

A Glimpse into Discrete Differential Geometry

Keenan Crane and Max Wardetzky

Communicated by Joel Hass

EDITOR'S NOTE. The organizers of the two-day AMS Short Course on Discrete Differential Geometry have kindly agreed to provide this introduction to the subject. The AMS Short Course runs in conjunction with the 2018 Joint Mathematics Meetings.

The emerging field of *discrete differential geometry* (DDG) studies discrete analogues of smooth geometric objects, providing an essential link between analytical descriptions and computation. 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




Figure 1. Discrete differential geometry reimagines 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.

Keenan Crane is assistant professor of computer science at Carnegie Mellon University. His e-mail address is kmcrane@cs.cmu.edu.

Max Wardetzky is professor of mathematics at University of Göttingen. His e-mail address is wardetzky@math.uni-goettingen.de.

For permission to reprint this article, please contact: reprint-permission@ams.org.

DOI: <http://dx.doi.org/10.1090/noti1578>

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. This article provides an overview of some of the themes in DDG.




Figure 2. A given geometric quantity from the smooth setting, like curvature κ , may have several reasonable definitions in the discrete setting. Discrete differential geometry seeks definitions that exactly replicate properties of their smooth counterparts.

The Game

Our article is organized around a “game” often played in discrete differential geometry in order to come up with a discrete analogue of a given smooth object or theory:

- (1) Write down several *equivalent* definitions in the smooth setting.
- (2) Apply each smooth definition to an object in the discrete setting.
- (3) Analyze trade-offs among the resulting discrete definitions, which are invariably *inequivalent*.

Most often, none of the resulting discrete objects preserve 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! To highlight some of the challenges and themes commonly encountered in DDG, we first consider the simple example of the curvature of a plane curve.

Discrete Curvature of Planar Curves

How do you define the curvature of a discrete curve? For a smooth arc-length parameterized curve $y(s) : [0, L] \rightarrow \mathbb{R}^2$, curvature κ is classically expressed in terms of second derivatives. In particular, if y has unit tangent $T := \frac{d}{ds}y$ and unit normal N (obtained by rotating T a quarter turn in the counter-clockwise direction), then

$$(1) \quad \kappa := \langle N, \frac{d^2}{ds^2}y \rangle = \langle N, \frac{d}{ds}T \rangle.$$

Suppose instead we have a polygonal curve with vertices $y_1, \dots, y_n \in \mathbb{R}^2$, as often used for numerical computation (see Figure 2, right). Here we hit upon the most elementary problem of discrete differential geometry: discrete geometric objects are often not sufficiently differentiable (in the classical sense) for standard definitions to apply. For instance, our curvature definition (Equation 1) causes trouble, since at vertices our discrete curve is not twice differentiable, nor does it have well defined normals. The

basic approach of DDG is to find alternative characterizations in the smooth setting that can be applied to discrete geometry in a natural way. With curvature, for instance, we can apply the fundamental theorem of calculus to Equation 1 to acquire a different statement: if φ is the angle from the horizontal line to the unit tangent T , then

$$\int_a^b \kappa \, ds = \varphi(b) - \varphi(a) \pmod{2\pi}.$$

Put more simply: curvature is the rate at which the tangent turns. This characterization can be applied naturally to our polygonal curve: along any edge the change in angle is clearly zero. At a vertex it is simply the *turning angle* $\theta_i := \varphi_{i,i+1} - \varphi_{i-1,i}$ between the directions $\varphi_{i-1,i}, \varphi_{i,i+1}$ of the two incident edges, yielding our first notion of discrete curvature:

$$(2) \quad \kappa_i^A := \theta_i \in (-\pi, \pi).$$

Are there other characterizations that also lead naturally to a discrete formulation? Yes: for instance we can consider the motion of y that most quickly reduces its length. In the smooth case it is well known that the change in length with respect to a smooth variation $\eta(s) : [0, L] \rightarrow \mathbb{R}^2$ that vanishes at the endpoints of the curve is given by integration against curvature:

$$\left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} \text{length}(y + \varepsilon\eta) = - \int_0^L \langle \eta(s), \kappa(s)N(s) \rangle \, ds.$$

