

CALCULATING VOXEL-POLYHEDRON INTERSECTIONS FOR MESHING IMAGES

Stephen A. Langer¹

Andrew C.E. Reid²

¹*Information Technology Laboratory, National Institute of Standards and Technology, Gaithersburg, MD, 20899, stephen.langer@nist.gov*

²*Material Measurement Laboratory, National Institute of Standards and Technology, Gaithersburg, MD, 20899, andrew.reid@nist.gov*

ABSTRACT

Finite element meshes constructed from 3D images are useful in materials science and medical applications when it is necessary to model the actual geometry of a sample, rather than an idealized approximation of it. Constructing the mesh involves computing the intersection of the mesh elements with the voxels (3D pixels) of the image. If done naively, this process is unstable, and small errors in the computed position of an intersection point can lead to large errors in the computed volume. We demonstrate the source of the instability and present a robust and efficient method of doing the computation, based on the r3d algorithm of Powell and Abel. The method was developed in the context of the OOF3D program at NIST, but the freely available code is completely independent of OOF3D.

Keywords: image meshing

1. INTRODUCTION

When computationally modeling a complex object, it is often convenient to begin with an image of the object. For example, in materials science, a material may comprise many different grains with different shapes, compositions, and orientations, as shown in Fig. 1. Simulating the behavior of a 3D region provides insight into the behavior of the material as a whole. Basing the simulation on a micrograph of the material ensures that it accurately reproduces the geometry of (at least) one particular instance of structure [1]. Similarly, a medical image might provide the real geometry of a bone, and a simulation based on the image could compute the bone's strength and other properties [2].

The Object-Oriented Finite element (OOF) software developed at the National Institute of Standards and Technology (NIST) uses finite element analysis to compute properties of complex materials, starting from experimental or simulated micrographs. When the project began, creating 3D micrographs was diffi-

cult, tedious, and uncommon, so OOF1 [3] and OOF2 [4, 5] worked only in 2D, using 2D physical approximations. 3D micrographs are now common, and the third major revision of OOF, OOF3D [6, 7], creates 3D meshes and uses 3D physics. This paper describes part of the mesh generation method in OOF3D. However, the method is applicable to any 3D technique for meshing images.

Mesh generation in OOF begins by first segmenting an image, assigning material properties to each pixel, and/or putting pixels into user-defined groups. (We will sometimes use the words “pixel” and “voxel” interchangeably, to avoid excess verbiage.) For the purposes of this paper, all that matters is that pixels can be grouped, and that pixels in one group are somehow different from pixels in another group. The difference may be one of color or material properties or something else. Each group can be assigned a unique integer, which we call its “category”. For a finite element mesh to be an accurate representation of the geometry of the image, one requirement is that each




Figure 1: A scanning electron microscope image (a) of a cast cobalt-chrome alloy with metal carbide inclusions, approximately $200 \mu\text{m}$ on a side. The light gray inclusions have a different composition than the darker background, which consists of two regions with the same composition but different crystalline orientations. The diagonal stripes and dark spots are artefacts of the sample preparation process. Image courtesy of Adam Creuziger at NIST. The region in the white rectangle has been meshed in (b) by OOF2.

element be as *homogeneous* as possible – that is, it should overlie pixels of only one category, considering that pixels are rectangles and voxels are rectangular prisms, not points. A typical inhomogeneous element is shown in Fig. 2. Note that it is always possible to create a completely homogeneous mesh by creating a single element per pixel, but this mesh would be over-refined – it would resolve unphysical pixel corners and contain far more elements than are usually required for a sufficiently accurate computation.

A simple way of computing the homogeneity of an element E would be to count the number of pixel centers within the element for each pixel category. If there are $n_i(E)$ such pixels of category i , then the homogeneity

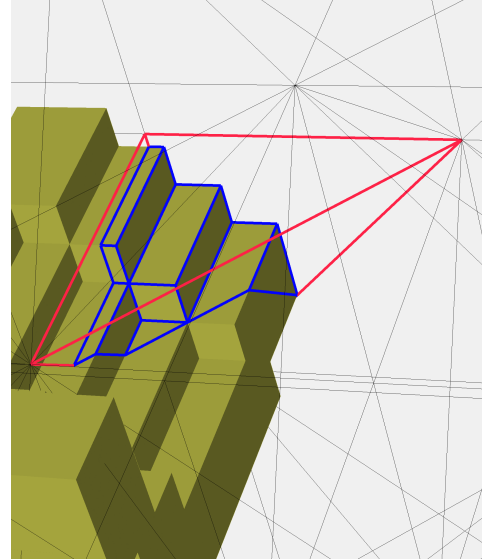


Figure 2: A screen shot from OOF3D. The yellow voxels have the same material type, which is different from the material of the the other, invisible, voxels. The thin lines are the edges of the elements in a uniform tetrahedral mesh. The element highlighted in red is inhomogeneous because it contains both yellow and invisible voxels. The blue edges have been computed by the method described here and outline the intersection of the yellow voxels and the red element.

could be defined as

$$h_{\text{simple}}(E) = \frac{\max_i n_i(E)}{\sum_i n_i(E)}. \quad (1)$$

This simple definition clearly fails when an element is small enough that very few (or zero) pixel centers lie within it. Furthermore, it produces a homogeneity that is a discontinuous and piece-wise constant function of the positions of the element’s nodes, which is undesirable for many mesh modification tools. For example, a tool which moves nodes will find that many small moves don’t change the homogeneity at all, providing no guidance on whether or not moving a node in a given direction will improve the mesh. (This simple method is discussed in more detail below, in Section 6.)

A more robust way of defining the homogeneity is

$$h(E) = \frac{\max_i V_i(E)}{V(E)} \quad (2)$$

where $h(E)$ is the homogeneity of element E , and V_i is the volume (in 3D) or area (in 2D) of the part of the element that intersects pixels of category i . The total volume (or area) of the element is $V = \sum_i V_i$. The definition (2) of the homogeneity has the advantage that it is a continuous function of the positions of the

element’s nodes. Its main disadvantage is that it is far more complicated to compute than the simple definition (1), because it uses the volume of intersecting polyhedra.

One way to compute V_i would be to loop over the voxels in category i , finding which ones intersect the element, and computing the volume of each intersection. Another way would be to aggregate the voxels of category i into a region, compute the boundary of the region, and then find the intersection of the element with the boundary. The boundary in general is a non-convex disconnected polyhedron requiring a more complicated storage scheme than a simple set of voxels. However, it can contain many fewer faces, edges, and vertices than the voxel set, and looping over them can be more efficient than looping over the voxels individually. This paper describes a way to compute the element homogeneity from the boundary of the voxel set.

There is little or no literature specifically on computing the intersections of tetrahedra or polyhedra with non-convex sets of voxels. In a method akin to marching cubes, Frey, *et al.*[8] mesh images by subdividing each voxel into many tetrahedra. This often creates far more elements than necessary, although it does avoid needing to compute intersecting volumes. The problem of intersecting convex polygons in a plane was addressed by O’Rourke, *et al.*[9], who comment on the difficulty of extending their method to polyhedra. Grandy[10] developed a method for polyhedra using “extremely complex geometric analysis[11]” which has been further developed by others[11, 12], but is not specifically applicable to sets of voxels. Ray, *et al.*[13] present an algorithm for clipping convex polyhedra to compute Voronoi cells on a GPU. Their algorithm is fast and parallel, but encounters rare special cases which must be handled separately. Arbitrary precision arithmetic[14] has also been used to avoid numerical issues in related problems.

In the remainder of this paper we first show how a naive way of computing Eq. 2 can run into trouble due to round-off error, and then introduce a method based on the r3d algorithm of Powell and Abel [15]. We show how the r3d method can fail in some cases, and how to fix it. This resulting algorithm applies to non-convex polyhedra, avoids the issue of round-off error completely, and entirely eliminates the need to handle special cases, being far simpler than Grandy’s method. In Section 4 we discuss a method for efficiently computing the polyhedral graphs required by the algorithm. This computation has a fairly large number of steps, but is straightforward, requires no floating point arithmetic, and only needs to be done once per image, not once per element. We conclude by comparing the performance of the r3d-based algorithm

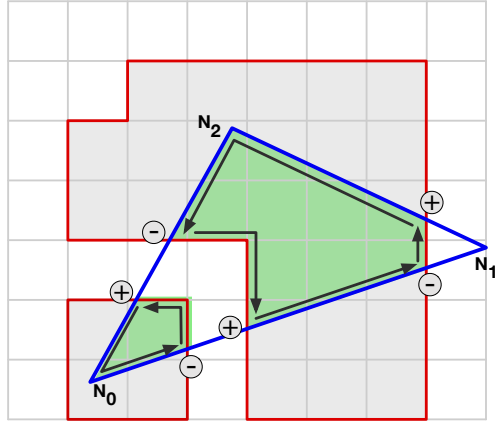


Figure 3: Finding the area of intersection of an element and a set of pixels in 2D. Intersections between the pixel set boundary (red) and element boundary (blue) are marked as entrances (+) or exits (-). Entrances are connected to exits counterclockwise along the element boundary and elements are connected to exits counterclockwise along the pixel set boundary (black arrows).

with the simple homogeneity calculation of Eq. 1. We find that our method is comparably fast in many cases, especially when the elements are large compared to the voxels.

