Discrete Differential Geometry

MA5210 Reading Report 2

Ang Yan Sheng
A0144836Y

Thanks to the contributions of Gauss, Riemann, Grassmann, Poincaré, Cartan, and many others, we now have a comprehensive classical theory of differential geometry. The most famous use of this theory might be in Einstein's theory of general relativity, but even today, differential geometry sees substantial applications in diverse areas such as computational biology, computer graphics, industrial design, and architecture—in short, any field which requires some form of digital geometry processing.

Of course, a computer does not directly work with smooth surfaces or differential forms, but with certain finite, discrete approximations of such objects, such as meshes. However, discretising the smooth objects and concepts of classical differential geometry poses some nontrivial problems. Crane and Wardetzky (2017) describe a "game" which is often played in this context:

1. Start with several *equivalent* definitions in the smooth setting;
2. Apply each definition to an object in the discrete setting;
3. Analyse the trade-offs among these (usually *inequivalent*) discrete definitions.

In this report, I will look at some instances of this phenomenon, also known as the *no free lunch scenario*, and highlight some of the results in the area of *discrete differential geometry*, which attempts to study the link between discrete geometric structures and their smooth counterparts.

Curvature

As a first example, we consider the curvature of a smooth plane curve $\gamma : [0, L] \to \mathbb{R}^2$. Assuming that γ is parametrised by arc length, its unit tangent vector is given by $T(t) := \gamma'(s)$, and its unit normal by $N(s) := \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} T(s)$, ie. obtained by rotating T counterclockwise by $90°$. Then the *(signed) curvature* of γ is

$$\kappa(s) := \langle N(s), T'(s) \rangle.$$

Here we see a problem with transferring this definition directly to a polygonal curve $\gamma_1\gamma_2 \cdots \gamma_n$: its arc length parametrisation is not twice-differentiable. However, there are other methods of defining discrete curvature, by using other characterisations of κ.

Tangential angle Note that $T : [0, L] \to \mathbb{S}^1 \cong \mathbb{R}/2\pi\mathbb{Z}$ can be lifted to a continuous function $\varphi : [0, L] \to \mathbb{R}$, the *tangential angle* that T makes with the positive x-axis, which is unique up to a multiple of 2π. Now it is easy to show that curvature is the rate of change of the tangential angle, so

$$\int_a^b \kappa(s)\, ds = \varphi(b) - \varphi(a).$$

The concept of tangential angle also makes sense for a polygonal curve, and so we have a first definition of discrete curvature at a vertex:

$$\boxed{\kappa_i^A := \theta_i} := \angle(\gamma_i - \gamma_{i-1}, \gamma_{i+1} - \gamma_i) \in (-\pi, \pi),$$

with angle measured counterclockwise from the first vector to the second. With this definition, we have a discrete analog of the Gauss-Bonnet theorem for planar curves, namely

$$\int_0^L \kappa(s)\, ds = 2\pi \qquad \longrightarrow \qquad \sum_{i=1}^{n} \kappa_i^A = 2\pi,$$

for simple positively-oriented smooth curves (left) or polygonal curves (right).

Length variation Another characterisation of κ is given by the following: for any smooth function $\eta : [0, L] \to \mathbb{R}^2$ such that $\eta(0) = \eta(L) = 0$, the variation of the arc length of γ is given by

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

Hence the variation that results in the fastest decrease in arc length is given by $\eta(s) = \kappa(s)N(s)$.

For polygonal curves, we can also look at the variation of $L = \sum_{i=1}^{n-1} \|\gamma_{i+1} - \gamma_i\|$:

$$\begin{aligned}\nabla_{\gamma_i} L &= \frac{\gamma_i - \gamma_{i-1}}{\|\gamma_i - \gamma_{i-1}\|} - \frac{\gamma_{i+1} - \gamma_i}{\|\gamma_{i+1} - \gamma_i\|} \\ &= -2\sin(\theta_i/2)N_i,\end{aligned}$$

where N_i is the unit vector in the direction of the internal angle bisector of $\angle\gamma_{i-1}\gamma_i\gamma_{i+1}$. This gives a second notion of discrete curvature,

$$\boxed{\kappa_i^B := 2\sin(\theta_i/2).}$$

Alternatively, we can use the corollary

$$\left.\frac{d}{d\varepsilon}\right|_{\varepsilon=0} \int_a^b \|(\gamma + \varepsilon N)'(s)\|\, ds = -\int_a^b \kappa(s)\, ds.$$

