



HAL
open science

Evocube: a Genetic Labeling Framework for Polycube-Maps

Corentin Dumery, François Protais, Sébastien Mestrallet, Christophe Bourcier, Franck Ledoux

► **To cite this version:**

Corentin Dumery, François Protais, Sébastien Mestrallet, Christophe Bourcier, Franck Ledoux. Evocube: a Genetic Labeling Framework for Polycube-Maps. Computer Graphics Forum, In press, 10.1111/cgf.14649 . hal-03657779v3

HAL Id: hal-03657779






<https://hal.science/hal-03657779v3>

Submitted on 8 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Evocube: a Genetic Labeling Framework for Polycube-Maps

C. Dumery^{1,4} , F. Protais² , S. Mestrallet¹ , C. Bourcier¹ , F. Ledoux³ 

¹CEA, Université Paris-Saclay

²Université de Lorraine, CNRS, Inria, LORIA

³CEA DAM, LIHPC, Université Paris-Saclay

⁴EPFL, School of Computer and Communication Sciences




Figure 1: Collection of labelings and polycube-maps of general shapes generated using the Evocube framework.

Abstract

Polycube-maps are used as base-complexes in various fields of computational geometry, including the generation of regular all-hexahedral meshes free of internal singularities. However, the strict alignment constraints behind polycube-based methods make their computation challenging for CAD models used in numerical simulation via Finite Element Method (FEM). We propose a novel approach based on an evolutionary algorithm to robustly compute polycube-maps in this context.

We address the labeling problem, which aims to precompute polycube alignment by assigning one of the base axes to each boundary face on the input. Previous research has described ways to initialize and improve a labeling via greedy local fixes. However, such algorithms lack robustness and often converge to inaccurate solutions for complex geometries. Our proposed framework alleviates this issue by embedding labeling operations in an evolutionary heuristic, defining fitness, crossover, and mutations in the context of labeling optimization. We evaluate our method on a thousand smooth and CAD meshes, showing Evocube converges to accurate labelings on a wide range of shapes. The limitations of our method are also discussed thoroughly.

CCS Concepts

• *Mathematics of computing* → *Mesh generation; Evolutionary algorithms;*

1. Introduction

Scientific computational analysis based on finite element method (FEM) or finite volume method (FVM) is increasingly used to model engineering problems. Recent developments incorporate coupled physical phenomena and require complex geometric shapes. However, FEM and FVM are limited by volumetric mesh generation. In practice, for many cases of interest, all-hexahedral meshes, or hex meshes, are preferred over tetrahedral

meshes [SJ08]. Their use drastically reduces computational cost and memory footprint while preserving numerical accuracy. Although tetrahedral meshing is now robustly performed on general 3D shapes, direct generation of good-quality hex meshes remains an open problem. To ensure high-accuracy and convergence speed, finite element analysis requires block-structured hex meshes with low cell distortion, as illustrated in Figure 2. As of today, such properties cannot be guaranteed for general shapes.




Figure 2: Block-structured hexahedral mesh of a sphere, structure on the left and resulting mesh on the right.

All-hexahedral meshing has been studied for several decades. Several approaches have emerged, relying on different techniques: advancing front [SCB92, BM93, OS00, TBM95, LW07], overlay-grid or octree-based [Sch96, Mar09, ZHB10, GSP19], frame field [HTWB11, LLX*12, RSL16, KLF16, PBS20], medial axis [SERB99, LMPS16, Qua14], or polycube-maps [THCM04, GSZ11, LVS*13]. As of today, none of these methods has been successful in automatically generating satisfying hex meshes for realistic 3D shapes. In practice, time-consuming domain partitioning is often performed interactively [PBSB07, Tak19, CHL19, LZS*21] instead. For the generation of block-structured meshes, we believe frame field and polycube-map methods are the most promising directions in state-of-the-art research. Frame fields are generated by optimizing a smoothness energy and yield the desired block-structure for simple geometries, but fail on more complex shapes. Polycube-maps, on the other hand, succeed on a broader spectrum of 3D shapes but do not provide a usable block structure since interior singularities are missing. Polycube-map generation is performed by deforming a 3D shape to obtain an orthogonal polyhedron (or *polycube* [THCM04]), and generating a volumetric map p between them. A hex mesh is then extracted from the polycube-map via subdivision following the integer grid, and morphed back onto the initial 3D shape following the inverse map p^{-1} . By construction, resulting meshes lack interior singularities and low-quality hexes are generated close to the boundary. In state-of-the-art pipelines, this issue is partially addressed with post-processing improvements that insert layers of cells along the boundary [KLSO12, CAS*19].

We observe that polycube alignment can be quickly estimated and evaluated, and propose a novel approach based on an evolutionary algorithm [Bä96] to robustly compute *polycube labelings*. Starting from a tetrahedral mesh T_Ω of the 3D domain Ω , polycube labeling consists in assigning one of the six base axis directions $\{\pm X, \pm Y, \pm Z\}$ to each face of the boundary ∂T_Ω . Adjacent triangles that are assigned the same direction form a **chart**. Following the terminology introduced by Livesu et al. [LVS*13], a labeling should define a **valid** polycube topology [EM10, SR15], and compromise between the following quality criteria illustrated in Fig. 3:

- **Fidelity:** angles between assigned directions and triangle normals remain low;
- **Compactness:** the number of charts is small;
- **Monotonicity:** boundaries between charts are exempt from *turning points*, i.e. significant changes in boundary direction.




Figure 3: Illustration of labeling properties and associated polycubes. We report the number of corners in the polycube N_c and the area distortion D_a [THCM04].

The nature of the polycube labeling problem makes genetic methods a good fit. In previous work, labelings were computed using deterministic operations which explored a limited set of solutions. We observe that generating and evaluating large quantities of labelings is significantly less costly than computing low-distortion polycube-maps for all possibilities. We thus propose an evolutionary framework that explores a wider range of solutions than previous methods. Starting from an initial solution, and over several iterations or *generations*, we generate new solutions referred to as *individuals*. At any given time, the current set of individuals being considered is referred to as *population* and is stored in a fixed-size *archive*. We measure the quality of an individual with a *fitness* function. Each generation, some individuals are sampled from the archive to undergo labeling modifications named *mutations*. Individuals with high fitness are then selected for *crossover*, where a pair of individuals generates a third one combining both of its parents traits. At the end of each generation, the fittest individuals are inserted back into the archive. This process aims to improve population quality over generations. If the content of the archive does not evolve for several consecutive generations, then we consider that the algorithm has converged and stop the process. For a large majority of inputs, *Evocube* converges towards a labeling associated with a low-distortion polycube within a dozen generations.

We validate our algorithm on a large collection of over one thousand CAD and smooth models, showcasing some in Figure 1. *Evocube* produced labelings leading to low-distortion polycubes on an overwhelming majority of inputs, as reported in Section 4. We also analyze failure cases and comment on the limitations of our method.