2. WHY IS THIS HARD?

The simplest way of computing the volume of the intersection between a set of voxels and an element is to consider each voxel independently, and sum over them at the end. As discussed above, this can be slow and inefficient, even in 2D. A more attractive algorithm is to extract the boundary of the set of voxels, and compute its intersection with the boundary of the element. In this way one can compute each facet of the intersection region, and therefore its volume. OOF1 and OOF2 use this method, although it is not without its subtleties.

For simplicity, we will demonstrate the method and one of its failure modes in 2D. Fig. 3 shows a triangular element intersecting a set of pixels (which need not be simply connected). We want to compute the area, \mathcal{A} , of the shaded intersection, \mathcal{I} (which need not be simply connected). We use a method based on O’Rourke’s[9] algorithm for intersecting polygons. First, all of the pixels of a given type are aggregated, and their boundary is extracted, as shown in red. Then the intersection points between the segments of the pixel boundary and the segments of the element boundary are detected. If the element boundary (shown in blue) is traversed counterclockwise, each intersection point can be labelled as an entrance or an exit, depending on

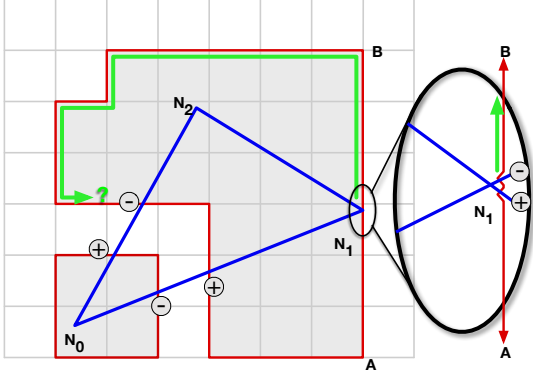


Figure 4: Failing to find the area of intersection of an element and a set of pixels in 2D. The entrances and exits at node N_1 are in the wrong order, so the green segments leaving the exit are incorrect.

whether the element boundary enters or exits the pixel set at that point. Then the new boundary segments of the intersection can be constructed by following the element boundary from an entrance to the next exit, and following a pixel set boundary from an exit to the next entrance. To compute the area, the segments don't even need to be linked into loops. If the segment i goes from point S_i^0 to S_i^1 , the area is

$$\mathcal{A} = \frac{1}{2} \sum_{i=0}^N (S_i^0 - C) \times (S_i^1 - C) \quad (3)$$

for some conveniently chosen center point C . To avoid loss of precision in the sum, C should not be too far away from the segments.

Now consider what happens if the rightmost node of the element in Fig. 3 coincides with the pixel set boundary at location N_1 , as shown in Fig. 4. Note that this is not at all an unlikely situation, since any algorithm that makes elements homogeneous will tend to put nodes on pixel boundaries. There are a number of things that can go wrong. Round-off error in the computations can lead to inconsistent information about whether point N_1 is inside or outside the pixel set, since it may make it seem that segment N_0N_1 intersects segment AB but N_1N_2 doesn't (or vice versa). Or if the program is smart enough to insist that N_1 must be either inside or outside, but not both, it can still misplace intersection points on segment AB , as shown in the inset. It's possible that node N_1 may be outside the element (by an infinitesimal amount) but that the computed entrance intersection of segment N_1N_2 with AB might *precede* the exit intersection of N_0N_1 on AB . In this case, following the pixel boundary from the the exit point to the entrance won't connect the two points, but will follow the green line, adding incorrect edges to the perimeter of \mathcal{I} . Thus a

microscopic error in the position of the points can lead to a macroscopic error in the area \mathcal{A} .

In 3D, the method and its failure modes are similar, except that both are more complex. We want to find the volume, \mathcal{V} , of the intersection, \mathcal{I} , of a set of voxels with a convex polyhedral element. First, the exterior of a set of voxels is divided into a set of facets. Each facet is a set of pixels in a plane. The intersection of that plane with the element is a polygon, and the intersection of that polygon with the pixels in the plane can be found using the 2D algorithm. These in-plane intersections form part of the exterior of \mathcal{I} . The remaining faces of \mathcal{I} lie on the faces of the element, and can be computed by finding the edges of the in-plane intersections that lie on the element faces, and linking them, if necessary, along the edges of the faces. All of the difficulties of the 2D method apply, with additional complications. Is a point actually on a face? Actually on an edge? Which face is it on? What if a segment appears to intersect only one face of the element but both of its endpoints appear to be outside?

In 2D, these difficulties can be overcome by careful bookkeeping and using simple geometry to resolve ambiguities, following a "topology first" scheme. For example, in the example of Fig. 4, the fact that two exits (-) appear consecutively is a clue that something is wrong.¹ In 3D, detecting and correcting errors devolves into a forest of special cases and the resulting code is fragile, inefficient, and untrustworthy.

3. POWELL AND ABEL'S R3D ALGORITHM

An algorithm by Powell and Abel [15], based on ideas from Sugihara [16], provides a robust way of computing the intersection volume. Powell and Abel's r3d algorithm is a way of conservatively transferring data from one mesh to another so that integrals of the data over the new mesh are the same as integrals over the old mesh. Their original paper [15] described a method that computed integrals of polynomials over the intersections of two convex polyhedra (the polyhedra being elements from the old and new meshes), but a modified algorithm [17] works even if one of the polyhedra is not convex. We can use the method to compute the volume of the intersections \mathcal{I} — our element is a convex polyhedron, our voxel set is a (possibly) non-convex polyhedron, and the polynomial that we are integrating is 1.

r3d works by constructing a planar graph of the non-convex polyhedron and clipping it successively by the

¹Actually, even in 2D OOF2 uses a method similar to that described in this paper because it is simpler and requires no consideration of special cases. Instead of examining the entire element, it trims the pixel set one element edge at a time.

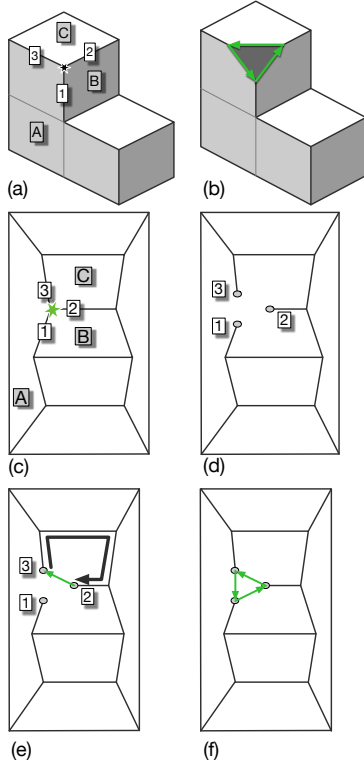


Figure 5: Clipping a single corner from a polyhedron composed of three voxels. (a) The unclipped polyhedron, with three faces and edges labelled. The clipping plane is not shown, but passes between the starred vertex and the rest of the polyhedron. (c) The graph of the polyhedron, with corresponding faces and edges labelled. (d) The clipped node has been removed from the graph, and new dangling nodes (circles) inserted. (e) The end-points of a missing edge are discovered by following graph edges from one dangling node to another, turning right at all intermediate nodes. (f) The completed graph, after all dangling nodes are completed. (b) The clipped polyhedron, reconstructed from the graph.

each of the planes of the faces of the convex polyhedron. The graph is isomorphic to the non-convex polyhedron — nodes of the graph correspond to vertices of the polyhedron, edges of the graph correspond to edges of the polyhedron, and they are connected identically. (We will use the word “vertex” to refer to corners of the polyhedron, and “node” to refer to the corresponding points in the graph.) If the polyhedron is sliced by a plane that divides it into two (possibly more, if nonconvex) regions, r3d constructs the graph of the new polyhedron on one side of the plane, using only information about which vertices are on which sides of the plane. Powell and Abel explain that by preserving topological consistency at all steps, and not relying on any arithmetic operations after deciding which ver-

tices to clip, the method is robust in the presence of numerical errors. Unlike the naive method discussed in Section 2, small errors in arithmetic lead only to small errors in the results.