We obtain the discrete analog of this operation by moving each edge of the polygonal curve outwards by ε. Depending on whether we connect the ends of the new edges by (A) circular arcs, (B) line segments, or (C) extending the edges to their intersection, we recover the two previous definitions of discrete curvature and obtain a third,

$$\kappa_i^A = \theta_i, \qquad \kappa_i^B = 2\sin(\theta_i/2), \qquad \boxed{\kappa_i^C := 2\tan(\theta_i/2).}$$

Osculating circles The usual geometric intuition behind curvature is that of the *osculating circle*, which agrees with the curve at a given point up to second order. In the discrete setting, we can consider the circumcircle of a given vertex and its two adjacent vertices, which gives rise to yet another discrete curvature,

$$\boxed{\kappa_i^D := \frac{1}{R_i} = \frac{2\sin\theta_i}{\|\gamma_{i+1} - \gamma_{i-1}\|}.}$$

We note that $\kappa^A, \kappa^B, \kappa^C$ are invariant under scaling, while κ^D varies inversely proportionally to the scaling factor, which is exactly the case for the usual curvature κ. This can be interpreted as saying that the first three discrete curvatures are "integrated" quantities, while the last one is a "pointwise" quantity.

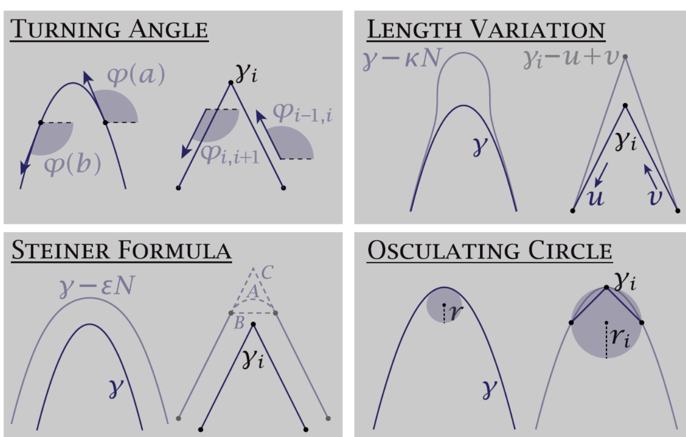

Figure 1: A summary of the four discrete curvatures.

Which of the above is the "best" discretisation? Given a polygonal curve that converges (in the appropriate sense) to a smooth curve, we can recover the usual smooth curvature from each of the four discrete curvatures (with appropriate scaling). Therefore we need to consider other properties, such as conserved quantities.

For instance, consider the problem of modelling the *curve-shortening flow*, where each point on the curve moves with velocity κN (Figure 2). It is known that if we start with a smooth simple closed curve, the curve-shortening flow satisfies the following properties:

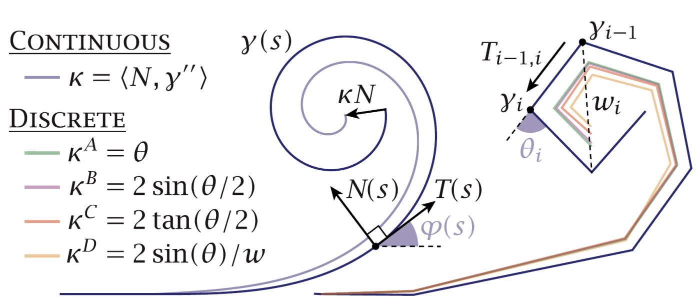

Figure 2: The smooth (left) and discrete (right) curve-shortening flows.

– the flow preserves the total curvature 2π;

– the flow preserves the centre of mass;
– the curve limits to a circle (after rescaling); and
– the curve remains simple (Gage-Grayson-Hamilton theorem).

To discretise this system, we can numerically solve $\frac{d}{dt}\gamma_i = \kappa_i N_i$, for instance by the iteration

$$\gamma_i(t+\tau) = \gamma_i(t) + \tau\kappa_i(t)N_i(t)$$

for a fixed time step $\tau > 0$. Here N_i is unit vector along the internal angle bisector of the two adjacent edges for κ^A, κ^B, or along the circumradius for κ^D. The results of such a simulation are summarised in Figure 3.

