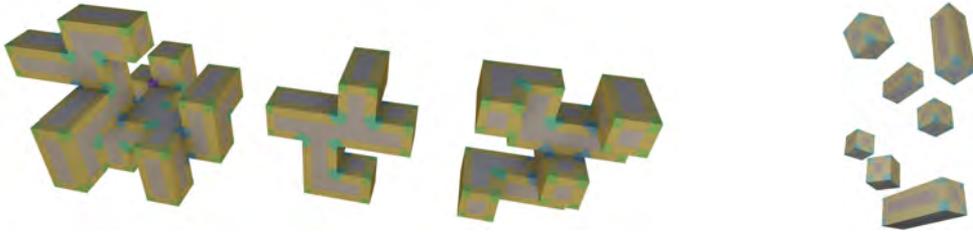


Figure 6: Simple models. Left: Simultaneous decomposition of three simple meshes in one scene, $S = S_1 \cup S_2 \cup S_3$. Right: decomposition of 7 boxes in one scene. Equivalent tiles in each figure were computed and colored automatically. The number of different microtiles are 6, 5, and 5 for S_1, S_2, S_3 . Each box on the right has 3 microtiles, as expected. The run-time for the decomposition of the first three meshes was around 10 min, the boxes on the right took approx. 1min.

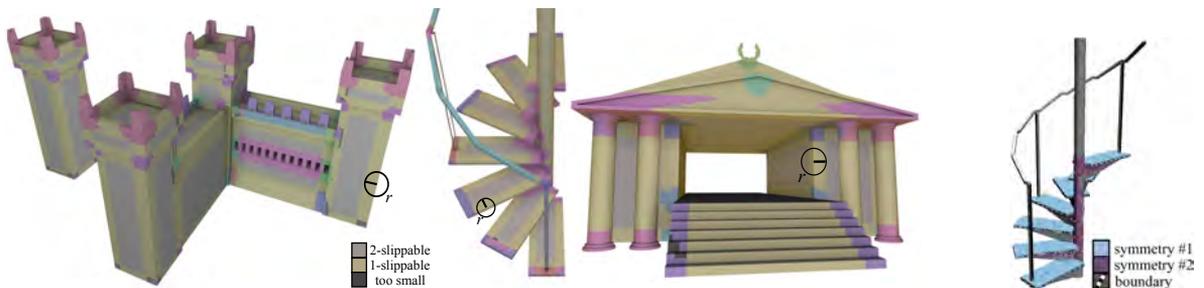


Figure 7: More complex models: Castle (left), staircase (middle), courthouse (right) models decomposed into r -microtiles. The classes can be computed reliably only for microtiles larger than 32 voxels. The diameter of this uncertain area is smaller than r . Run-time: around 1 hour (castle), 18 min (staircase), and 40 min (courthouse). Far right: comparison to [BBW*09].