1.1. Related Work

Polycube labeling. Polycubes were first used in computer graphics for seamless texturing of triangulated surfaces [THCM04]. They rely on a polyhedral structure and a volumetric map, which can either be computed one after the other, or together via mesh deformation towards the orthogonal polyhedron closest to the input 3D shape. The polyhedral structure is usually defined by labeling each element of ∂T_Ω with a value that represents one of the six base axes $\{\pm X, \pm Y, \pm Z\}$. A naive labeling can be computed by assigning to each surface triangle the label closest to its normal [GSZ11]. However, this does not produce a valid structure in general and ad-




Figure 4: Starting from a triangular mesh representation of a 3D shape, (a) we generate an initial non-optimal and potentially invalid labeling using a graph-cut method [LVS*13]; (b) this labeling is then refined within our Evocube framework; (c) we generate a polycube with topology matching our labeling, and (d) a grid-like mesh is extracted, padded and morphed back onto the initial 3D shape.

ditional labeling refinement is necessary. This refinement is guided by a set of sufficient topological conditions for the existence of a valid polyhedron defined by Eppstein et al. [EM10] and used by Livesu et al. [LVS*13]. Similarly, Hu et al. [HZ16] use a modified Centroidal Voronoi Tesselation labeling in the space of normals, followed by a post-processing stage to sanitize topological inconsistencies.

Polycube deformation. In previous work, mesh deformation is used to compute an approximate polycube by minimizing an energy that penalizes surface normals poorly aligned with base axes, until the polycube structure is revealed [YZWL14, HJS*14, FXBH16]. Some methods also start from an initial labeling [GSZ11, LVS*13] to assist the deformation, or interleave both approaches by updating a target labeling between each deformation iteration [FBL16]. Mesh deformation may lead to strongly distorted or flipped elements [GSZ11, HJS*14]. In recent work [GLYL20], the AMIPS [FLG15] energy guarantees an inversion-free deformation by diverging in the presence of inverted cells. However, in some degenerate cases and due to inconsistent polyhedron topology, additional post-processing stages are necessary. Unfortunately, Sokolov et al. [SR15] show that degenerate conditions are not merely local and the computation of a globally valid structure remains a challenging problem.

Hex quality improvements. Hex meshes generated via polycube-maps do not have inner singularities. As a consequence, state-of-the-art polycube-based pipelines insert one or more layers of hexes along the whole boundary [GSZ11]. Such a process gives more degrees of freedom to smooth the mesh by pushing singularities inside. Unfortunately, naive global insertion can also locally decrease mesh quality. Kowalski et al. [KLSO12] define three types of fundamental layers that can be added locally to better capture boundary curves and surfaces solving an integer linear program. Similarly, Cherchi et al. [CAS*19] define selective padding which is able to add hex layers with a quality improvement guarantee. Recently, Guo et al. [GLYL20] enhance polycube-maps through clever cuts directly on the polycube, achieving similar results. Authors of [GPW*17] propose an approach to simplify a mesh using its base complex structure. While this approach is more dedicated to optimize a mesh generated by an overlay-grid approach, it could be used to post process Polycube-like meshes in order to extract coarse structure.

Machine learning and mesh generation. A variety of ge-

ometry and mesh generation problems have seen great advancements in the past decade via machine learning techniques [XZCOC12, MTP*15, LYZ*20, DLLK21]. Marcias et al. [MTP*15] describe a novel interactive method to suggest quad surface meshes to final users. Lim et al. [LYZ*20] designed an evolutionary algorithm that performs automatic blocking of a 2D manifold. They define a set of simple genetic operators on a set of points from which they robustly extract a quad layout. This layout is then evaluated and used to rank a population of sets of points. Their work achieves near-optimal blocking configurations after a large number of generations. Starting from an initial population of 3D models, Xu et al. [XZCOC12] use an evolutionary algorithm to generate novel shapes. The process is interactively driven and user preferences define the fitness function. Shapes are described as an assembly of parts and the crossover operation swaps different parts between parent shapes. Similarly, we use an evolutionary algorithm to perform polycube labeling. While the main principle and terminology are similar, the nature of the problems induces very different choices. In particular, the highly non-local nature of polycube labeling prevents the use of a naive per-triangle crossover operator.

1.2. Main contributions and pipeline overview

Our contributions are embedded in a full hex meshing pipeline as illustrated in Figure 4.

- In Section 2, we define *Evocube*, an extensible evolutionary-based framework for polycube labeling;
- In Section 3, we compute polycube-maps matching a desired labeling, and derive low distortion all-hex meshes.

Finally, in Section 4, we evaluate our method on over a thousand natural and CAD models, and discuss its limitations. To foster future research on polycube labeling, our implementation of *Evocube* is open-source and can easily be extended: <https://github.com/LIHPC-Computational-Geometry/evocube>

2. Genetic labeling optimization

Our method aims to generate a low-distortion polycube from an input domain. Precisely, the input to our algorithm is:

- a set of three-dimensional vertices \mathcal{V} ;
- a set of tetrahedral cells \mathcal{T}_Ω connecting points in \mathcal{V} , along with its boundary triangular mesh $\partial\mathcal{T}_\Omega$;

- the desired edge length l_e that defines the size of a cube in our polycube-map.

We aim to compute new vertex positions \mathcal{V}' such that all the normals in $(\mathcal{V}', \partial\mathcal{T}_\Omega)$ are aligned with one of the three main axes while $(\mathcal{V}', \mathcal{T}_\Omega)$ remains a low-distortion deformation of $(\mathcal{V}, \mathcal{T}_\Omega)$. For such a problem to be valid, we assume that \mathcal{T}_Ω defines a volume, is free of self-intersections, and $\partial\mathcal{T}_\Omega$ is manifold. Specifically, our work addresses the labeling task, which consists in precomputing polycube topology directly on $\partial\mathcal{T}_\Omega$ with a labeling vector $\ell \in \{\pm X, \pm Y, \pm Z\}^{|\partial\mathcal{T}_\Omega|}$ assigning any of the six possible orientations to each boundary triangle.

Intuitively, ℓ subdivides $\partial\mathcal{T}_\Omega$ in a set of charts \mathcal{C} comprised of neighboring triangles sharing target orientation. Charts are separated by a set of edge boundaries \mathcal{B} . In the event that a boundary $b \in \mathcal{B}$ is not straight, we use the method of Livesu et al. [LVS*13] described in Section 2.2 to identify turning points.

Previous work [GSZ11, LVS*13] has shown the feasibility of precomputing such a labeling of boundary faces using a set of modification operations that greedily improve an initial labeling. Our key insight is to embed local labeling fixes and associated quality criteria in a novel genetic framework. We recast the labeling problem as one of optimization with an objective function, and define genetic operations in the context of polycube labeling. Stochastic selection and crossover of candidate solutions allow *Evocube* to simultaneously consider several search directions and reduce susceptibility to local minima.

The main challenge with an approach based on optimization using a heuristic is the definition of an appropriate fitness function. In the case of a genetic framework, fitness evaluation is often a time bottleneck as every individual generated needs to be ranked and assessing the quality of a solution is not straightforward. Ideally, polycube distortion with regards to the input mesh could be used as a labeling quality metric. However, computing a satisfying volumetric polycube from a labeling is time-consuming and cannot be performed repeatedly. Hence, we define distortion proxies and labeling modifications that guide us towards interesting solutions in the set of all possible labelings.