To be specific, when intersecting polyhedron \mathcal{P} with convex polyhedron \mathcal{E} , the process is [17]:

1. Construct a planar graph of \mathcal{P} , making sure that the edges at each node are in the same order as the edges of the polyhedron at the corresponding vertex, when viewed from outside the polyhedron. Figs. 5a and c show a polyhedron and its planar graph. A planar graph is one that can be drawn in a plane with no crossing edges. Each node must link to exactly three edges. (If a vertex of the polyhedron is on more than three edges, that vertex must be represented by more than one node in the graph, possibly connected by edges of zero length.) The edges of the graph divide the plane into regions that correspond to the faces of the polyhedron, including the part of the plane that lies outside of the graph. (Every convex polyhedron can be represented by a planar graph, but sets of pixels can form non-convex polyhedra.)
2. Choose one of the faces of \mathcal{E} . The clipping plane is the plane containing this face.
3. Remove the vertices of the graph that are outside the clipping plane, leaving dangling edges, as shown in Fig. 5d. Put a new node at the end of each dangling edge. The position of the corresponding vertex on the polyhedron can be found by interpolating along the original unclipped polyhedron edge. If rounding error puts an intersection point just past the end of the edge, it can be assumed to coincide with the endpoint. The graph itself contains no position information, so in the graph the position of the node is immaterial.
4. Choose one of the new nodes, \mathbf{N}' . In Fig. 5e, this is node 3.
5. From \mathbf{N}' , follow existing edges from node to node, always picking the right hand branch at pre-existing nodes (all of which have three edges) until another new vertex, \mathbf{N}'' , is encountered. In Fig. 5e, this is vertex 2. It makes sense to refer to the “right hand branch” because the graph is planar.
6. Add a directed edge from \mathbf{N}'' to \mathbf{N}' , as shown in green in Fig. 5. The direction can be used later to determine the outward normal of the face, but is not strictly necessary.

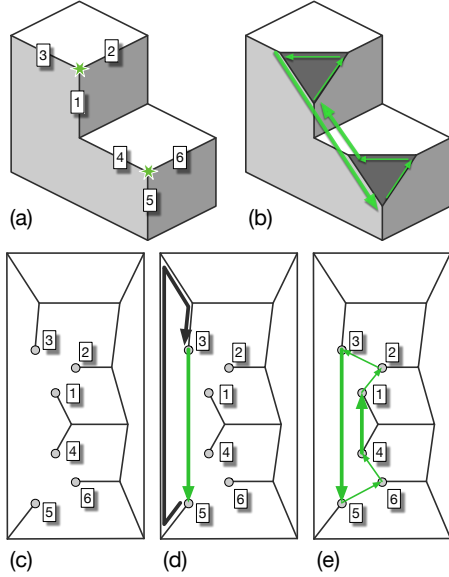


Figure 6: Clipping two corners from the polyhedron of Fig. 5. (a) The unclipped polyhedron. The starred vertices are outside the clipping plane and will be removed. (c) The polyhedral graph with the clipped nodes removed. The unclipped graph is the same as Fig. 5c. (d) Following the graph from new node 5 to new node 3 along the heavy black line creates an edge (green arrow) that appears to span too much of the graph. (e) The completed graph. (b) The clipped polyhedron. The heavy arrows indicate collinear edges of the new facets. The edges cancel each other out in the region between the new triangular facets (dark gray).

7. Go back to step 4 and repeat until all nodes have three edges, as in Fig. 5f.
8. Go back to step 2 and repeat for each face of the element \mathcal{E} .
9. Find the facets of the clipped polyhedron by picking an edge of the graph and moving from edge to edge, turning left at each node, until the original edge is reached. From the edges, compute the area and normal vector of each facet, and from that the volume.

Notice in Fig. 5 what happens if the vertex is within round-off error of the clipping plane, and the decision of whether or not it should be clipped is ambiguous. If it is clipped, then the topology of the new polyhedron is exactly what is shown in the figure, but the new facet will be small, and the volume removed will be insignificant, on the order of the error in position cubed. If it is not clipped, then of course the volume is unchanged. This is what makes the calculation robust. The volume calculation does not depend on how the position of an ambiguous point is resolved.

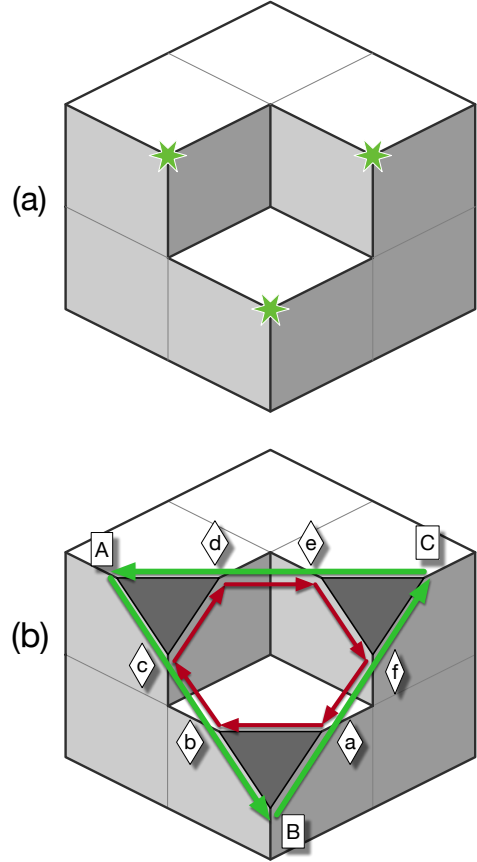


Figure 7: Clipping another concave set of voxels. (a) A set of 7 voxels, arranged in a cube with one corner missing. The corners marked with stars are outside the clipping plane and will be removed. (b) The result of applying the r3d algorithm. There are two new faces denoted by the green and red arrows. The faces are coplanar and oppositely directed, so the inner hexagonal face ($abcdef$) cancels the interior of the outer triangular face (ABC).

When the voxel set has concavities, it is possible that a clipping plane will intersect it in more than one region. Fig. 6 illustrates one such situation and the application of the algorithm to it. When both of the starred vertices in Fig. 6a are clipped, the intuitive solution is to replace each with a triangular facet. The algorithm instead generates two linked triangular facets, shown in Fig. 6b. However, the new vertices at points 3, 1, 4, and 5 are collinear (on the line formed by the intersection of the left face of the voxels and the clipping plane), so between points 1 and 4 the bridge between the triangles comprises two oppositely directed collinear segments. These make no net contribution to the area of the facet or volume \mathcal{I} of the clipped polyhedron, so this somewhat convoluted geometry is harmless. (If the vertices were *not* collinear, the left face

would necessarily be split into multiple facets, which would introduce new edges in the graph, and the algorithm would no longer connect vertices 3 and 5.) Similarly, in Fig. 7 a concave set of voxels is clipped so that three corners are removed. In this case, there are two new coplanar facets, a triangle and a hexagon, with oppositely directed normals. The triangular facet is indicated by the heavy green arrows in the figure. The hexagonal facet, indicated by the thin red arrows, cancels the area of the interior of the triangular facet, so that net result is the three dark triangular facets. Using directed edges when reconnecting the clipped graph makes it clear that the hexagonal and triangular faces have opposing normals. Again, the result is correct for the purpose of computing the volume of the clipped polyhedron.

To maintain topological consistency, r3d relies on the graph of the polyhedron being 3-vertex-connected, that is, not separable into two pieces by removing any two nodes. Polyhedra created from voxel sets can easily fail this test, but the algorithm still works with a slight modification. Consider the polyhedron shown in Fig. 8, built from a $2 \times 2 \times 4$ brick of voxels, with 2 voxels removed from one side. The graph of this polyhedron, Fig. 8b, is not 3-vertex-connected, because it falls into two pieces if nodes 2 and 3 are removed. Now clip the polyhedron with a plane that removes only vertices 2 and 3. This can only occur due to numerical error, because a plane that clips vertices 2 and 3 must necessarily also clip either vertex 1 or vertex 4 or both, but this is exactly the sort of numerical error that is expected and must be handled. Applying r3d to Fig. 8b results in the disconnected pair of graphs shown in Fig. 8c, with the new edges shown in green. The lower part, containing face C, is an inverted (negative volume) pentagonal prism. It is difficult to interpret the upper graph. Do the two oppositely directed edges cancel each other out? Does that leave dangling nodes at their ends?