- [BBW*09] BOKELOH M., BERNER A., WAND M., SEIDEL H. P., SCHILLING A.: Symmetry detection using feature lines. *Computer Graphics Forum* 28, 2 (2009), 697–706. 2, 7, 8, 9, 10
- [BIT04] BHAT P., INGRAM S., TURK G.: Geometric texture synthesis by example. In *Symp. Geometry Processing* (2004). 3
- [BWKS11] BOKELOH M., WAND M., KOLTUN V., SEIDEL H.-P.: Pattern-aware shape deformation using sliding dockers. *ACM Trans. Graph.* 30 (Dec. 2011), 123:1–123:10. 2
- [BWS10] BOKELOH M., WAND M., SEIDEL H.-P.: A connection between partial symmetry and inverse procedural modeling. *ACM Trans. Graph.* 29 (July 2010), 104:1–104:10. 1, 2, 4, 7, 8
- [EL99] EFROS A. A., LEUNG T. K.: Texture synthesis by non-parametric sampling. In *Proc. Int. Conf. Comp. Vision* (1999). 3
- [GCO06] GAL R., COHEN-OR D.: Salient geometric features for partial shape matching and similarity. *ACM Trans. Graph.* 25, 1 (2006), 130–150. 2
- [GG04] GELFAND N., GUIBAS L.: Shape segmentation using local slippage analysis. In *Proc. Symp. Geom. Processing* (2004). 5, 7
- [GSH*07] GAL R., SHAMIR A., HASSNER T., PAULY M., COHEN-OR D.: Surface reconstruction using local shape priors. In *Proc. Symp. Geometry Processing* (2007). 2
- [GSMCO09] GAL R., SORKINE O., MITRA N., COHEN-OR D.: iwires: An analyze-and-edit approach to shape manipulation. *ACM Trans. Graph.* 28, 3 (2009). 2
- [HF97] HART J., FLYNN O. C. P.: Similarity hashing: A computer vision solution to the inverse problem of linear fractals. *Fractals* 5 (1997), 35–50. 1
- [LCDF10] LIPMAN Y., CHEN X., DAUBECHIES I., FUNKHOUSER T.: Symmetry factored embedding and distance. *ACM Transactions on Graphics (SIGGRAPH 2010)* (July 2010). 2, 3, 8
- [Mar07] MARTINET A.: *Structuring 3D Geometry based on Symmetry and Instancing Information*. PhD thesis, INP Grenoble, 2007. 2
- [MBB10] MITRA N. J., BRONSTEIN A., BRONSTEIN M.: Intrinsic regularity detection in 3d geometry. In *ECCV* (2010). 2
- [Mer07] MERRELL P.: Example-based model synthesis. In *Symp. Interactive 3D Graphics and Games* (2007), pp. 105–112. 3
- [MGP06] MITRA N. J., GUIBAS L. J., PAULY M.: Partial and approximate symmetry detection for 3d geometry. *ACM Trans. Graph.* 25, 3 (2006), 560–568. 1, 2
- [MPWC12] MITRA N. J., PAULY M., WAND M., CEYLAN D.: Symmetry in 3d geometry: Extraction and applications. In *EUROGRAPHICS State-of-the-art Report* (2012). 1, 2, 7
- [MY*10] MITRA N. J., YANG Y.-L., YAN D.-M., LI W., AGRAWALA M.: Illustrating how mechanical assemblies work. *ACM Transactions on Graphics* 29, 3 (2010). 2
- [OSG08] OVSJANIKOV M., SUN J., GUIBAS L.: Global intrinsic symmetries of shapes. In *Eurographics Symposium on Geometry Processing (SGP)* (2008). 2
- [PMW*08] PAULY M., MITRA N. J., WALLNER J., POTTMANN H., GUIBAS L.: Discovering structural regularity in 3D geometry. *ACM Trans. Graph.* 27, 3 (2008). 2

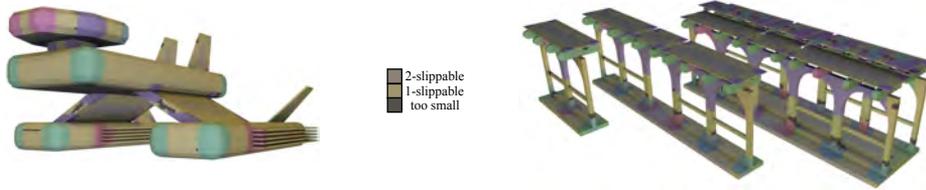


Figure 8: A spaceship model (left) - the tiles are recognized correctly, but again, some unassigned area remains (gray). A cascade of models with increasing complexity (right) - the newly added parts create new tiles.