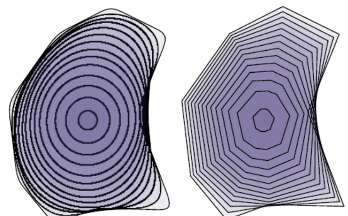

	Total	Drift	Round
κ^A	✓	×	×
κ^B	×	✓	×
κ^D	×	×	✓

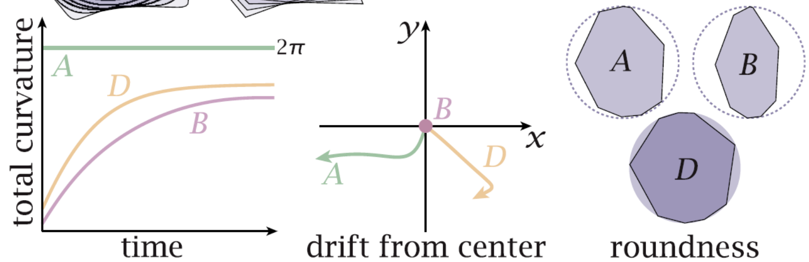

Figure 3: Three discrete curve-shortening flows, and their conserved quantities.

We have already seen that the flow under κ^A preserves total curvature, and it is easy to see (by summing the definition over all vertices) that the flow under κ^B preserves the centre of mass. It can also be shown that the limiting shape of the flow under κ^D is a cyclic polygon. However, no single discrete curvature satisfies all three properties. Moreover, for a given time step τ, none of the flows can guarantee that the polygonal curve stays simple. This is one example of the no free lunch scenario, and so the correct choice of discretisation depends on the most relevant conserved quantity for the problem.

Exterior calculus

To state our next example in a coordinate-free formulation, we will need to introduce some notions from exterior calculus.

Musical isomorphisms Recall that in $\mathbb{R}^n$, there are canonical isomorphisms (the *musical isomorphisms*) between $T_p\mathbb{R}^n$ and $T_p^*\mathbb{R}^n$, given by

$$\sum_i \alpha^i \frac{\partial}{\partial x^i} \underset{\sharp}{\overset{\flat}{\rightleftharpoons}} \sum_i \alpha_i dx^i.$$

Consider a surface S embedded in $\mathbb{R}^3$, with chart $f : U \subseteq S \to \mathbb{R}^2$. The embedding in $\mathbb{R}^3$ induces a metric tensor on $\mathbb{R}^2$ by $g_p(u, v) = \langle df^{-1}(u), df^{-1}(v)\rangle$. By the uniformisation theorem for Riemann surfaces, we may assume that f is conformal; hence we may transfer the musical isomorphisms to S via

$$u^\flat \coloneqq \det(g_p) f^*(df(u)^\flat), \qquad \alpha^\sharp \coloneqq \det(g_p)^{-1} df^{-1}((f^{-1})^*(\alpha)^\sharp)$$

for all $u \in T_pS$, $\alpha \in T_p^*S$. We can check that the above expressions are well-defined under conformal transition maps. In particular, we have identities corresponding to the $\mathbb{R}^n$ case, such as $u^\flat(v) = \langle u, v\rangle$.

Hodge duality Since the vector space of k-forms and $(n-k)$-forms on $\mathbb{R}^n$ have the same dimension, one might suspect that there is some duality relation between them. Indeed, we can define the *Hodge star* operator by the relation

$$\alpha \wedge *\beta = \langle\langle \alpha, \beta\rangle\rangle \omega$$

for any two k-forms α, β, where $\omega = dx^1 \wedge \cdots \wedge dx^n$ is the volume form, and $\langle\langle \alpha, \beta\rangle\rangle \coloneqq \sum_i \alpha_i \beta_i$ is the inner product on k-forms. This gives the relation

$$*(dx^{i_1} \wedge \cdots \wedge dx^{i_k}) = dx^{i_{k+1}} \wedge \cdots \wedge dx^{i_n},$$

for any even permutation $(i_1, \ldots, i_n)$ of $(1, \ldots, n)$.