Hence, the velocity that most quickly reduces length is κN . For a polygonal curve, we can simply differentiate the sum of the edge lengths $L := \sum_{i=1}^{n-1} |y_{i+1} - y_i|$ with respect to any vertex position. At a vertex i we obtain

$$(3) \quad \partial_{y_i} L = \frac{y_i - y_{i-1}}{|y_i - y_{i-1}|} - \frac{y_{i+1} - y_i}{|y_{i+1} - y_i|} =: T_{i-1,i} - T_{i,i+1},$$

i.e., just a difference of unit tangent vectors $T_{i,i+1}$ along consecutive edges. If $N_i \in \mathbb{R}^2$ is the unit angle bisector at vertex i , this difference can also be expressed as

$$(4) \quad \kappa_i^B N_i := 2 \sin(\theta_i/2) N_i,$$




Figure 3. Different characterizations of curvature in the smooth setting naturally lead to different notions of discrete curvature. (Here we abbreviate $T_{i,i+1}$ and $T_{i-1,i}$ by u and v , respectively.)

providing a discretization of the curvature *normal* κN . A closely-related idea is to consider how the length of a curve changes if we displace it by a small constant amount in the normal direction. As observed by Steiner, the new length of a smooth curve can be expressed as

$$(5) \quad \text{length}(\gamma + \varepsilon N) = \text{length}(\gamma) - \varepsilon \int_0^L \kappa(s) \, ds.$$

Since this formula holds for any small piece of the curve, it can be used to obtain a notion of curvature at each point. How do we define normal offsets in the polygonal case? At vertices we again encounter the issue that we have no notion of normals. One idea is to break the curve into individual edges which can then be translated by ε along their respective normal directions. We can then close the gaps between edges in a variety of ways: using (A) a circular arc of radius ε , (B) a straight line, or by (C) extending the edges until they intersect (see Figure 3, bottom left). If we then calculate the lengths for these new curves, we get

$$\begin{aligned} \text{length}_A &= \text{length}(\gamma) - \varepsilon \sum_{i=2}^{n-1} \theta_i, \\ \text{length}_B &= \text{length}(\gamma) - \varepsilon \sum_{i=2}^{n-1} 2 \sin(\theta_i/2), \\ \text{length}_C &= \text{length}(\gamma) - \varepsilon \sum_{i=2}^{n-1} 2 \tan(\theta_i/2). \end{aligned}$$

Mirroring the observation in the smooth setting, we can now say that whatever change we observe in the length provides a definition for discrete curvature. The first two are the same as ones we have seen already: the circular arc yielding the expression from Equation 2, and the straight line corresponding to Equation 4. The third one provides yet another notion of discrete curvature

$$\kappa_i^C := 2 \tan(\theta_i/2).$$

Finally, in the smooth case it is also well known that curvature has magnitude equal to the inverse of the radius of the so-called *osculating circle*, which agrees with the curve up to second order. A natural way to define an osculating circle for a polygon is to take the circle passing through a vertex and its two neighbors. From the formula for the radius R_i of a circumcircle in terms of the side lengths of the corresponding triangle, one easily gets a discrete curvature that is different from the ones we saw before:

$$(6) \quad \kappa_i^D := 1/R_i = 2 \sin(\theta_i)/w_i,$$

where $w_i := |y_{i+1} - y_{i-1}|$. Apart from merely being different expressions, we can notice that κ^A , κ^B , and κ^C are all invariant under a uniform scaling of the curve, whereas κ^D scales like the smooth curvature κ . This situation demonstrates another common phenomenon in discrete differential geometry, namely that depending on which smooth characterization is used as a starting point, one

The fundamental behavior of geometry is neither inherently smooth nor discrete.

may end up with pointwise or integrated quantities in the discrete case.