2.1. Fitness of a labeling

Evaluating labeling quality is a complex problem that encompasses several aspects of a labeling and therefore requires the definition of several metrics. In the following, we distinguish between validity conditions - which must be fulfilled for a solution to be considered feasible - and optimization criteria - which we aim to minimize.

Validity proxy. A labeling is said to be *valid* if there exists a matching polycube polyhedron. Unfortunately, evaluating the validity of a labeling for a general shape remains an open problem, and the partial solutions provided in previous work [EM10, ZLW*19] are too computationally costly to be used repeatedly in a heuristic. Instead, we define a proxy validity criterion which offers no theoretical guarantee but can be efficiently evaluated.

Considering a set of charts and boundaries, the number of boundaries intersecting at a given vertex is referred to as the vertex's valency. Let us consider the following sets:

- invalid corners \mathcal{V}_{inv} with valency at least 4;
- invalid boundaries \mathcal{B}_{inv} between charts with opposite labels;
- invalid charts \mathcal{C}_{inv} with strictly fewer than 4 neighbors.

We thus define the **validity proxy** of a labeling $V_p(\ell)$ as:

$$V_p(\ell) = |\mathcal{V}_{\text{inv}}| + |\mathcal{B}_{\text{inv}}| + \sum_{c \in \mathcal{C}_{\text{inv}}} (4 - N_c) \quad (1)$$

where N_c is the number of neighbors of chart c . Examples of labelings with non-empty invalid sets are shown in Figure 5.




Figure 5: Example labelings that do not lead to a valid polycube.

A labeling ℓ will be considered *pseudo-valid* if and only if $V_p(\ell)$ is equal to 0, or equivalently, the three invalid sets are empty. We provide some additional insight into the limitations of our validity proxy in Appendix A.

Optimization criteria. Since our objective is the minimization of volumetric polycube parameterization distortion, which is too time-consuming to evaluate directly, our method requires a set of metrics on ℓ that approximate the distortion on $(\mathcal{V}', \mathcal{T}_\Omega)$.

We observe that a surface polycube of the boundary $\partial\mathcal{T}_\Omega$ can be computed significantly faster than its volumetric counterpart. We thus define the **workability** of a labeling $E_w(\ell)$ using the distortion of the mapping from $(\mathcal{V}, \partial\mathcal{T}_\Omega)$ to a fast polycube $(\mathcal{V}_f, \partial\mathcal{T}_\Omega)$. The term workability refers to the ability of some material to be easily deformed into a different shape. We compute our fast surface polycube as follows. Given a labeling ℓ defining a set of charts \mathcal{C} , we use a change of variables to constrain all vertices on a chart to the same value on the chart's label axis. We then minimize a least-squares problem aiming to preserve edge lengths on the two other axes. The solution yields a set of new vertex positions \mathcal{V}_f . The per-triangle distortion is measured using the singular values σ_1 and σ_2 of the Jacobian of the mapping from the initial triangle to its equivalent in $(\mathcal{V}_f, \partial\mathcal{T}_\Omega)$ [THCM04]:

$$e_w = \sigma_1 + \sigma_2 + \frac{1}{\sigma_1 \sigma_2} + \frac{\sigma_1}{\sigma_2} + \frac{\sigma_2}{\sigma_1} - 4$$

For some invalid labelings, some triangles may be degenerate and have infinite e_w . We clamp e_w to some arbitrarily large value, greatly penalizing invalid labelings. Finally, we compute the overall workability $E_w(\ell)$ by integrating e_w^2 over the boundary $\partial\mathcal{T}_\Omega$. The resulting surface polycube may contain inverted triangles which would require more time-consuming optimization to resolve, but in practice it remains a good proxy of the final polycube quality issued from ℓ .

We complement our novel metric with additional ones used in

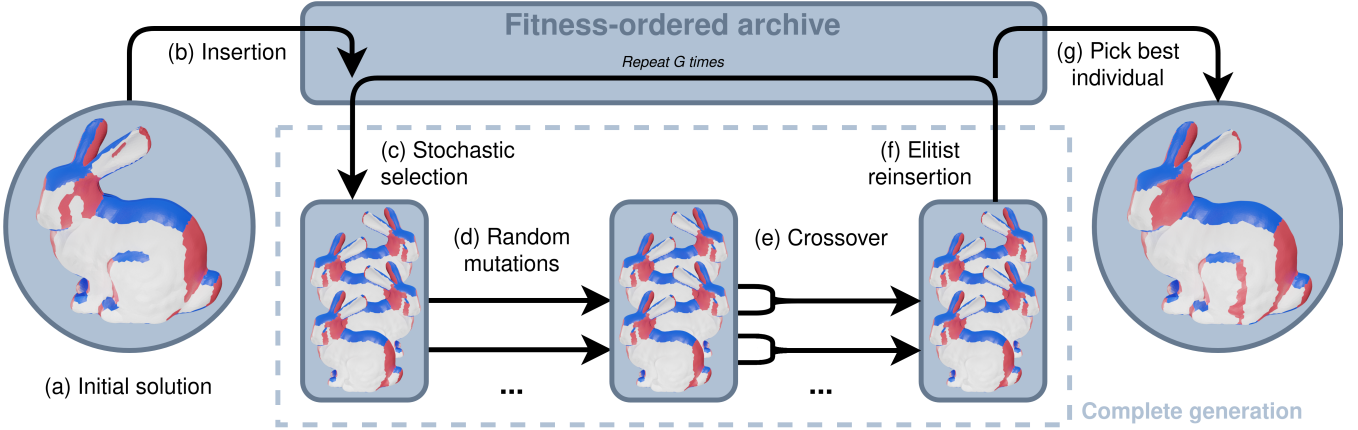


Figure 6: Architecture of our genetic optimization framework.

previous work [LVS*13]. First, the **fidelity** of a given face measures how well its label fits its orientation. It is defined as the dot product of its normal vector and its assigned direction. The fidelity of a labeling E_F is then computed by integrating over all boundary triangles. Finally, **compactness** E_C is the number of corner vertices in the polycube graph associated with ℓ .

By combining the above metrics, we define a *fitness* function that can be embedded in our genetic framework:

$$fitness(\ell) = V_p(\ell) + \omega_1 E_W(\ell) + \omega_2 E_F(\ell) + \omega_3 E_C(\ell)$$

The coefficients vector $\omega = (\omega_1, \omega_2, \omega_3)$ is set to $(10^2, 10^{-2}, 10^{-2})$ in our experiments. This *fitness* function is able to discriminate between candidate solutions and favor search directions reducing parameterization distortion as evaluated by our metrics.

2.2. Initial solution

Our algorithm requires an initial labeling to improve upon. We use the graph-cut initialization method presented by Livesu et al. [LVS*13]. Considering the dual graph of a triangle mesh, its main advantage is the trade-off between:

- a unary cost for assigning a given label to a face, penalizing label directions poorly aligned with triangle normals;
- and a binary cost for two neighboring triangles labeled differently, with a greater cost for neighbors that are close to be coplanar.