The fact that this situation must be the result of numerical error indicates how it should be resolved. Vertices 1, 2, 3, and 4 must actually be collinear. (Any modification to the polyhedron that makes them non-collinear without adding extra vertices between them also restores 3-vertex-connectivity.) Infinitesimally perturbing the polyhedron by shaving a strip off of the corner, as shown in Fig. 9, will not change its volume. The perturbed polyhedron has two new faces, labelled ‘a’ and ‘b’, and its graph has 3-vertex-connectivity. Using r3d to clip the *four* starred vertices produces the disjoint pair of graphs in Fig. 8c. The lower graph is again an inverted pentagonal prism, one face of which has two infinitesimal edges. The upper graph is a pentagonal prism with an infinitesimally thin facet that is broken into three sub facets labelled a, b, and c, all positively oriented. Collapsing

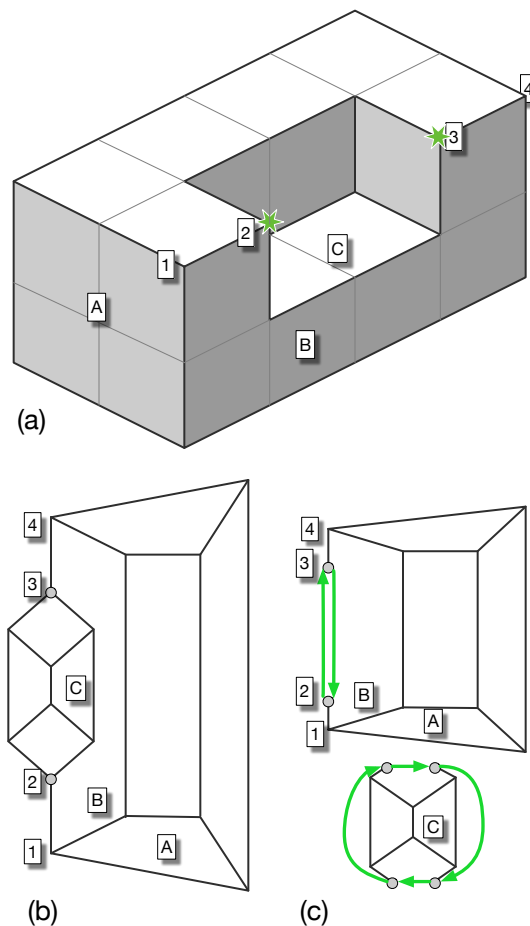


Figure 8: Applying r3d blindly to an insufficiently connected polyhedron. (a) The voxels forming the polyhedron. The vertices marked with stars are outside the clipping plane. (b) The graph of the polyhedron, with some nodes and faces labelled for comparison with (a). (c) The result of applying r3d to (b). The section including face C has become a disconnected graph.

the infinitesimal edges and faces leads to the graph and polyhedron in Fig. 10. (The arrows have been dropped from the new edges, because they no longer belong to new faces whose orientation must be determined.) The eight vertices on the thin facet in Fig. 9c are all collinear so the facets collapse into a single line from vertex 1 to vertex 4. In other words, the correct stratagem is not to connect nodes 2 and 3 in Fig. 8c with a double edge, but to connect nodes 1 and 4 with a single edge. The final graph is two separate quadrilateral prisms, one of which is inverted. The polyhedron is formed by the subtraction of the smaller polyhedron from the larger, and its net volume is the same as the initial set of voxels in Fig. 8, as it should be.

Computationally, it is difficult and expensive to verify

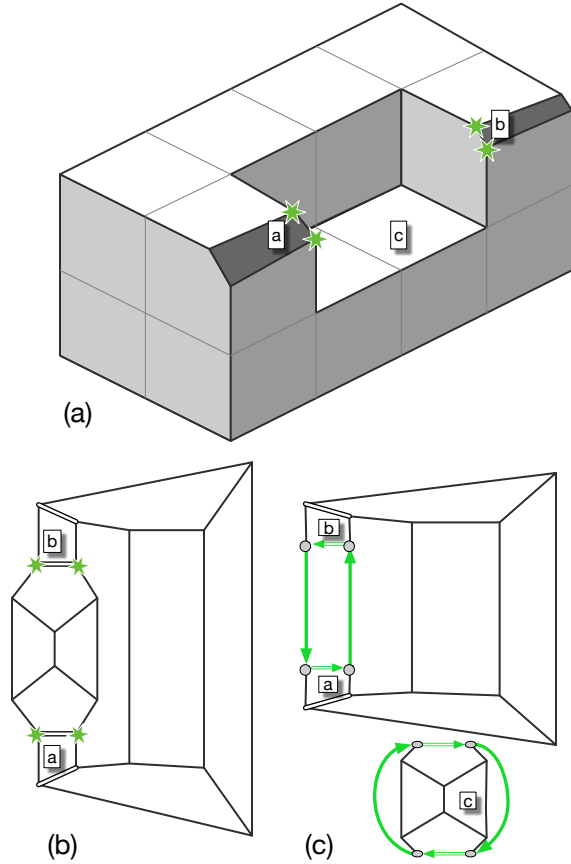


Figure 9: A perturbed version of Fig. 8. (a) The perturbed polyhedron, with two new infinitesimal faces. Stars mark the vertices to be clipped. (b) The graph of the polyhedron. The doubled edges are infinitesimal. (c) The result of applying `r3d` to (b).

3-vertex-connectivity and to know when this procedure will need to be imposed. However, it is fast and easy to check for doubled edges in the graph, and to repair them by replacing them and the single edges on either side (*i.e.*, 1234 in Fig 8c) with one single edge (14 in Fig 10a).

4. CONSTRUCTING GRAPHS OF VOXEL SET BOUNDARIES

Before `r3d` can be applied to a segmented image, the graph of each voxel set must be found. This can be a time consuming procedure, but it can be done just once, and doesn't need to be repeated unless the voxel categories change. Here we discuss how this can be done using C++.

Although the code described here was developed for OOF3D, it was written to be independent of it us-




Figure 10: Resolution of the problems in Fig. 8. Vertices 1 and 4 are joined directly (thin green line) as are vertices 2 and 3 (thick green line). The inner polyhedron is subtracted from the outer one.

ing C++ templates. The template parameters are an `IMAGE` (a three dimensional array class), an `IMAGEVAL` (the value of a voxel), a `COORD` (a floating point coordinate), an `ICOORD` (an integer coordinate). The simple requirements for those classes are described in the source code header files (see Section 8).

4.1 Algorithm Components

It's convenient to first describe the classes that we use to represent the voxel set boundary and its graphs.

A `VoxelSetBdy` describes a voxel set boundary. It creates a graph from an image and computes intersections.

A `VSBGraph` is the graph of the boundary edges of a voxel set. It is contained in a `VoxelSetBdy`.

A `VSBNode` is a node in a `VSBGraph`. Each `VSBNode` knows its neighboring nodes, and its position in space.

A `ProtoVSBNode` is a point in the image at a corner where voxels meet, and encodes the geometry of the voxel set at the corner.

The algorithm constructs a `VSBGraph` of the voxel set boundary. The first step of the algorithm is to identify relevant voxel corners, which are associated with `ProtoVSBNode` objects. The `ProtoVSBNode` object subtypes reflect the local geometry, which in turn determines how many (zero or more) `VSBNodes` are required for each `ProtoVSBNode`. Finally, using the geometry encoded in the `ProtoVSBNodes`, the appropriate edges are inserted between the corresponding `VSBNodes`, the `ProtoVSBNodes` are discarded, and the graph is now complete.

4.2 Creating Proto Nodes

The first step in building a graph is constructing an array of `ProtoVSBNode` (Proto Voxel Set Boundary Node) pointers, one for every voxel corner in the image. If the image is $k \times l \times m$ voxels, the array size is $(k + 1) \times (l + 1) \times (m + 1)$ because the `ProtoVSBNodes` live at the corners of the voxels, not the centers. Each `ProtoVSBNode` contains information about the eight voxels that meet at that point in the image. A voxel is *occupied* if it's in the current category. (On the image edges, the non-existent voxels outside the image are considered to be unoccupied.) Assigning a 1 to an occupied voxel and 0 to an unoccupied voxel and arranging the bits in an arbitrary but consistent order reduces the local configuration of eight voxels to a one byte signature. (See Fig. 11.) There are different `ProtoVSBNode` subclasses for different signatures. Although there are 256 signatures, only 17 `ProtoVSBNode` subclasses are needed, because voxel configurations that can be rotated into one another are represented by a single class. For example, one `ProtoVSBNode` subclass covers all of the eight configurations that have only one occupied voxel. Another subclass covers the 24 orientations of the configuration of three occupied voxels in Fig. 12. A `ProtoVSBNode` of the correct subclass is created at each voxel corner using a signature-based lookup table, storing the rotation required to bring the actual voxel configuration into alignment with the subclass's reference orientation. A `ProtoVSBNode` is not allocated for the 38 configurations of voxels that don't correspond to a corner in the aggregated voxel set.