As one might imagine, there are many other possible starting points for obtaining a discrete analogue of curvature. Eventually, however, all starting points end up leading back to the same definitions, suggesting that there may be only so many possibilities. For example, if $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$ is the signed distance from a smooth closed curve γ , then applying the Laplacian Δ yields the curvature of its level curves; in particular, $\Delta\phi|_{\phi=0}$ yields the curvature of γ . Likewise, if we apply the Laplacian to the signed distance function for a discrete curve, we recover κ^A on one side and κ^B on the other. Yet another approach is the theory of *normal cycles* (as discussed by Morvan), related to the Steiner formula from Equation 5. Here, rather than settle on a single normal N_i at each vertex we consider *all* unit vectors between the unit normals of the two incident edges, ultimately leading back to the first discrete curvature κ^A . The theory of normal cycles applies equally well to both smooth and polygonal curves, again reinforcing the perspective that the fundamental behavior of geometry is neither inherently smooth nor discrete, but can be well captured in both settings by picking the appropriate ansatz. More broadly, the fact that equivalent characterizations in the smooth setting lead to different inequivalent definitions in the discrete setting is not special to the case of curves, but is one of the central themes in discrete differential geometry.

From here, a natural question arises: *which discrete curvature is “best”*? A traditional criterion for discriminating among different discrete versions is the question of convergence: if we consider finer and finer approximating polygons, will our discrete curvatures converge to the classical smooth one? However, convergence does not always single out a best version: treated appropriately, all four of our discrete curvatures will converge. We must




Figure 4. Typically, not all properties of a smooth object can be preserved exactly at the discrete level. For curve-shortening flow, for example, κ^A exactly preserves the total curvature, κ^B exactly preserves the center of mass, and with κ^D the flow remains stationary (up to rescaling) for any circular solution. However, no local definition of discrete curvature can provide all three properties simultaneously.

therefore look beyond convergence, toward exact preservation of properties and relationships from the smooth setting. Which properties should we try to preserve? The answer of course depends on what we aim to use these curvatures for.

As a toy example, consider the *curve-shortening flow* (depicted in Figure 4, top left), where a curve evolves according to the velocity that most quickly reduces its length. As discussed above, this velocity is equal to the curvature normal κN . A smooth, simple curve evolving under this flow exhibits several basic properties: it has at all times total curvature 2π , its center of mass remains fixed, it tends toward a circle of vanishing radius, and remains embedded for all time, *i.e.*, no self-crossings arise (Gage-Grayson-Hamilton). Do our discrete curvatures furnish these same properties? A numerical experiment is shown in Figure 4. Here we evolve our polygon by a simple time-discrete flow $y_i \leftarrow y_i + \tau \kappa_i N_i$ with a fixed time step $\tau > 0$. For κ^D , N_i is the unit vector along the circumradius; otherwise it is the unit angle bisector. Not surprisingly, κ^A preserves total curvature (due to the fundamental theorem of calculus); κ^B does not drift (consider summing Equation 3 over all vertices); and κ^D has circular polygons as limit points (since all velocities point toward the center of a common circle). However, no discrete curvature satisfies all three properties simultaneously. Moreover, for a constant time step τ no such flow can guarantee that new crossings do not occur. This situation illustrates the no free lunch idea: no matter how hard we try, we cannot find a single discrete object that preserves all the properties of its smooth counterpart. Instead, we have to pick and choose the properties best suited to the task at hand.

Suppose that instead of curvature flow, we consider two other beautiful topics in the geometry of plane curves: the Whitney-Graustein theorem, which classifies regular homotopy classes of curves by their total curvature, and Kirchhoff's famous analogy between motions of a spherical pendulum and *elastic curves*, *i.e.*, curves that extremize the bending energy $\int_0^L \kappa^2 ds$ subject to boundary conditions. Among the curvatures discussed above, only κ^A provides a discrete version of Whitney-Graustein, but does not provide an exact discrete analogue of Kirchhoff. Likewise, κ^C preserves the structure of the Kirchhoff analogy, but not Whitney-Graustein. This kind of no free lunch situation is a characteristic feature of DDG. A similar obstacle is encountered in the theory of ordinary differential equations, where it is known that there are no numerical integrators for Hamiltonian systems that simultaneously conserve energy, momentum, and the symplectic form. From a computational point of view, making judicious choices about which quantities to preserve for which applications goes hand-in-hand with providing formal guarantees on the reliability and robustness of algorithms.

