

Preface




FIGURE 1. Discrete differential geometry re-imagines classical ideas from differential geometry without reference to differential calculus. For instance, surfaces parameterized by principal curvature lines are replaced by meshes made of circular quadrilaterals (top left), the maximum principle obeyed by harmonic functions is expressed via conditions on the geometry of a triangulation (top right), and complex-analytic functions can be replaced by so-called *circle packings* that preserve tangency relationships (bottom). These discrete surrogates provide a bridge between geometry and computation, while at the same time preserving important structural properties and theorems.

1. Overview

The emerging field of discrete differential geometry (DDG) studies discrete analogues of smooth geometric objects, providing an essential link between analytical descriptions and computation¹ (see Figure 1). In recent years it has unearthed a rich variety of new perspectives on applied problems in computational anatomy/biology, computational mechanics, industrial design, computational architecture, and digital geometry processing at large.

The basic philosophy of discrete differential geometry is that a discrete object like a polyhedron is not merely an approximation of a smooth one, but rather a differential geometric object in its own right. In contrast to traditional numerical analysis, which focuses on eliminating approximation error in the limit of refinement (e.g., by taking smaller and smaller finite differences), DDG places an emphasis on the so-called “mimetic” viewpoint, where key properties of a system are preserved exactly, independent of how large or small the elements of a mesh might be. Just as algorithms for simulating mechanical systems might seek to exactly preserve physical invariants such as total energy or momentum, structure-preserving models of discrete geometry seek to exactly preserve global geometric invariants such as total curvature. More broadly, DDG focuses on the discretization of objects that do not naturally fall under the umbrella of traditional numerical analysis.

The Game. The spirit of discrete differential geometry is well-illustrated by a “game” often used to develop discrete analogs of a given smooth object:

- (1) Write down several equivalent definitions in the smooth setting.
- (2) Apply each smooth definition to an object in the discrete setting.
- (3) See which properties of the original smooth object are preserved by each of the resulting discrete objects, which are invariably inequivalent.

Most often, no discrete object preserves all the properties of the original smooth one—a so-called *no free lunch* scenario. Nonetheless, the properties that are preserved often prove invaluable for particular applications and algorithms. Moreover, this activity yields some beautiful and unexpected consequences, such as a connection between conformal geometry and pure combinatorics, or a description of constant-curvature surfaces that requires no definition of curvature! These notes provide an incomplete overview of several contemporary topics in DDG; a broad overview can also be found in the recent *Notices* article, “A Glimpse into Discrete Differential Geometry” by Crane and Wardetzky.

1.1. Geometry and finitism. There is a long debate in mathematics about the meaning and implications of various notions of infinity dating back at least to Aristotle and brought to a head in the late 19th century by Cantor. *Finitists* take the view that (in one way or another) the only “real” objects are those that have finite descriptions (perhaps allowing for “potential infinities,” like the natural numbers). For some, limiting mathematics to discrete or finite objects is a deep philosophical conviction; for others, it is an entertaining game. In the context of computation, the search for finite descriptions has a more pragmatic motivation: Turing’s abstract model of computation is discrete; real physical computing devices

¹This preface is adapted from Crane and Wardetzky, “A Glimpse into Discrete Differential Geometry,” *Notices of the American Mathematical Society*, Vol. 64, No. 10, pp. 1153–1159, 2017.

are necessarily finite. If one wishes to use machines to analyze geometry, one must therefore have good finite models².

Since one knows a priori that geometry must have a finite description, it seems worthwhile to consider geometric models that directly represent the way geometry behaves in nature, rather those that converge to the true behavior only in an unattainable limit. In fact, this situation leads to a natural question: Is some notion of infinity strictly required to faithfully model nature? Or can the basic things we want to say about natural geometry be captured by purely discrete models? The differential calculus of Leibniz and Newton has proven to be an unreasonably effective tool for modeling the natural world. For this reason it is tempting to feel that models based on differential equations are the true, canonical Platonic forms, and anything discrete and finite is a mere approximation. Yet nobody denies that even beloved continuum models (say, the Navier-Stokes equations) break down at very small scales, where nature looks increasingly like quantized packets of matter and energy. From this point of view, one might also be justified in feeling like nature is inherently discrete and that continuum models are perceived as “correct” only because they closely approximate observed behavior at the macroscopic scale. Moreover, infinity can creep into geometry in ways that betray our experience of nature. For instance, the *Banach-Tarski paradox* allows (by sly use of infinity) a solid ball to be decomposed into finitely many pieces, then reassembled into two balls—each of the same volume as the original one!

The purpose of mentioning such examples is not to ridicule the use of infinity, but rather to say: in order for a discrete model to be a “first class” description of an object, it need not duplicate all the strange behavior found in the continuous setting. One should instead ask, on both the smooth and discrete side: which phenomena are fundamental, and which are purely artifacts of the language in which the model is expressed? In fact, the dual perspectives offered by smooth and discrete differential geometry help to distill which phenomena are fundamental and which are superficial, since the most fundamental properties tend to arise naturally and easily on both sides (say, conservation laws in classical mechanics or global topological invariants in Riemannian geometry). In other words, geometry is neither inherently smooth nor discrete; rather, it is comprised of all the ideas that persist independent of the particular language used to write them down.

One can also tell this story from an analytical point of view: by considering only finite, discrete models, one obtains a degree of regularity that naturally excludes certain strange examples and behaviors. A nice example is the question of whether a 2-sphere embedded in \mathbb{R}^3 always divides space into two simply-connected regions. Intuitively, the answer is “yes,” since the 2-sphere is simply connected, and embeddings preclude any change in global topology. Yet one can construct a topological embedding (again by sly use of infinity) called *Alexander’s horned sphere*, where a loop around one of the “horns” can never be contracted. On the other hand, Alexander himself showed that if one restricts the question to discrete, or more precisely, *piecewise linear* embeddings, then no such example is possible. In other words, working in the simplicial category yields a result that is perhaps more in line with the kind of behavior that we would expect to encounter in nature. (Likewise, Banach-Tarski has no analogue in the finite, simplicial setting.)

²Note that even the symbolic expressions used to reason about geometry in a computer algebra system—or with pen and paper—are finite expressions in a formal grammar!

A more negative point of view—encapsulated by the *no free lunch* situation—is that discrete models are ultimately too rigid to describe all the geometric behavior observed in nature. The question of rigidity is a central and ongoing question in discrete differential geometry and a major theme of these notes. On the one hand, it is often true that no single discrete object in a certain class can faithfully reproduce a given list of properties from the smooth setting (see for instance our discussion of discrete Laplace operators). Are such results the end of the story, or does one simply need to consider an alternative approach to discretization? By adopting a different point of view, as outlined in *The Game*, situations that initially look hopeless sometimes give way to rich and flexible theories that capture much of the structure found in the smooth setting (such as conformal equivalence of triangle meshes). A major goal of discrete differential geometry is to see how far we can push the “finitist” point of view, not only to enable practical computation, but also to obtain clearer descriptions and a deeper understanding of the way shape behaves in nature.

Keenan Crane