4.3 Creating Graph Nodes

The second step is to create a new empty `VSBGraph` object, and for each `ProtoVSBNode` to create one or more `VSBNodes` in it. The `VSBGraph` is mostly just




Figure 11: The geometry of a `ProtoVSBNode`. The green circle is the position of the node. The binary digits are the labels of the eight neighboring voxels. A voxel is *occupied* if it's in the category whose graph is being computed. The *signature* of the `ProtoVSBNode` is the bitwise-or of the labels of the occupied voxels.




Figure 12: Two of the 24 voxel configurations in one of the `ProtoVSBNode` subclasses. The green circle is the position of the `ProtoVSBNode`, and the occupied voxels are shaded. The green lines show the directions in which this `ProtoVSBNode` needs to look for neighboring `ProtoVSBNodes` when connecting the graph.

a container for `VSBNodes`. A `VSBNode` stores its position in the image and pointers to its three neighboring `VSBNodes` in the graph (thereby mixing information from the real-space and graph-space representations of the polyhedron). In an unclipped graph, a `VSBNode`'s position is at integer voxel coordinates, but node positions in a clipped graph are not constrained and must be stored as floating point numbers. The list of pointers to neighbors is ordered so that neighbor $(i + 1) \bmod 3$ is clockwise from neighbor i , which facilitates graph traversal in the `r3d` algorithm.

A single `ProtoVSBNode` must create more than one `VSBNode` when there are more than three edges that meet at a point. A simple example is the case of two voxels (or any two rectangular blocks of voxels) that touch at a corner, as shown in Fig. 13. The graph of this polyhedron has two independent parts, although one node in each part shares its position with a node




Figure 13: A polyhedron formed from two voxels that touch at a corner, and its graph, which consists of two independent sections. The central vertex of the polyhedron is represented by two nodes (drawn inside the shaded circle) in the graph.

in the other. The `ProtoVSBNode` at that point creates two `VSBNodes` and links to six edges.

The colocated graph nodes in Fig.13 are not connected directly to one another, and should not be connected because they don't share faces of the polyhedron. Some `ProtoVSBNodes`, however, require the addition of multiple `VSBNodes` connected to one another by edges of length zero. One example is shown in Fig. 14.

4.4 Connecting Graph Nodes

The third step in constructing the graph is to add edges connecting the `VSBNodes`. The `ProtoVSBNodes` at each point know in which directions they connect, so they can search in those directions for their nearest neighbors. The two `ProtoVSBNodes` at either end of an edge then collaborate to decide which of their `VSBNodes` should be connected to one another, and, importantly, how the edges must be inserted into the `VSBNodes` so that they are in the correct order for `r3d`.

The `ProtoVSBNodes` do all of the work in determining how `VSBNodes` should be connected, because they have complete information about the local voxel structure. Each `ProtoVSBNode` subclass knows which of its `VSBNodes` connect to the graph edges in each direction in its reference orientation. For example, Fig. 14b shows the configuration of nodes and edges for the `ProtoVSBNode` shown in its reference orientation in Fig. 14a. The `VSBNode` numbered 0 connects to the `VSBNodes` numbered 1 and 5 and to a `VSBNode` created by the next `ProtoVSBNode` in the `+x` direction.

The actual connection process falls into one of three categories, described in the following subsections.

4.4.1 Connecting a Single Pair of Nodes

Connecting two `VSBNodes` in the same `ProtoVSBNode` is trivial, because the `ProtoVSBNode` contains all of




Figure 14: (a) A `ProtoVSBNode` for a configuration of four voxels (one is hidden). The green edges are the ones that connect to neighboring nodes. (b) A section of the graph created by the `ProtoVSBNode`. Six `VSBNodes` are required. The edges between them have zero length. The numbers indicate the indexing used internally to ensure that the correct edges are connected to the nodes.

the information that it needs. The simplest non-trivial case is when two `ProtoVSBNodes` must connect their `VSBNodes` along a direction that requires a single edge (ie, a situation like that in Figs. 12 or 14 but not Figs. 15 or 16). For two `ProtoVSBNode` objects `A` and `B` to figure out which of their `VSBNodes` connect along which edges, `A` calls `B.connect(A)`. `B` finds the spatial direction from itself to `A`, rotates the direction into its (`B`'s) reference frame, and thereby knows which connection is being made. For example, if `B` is the `ProtoVSBNode` in Fig. 14, and `A` is in direction `+y` in `B`'s reference orientation, then it's connecting to `B`'s node number 2. `B.connect` knows which of its `VSBNodes`, `bnode`, needs to be connected, but it doesn't know which of `A`'s to connect to. To find out, it calls `A.connectBack(B, bnode)`, which finds the appropriate `VSBNode` (`anode`) in `A`, inserts `bnode` in the correct slot, and returns `anode`, which `B.connect` can now insert in the correct neighbor slot in `bnode`. (Simple!) `ProtoVSBNode.connect` and `ProtoVSBNode.connectBack` are virtual functions, defined independently in each `ProtoVSBNode` subclass, enabling the connection process to depend on the local geometries of both `ProtoVSBNodes`.




Figure 15: A situation requiring both coincident nodes and coincident edges in the graph. The two blocks of the polyhedron (a) could be single voxels or blocks of voxels. They share the green edge. The graph (b) is the same as the graph of two cubes, although the two nodes at point 1 coincide, as do the nodes at point 2. Both green edges correspond to the green edge in (a). When constructing the graph, it is important that the nodes are connected correctly. Misconnecting nodes along the doubled edge as in (c) results in a graph that is not 3-vertex-connected.

4.4.2 Connecting a Double Pair of Nodes

Some varieties of `ProtoVSBNode` need to create two graph edges in the same spatial direction. Fig. 15 shows a simple example. There are 5 edges at each of the vertices marked 1 and 2, which are both represented by two graph nodes. The vertices share the same location in physical space, but the nodes are distinct in the graph topology. If the wires are crossed, so to speak, the resulting graph may not have the required connectivity and may produce incorrect results when clipped. Fig. 15c is not 3-vertex-connected, but as it happens this configuration doesn't cause a problem for `r3d`. Fig. 16, on the other hand, shows a case in which misconnection of the edges produces a graph

for which `r3d` results in nonsense, as seen in Fig. 17. We should not expect `r3d` to work on Fig. 17a, because the graph is not planar, but this illustrates the importance of ensuring that the connection procedure guarantees planarity.

Note that the two graph edges that leave the two `VSBNodes` in each of the `ProtoVSBNodes` always lie along the edges of two voxels that share an edge, as in Figs. 15 and 16. Call these the *key* voxels. Thinking of the key voxels as having an infinitesimal gap between them, each of the two graph edges follows the edge of its own voxel, and connects to the graph node associated with the same key voxel at both `ProtoVSBNodes`. The key voxels must have the same relative spatial positions in both `ProtoVSBNodes` (in real space, not necessarily in the reference orientation of the `ProtoVSBNode`), or else there would be another `ProtoVSBNode` intervening between A and B. Therefore, the correct way to connect the `VSBNodes` is to find the corresponding key voxels in each `ProtoVSBNode` and connect the `VSBNodes` that belong to the edges of those voxels. However, A and B may be in different `ProtoVSBNode` subclasses, with different orientations, and do not know the correspondence between their sets of key voxels.

If the key voxels at A are i and j , let p_i^A and p_j^A be their real space coordinates, and let N_i^A and N_j^A be the corresponding `VSBNodes`. p_i^A and p_j^A are vectors with integer coordinates, and we define an arbitrary but self-consistent greater-than operator for such vectors. (We are guaranteed that $p_i^A \neq p_j^A$.) Because the key voxels in B must have the same relative position as the key voxels in A, $p_i^A > p_j^A \iff p_k^B > p_l^B$ if k and l are the key voxels in B that correspond with i and j in A. This tells us how to connect the nodes. In the program, `A.connect(B)` finds A's key voxels and passes the ordered pair to a virtual function, `B.connectDoubleBack(A, N0, N1)`, where `N0` is a pointer to the `VSBNode` whose key voxel has the larger position and `N1` is a pointer to the other. `B.connectDoubleBack` figures out which of its key voxels has the larger position, sets the neighbor pointers in its `VSBNodes`, and returns them as an ordered pair. `A.connect(B)` can then set the neighbor pointers in its own nodes correctly.