We now give a few glimpses into recent topics and trends in DDG.




Figure 5. What is the simplicial analogue of a conformal map? Requiring all angles to be preserved is too rigid, forcing a global similarity (left). Asking only for preservation of so-called *length cross ratios* provides just the right amount of flexibility, maintaining much of the structure found in the smooth setting such as invariance under Möbius transformations (right).

Discrete Conformal Geometry

A *conformal map* is, roughly speaking, a map that preserves angles (see Figure 1, bottom left). A good example is Mercator's projection of the globe: even though area gets stretched out badly—making Greenland look much bigger than Australia!—the directions “north” and “east” remain at right angles, which is very helpful if you are trying to navigate the sea. A beautiful fact about conformal maps is that any surface can be conformally mapped to a space of constant curvature (“uniformization”), providing it with a canonical geometry. This fact, plus the fact that conformal maps can be efficiently computed (*e.g.*, by solving sparse linear systems), have led in recent years to widespread development of conformal mapping algorithms as a basic building block for digital geometry processing algorithms. In applications, discrete conformal maps are used for everything from sensor network layout to comparative analysis of medical or anatomical data. Of course, to process real data one must be able to compute conformal maps on discrete geometry.

What does it mean for a discrete map to be conformal? As with curvature, one can play the game of enumerating several equivalent characterizations in the smooth setting. Consider for instance a map $f : M \rightarrow D^2 \subset \mathbb{C}$ from a disk-like surface M with Riemannian metric g to the unit disk D^2 in the complex plane. This map is conformal if it

preserves angles, if it preserves infinitesimal circles, if it can be expressed as a pair of real conjugate harmonic functions $f = a + bi$, if it is a critical point of the Dirichlet energy $\int_M |df|^2 dA$, or if it induces a new metric $\tilde{g} := df \otimes df$ that at each point is a positive rescaling of the original one: $\tilde{g} = e^{2u} g$. Each starting point leads down a path toward different

too few degrees of freedom relative to the number of constraints

consequences in the discrete setting, and to algorithms with different computational tradeoffs.

Oddly enough, the most elementary characterization of conformal maps, angle preservation, does not translate very well to the discrete setting (see Figure 5). Consider for instance a simplicial map that takes a triangulated disk $K = (V, E, F)$ to a triangulation in the plane. Any map that preserves interior angles will be a *similarity* on each triangle, *i.e.*, it can only rigidly rotate and scale. But since adjacent triangles share edges, the scale factor for all triangles must be identical. Hence, the only discrete surfaces that can be conformally flattened in this sense are those that are (up to global scale) *developable*, *i.e.*, that can be rigidly unfolded into the plane. This outcome is in stark contrast to the smooth setting, where *any* disk can be conformally flattened. This situation reflects a common scenario in DDG: *rigidity*, or what in finite element analysis is sometimes called *locking*. There are simply too few degrees of freedom relative to the number of constraints: we want to match angles at all $3F$ corners, but have only $2V < 3F$ degrees of freedom. Hence, if we insist on angle preservation we have no chance of capturing the flexibility of smooth conformal maps.

Other characterizations provide greater flexibility. One idea is to associate each vertex of our discrete disk K with a circle in the plane. A theorem of Koebe implies that one can always arrange these circles such that two circles are tangent if they belong to a shared edge and all boundary circles are tangent to a common circle bounding the rest. For a regular triangular lattice approximating a region $U \subset \mathbb{C}$, Thurston noticed that this map approximates a smooth conformal map $f : U \rightarrow D^2$ as the region is filled by smaller and smaller circles (see Figure 1, bottom), as later proved by Rodin and Sullivan. Unlike a traditional finite element discretization, these so-called *circle packings* also preserve many of the basic structural properties of conformal maps. For instance, composition with a Möbius transformation of the disk yields another uniformization map, as in the smooth setting. More broadly, circle packings provide an unexpected bridge between geometry and combinatorics, since the geometry of a map is determined entirely by incidence relationships.¹ On the flip side, this means a different theory is needed to account for the geometry of *irregular* triangulations, as more commonly used in applications.