The unary cost helps find a labeling with low fidelity error, while the binary improves compactness and reduces boundary size. This method provides us with a satisfying initial solution that encompasses important fidelity information. While it is a good starting point, in general this solution has invalid patches. It cannot be used as is, and will be enhanced by our proposed genetic framework. We use a graph-cut optimization library by Boykov et al. [BVZ01] which implements the improvements described in additional research [BK04, KZ04]. We allow neighbors with opposite orientations in our initial solution and propose our own fix in Section 2.5.

We compute **turning points** following a graph-cut approach on boundary edges, similarly to Livesu et al. [LVS*13]. Considering a directed boundary (Fig 7b) and its corresponding axis (Fig 7a), our graph-cut approach partitions the boundary with two labels (Fig 7e) and we define turning points as the vertices where the label switches. We use the unary cost u (Fig 7c) to infer a bias based on the dot product of the axis and a given edge. Then, the binary cost b encourages cuts between poorly aligned consecutive edges. Specifically, for a normalized edge \vec{e} and its label l_e , we define these costs as follows:

$$\begin{cases} u(\vec{e}, 0) = 1 - e^{-\frac{1}{2} \left(\frac{\vec{e} \cdot \vec{axis}}{0.9} \right)^2}, & \text{if } \vec{e} \cdot \vec{axis} < 0, \text{ and } 0 \text{ otherwise;} \\ u(\vec{e}, 1) = 1 - e^{-\frac{1}{2} \left(\frac{\vec{e} \cdot \vec{axis}}{0.9} \right)^2}, & \text{if } \vec{e} \cdot \vec{axis} > 0, \text{ and } 0 \text{ otherwise;} \\ b(\vec{e}_1, \vec{e}_2, l_{e_1}, l_{e_2}) = e^{-\frac{(\vec{e}_1 \cdot \vec{e}_2 - 1)^2}{2}}, & \text{if } l_{e_1} \neq l_{e_2}, \text{ and } 0 \text{ otherwise.} \end{cases}$$




Figure 7: Turning points identification.

2.3. Genetic framework

From an initial labeling, our framework generates new solutions and selects promising search directions. The overall architecture of our Evocube framework is illustrated in Figure 6. We define the key elements to select and cross individuals in this section, and labeling mutations in section 2.4.

Crossover. The crossover operator takes two solutions ℓ_1 and ℓ_2 as input, and generates a new individual including both of its parents' mutations. To achieve this, when both parent solutions agree on a face label, the child is assigned the same label. When labels

conflict, the label that was changed during the most recent generation is kept. In case of a draw, we consistently pick ℓ_1 's label.

Archive. We use an archive system to keep track of the best solutions to date. The archive has a fixed maximum size, and when an element is submitted, sufficient *fitness* greater than that of the worst element is required. The latter is then discarded. When selecting an element from the archive, we use a stochastic model favoring higher ranked solutions, similarly to Lim et al. [LYZ*20]. For an archive containing ranked solutions from ℓ_1 to ℓ_n , the probability of randomly picking ℓ_i is as follows:

$$P(\ell_i) = \frac{n - i + 1}{1 + \dots + n}$$

Generations. Here, we describe the core architecture of our genetic framework illustrated in Figure 6. The initial solution is generated using the graph-cut approach described previously. At the beginning of each generation, N candidate solutions are selected from the archive. Following our stochastic model, the highest-ranked solutions are picked several times. The selected individuals then undergo random mutations. Since individual mutations are independent, they are computed in parallel and a greater N can be used while maintaining reasonable time complexity.

Afterwards, pairs of individuals are stochastically selected and combined using the crossover operator. This is repeated C times, leading to a population of size $N + C$. Finally, we insert all new solutions with sufficient score in the archive, and discard all unfit individuals, ending the generation.

This process ends after $G = 40$ generations, or after three consecutive generations without change in the archive's best solution. In our experiments, we found that $N = 100$ and $C = 10$ were sufficient to reach convergence for most input meshes.

2.4. Labeling mutations

We propose a set of mutations that aim to improve an initial labeling. Given that the search space of all possible labelings is of size $6^{|\partial\mathcal{T}_\Omega|}$, we restrict our search to meaningful solutions. Our mutations specifically focus on fixing invalid patches and identified turning points. These operations may yield poor results under some circumstances, but by virtue of our genetic framework, our method identifies and discards detrimental modifications. Our labeling mutations are illustrated in Figure 8. Future work may complement *Evocube* with additional operations to account for specific configurations, but we found these to be sufficient for most geometries.

Directional path. In Figure 8a, a turning point is chosen randomly, along with one of the 4 orthogonal directions on the chart. We then greedily select a path matching that desired direction until we reach another boundary. Once the path is defined, we apply the new label to the triangles around the path and propagate to their neighbors. The introduced label is picked from only two possibilities, since the chart's label and its opposite cannot be picked, and the path is highly non-constant on its propagation axis, effectively forbidding two more labels. This operation is similar to the orthogonal path described by Gregson et al. [GSZ11].

Chart removal. When a chart is invalid as in Figure 8b, it can sometimes simply be removed. To perform this, we use the same

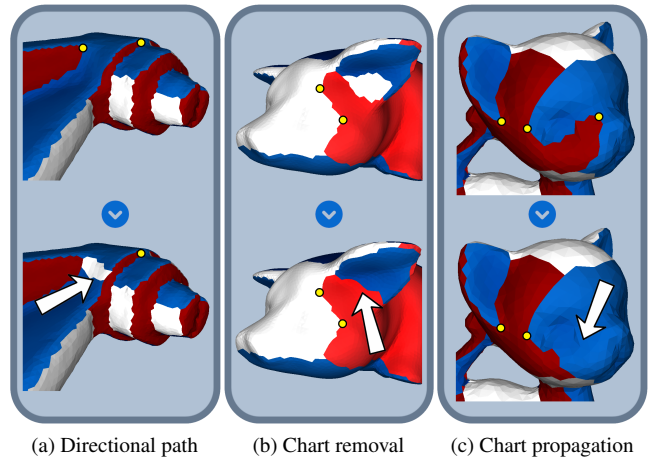


Figure 8: Base mutations in our genetic framework. Turning points are highlighted in yellow.

graph-cut method used to compute an initial solution, with two modifications. Firstly, the rest of the labeling is locked to prevent any unwanted modification, and secondly we forbid the initial label to be assigned again to the chart that is being removed.

Chart Propagation. Given a random border with a turning point, the label of one side is applied to the other side around the selected turning point, as illustrated in Fig 8c. If the border is already monotone, then we propagate along the whole border instead.

In our genetic framework, mutations are chosen randomly and applied in random locations. In order to help speedup convergence, we select among invalid charts if some remain for chart removal, or among turning points for directional path and chart propagation. If none remains, the location is randomly selected among all charts, or boundary vertices, respectively. In our experiments, the distance for propagation mutations is sampled from $[l_{avg}, 5 l_{avg}]$, where l_{avg} is the average edge length in the mesh.