4.4.3 A Special Case

There is one more subtlety to the node connection procedure. The `ProtoVSBNode` configuration shown in Fig. 18a contains two stacked voxels with a third sharing an edge with one of the other two. The `ProtoVSBNode` needs to create one `VSBNode` connecting along all three edges of the third voxel. Another graph edge passes through the `ProtoVSBNode` along the edge of the stacked voxels, but doesn't connect to a

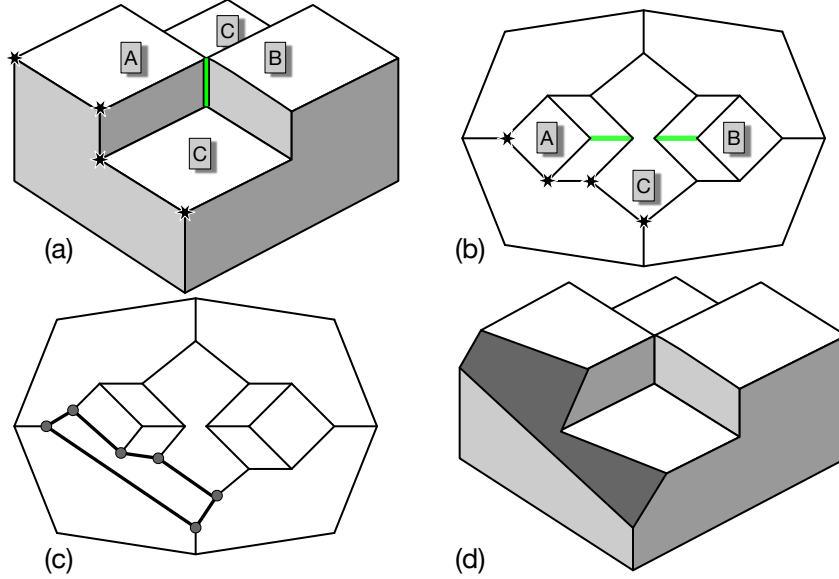


Figure 16: (a) A polyhedron that yields a doubled graph edge, as in Fig. 15, but in a non-trivial environment. The green edge in the polyhedron, where two voxels touch edge-wise, will become two edges in the graph. The two facets labeled C are coplanar. (b) The graph of the polyhedron, with the green edges from (a) highlighted. (c) The clipped graph, after the vertices marked with stars in (a) and (b) are removed. The heavy lines and circles mark the new vertices and edges. (d) The clipped polyhedron.

VSBNode there. This causes a problem, because that extra edge means that this ProtoVSBNode sits between two other ProtoVSBNodes that need to connect to one another, but ProtoVSBNodes always connect to their nearest neighbors. The expedient solution is to allow it to create a VSBNode with only two edges, as shown in Fig. 18c. Then after all connections are made, those nodes and their edges can be removed and replaced by a single edge.

4.5 Cleaning Up

Finally, after all connections are made, all of the ProtoVSBNodes can be deleted.

5. ACCURACY CHECKS

To check the accuracy of the algorithm, we used OOF3D to create artificial microstructures containing two categories of voxels and superimposed finite element meshes on the microstructure. We calculated the volumes of the elements in two ways: V_g is the geometric volume, using just the positions of the nodes, while V_r uses r3d to find the intersection of the element with each voxel category, and sums over categories. For each element we found the relative error, $E = |V_r - V_g|/V_g$, and computed its average, \bar{E} , standard deviation, δE , and maximum over the entire mesh, E_{\max} . In addition, we computed the total vol-

ume, V_i^i of each voxel category i by summing the r3d results in each element and also by simply counting the number of voxels N_i in each category in the entire microstructure, multiplied by the volume of a voxel. The relative discrepancy between these two results, $E_i = |V_i^i - N_i V_{\text{vox}}|/(N_i V_{\text{vox}})$ gives another error estimate for each category.

All tests were performed on microstructures that were cubes of side 1 (in arbitrary units). The cubes were divided into $M \times M \times M$ voxels. A uniform finite element mesh was created by dividing the microstructure into $N_x \times N_y \times N_z$ cubes, and splitting each cube into 5 tetrahedral elements. Usually, we took $N_x = N_y = N_z \equiv N$.

As a baseline check, we used only one category of voxel, making all intersection calculations trivial. With $M = 50$ and N ranging from 1 to 50, \bar{E} increased roughly linearly from 6.3×10^{-16} to 1.5×10^{-14} , with $\delta E \approx \bar{E}$ at each N . The largest value of E_{\max} was 1.6×10^{-13} . The error in the total volume of the category, E_0 , was larger and noisier as a function of N , but still always below 7×10^{-12} . (Machine epsilon was 2×10^{-16} .)

Nontrivial tests were done with two different microstructure geometries, both using two categories of voxels. In the first geometry, voxels were randomly assigned to one of two categories with probability p . In the second geometry, all voxels within a set of ran-




Figure 17: (a) The nodes on the doubled edge in the polyhedron in Fig. 16 are connected incorrectly. The graph is nonplanar. (b) The result of using r3d to remove the starred nodes from (a). The circles represent new nodes. Green lines show one plausible but incorrect facet. The red lines show r3d’s path through the graph when starting at the red node. The path returns to its starting node, indicating that the algorithm has failed. (c) The polyhedron represented by (b), drawn on top of the correctly clipped polyhedron from Fig. 16. The incorrect path is marked in red.

domly placed and possibly overlapping spheres were in one category, and voxels outside the spheres were in another. The number of spheres, their mean radius, and the width of the radius distribution were adjustable. All randomized configurations were repeated at least 20 times. Using M of 5, 50, and 100, and N of 4, 10, 11, no errors greater than 10^{-14} were observed in any runs of either geometry. The error magnitude increased when the elements’ aspect ratio was allowed to vary, but even setting $N_x = 10000$, $N_y = N_z = 1$ only raised the average error \bar{E} to 1.3×10^{-12} and E_{\max} to 1.2×10^{-11} (in a run with $M = 100$ and 60 spheres of mean radius 0.15).

The largest errors observed in any run occurred when the mesh was modified by moving nodes at random (while ensuring that elements were still well-formed). In one case, an element with a y -dimension that was 10^{-3} times its x and z dimensions had an error E of 3×10^{-10} (using a microstructure of 100 spheres with mean radius 0.1, $M = 100$, and $N = 4$). No larger errors were observed.

Given these results, it is reasonable to assume that errors in the calculation of element/voxel intersections will be negligible compared to other sources of error, such as finite element discretization or numerical error in the solution of PDEs.

6. COMPARISON TO THE SIMPLE METHOD

The simple homogeneity calculation given by Eq. 1 is substantially faster than the robust r3d method (up to 80 times faster in some observations), so it is natural to ask if the r3d method is actually necessary. To make Eq. 1 usable, we add the stipulation that the homogeneity of an element that contains *zero* voxel

centers is 1.0, and that it belongs to the voxel category at the location of the geometric center of the element.

A useful measure of how well a mesh approximates an image is what we call the “homogeneity index”, H . Recall that the homogeneity, $h(E)$, or its approximation $h_{\text{simple}}(E)$, is the largest fraction of an element’s volume that is occupied by a single voxel category. Assuming that the element’s properties are determined by that maximal voxel category, then the homogeneity index

$$H = \frac{\sum V(E)h(E)}{\sum V(E)} \quad (4)$$

is the volume fraction of the entire image in which the voxel category matches the element category. The sums in Eq. 4 are over elements. A mesh that perfectly matches its image will have $H = 1$. We can compute an approximate homogeneity index

$$H_{\text{simple}} = \frac{\sum V(E)h_{\text{simple}}(E)}{\sum V(E)} \quad (5)$$

similarly using by using the simple homogeneity calculation.