An alternative theory starts from the idea that under a conformal map the Riemannian metric g experiences a uniform scaling at each point: $\tilde{g} = e^{2u}g$. In other words, vectors tangent to a given point $p \in M$ shrink or grow by a positive factor e^u . In the simplicial setting g is replaced by a piecewise Euclidean metric, *i.e.*, a collection of positive edge lengths $\ell : E \rightarrow \mathbb{R}_{>0}$ that satisfy the triangle inequality in each face. Two such metrics $\ell, \tilde{\ell}$ are then said to be *discretely conformally equivalent* if they are related by $\tilde{\ell}_{ij} = e^{(u_i+u_j)/2} \ell_{ij}$ for any collection of discrete scale factors $u : V \rightarrow \mathbb{R}$. Though at first glance this

relationship looks like a simple numerical approximation, it turns out to provide a complete discrete theory that preserves much of the structure found in the smooth setting, with close ties to theories based on circles. An equivalent characterization is the preservation of *length cross ratios* $\epsilon_{ijkl} := \ell_{ij}\ell_{kl}/\ell_{jk}\ell_{li}$ associated with each edge $e \in E$; for a mesh embedded in \mathbb{R}^n these ratios are invariant under Möbius transformations, again mimicking the smooth theory. This theory also leads to efficient, convex algorithms for discrete *Ricci flow*, which is a starting point for many applications in digital geometry processing.

More broadly, discrete conformal geometry and discrete complex analysis is an active area of research, with elegant theories not only for triangulations but also for lattice-based discretizations, which make contact with the topic of (discrete) integrable systems, discussed below. Yet basic questions about properties like convergence, or descriptions that are compatible with *extrinsic* geometry, are still only starting to be understood.

Discrete Differential Operators

Differential geometry and in particular Riemannian manifolds can be studied from many different perspectives. In contrast to the purely geometric perspective (based on, say, notions of distance or curvature), differential operators provide a very different point of view. One of the most fundamental operators in both physics and geometry is the Laplace–Beltrami operator Δ (or Laplacian for short) acting on differential k -forms. It describes, for example, heat diffusion, wave propagation, and steady state fluid flow, and is key to the Schrödinger equation in quantum mechanics. It also provides a link between analytical and topological information: for instance, on closed Riemannian manifolds the dimension of harmonic k -forms (*i.e.*, those in the kernel of Δ) equals the dimension of the k th cohomology—a purely topological quantity. The spectrum of the Laplacian (*i.e.*, the list of eigenvalues) likewise reveals a great deal about the geometry of the manifold. For example, the first nonzero eigenvalue of the 0-form Laplacian provides an upper and a lower bound on optimally cutting a compact Riemannian manifold M into two disjoint pieces of, loosely speaking, maximal volume and minimal perimeter (Cheeger–Buser). These so-called *Cheeger cuts* have a wide range of applications across machine learning and computer vision; more broadly, eigenvalues and eigenfunctions of Δ help to generalize traditional Fourier-based simulation and signal processing to more general manifolds.

These observations motivate the study of discrete Laplacians, which can be defined even in the purely combinatorial setting of graphs. Here we briefly outline their definition for orientable finite simplicial n -manifolds, such as polyhedral surfaces, without boundary. Our exposition is similar to what has become known as *discrete exterior calculus*. To this end, consider the simplicial boundary operators $\partial_k : C_k \rightarrow C_{k-1}$ acting on k -chains (*i.e.*, formal linear combinations of k -simplices). The corresponding dual spaces (cochains) $C^k := \text{Hom}(C_k, \mathbb{R})$ and

¹See “Circle Packing” in the December 2003 Notices www.ams.org/notices/200311/fea-stephenston.pdf.