2.5. Labeling repairs

We complement our labeling mutations with a set of fixes illustrated in Figure 9 which rectify some easily identifiable labeling problems. Unlike mutations, these operations are applied deterministically, as specified below.

In a valid labeling, no pair of neighboring charts can have opposite labels. We repair any boundary separating two such charts by introducing a new chart with one of the remaining labels around the boundary (Fig 9a). The new chart can be inserted on both sides of the boundary or on either side alone, and its size is taken to be a multiple of the mesh's average edge length. We pick the preferred option by measuring *fitness* for all possibilities. We apply this repair only on the initial and final solutions.

Similarly, we propose a simple fix for corners with invalid valency in the polycube graph (Fig 9b). We introduce a new chart with one of the remaining labels around the problematic vertex,




Figure 9: Labeling repairs.

and pick the optimal size using our *fitness* evaluator. Similarly, this operation is used only on the initial and final solutions.

Finally, any triangle with two edges on the same chart boundary will be flattened in parameterization space. We consider pairs of triangles to be neighbors only if they share an edge on a chart boundary. We then smooth boundaries by relabeling any triangle surrounded by two neighbors sharing a label different of its own (Figure 9c). In some cases, several iterations are needed to obtain a smooth labeling. This operation is fast and we apply it to any individual before evaluation.

3. Hexahedral meshing

Given a pseudo-valid polycube labeling ℓ of the boundary mesh $(\mathcal{V}, \partial\mathcal{T}_\Omega)$, we now seek to compute a volumetric polycube-map matching this labeling and an all-hex mesh. Note that previous work [GSZ11, LVS*13] has also described effective methods which could be used in combination with Evocube.

We compute a low-distortion volumetric polycube of $(\mathcal{V}, \mathcal{T}_\Omega)$. This differs from the polycube-based fitness evaluator described in Section 2.1, which is only a fast surface mapping of $(\mathcal{V}, \partial\mathcal{T}_\Omega)$. Similarly, we first proceed to a change of variable along chart vertices to enforce the polycube topology induced by our labeling ℓ . Consequently, vertices on the same chart will share the same coordinate on the axis associated with the chart's label. We then minimize a Laplacian energy penalizing edge distortion over both interior and exterior vertices. After this step, the resulting polycube may still contain inverted tetrahedra. Subsequently, we minimize the energy proposed by Garanzha et al. [GKK*21] under the same change of variable to restore inverted cells and further improve the quality of the volumetric polycube-map.

For all-hex meshing, we then quantize our polycube-map using the robust mixed-integer method described by Protais et al. [PRR*20], and extract our initial hexahedral mesh by applying the inverse mapping from our quantized polycube to input space. While polycube-based hexahedral meshing is appreciated for its regularity, it also suffers from the lack of inside singularities. We introduce

a layer of elements on the boundary, and smooth the resulting hex mesh by minimizing a combination of the energy defined in [BG02] and the mesh distance to the input boundary and features.

There exists various methods to extract a hexahedral mesh from a valid labeling. We include our method for completeness. In practice, it has proven to be robust and effective during our experiments. To further improve on the quality of hex meshes extracted from polycubes and push singularities towards the interior, our simple pillowing can be replaced with more advanced methods such as selective-padding [CAS*19] or cut-enhancements [GLYL20].

4. Experiments

We tested our method on a collection of 1315 meshes from the MAMBO dataset [Led20], the natural and CAD shapes used by Gao et al. [GSP19], and a subset of ABC [KMJ*19] consisting of the inputs sampled by Reberol et al. [RGR21]. Our results on these datasets are synthesized in Table 1. We generate initial solutions with a ratio of $\omega_{\text{unary}}/\omega_{\text{binary}} = 3$, as recommended by Livesu et al. [LVS*13]. If no pseudo-valid labeling is found on our first attempt, we divide this ratio by 3 and apply our labeling optimization starting from this new initial solution.

We generate tetrahedral meshes with *netgen* [Sch97], and use *libigl* [JP*18] for common geometry processing operations. In the following, all-hex meshes are visualized in *HexaLab* [BTP*19]. In our supplemental material, we provide renderings of all labelings and polycubes, as well as per-model statistics.

Dataset	size	mesh	$V_p(\ell) = 0$	$D_a < 2$	$\text{minSJ} \geq 0$
Mambo	113	100%	100%	97.4%	96.5%
Smooth	93	100%	100%	75.2%	89.2%
CAD	109	100%	94.5%	87.1%	85.3%
ABC	1000	99.4%	95.5%	82.3%*	61.9%*

Table 1: Results on CAD and smooth datasets. For each entry we report the number of inputs, the ratios of (a) tetrahedral meshes successfully generated, (b) pseudo-valid labelings, (c) polycubes with area distortion D_a [THCM04] below 2 (the ideal value being 1), and (d) hex meshes with $\text{minSJ} \geq 0$. *Given the size of the ABC dataset, these results were simply computed using our fast polycube estimator and *libHexEx* [LBK16] without any post-processing.

4.1. Comparison

We compare our work with PolyCut [LVS*13] using the binaries provided by the authors incorporating subsequent work [LSVT15]. On some inputs, we had to stop labeling optimization after one hour without progress. We also report on the initial solution computed with graph-cut, showing the importance of our labeling optimization. The results, detailed in Table 2, show that our method converges to a pseudo-valid labeling in an overwhelming majority of cases. When compared with another labeling based method [LVS*13], our method is more robust and our labelings subsequently lead to high quality hexahedral meshes in a greater range of cases. Our ability to compute satisfying polycubes on more models than previous work can be used to improve the robustness of




Figure 10: Comparison with previous methods on *moai*, *camil*, *pinion* and *venus*. Mesh quality is reported in Table 2.

previous polycube enhancement methods [CAS*19, GLYL20] and further enhance our resulting hex meshes.

As explained in Section 3, we compute volumetric polycubes and as such we do not optimize for surface polycube distortion [THCM04]. Our labelings and volumetric polycubes aim to pre-determine polycube topology and the number of cells subdividing the input volume. The exact position of these cells in the final hex mesh is determined by our all-hex smoothing, effectively removing the need for surface polycubes in comparison with previous work. To compute a low-distortion surface parameterization, our polycubes can be used as an initial solution for an inverse optimization with remeshing such as the one used by Livesu et al. [LVS*13].

We also compare our results with the feature-aware octree-based method proposed by Gao et al. [GSP19], which is known to be more robust than polycubes but yield less regular meshes. As expected, our method produces satisfying hex meshes on a smaller range of models but still results in better average cell quality on the whole dataset. As illustrated in Figure 10, our method produces highly regular hex-meshes with cells of similar sizes, a property which is greatly appreciated for the interpretation of numerical simulation results. In comparison, the results of Gao et al. [GSP19] include

cells of varying sizes and are able to capture complex shapes with more accuracy, at the expense of interpretability and singularities.