We can also transfer this concept to a Riemannian n-manifold, by replacing the volume form with $\omega \coloneqq \sqrt{\det(g_p)} dx_1 \wedge \cdots \wedge dx_n$. In particular, the Hodge star operator is defined on surfaces embedded in $\mathbb{R}^3$.

Example: Vector calculus on $\mathbb{R}^3$ To showcase the notions that we have introduced, let us derive coordinate-free definitions of the gradient, curl and divergence operators from vector calculus on $\mathbb{R}^3$.

Let $\varphi : \mathbb{R}^3 \to \mathbb{R}$ be a scalar field (ie. 0-form). Its exterior derivative is just the differential $d\varphi = \sum_i \frac{\partial \varphi}{\partial x^i} dx^i$. Comparing with the gradient $\nabla\varphi = \sum_i \frac{\partial \varphi}{\partial x^i} \frac{\partial}{\partial x^i}$, we see that the only difference is that the first is a 1-form while the second is a vector field. Hence $\boxed{\nabla\varphi = (d\varphi)^\sharp.}$

Let $X = \sum_i \alpha^i \frac{\partial}{\partial x^i}$ be a vector field on $\mathbb{R}^3$, with associated 1-form $X^\flat = \alpha = \sum_i \alpha_i dx^i$. Then the exterior derivative of α is

$$\begin{aligned} d\alpha &= \sum_{i,j} \frac{\partial \alpha_j}{\partial x_i} dx^i \wedge dx^j \\ &= \left(\frac{\partial \alpha_3}{\partial x^2} - \frac{\partial \alpha_2}{\partial x^3}\right) dx^2 \wedge dx^3 \\ &\quad + \left(\frac{\partial \alpha_1}{\partial x^3} - \frac{\partial \alpha_3}{\partial x^1}\right) dx^3 \wedge dx^1 \\ &\quad + \left(\frac{\partial \alpha_2}{\partial x^1} - \frac{\partial \alpha_1}{\partial x^2}\right) dx^1 \wedge dx^2. \end{aligned}$$

Note that the components are exactly the same as the curl of X, but since $d\alpha$ is a 2-form while $\nabla \times X$ is a vector field, we have $\boxed{\nabla \times X = (*dX^\flat)^\sharp.}$

Also, note that

$$*X^\flat = X_1 dx^2 \wedge dx^3 + X_2 dx^3 \wedge dx^1 + X_3 dx^1 \wedge dx^2,$$

which has exterior derivative

$$d(*X^\flat) = \sum_i \frac{\partial X_i}{\partial x^i} dx^1 \wedge dx^2 \wedge dx^3.$$

Hence the divergence of X is $\boxed{\nabla \cdot X = *d * X^\flat.}$

The Laplacian For a scalar field φ, the *Laplacian* of φ is defined as $\Delta\varphi \coloneqq \nabla \cdot \nabla\varphi$, ie. the divergence of the gradient. By the discussion above, we may write this as $\boxed{\Delta\varphi = *d * d\varphi.}$

The Laplace-Beltrami operator

Note that the expression $\Delta\varphi = *d * d\varphi$ still makes sense for arbitrary Riemannian manifolds, where it defines the (scalar) *Laplace-Beltrami operator* on scalar fields*.

Proposition (Green's first identity). *Let M be a Riemannian n-manifold, and $\varphi, \psi : M \to \mathbb{R}$ be scalar fields. Then*

$$\int_M \langle \nabla\varphi, \nabla\psi \rangle = \int_{\partial M} \langle N, \nabla\varphi \rangle \psi - \int_M \psi \Delta\varphi,$$

where N is the outward unit normal of M.

Proof. By Stokes' theorem on the $(n-1)$-form $\psi * d\varphi$, we have

$$\begin{aligned}
\int_{\partial M} \psi * d\varphi &= \int_M d(\psi * d\varphi) \\
&= \int_M d\psi \wedge *d\varphi + \int_M \psi d * d\varphi \\
&= \int_M d\psi \wedge *d\varphi + \int_M \psi * \Delta\varphi \\
&= \int_M \langle\langle d\psi, d\varphi \rangle\rangle \omega + \int_M \langle\langle \psi, \Delta\varphi \rangle\rangle \omega \\
&= \int_M \langle \nabla\varphi, \nabla\psi \rangle + \int_M \psi\Delta\varphi.
\end{aligned}$$

*More generally, the k-form Laplacian is defined as $\Delta \coloneqq *d * d + d * d*$, but the second term is zero for scalar fields (since every $(n+1)$-form on an n-manifold is zero).