respective dual operators $\delta_k : C^k \rightarrow C^{k+1}$ give rise to the chain complex

$$\{0\} \rightarrow C^0 \rightarrow C^1 \rightarrow \dots \rightarrow C^n \rightarrow \{0\}.$$

The chain property says that $\delta_k \circ \delta_{k-1} = 0$, and one hence obtains simplicial cohomology $H^k := \ker(\delta_k)/\text{im}(\delta_{k-1})$. To define a Laplacian in this setting, we equip each C^k with a positive definite inner product $(\cdot, \cdot)_k$, and let δ_k^* be the adjoint operator with respect to these inner products, i.e., $(\delta_k \alpha, \beta)_{k+1} = (\alpha, \delta_k^* \beta)_k$ for all α, β . The Laplacian on k -cochains is then defined as

$$\Delta_k := \delta_k^* \delta_k + \delta_{k-1} \delta_{k-1}^*.$$

The resulting space of harmonic k -cochains, $\{\alpha \in C^k \mid \Delta_k \alpha = 0\}$ is then isomorphic to H^k —just as in the smooth setting. This fact is independent of the choice of inner product, mirroring the fact that cohomology depends only on topological structure. Likewise, for any inner product one obtains a discrete Hodge decomposition

$$C^k = \ker(\Delta_k) \oplus \text{im}(\delta_{k-1}) \oplus \text{im}(\delta_k^*),$$

where here the subspaces *do* depend on the choice of inner product.

At this point we return again to the game of DDG: which choice of inner product is *best*? A trivial inner product leads to purely combinatorial *graph Laplacians*, which do not (in general) converge to their smooth counterparts (e.g., when approximating a smooth manifold by a polyhedral one). Another choice is to consider linear interpolation of k -cochains over n -dimensional simplices, resulting in what is known as *Whitney elements*. For $n = 2$, we get the so-called *cotan Laplacian* (Pinkall and Polthier), which is widely used in digital geometry processing. Though other choices are possible, we again encounter a no free lunch situation: no choice of inner product can preserve all the properties of the smooth Laplacian. Which properties do we care about? Beyond convergence, perhaps the most desirable properties are the *maximum principle* (which ensures, for instance, proper behavior for heat flow), and the property that, for flat domains, linear functions are in the kernel (leading to a proper definition of barycentric coordinates). For general unstructured meshes there are no discrete Laplacians with all of these properties. However, certain types of meshes (such as weighed *Delaunay triangulations*) *do* indeed allow for “perfect” discrete Laplacians, offering a connection between geometry and (discrete) differential operators.

Discrete Integrable Systems

Another topic that has provided inspiration for many ideas in DDG is parameterized surface theory. Consider for instance the problem of dressing a given surface by a fishnet stocking, i.e., a woven material composed of inextensible yarns following transversal “warp” and “weft” directions (see Figure 6, right). This task corresponds to decorating a surface with a tiling where each vertex is incident to four parallelograms. Infinitesimally, such a tiling is known as a weak *Chebyshev net* (Chebyshev 1878), and locally corresponds to a regularly parameterized

surface $f : U \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$ where the directional derivatives f_u and f_v along the coordinate directions satisfy $|f_u|_v = |f_v|_u = 0$, i.e., partial derivatives with respect to one parameter have constant length along the parameter lines of the other parameter. The special case of rhombic tilings ($|f_u| = |f_v| = 1$) are known as (strong) Chebyshev nets. Can every smooth surface be wrapped in a stocking? Locally (i.e., in a small patch around any given point) the answer is “yes.” Globally, however, there are severe obstructions to doing so, which provide some fascinating connections to physics.