4.2. Analysis

Owing to its genetic nature, Evocube is highly parallelizable and time-efficient. Each generation, individuals mutate and are evaluated independently. In our experiments, this led to a speedup by a factor of 14, as illustrated in Figure 11. Despite our fast surface polycube evaluator, which is several magnitudes faster than computing the final polycube, the main time bottleneck of our framework remains fitness evaluation. On average, our method required around 15 seconds per input for labeling optimization and was able to compute the labelings leading to Table 1 in under six hours, as detailed in Appendix B.

Nonetheless, our method failed to find valid labelings on some inputs. Similarly to previous research on polycube-maps, our method is unable to deal with complex shapes when they are poorly aligned with principal axes, as illustrated in Figure 12. There are also cases in which a pseudo-valid labeling did not lead to a satisfying hex mesh, as reported in Tables 1 and 2, where the ratio of such labelings is higher than the ratio of successful hex-meshes.

Method	pseudo-valid	minSJ ≥ 0	$\overline{\text{minSJ}}$	$\overline{\text{avgSJ}}$
Mambo Basic				
Graph-cut	77.0%			
[LVS*13]	94.6%	73.0%	-0.00	0.60
Ours	100%	100%	0.21	0.90
Mambo Simple				
Graph-cut	66.7%			
[LVS*13]	96.7%	53.3%	-0.24	0.40
Ours	100%	90.0%	0.08	0.84
Mambo Medium				
Graph-cut	33.3%			
[LVS*13]	88.9%	11.1%	-0.54	0.06
Ours	100%	88.9%	-0.00	0.75
Smooth [GSP19]				
Graph-cut	18.3%			
[LVS*13]	76.1%	55.7%	-0.20	0.37
[GSP19]	N/A	100%	0.19	0.81
Ours	100%	89.2%	0.15	0.86
CAD [GSP19]				
Graph-cut	72.5%			
[LVS*13]	85.8%	52.8%	-0.23	0.39
[GSP19]	N/A	100%	0.17	0.83
Ours	94.5%	85.3%	0.19	0.87

Table 2: Comparison in terms of ratio of pseudo-valid labelings, hex meshes with $\text{minSJ} \geq 0$, and hex quality. We denote \bar{X} the average of X over a given dataset. When no output is produced, the worse possible value of -1 is assumed instead.




Figure 11: Time plot (in hours of real and CPU time) of our labeling optimization over all 1315 tested models. We provide a detailed time report in Appendix B.

By construction, our labeling corresponds to base axes and is dependent on the input’s orientation. For some shapes, this assumption leads to non-optimal labelings. For example, a U-shaped cylinder will be mapped to a U-shaped polycube, as illustrated in Figure 14, whereas a simple cuboid would be preferable and lead to fewer singularities in the resulting hex-mesh. This limitation could be addressed in our framework with recent work focusing on removing undesired polycube corners [MCBC22].

The consistency and reproducibility of our results are hindered by the non-deterministic nature of our approach. This drawback can be mitigated by increasing the number of individuals N per generation. More mutations are considered each generation, leading to a greater pool to choose from for the next and an improved breadth of search overall. In contrast, the number of generations G relates to the total number of mutations that can be applied, analogous to the

Input	N_c	angle/area d.	minSJ	avgSJ	#hex
Rocker Fig 4					
[GSP19]			0.165	0.813	30k
[LVS*13]	62	1.066/1.051	0.370	0.890	57k
Ours	74	1.315/1.523	0.241	0.938	60k
Moai Fig 10					
[GSP19]			0.263	0.824	25k
[LVS*13]	8		0.497	0.950	3k
Ours	28	1.073/1.057	0.530	0.968	18k
cam1 Fig 10					
[GSP19]			0.056	0.829	25k
Ours	72	1.164/1.221	0.108	0.802	3k
pinion Fig 10					
[GSP19]			0.026	0.783	89k
[LVS*13]			0.106	0.863	5k
Ours	112	1.122/1.096	0.142	0.816	5k
venus Fig 10					
[GSP19]			0.131	0.800	20k
Ours	36	1.131/1.122	0.326	0.950	49k
S22 Fig 13	64	1.053/1.032	0.023	0.910	23k
S26 Fig 14	80	1.044/1.027	0.016	0.899	21k
B16 Fig	16	1.160/1.161	0.098	0.917	10k
B51 Fig	32	1.115/1.101	0.085	0.951	20k
B76 Fig	26	1.098/1.058	0.059	0.934	8k
B38 Fig	24	1.303/1.149	0.058	0.845	9k
B49 Fig	12	1.020/1.021	0.083	0.989	49k

Table 3: Statistics of presented polycubes and hex meshes. We report the number of corners N_c of polycubes, angle and area distortion, as well as quality metrics and number of cells for hex meshes.




Figure 12: Example meshes for which our Evocube implementation did not converge to a pseudo-valid labeling. The majority of failures fell in one of these categories: complex geometries with features poorly aligned with the base axes (left); shapes requiring a polycube with valency 4 corners (middle); and inputs with tiny holes that are too coarse to be satisfactorily labeled (right).

depth of search. In Figure 15, we observe that a very high number of generations does not necessarily imply better results.

In complex cases and similarly to previous work [LVS*13], the final labeling may assign the same label across some feature-edges. In such cases, the feature-edge does not appear in the polycube and is most often deteriorated in the resulting hex-mesh, as illustrated in Figure 13. For CAD models, this issue could be alleviated in future work through feature-aware mutations and fitness.

In our experiments, we further observe that when a vastly superior mutation is found, a single individual and its offspring tend to overwhelm the archive, reducing population diversity. For the most




Figure 13: In some cases, Evocube finds a low-distortion labeling that assigns the same label across a feature edge. In our current post-processing pipeline and similarly to other polycube-based methods, such feature edges are not preserved.




Figure 14: Additional results on CAD inputs from the MAMBO dataset. From left to right, top to bottom: S26, B16, B51, B76, B38, B49.

complex shapes, advanced evolutionary mechanisms such as *speciation* [DS97] could be added to favor solution diversity.

5. Conclusion and future work

We have presented *Evocube*, a novel evolutionary approach to compute polycube labelings. Our method generates pseudo-valid labelings on a wide range of CAD and smooth models, paving the way for future work on labeling enhancements within this framework. Resulting polycubes have meaningful complexity and produce high-quality hex meshes on a wide range of input geometries.

Our method has a few limitations which may be addressed in

future work. Additional mutations are required to further improve our pseudo-valid labelings. In particular, boundary displacement and feature-aware operations would integrate well in *Evocube* and produce more accurate hex meshes while greatly reducing parameterization distortion. Tailor-made mutations may also be added to systematically deal with specific CAD features, such as slopes or sharp wedges. Finally, we hope our use of an evolutionary algorithm applied to polycube labeling can inspire similar work on other problems in mesh generation and computer graphics in general.




Figure 15: In complex cases, a large number of generations may not be sufficient to reach a satisfying labeling. On this model (ABC #4035), sharp wedges are numerous and poorly aligned with base axes. After 500 generations, $V_p(l)$ is reduced from 45 to 26, but many invalidities remain and some of the solutions found are not satisfying.