- [PSG*06] PODOLAK J., SHILANE P., GOLOVINSKIY A., RUSINKIEWICZ S., FUNKHOUSER T.: A planar-reflective symmetry transform for 3D shapes. *ACM Trans. Graph.* 25, 3 (2006). 2
- [RBBK10] RAVIV D., BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Full and partial symmetries of non-rigid shapes. *Intl. Journal of Computer Vision (IJCV)* 89, 1 (August 2010), 18–39. 2
- [ŠBM*10] ŠT’AVA O., BENEŠ B., MĚCH R., ALIAGA D., KRIŠTOF P.: Inverse procedural modeling by automatic generation of l-systems. *Computer Graphics Forum* (2010). 1
- [SKS06] SIMARI P., KALOGERAKIS E., SINGH K.: Folding meshes: hierarchical mesh segmentation based on planar symmetry. In *Proc. Symp. Geometry Processing* (2006), pp. 111–119. 2
- [WL00] WEI L., LEVOY M.: Fast texture synthesis using tree-structured vector quantization. In *Proc. Siggraph 2000* (2000). 3
- [WXL*11] WANG Y., XU K., LI J., ZHANG H., SHAMIR A., LIU L., CHENG Z., XIONG Y.: Symmetry hierarchy of man-made objects. In *Proc. Eurographics* (2011). 2
- [ZHW*06] ZHOU K., HUANG X., WANG X., TONG Y., DESBRUN M., BAINING GUO H.-Y. S.: Mesh quilting for geometric texture synthesis. *ACM Trans. Graph.* 25, 3 (2006), 690–697. 3

Appendix A: Correctness of the Tile Extraction Algorithm

The algorithm outlined in Section 7 computes a correct microtile decomposition if there is no rigid transformation $T \in \mathcal{T}$ such that points on different microtiles of the output are r -symmetric under T . In other words we need to compute all transformations that can map two r -regions on \mathcal{S} symmetrically.

There are three cases of r -neighborhoods, that we need to consider. They can be 2-slippable, 1-slippable or non-slippable. The 2-slippable surfaces on a triangle mesh can only be planes. They are characterized by a single microtile, and we check for its existence in the input model during slippage analysis. If a point on a triangle mesh is 1-slippable, then its r neighborhood contains one or more edges, that are parallel to the line features of \mathcal{S} . Even though 1-slippable microtiles are mapped by infinitely many transformations to symmetric tiles, it is possible to consider segments of non-zero lengths instead of the infinitely small microtiles. Detecting discrete symmetries between such segments allows to decompose the 1-slippable microtiles. To this end, we need to compute all transformations that map pairs of line features to each other (s. [BBW*09]). It remains to find all transformations that map non-slippable r -neighborhoods symmetrically. Observe that any such region has to contain at least two non-parallel edges (line features).

Transformations that map pairs of non-parallel line features have to align the center of the shortest line segment between the two lines and the line directions. These transformations can be computed by exhaustively checking all possible pairs of features at distance no more than r from each other.

Appendix B: Implementation Details

A straightforward implementation of the abstract algorithm is quite slow because of the large amount of initial candidate transformations that need to be processed. Therefore, we try to discard candidate transformations as early as possible. We ignore transformations between corner features that fail to align all edges meeting at the corner. We discard duplicate transformations, and those that map corner points that are not r -symmetric. If a transformation matches two line features, we only consider it as a valid candidate if all 1-slippable voxels along the shorter feature are mapped to symmetric (1-slippable) voxels along the longer feature.

The voxel- and feature-based approach has some drawbacks we need to address: The precision of the decomposition is limited by the discretization of the scene. Because of that, equivalent microtiles will be decomposed in different ways, and will have a slightly different voxel representations. We are therefore unable to directly compute the equivalence classes or cliques on per-voxel basis. After an initial segmentation based on pairwise matching, we compute the classes of equivalent microtiles. We transform the microtiles with each of the r -symmetry transformations to find their equivalent counterparts. Because a single tile will overlap multiple other parts, we gather votes from the overlapping voxels and select the tile that is most similar in terms of size and set of symmetry transformations. This approach relies on a fine discretization. Another practical issue is related to the precision with which we can compute the matrices for the actual transformations. Because we align each of them at a single corner feature, the mapping becomes imprecise with the distance to the feature. In combination with the voxel quantization, the result is that near tile borders, r -symmetry detection becomes inconsistent and the set of transformations for many of the voxels is incomplete. This shows up as a large amount of small microtiles, that make further extraction of the equivalence classes virtually impossible. We address this problem by a filter: near tile boundaries, we merge small tiles to neighboring larger one, whenever the voxel distribution and set of transformations of the smaller one suggest that it can be a part of the larger tile. We never discard a tile that has at least one voxel completely surrounded by voxels on the same tile. This ensures that the area we filter will converge to zero if the voxel size does so.