Consider for instance the special case of so-called *K-surfaces*, characterized by constant Gauß curvature $K = -1$. Every K-surface admits a parameterization $f : U \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$ aligned with the two transversal *asymptotic directions* along which normal curvature vanishes. Hence, if N is the unit surface normal then

$$(7) \quad \langle f_{uu}, N \rangle = \langle f_{vv}, N \rangle = 0.$$

Asymptotic parameterizations are weak Chebyshev nets since

$$(8) \quad a := |f_u|, \quad b := |f_v| \quad \text{satisfy} \quad a_v = b_u = 0.$$

Moreover, one can show that the angle ϕ between asymptotic lines satisfies the *sine-Gordon equation*

$$(9) \quad \phi_{uv} - ab \sin \phi = 0,$$

and conversely, every solution to the sine-Gordon equation describes a parameterized K-surface. Hilbert used this equation (and Chebyshev nets) to prove that the complete hyperbolic plane cannot be embedded isometrically into \mathbb{R}^3 . More generally, the sine-Gordon equation has attracted much interest both in mathematics as an example of an infinite-dimensional *integrable system*, and in physics as an example of a system that admits remarkably stable *soliton solutions*, akin to waves that travel uninterrupted all the way across the ocean. Another key property of the sine-Gordon equation is the existence of a so-called *spectral parameter* $\lambda > 0$: Equation 9 is invariant under a rescaling $a \rightarrow \lambda a$ and $b \rightarrow \lambda^{-1} b$, giving rise to a one-parameter *associated family* of K-surfaces. Geometrically, the parameter λ rescales the edges of parallelograms while preserving the angle between asymptotic lines.

Do these properties depend critically on the smooth nature of the solutions, or can they also be faithfully captured in the discrete setting? Hirota derived such a discrete version without any reference to geometry. Later Bobenko and Pinkall suggested a geometric definition of discrete K-surfaces that recovers Hirota's equation. In their setting, discrete K-surfaces are defined as discrete (weak) Chebyshev nets with the additional property that all four edges incident to any vertex lie in a common plane. The last requirement is a natural discrete analogue of Equation 7. This definition of discrete K-surfaces also comes with a spectral parameter λ and results in Hirota's discrete sine-Gordon equation—without requiring any notion of discrete Gauß curvature. Only recently has a discrete version of Gauß curvature been suggested that results in discrete K-surfaces indeed having constant negative Gauß curvature.




Figure 6. Left: two discrete parameterizations of a pseudosphere (constant Gauß curvature $K = -1$), one with a Chebychev net along asymptotic directions (left) and another along principal curvature lines (right). Right: a discrete Chebyshev net on a surface of varying curvature, resembling the weft and warp directions of a woven material.



Figure 7. Discrete parameterized surfaces play a role in architectural geometry, where special incidence relationships on quadrilaterals translate to manufacturing constraints like zero nodal torsion, or offset surfaces of constant thickness. Here a curvature line parameterized surface discretized by a conical net is used in the design of a railway station.

For discrete K -surfaces with all equal edge lengths (i.e., the rhombic case) the four neighboring vertices of a given vertex must lie on a common circle. By considering a subset of the diagonals of the quadrilaterals, one obtains another quad mesh with the property that all quads have a circumscribed circle, resulting in so-called cK-nets (see Figure 6, left). In the discrete setting, regular networks of circular quadrilaterals play the role of curvature line parameterized surfaces (as in Figure 1, top left). This transformation therefore mimics the smooth setting, where the angle bisectors of asymptotic lines are lines of principal curvature. More broadly, the theory of quad nets with special incidence relationships is closely linked to physical manufacturing considerations in the field of architectural geometry. For example, a quad net is *conical* if the four quads around each vertex are tangent to a common cone—such surfaces admit face offsets of constant width, making them attractive for the construction of (for instance) glass-paneled structures, as in Figure 7.

For further reading, see *Discrete Differential Geometry* (2008, Alexander Bobenko ed.).

Image Credits

Figures 1–6 courtesy of Keenan Crane and Max Wardetzky.
Figure 7 courtesy of B. Schneider.
Author photo of Keenan Crane courtesy of Keenan Crane.
Author photo of Max Wardetzky courtesy of Max Wardetzky.

ABOUT THE AUTHORS

Keenan Crane works at the interface between differential geometry and geometric algorithms. His scientific career started out studying Pluto—back when it was still a planet!



Keenan Crane

Max Wardetzky's interest in discrete differential geometry comes from his outstanding and truly inspiring teachers in differential geometry and his exposure to marvelous people in computer graphics.



Max Wardetzky