References

- [BG02] BRANETS L. V., GARANZHA V. A.: Distortion measure of trilinear mapping. Application to 3-D grid generation. *Numerical Linear Algebra with Applications* 9, 6-7 (2002), 511–526. 7
- [BK04] BOYKOV Y., KOLMOGOROV V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 9 (2004), 1124–1137. 5
- [BM93] BLACKER T. D., MEYERS R. J.: Seams and wedges in plastering: A 3-d hexahedral mesh generation algorithm. *Eng. with Comput.* 9, 2 (June 1993), 83–93. 2
- [BTP*19] BRACCI M., TARINI M., PIETRONI N., LIVESU M., CIGNONI P.: Hexalab.net: An online viewer for hexahedral meshes. *Computer-Aided Design* 110 (2019), 24–36. 7
- [BVZ01] BOYKOV Y., VEKSLER O., ZABIH R.: Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 11 (2001), 1222–1239. 5
- [Bä96] BÄCK T.: *Evolutionary algorithms in theory and practice - evolution strategies, evolutionary programming, genetic algorithms*. Oxford University Press, 1996. 2
- [CAS*19] CHERCHI G., ALLIEZ P., SCATENI R., LYON M., BOMMES D.: Selective padding for polycube-based hexahedral meshing. *Computer Graphics Forum* 38, 1 (2019), 580–591. 2, 3, 7, 8
- [CHL19] CALDÉRAN S., HUTZLER G., LEDOUX F.: Dual-based user-guided hexahedral block generation using frame fields. In *Proceedings of 28th International Meshing Roundtable* (2019). 2
- [DLLK21] DIELEN A., LIM I., LYON M., KOBELT L.: Learning direction fields for quad mesh generation. *Computer Graphics Forum* 40, 5 (2021). 3
- [DS97] DEB K., SPEARS W.: Population structures, c2.6 speciation methods. In *Handbook of Evolutionary Computation*. (1997). 10
- [EM10] EPPSTEIN D., MUMFORD E.: Steinitz theorems for orthogonal polyhedra. In *Proceedings Symposium on Computational Geometry* (2010), p. 429–438. 2, 3, 4, 12
- [FBL16] FU X.-M., BAI C.-Y., LIU Y.: Efficient volumetric polycube-map construction. *Computer Graphics Forum* 35, 7 (2016), 97–106. 3
- [FLG15] FU X.-M., LIU Y., GUO B.: Computing locally injective mappings by advanced mips. *ACM Trans. Graph.* 34, 4 (July 2015). 3
- [FXBH16] FANG X., XU W., BAO H., HUANG J.: All-hex meshing using closed-form induced polycube. *ACM Trans. Graph.* 35, 4 (2016), 124:1–124:9. 3
- [GKK*21] GARANZHA V., KAPORIN I., KUDRYAVTSEVA L., PROTAIS F., RAY N., SOKOLOV D.: Foldover-free maps in 50 lines of code. *ACM Trans. Graph.* 40, 4 (2021). 7
- [GLYL20] GUO H.-X., LIU X., YAN D.-M., LIU Y.: Cut-enhanced polycube-maps for feature-aware all-hex meshing. *ACM Trans. Graph.* 39, 4 (2020), 106–1. 3, 7, 8
- [GPW*17] GAO X., PANOZZO D., WANG W., DENG Z., CHEN G.: Robust structure simplification for hex re-meshing. *ACM Trans. Graph.* 36, 6 (nov 2017). 3
- [GSP19] GAO X., SHEN H., PANOZZO D.: Feature preserving octree-based hexahedral meshing. *Computer Graphics Forum* 38, 5 (2019), 135–149. 2, 7, 8, 9
- [GSZ11] GREGSON J., SHEFFER A., ZHANG E.: All-hex mesh generation via volumetric polycube deformation. *Computer Graphics Forum* 30, 5 (2011), 1407–1416. 2, 3, 4, 6, 7
- [HJS*14] HUANG J., JIANG T., SHI Z., TONG Y., BAO H., DESBRUN M.: ℓ_1 -based construction of polycube maps from complex shapes. *ACM Trans. Graph.* 33, 3 (2014), 25:1–25:11. 3
- [HTWB11] HUANG J., TONG Y., WEI H., BAO H.: Boundary aligned smooth 3d cross-frame field. *ACM Trans. Graph.* 30, 6 (2011), 1–8. 2
- [HZ16] HU K., ZHANG Y. J.: Centroidal voronoi tessellation based polycube construction for adaptive all-hexahedral mesh generation. *Computer Methods in Applied Mechanics and Engineering* 305 (2016), 405–421. 3
- [JP*18] JACOBSON A., PANOZZO D., ET AL.: libigl: A simple C++ geometry processing library, 2018. <https://libigl.github.io/>. 7
- [KLF16] KOWALSKI N., LEDOUX F., FREY P.: Smoothness driven frame field generation for hexahedral meshing. *Computer-Aided Design* 72 (2016), 65–77. 2
- [KLSO12] KOWALSKI N., LEDOUX F., STATEN M. L., OWEN S. J.: Fun sheet matching: towards automatic block decomposition for hexahedral meshes. *Engineering with Computers* 28, 3 (2012), 241–253. 2, 3
- [KJM*19] KOCH S., MATVEEV A., JIANG Z., WILLIAMS F., ARTEMOV A., BURNAEV E., ALEXA M., ZORIN D., PANOZZO D.: Abc: A big cad model dataset for geometric deep learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019). 7
- [KZ04] KOLMOGOROV V., ZABIH R.: What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 2 (2004), 147–159. 5
- [LBK16] LYON M., BOMMES D., KOBELT L.: Hexex: robust hexahedral mesh extraction. *ACM Trans. Graph.* 35, 4 (2016), 123. 7
- [Led20] LEDOUX F.: The MAMBO dataset. <https://gitlab.com/franck.ledoux/mambo>, 2020. 7
- [LLX*12] LI Y., LIU Y., XU W., WANG W., GUO B.: All-hex meshing using singularity-restricted field. *ACM Trans. Graph.* 31, 6 (2012). 2
- [LMP16] LIVESU M., MUNTONI A., PUPPO E., SCATENI R.: Skeleton-driven adaptive hexahedral meshing of tubular shapes. *Computer Graphics Forum* 35, 7 (2016), 237–246. 2
- [LSVT15] LIVESU M., SHEFFER A., VINING N., TARINI M.: Practical hex-mesh optimization via edge-cone rectification. *Transactions on Graphics (Proc. SIGGRAPH 2015)* 34, 4 (2015). 7
- [LVS*13] LIVESU M., VINING N., SHEFFER A., GREGSON J., SCATENI R.: Polycut: Monotone graph-cuts for polycube base-complex construction. *ACM Trans. Graph.* 32, 6 (2013). 2, 3, 4, 5, 7, 8, 9
- [LW07] LEDOUX F., WEILL J.: An extension of the reliable whisker weaving algorithm. In *Proceedings of the 16th International Meshing Roundtable, October 14-17, 2007, Seattle, Washington, USA, Proceedings* (2007), Springer, pp. 215–232. 2
- [LYZ*20] LIM C. W., YIN X., ZHANG T., SELVARAJ S. K., SU Y., GOH C.-K., MORENO A., SHAHPAR S.: Towards automatic blocking of shapes using evolutionary algorithm. *Computer-Aided Design* 120 (2020), 102798. 3, 6


