Maximal Convexity, Flexible Sheets, and Segmented Medical Images

Jason D. Williams

Overview My research lies in the intersection of geometry and computing. I work on problems that combine solid modeling, differential geometry, computer graphics, computational geometry, and mathematical morphology. I have professional and academic experience in medical imaging, and I intend to use that experience to apply geometric computing to biomedical engineering. Given the increasing availability of 3D printers and other computer-controlled tools, I also intend to develop computer-aided design techniques targeted at custom, on-demand digital fabrication.

Below I pose three research questions that I have addressed, but which merit further investments of time, skill, and material resources. The first asks how we can design maximally convex shapes; the second asks how we can provide software tools for the design of objects made from flexible, inelastic sheets; and the third asks how we can estimate the boundary of an anatomical structure from a segmented medical image.

How can we maximize convexity? Convexity is fundamental in both pure and applied mathematics. Optimizing a convex function over a convex domain is far easier than more general kinds of optimization, and convex hulls are routinely used in problems involving collision detection and occlusion. Classically, a set is either convex or it is not. What would it mean to characterize a set as "nearly convex?" We might easily face a situation where we would like to take advantage of the properties of a convex set, but no convex set exists that satisfies our problem's constraints. Alternatively, we might want to control the tradeoff between the benefits of convexity and geometric fidelity: "What is the most convex approximation of this set, given a limited error budget?"

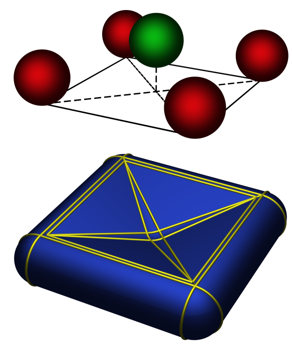

My approach to convexity measurement is to adapt a strategy from the differential geometry literature on tight embeddings, which defines a set as maximally convex if its expected number of height function critical points over all random height functions is minimal. The boundary of a maximally convex set then exhibits a minimum amount of oscillation, making it *tight*. My contribution is to express tightness in terms of normal field behavior. I define normal fields for arbitrary subsets of Euclidean space, along with a variety of means for describing a set's normal field variation. Under my approach, convexity is a linearly decreasing function of total normal variation.

With tightness as a measure of convexity, we can define a variety of maximally convex sets, such as the *tight hull* I describe in my thesis. In the image above, for instance, the blue tight hull boundary wraps around the red balls, remaining maximally convex while separating the red balls from the green ball. The tight hull is a simple generalization of the convex hull: if there were no green ball, the blue surface would bound the red balls' convex hull. Its simplicity gives the tight hull generality, making it applicable in diverse shape design and surface construction scenarios.

How can we construct shapes from flexible sheets? Many shapes incorporate surfaces that can be produced by bending, folding, and joining flexible but inelastic sheets. Bending and folding, illustrated in a sculpture by Huffman, have proved popular in manufacturing because they are comparatively inexpensive deformation techniques. As a result, some of the largest freeform surfaces in architecture, such as the facade of the Guggenheim at Bilbao, are made from bent metal sheets. Robots can now introduce curved folds into metal, with applications in the automotive industry. Product packaging, display cases, pop-up books, and greeting cards are routinely manufactured by cutting and folding paper, and artists, designers, and hobbyists produce paper artworks from elaborate designs.

Computer-aided design tools offer limited support for working with flexible sheets. We can formalize flexible sheets as developable surfaces, which are precisely those surfaces that can be unfolded or unrolled onto a plane without distortion. Researchers continue to investigate computational techniques for manipulating developable surfaces. For instance, several investigators have asked how to compute the "best" developable patch interpolating a given boundary curve. My work on tight hulls provides a well-motivated answer for what the best patch should be, and it further provides a mechanism for constructing developable surfaces using volumetric constraints rather than interpolatory constraints. Constraining a set to contain or exclude another set is more difficult that constraining its boundary to pass through certain points, but handling volumetric constraints expands our possible applications. For instance, it lets us deform a developable surface by poking it with a virtual finger.

Objects composed of developable surfaces are amenable to custom, on-demand manufacturing with computer-controlled tools, because the primary tool needed to fabricate them is a cutter. Tools like laser cutters, vinyl cutters, and water jet cutters have limited 3D sculpting capabilities, but they can process large pieces of diverse materials. Exploiting these properties, designers like Ronen Kadushin distribute patterns for laser-cut consumer goods over the internet. In time, the economic importance of folding a custom-cut sheet of steel into a coffee table at home may rival that of assembling a comparable table bought at Ikea. Custom, on-demand manufacturing can reduce transportation and warehousing costs, and it offers consumers access to higher-quality materials and more personalized designs. For that mode of production to reach its full potential, however, designers need software tools for manipulating developable surfaces. My research provides such tools.

How can we estimate organ boundaries from segmented medical images?

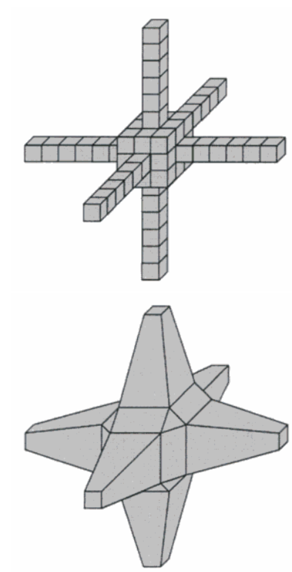

Given a grayscale medical image, we can extract an isosurface from it using an algorithm like Marching Cubes. Marching Cubes scans a volumetric image and builds an isosurface by examining the image sample values in a small moving window. While fast, this approach produces jagged, low-quality surfaces when applied to binary image data. That is the kind of data we obtain when we segment a medical image by labeling each sample as inside or outside of a particular anatomical structure. Formally, the isosurfaces produced by Marching Cubes are nonconvergent normal field and boundary measure estimators: no matter how high our input image resolution may be, the normal field and boundary measure of its Marching Cubes isosurface remain inaccurate.

Tight hulls can help us construct a boundary estimator with better convergence properties. The relative convex hull, which is a convex hull generalization defined in prior art, was proved by Sloboda and Zatko to provide a convergent estimate of a boundary's surface area. Heuristically, this is because relative convex hulls exhibit slanted surfaces where Marching Cubes produces staircase patterns. Tight hulls share this behavior, and because tight hulls minimize normal variation, their normal fields are "simple." This simplicity suggests that a tight hull's normal field is a good estimate of the normal field of the imaged object. A tight hull's shading will then be accurate when we render it on a computer — much more accurate than the shading of a jagged isosurface.

Although the relative convex hull is widely used for two-dimensional boundary estimation, efficient three-dimensional relative convex hull construction remains an open problem. My work to date suggests that tight hulls involving polyhedral sets can be constructed in low-order polynomial time, and I intend to exploit the properties of segmented image data to make their construction practical. I can then build on the tight hull algorithm for segmented images by developing a general tight hull algorithm that separates polyhedral sets with a simple polygonal surface. This tight separating surface has applications to statistics, classification, and other aspect of applied mathematics.

Summary Using the concept of convexity maximization, I have outlined approaches to two problems: designing shapes composed of developable surfaces and estimating the boundaries of anatomical structures from segmented medical image data. These problems illustrate two areas where geometric computing can grow and contribute: digital fabrication and medical modeling. My work to date offers clear paths for fostering that growth.