In Fig. 19 a uniform mesh has been imposed on a 3D image that is split 40:60 between two colors along a plane perpendicular to one axis. Meshes of various sizes were created by splitting the image into $N \times N \times N$ cubes, and dividing each cube into 5 tetrahedra. (In the figure, $N = 7$.) Then each mesh was annealed to increase its overall homogeneity – each node was displaced in x , y , and z by a distance chosen from a Gaussian distribution of width $\delta = 1$ (in voxel units) and the new position was accepted or rejected according to whether or not it increased the homogeneity of the element, using either $h(E)$ or $h_{\text{simple}}(E)$. A shape quality factor ensured that moves that created badly shaped elements were rejected[6], and nodes on





Figure 18: A special case. The ProtoVSBNode (a) needs to create a single VSBNode with three edges (b) and also allow neighboring VSBNodes in the $\pm z$ directions to connect directly to one another. Creating a temporary VSBNode (c) with only two edges allows the connection process to involve only nearest neighbor ProtoVSBNodes.

the corners, edges, and faces were constrained to stay on those corners, edges, and faces. In a single annealing iteration, each node was addressed once, and the process was repeated for 50 iterations after which the homogeneity index H was recorded. Then the mesh was reset to its initial state and annealed again. The results shown in Table 1 are averages over 10 runs. (Please note that this is *not* an efficient method of meshing this particular image.)

The $N = 1$ mesh contains no moveable nodes and can't be annealed. For $N = 5$ and $N = 10$ the boundary between green and blue pixels falls on a plane of element faces. All elements in the initial mesh are completely homogeneous with either calculation, so the annealing procedure does nothing. For $N = 7, 8,$ and 9 the simple algorithm falsely concludes that the initial mesh is completely homogeneous, and does nothing (because any node motion can at best decrease the homogeneity), while the r3d algorithm improves the homogeneity. Notice that annealing with r3d achieves a larger homogeneity index than the simple method for all N except 4.




Figure 19: (a) A uniform mesh on a simple two color $10 \times 10 \times 10$ voxel image. The inhomogeneous elements in the third layer give the whole mesh a homogeneity index $H = 0.971$, although the simple algorithm finds $H_{\text{simple}} = 1.0$. (b) After annealing, the homogeneity index has increased to $H = 0.992$.

The simple homogeneity tends to overestimate the homogeneity of the elements, causing the annealing process to accept moves that don't actually increase the homogeneity (when the the homogeneity after the move was erroneously large) or to reject moves that do (when the homogeneity before the move was erroneously large). To see if this overestimation occurs in a more realistic scenario, we generated artificial micrographs by selecting spherical regions of voxels within a $100 \times 100 \times 100$ image, as shown in Fig. 20, and compared H and H_{simple} for uniform meshes on those images. (For each image, 250 spherical regions were selected by choosing a center point from a uniform random distribution and choosing a radius from a Gaussian distribution with mean $R = 0.08$ and standard deviation $W = 0.03$, in units of the length of the image side. No effort was made to avoid overlapping spheres.) Voxels within spheres make up about 50% of each image, and for the purposes of the homogeneity calculation are assumed to be different from the rest of the voxels.

Table 2 contains the difference δ between H and H_{simple} for various mesh sizes, averaged over 50 different randomly generated microstructures. A positive difference indicates that the simple method overestimated the homogeneity on average. (When the difference is small, the simple calculation is occasionally a slight underestimate.) Although the differences are in general small, they are not all negligible. A fraction $1 - H$ of the voxels in a mesh are miscategorized by the mesh. These voxels are naturally concentrated at the internal boundaries in the image, which define the features that the mesh should reproduce. A better measure of the quality of a mesh (in this regard) is the excess inhomogeneity, shown in the last column of the table. This normalizes the difference δ by the *inhomogeneity* $1 - H$. The excess inhomogeneity can easily be a percent or two, and since this can arise from a small subset of the elements, it indicates that the simple al-

N	Initial		H after Annealing with...	
	H_{simple}	H	... H_{simple}	... H
1	0.7914	0.7893		
2	0.8981	0.9000	0.968 ± 0.006	0.9966 ± 0.0007
3	0.9537	0.9333	0.975 ± 0.004	0.996 ± 0.001
4	0.9328	0.9473	0.968 ± 0.003	0.966 ± 0.003
5	1.0	1.0	1.0	1.0
6	0.9892	0.9649	0.9686 ± 0.001	0.9700 ± 0.0005
7	1.0	0.9714	0.9714	0.9937 ± 0.0006
8	1.0	0.9475	0.9475	0.9935 ± 0.0009
9	1.0	0.9766	0.9766	0.9781 ± 0.0002
10	1.0	1.0	1.0	1.0

Table 1: Homogeneity indices computed before and after annealing a uniform mesh on the $10 \times 10 \times 10$ image shown in Fig. 19. N is the number of elements along the side of the image. The columns marked “Initial” contain the initial indices computed using the fast simple method of Section 6 (H_{simple}) and the true index computed with the r3d method (H). Note that the simple method generally overestimates the homogeneity. The last two columns contain the true homogeneity indices after annealing, where the annealing was conducted using either the simple or r3d homogeneity. The values are the mean and standard deviation of 10 different annealing runs on each mesh. For $N = 7, 8,$ and 9 the final homogeneity index after annealing with H_{simple} is the same as the initial homogeneity index because H_{simple} could not be improved by any node motion, so the mesh did not change.

gorithm may be performing badly in those cases.

The last column in Table 2 is the average ratio of the computation time² for the r3d method to the computation time for the simple point counting method. Two things should be noted. First, in many cases the r3d method is nearly as fast, or faster, than the simple method. Second, there are some cases in which r3d performs much worse. However, those are also the cases in which the simple method is least accurate.

The actual difference in execution time between the two algorithms depends in a complicated way upon the geometry of the image and the relative size of the elements and voxels. It may be possible to use a hybrid method in which the simple algorithm is used for large node motions on large, well formed elements, and the r3d algorithm is used for small elements, high aspect ratio elements, small node motions, and any other situations in which counting voxel centers is insufficiently descriptive.

As a robustness check, while computing Table 2 the total volume of each element was computed by summing the subvolumes computed by r3d and comparing them to the volume computed geometrically from the corners of the tetrahedron. A message was printed if the relative error was greater than 10^{-10} . This occurred only for the $N \times 1 \times 1$ meshes, and in all of

²The r3d times include the time taken to clip the graphs and compute the homogeneity, but not the time involved in constructing the graphs. This is because in many situations the graphs only need to be constructed once per image, but the clipping operation will be performed often on each element. The graph construction time depends on a number of factors, but for the examples used in the tables was roughly 10%-50% of the total time.

those cases the maximum relative error was less than 3×10^{-9} .

7. OPTIMIZATIONS

The elements of a mesh representing an image may be large on the scale of the voxels in the image, but are generally going to be small on the scale of the image itself. This means that if a graph is constructed from the whole image, a lot of time in the graph traversal process will be spent eliminating voxels that are far from the current element. The process can be made much more efficient if the image is first sliced up into bins that are a bit larger than the average size of an element, and separate graphs are created for each bin. Then only graphs whose bins intersect with the bounding box of the element need to be considered. The data presented in Tables 2 and 1 were computed in this way, with a minimum bin size of $5 \times 5 \times 5$ voxels. A further optimization, not implemented yet, would be to maintain sets of graphs for a variety of bin sizes and to use the appropriate set for each element.

If the image is very large, it may be inconvenient to create and store the full array of `ProtoVSBNodes` during the graph construction step. It may be more efficient to create the `ProtoVSBNodes` on demand, and to delete them when they’re known not to be needed. `ProtoVSBNodes` only need to look for neighbors in the positive $x, y,$ and z directions. If the loop over nodes is done in the positive direction, all nodes with x positions less than the current value of x can safely be deleted. (This optimization has not yet been implemented in OOF3D.)




Figure 20: An artificial microstructure consisting of randomly generated spheres of one material embedded in another material, typical of the microstructures used to generate Table 2.

It may be possible to speed up our code by reducing its object-oriented nature, at the loss of some generality. Powell and Abel’s r3d code in C uses lists and list indices, whereas ours uses objects and pointers.

8. OBTAINING THE CODE

The method described here is used in OOF3D but could be useful in other contexts in which convex polygons intersect image data. The implementation used in OOF3D is independent of other aspects of OOF3D so that it can be easily incorporated in other projects. It may be downloaded from <http://www.ctcms.nist.gov/oof/vsb>. OOF3D itself may be downloaded from <http://www.ctcms.nist.gov/oof/oof3d>. The voxel set boundary code is in OOF3D’s SRC/common/VSB subdirectory. Both downloads are free.

9. ACKNOWLEDGMENTS

The authors would like to thank Devon Powell for helpful conversations and for sharing his software.