- [LZS*21] LI L., ZHANG P., SMIRNOV D., ABULNAGA S. M., SOLOMON J.: Interactive all-hex meshing via cuboid decomposition. *ACM Trans. Graph.* 40, 6 (dec 2021). 2
- [Mar09] MARÉCHAL L.: Advances in octree-based all-hexahedral mesh generation: Handling sharp features. In *Proceedings of International Meshing Roundtable* (2009), pp. 65–84. 2
- [MCBC22] MANDAD M., CHEN R., BOMMES D., CAMPEN M.: Intrinsic mixed-integer polycubes for hexahedral meshing. *Computer Aided Geometric Design* 94 (2022), 102078. 9
- [MTP*15] MARCIAS G., TAKAYAMA K., PIETRONI N., PANOZZO D., SORKINE-HORNUNG O., PUPPO E., CIGNONI P.: Data-driven interactive quadrangulation. *ACM Trans. Graph.* 34, 4 (Aug. 2015). 3
- [OS00] OWEN S. J., SAIGAL S.: H-morph: An indirect approach to advancing front hex meshing. *International Journal for Numerical Methods in Engineering* (5 2000). 2
- [PBS20] PALMER D., BOMMES D., SOLOMON J.: Algebraic representations for volumetric frame fields. *ACM TOG* 39, 2 (2020). 2
- [PBSB07] PARRISH M., BORDEN M. J., STATEN M. L., BENZLEY S. E.: A selective approach to conformal refinement of unstructured hexahedral finite element meshes. In *Proceedings of International Meshing Roundtable* (2007), pp. 251–268. 2
- [PRR*20] PROTAIS F., REBEROL M., RAY N., CORMAN E., LEDOUX F., SOKOLOV D.: Robust Quantization for Polycube Maps. preprint, Dec. 2020. 7
- [Qua14] QUADROS W. R.: Laytracks3d: A new approach to meshing general solids using medial axis transform. *Procedia Engineering* 82 (2014), 72 – 87. 2
- [RGR21] REBEROL M., GEORGIADIS C., REMACLE J.-F.: Quasi-structured quadrilateral meshing in gmsh – a robust pipeline for complex cad models, 2021. 7
- [RSL16] RAY N., SOKOLOV D., LÉVY B.: Practical 3d frame field generation. *ACM Trans. Graph.* 35, 6 (2016), 233. 2
- [SCB92] STEPHENSON M. B., CANANN S. A., BLACKER T. D.: Plastering: A new approach to automated 3d hexahedral mesh generation, 2 1992. 2
- [Sch96] SCHNEIDERS R.: Refining quadrilateral and hexahedral element meshes. *transition* 2 (1996), 1. 2
- [Sch97] SCHOEBERL J.: Netgen an advancing front 2d/3d-mesh generator based on abstract rules. *Computing and Visualization in Science* 1 (07 1997), 41–52. 7
- [SERB99] SHEFFER A., ETZION M., RAPPOPORT A., BERCOVIER M.: Hexahedral mesh generation using the embedded voronoi graph. *Engineering with Computers* 15, 3 (1999), 248–262. 2
- [SJ08] SHEPHERD J. F., JOHNSON C. R.: Hexahedral mesh generation constraints. *Engineering with Computers* 24, 3 (2008), 195–213. 1
- [SR15] SOKOLOV D., RAY N.: *Fixing normal constraints for generation of polycubes*. Tech. rep., LORIA, 2015. 2, 3, 12
- [Ste22] STEINITZ E.: Polyeder und raumeinteilungen, 1922. 12
- [Tak19] TAKAYAMA K.: Dual sheet meshing: An interactive approach to robust hexahedralization. *Computer Graphics Forum* 38, 2 (2019). 2
- [TBM95] TAUTGES T. J., BLACKER T., MITCHELL S. A.: The whisker weaving algorithm: A connectivitybased method for constructing all-hexahedral finite element meshes, 1995. 2
- [THCM04] TARINI M., HORMANN K., CIGNONI P., MONTANI C.: Polycube-maps. *ACM Trans. Graph.* 23, 3 (2004). 2, 4, 7, 8
- [XZCOC12] XU K., ZHANG H., COHEN-OR D., CHEN B.: Fit and diverse: Set evolution for inspiring 3d shape galleries. *ACM Trans. Graph.* 31, 4 (2012). 3
- [YZWL14] YU W., ZHANG K., WAN S., LI X.: Optimizing polycube domain construction for hexahedral remeshing. *Computer-Aided Design* 46 (2014), 58 – 68. 3

- [ZHB10] ZHANG Y., HUGHES T. J. R., BAJAJ C. L.: An automatic 3d mesh generation method for domains with multiple materials. *Computer Methods in Applied Mechanics and Engineering* 199, 5-8 (2010). 2
- [ZLW*19] ZHAO H., LI X., WANG W., WANG X., WANG S., LEI N., GU X.: Polycube shape space. *Computer Graphics Forum* 38 (2019). 4, 12


Appendix A: Validity limitations

The pseudo-validity conditions adopted by our method are not necessary for the existence of a corresponding polycube polyhedron. A counterexample is shown in Figure 16a, where a labeling not fulfilling the conditions still corresponds to a valid polycube polyhedron. The conditions are not sufficient either, as illustrated by the counterexample in Figure 16b [SR15], where a labeling fulfills the conditions but does not correspond to a valid polycube polyhedron.

Our pseudo-validity conditions are a simplification of the ones described by Eppstein and Mumford [EM10], adapting the Steinitz criteria [Ste22] for convex polyhedra. Zhao et al. [ZLW*19] managed to extend their analysis to orthogonal polyhedra with higher genus. The main advantage of our validity proxy is that it can efficiently be evaluated, allowing for optimization via a heuristic.



(a) Not necessary



(b) Not sufficient

Figure 16: Failure cases of pseudo-validity conditions.

Appendix B: Time measurements

In Table 4, we report detailed time measurements illustrated in Figure 11. We measure time both in CPU and real time, and observe a 14 fold speedup. All computations are performed on a 16-core AMD Ryzen Threadripper 1950X 2.2 GHz and 128 Gb RAM.

Operation	CPU time (h)	Real time (h)
Pre-optimization computation	0.33	0.33
Individual selection	3.20	0.21
Individual mutations	7.63	0.51
Charts and turning points	6.25	0.42
Fitness evaluation	55.2	3.66
Crossing	10^{-3}	10^{-3}
Insertion in archive	10^{-4}	10^{-4}
Post-optimization	0.21	0.21
Total	72.82	5.34

Table 4: Labeling optimization timings over all 1315 models.