References

- [1] Garcia R., Carter W., Langer S. “The Effect of Texture and Microstructure on the Macroscopic Properties of Polycrystalline Piezoelectrics: Application to Barium Titanate and PZN–PT.” *Journal of the American Ceramic Society*, vol. 88, no. 3, 750–757, 2005
- [2] Badilatti S., Kuhn G., Ferguson S., Müller R. “Computational modelling of bone augmentation in the spine.” *Journal of Orthopaedic Translation*, vol. 3, no. 4, 185–196, 2015
- [3] Langer S., Fuller E., Carter W. “OOF: An Image-Based Finite-Element Analysis of Material Microstructures.” *Computing in Science and Engineering*, vol. 3, no. 3, 15, 2001
- [4] Reid A., Langer S., Lua R., Coffman V., Haan S.I., García R. “Image-based finite element mesh construction for material microstructures.” *Computational Materials Science*, vol. 43, no. 4, 989–999, 2008
- [5] URL <http://www.ctcms.nist.gov/oof/oof2>
- [6] Coffman V., Reid A., Langer S., Dogan G. “OOF3D: An image-based finite element solver for materials science.” *Mathematics and Computers in Simulation*, vol. 82, 2951, 2012
- [7] URL <http://www.ctcms.nist.gov/oof/oof3d>
- [8] Frey P., Sarter B., Gautherie M. “Fully automatic mesh generation for 3-D domains based upon voxel sets.” *International Journal for Numerical Methods in Engineering*, vol. 37, 2735–2753, 1994
- [9] O’Rourke J., Chien C.B., Olson T., Naddor D. “A New Linear Algorithm for Intersecting Convex Polygons.” *Computer Graphics and Image Processing*, vol. 19, 384–391, 1982
- [10] Grandy J. “Conservative Remapping and Region Overlays by Intersecting Arbitrary Polyhedra.” *Journal of Computational Physics*, vol. 148, 433–466, 1999
- [11] Chen X., Zhang X., Jia Z. “A robust and efficient polyhedron subdivision and intersection algorithm for three-dimensional MMALE remapping.” *Journal of Computational Physics*, vol. 338, 1–17, 2017
- [12] Jia Z., Liu J., Zhang S. “An effective integration of methods for second-order three-dimensional multi-material ALE method on unstructured hexahedral meshes using MOF interface reconstruction.” *Journal of Computational Physics*, vol. 236, no. 513–562, 2013

- [13] Ray N., Sololov D., Lefebvre S., Lévy B. “Meshless Voronoi on the GPU.” *ACM Trans. Graph.*, vol. 37, no. 6, 265:1–265:12, 2018
- [14] Shewchuk J. “Adaptive Precision Floating-Point Arithmetic and Fast Robust Geometric Predicates.” *Discrete and Computational Geometry*, vol. 18, 305–363, 1997
- [15] Powell D., Abel T. “An exact general remeshing scheme applied to physically conservative voxelization.” *Journal of Computational Physics*, vol. 297, 340–356, 2015
- [16] Sugihara K. “A robust and consistent algorithm for intersecting convex polyhedra.” *Computational Graphics Forum*, vol. 13, no. 3, 45–54, 1994
- [17] Powell D. “r3d: Software for fast, robust geometric operations in 3D and 2D.” URL <https://github.com/devonmpowell/r3d/blob/master/1a-ur-15-26964.pdf>

$N_x \times N_y \times N_z$	δ		Excess inhomogeneity		$T_{\text{r3d}}/T_{\text{simple}}$
Equiaxed elements, decreasing size					
$1 \times 1 \times 1$	1.04×10^{-5}	$\pm 2. \times 10^{-5}$	2.37×10^{-5}	$\pm 4. \times 10^{-5}$	0.416
$5 \times 5 \times 5$	0.000318	$\pm 4. \times 10^{-5}$	0.00133	± 0.0002	1.7
$10 \times 10 \times 10$	0.000835	$\pm 5. \times 10^{-5}$	0.00602	± 0.0004	2.43
$20 \times 20 \times 20$	0.00205	$\pm 8. \times 10^{-5}$	0.0274	± 0.0008	5.24
$30 \times 30 \times 30$	0.0106	± 0.0004	0.195	± 0.003	7.66
$40 \times 40 \times 40$	0.0138	± 0.0005	0.327	± 0.005	10.3
Long, increasingly prolate elements					
$5 \times 5 \times 1$	9.16×10^{-5}	$\pm 6. \times 10^{-5}$	0.000259	± 0.0002	1.03
$10 \times 10 \times 1$	3.29×10^{-5}	± 0.0002	0.000106	± 0.0005	1.13
$20 \times 20 \times 1$	-3.03×10^{-5}	± 0.0005	-8.53×10^{-5}	± 0.002	1.27
$30 \times 30 \times 1$	0.00573	± 0.0006	0.0183	± 0.002	1.65
Prolate elements, fixed 10:1 aspect ratio					
$10 \times 10 \times 1$	3.29×10^{-5}	± 0.0002	0.000106	± 0.0005	1.13
$20 \times 20 \times 2$	1.02×10^{-5}	± 0.0004	5.67×10^{-5}	± 0.002	2.46
$30 \times 30 \times 3$	0.0065	± 0.0005	0.0328	± 0.002	2.86
$40 \times 40 \times 4$	0.00835	± 0.0008	0.0512	± 0.005	5.5
Prolate elements, fixed 5:1 aspect ratio					
$5 \times 5 \times 1$	9.16×10^{-5}	$\pm 6. \times 10^{-5}$	0.000259	± 0.0002	1.03
$10 \times 10 \times 2$	0.000118	± 0.0001	0.000444	± 0.0004	2.14
$15 \times 15 \times 3$	0.00214	± 0.0004	0.0102	± 0.002	1.68
$20 \times 20 \times 4$	0.000215	± 0.0002	0.00125	± 0.001	3.33
$25 \times 25 \times 5$	0.000589	± 0.0002	0.0041	± 0.001	2.96
Large, increasingly oblate elements					
$5 \times 1 \times 1$	6.28×10^{-5}	$\pm 5. \times 10^{-5}$	0.000152	± 0.0001	0.765
$10 \times 1 \times 1$	9.24×10^{-5}	± 0.0001	0.000225	± 0.0004	0.756
$20 \times 1 \times 1$	0.000514	± 0.0003	0.00126	± 0.0008	0.751
$30 \times 1 \times 1$	0.00217	± 0.0008	0.00534	± 0.002	0.883
Oblate elements, fixed 1:5 aspect ratio					
$5 \times 1 \times 1$	6.28×10^{-5}	$\pm 5. \times 10^{-5}$	0.000152	± 0.0001	0.765
$10 \times 2 \times 2$	0.00028	$\pm 6. \times 10^{-5}$	0.000827	± 0.0002	2.3
$15 \times 3 \times 3$	0.00126	± 0.0005	0.00454	± 0.002	1.45
$20 \times 4 \times 4$	0.00118	± 0.0001	0.00505	± 0.0005	3.39
$25 \times 5 \times 5$	0.00173	± 0.0001	0.00864	± 0.0005	1.93
$30 \times 6 \times 6$	0.00409	± 0.0003	0.0234	± 0.002	2.47
Oblate elements, fixed 1:10 aspect ratio					
$10 \times 1 \times 1$	9.24×10^{-5}	± 0.0001	0.000225	± 0.0004	0.756
$20 \times 2 \times 2$	0.000689	± 0.0002	0.00205	± 0.0006	2.23
$30 \times 3 \times 3$	0.00372	± 0.0005	0.0135	± 0.002	1.89
$40 \times 4 \times 4$	0.00449	± 0.0006	0.0195	± 0.003	4.36
$50 \times 5 \times 5$	0.00707	± 0.0006	0.0355	± 0.003	2.6

Table 2: The difference δ between the homogeneity index H_{simple} calculated using the simple algorithm and the index H calculated using the r3d algorithm. The average and standard deviation were computed from a set of 50 random microstructures, one of which is shown in Fig 20. The meshes were created by dividing the microstructure into $N_x \times N_y \times N_z$ rectangular prisms, and splitting each prism into 5 tetrahedral elements. For all geometries, $H_{\text{simple}} > H$ (i.e., $\delta > 0$), although the difference is insignificant when the elements are large compared to the voxels. The excess inhomogeneity, $(H_{\text{simple}} - H)/(1 - H)$, is scaled by the amount of inhomogeneity in the mesh and is a better measure of how well the simple algorithm approximates the more accurate r3d algorithm. The last column is the average ratio of the running time for r3d relative to the time for the simple algorithm. Note that the times are comparable, except when the elements are small, which is also where the simple method is least accurate.