Now if ω_∂ is the volume form on ∂M, then $\omega = N^\flat \wedge \omega_\partial$ is the volume form on M. But

$$N^\flat \wedge *d\varphi = \langle\langle N^\flat, d\varphi\rangle\rangle\omega = \langle N, \nabla\varphi\rangle\omega,$$

so $*d\varphi = \langle N, \nabla\varphi\rangle\omega_\partial$ since $*d\varphi$ is a multiple of ω_∂. Hence the LHS of the first equation is $\int_{\partial M}\langle N, \nabla\varphi\rangle\psi$, and we are done. □

Suppose that S is a surface embedded in $\mathbb{R}^3$, and we are given $\varphi_0 : \partial S \to \mathbb{R}$. For many applications, we would like to find the "smoothest" interpolating function $\varphi : S \to \mathbb{R}$ which extends φ_0. One natural measure of such "smoothness" is the *Dirichlet energy*

$$E(\varphi) = \int_S \|\nabla\varphi\|^2.$$

To find the minimiser of this functional, we calculate the gradient of E. If $\psi : S \to \mathbb{R}$ is a variation function, so that $\psi = 0$ on ∂S, then

$$\|\nabla(\varphi + \varepsilon\psi)\|^2 = \|\nabla\varphi\|^2 + 2\varepsilon\langle\nabla f, \nabla g\rangle + O(\varepsilon^2),$$

so

$$E(\varphi + \varepsilon\psi) = E(\varphi) - 2\varepsilon\int_S \psi\Delta\varphi + O(\varepsilon^2).$$

Hence if φ is a minimiser for E, then $\Delta\varphi = 0$ on S. Such functions are called *harmonic functions* on S.

More generally, if we want to minimise the functional

$$E_X(\varphi) = \int_S \|\nabla\varphi - X\|^2,$$

where X is a vector field on S, then proceeding as above, we get

$$\begin{aligned} E_X(\varphi + \varepsilon\psi) &= E_X(\varphi) - 2\varepsilon\int_S (\psi\Delta\varphi + \langle X, \nabla\psi\rangle) + O(\varepsilon^2) \\ &= E_X(\varphi) - 2\varepsilon\int_S \psi(\Delta\varphi - \nabla\cdot X) + O(\varepsilon^2), \end{aligned}$$

by the divergence theorem. Hence φ is a solution to $\boxed{\Delta\varphi = \rho,}$ with $\rho = \nabla\cdot X$. This is the ubiquitous *Poisson problem*.

In practice, interpolation with the Laplace-Beltrami operator can be used on positions, displacements, vector fields, and other functions. These can be incorporated into more advanced applications, such as surface parametrisation, vector field design, shape matching and reconstruction, geodesic distance computation, and fluid dynamics. The eigenfunctions of this operator also play an important role in geometric processing—in fact, Fourier analysis on $\mathbb{R}^n$ and spherical harmonics are just two special cases of this powerful tool.

On the other hand, for a discretisation of a surface, such as a polygonal mesh, it is far from clear how to obtain an analogue of the Laplace-Beltrami operator with the corresponding properties. We now present two different ways of achieving this discretisation; surprisingly, they yield the same result!

Finite element method For scalar fields $\varphi, \psi : S \to \mathbb{R}$, define the L^2 inner product $\langle \varphi, \psi \rangle \coloneqq \int_S \varphi\psi$, and write $\|\varphi\| = \langle \varphi, \varphi \rangle^{1/2}$.

Consider a Poisson problem $\Delta\varphi = \rho$. Let f_i be a set of basis functions, and suppose that φ is a linear combination of the f_i. If we want φ to be a good approximation to the solution, in the sense that $\|\Delta\varphi - \rho\|$ is small, then a natural condition to consider is

$$\langle \Delta\varphi - \rho, f_i \rangle = 0 \qquad \text{for all } i.$$

Assume that S has no boundary. Then for a triangular mesh of S, we can split the first term of the above expression into a sum over triangles σ_k, then apply Green's identity:

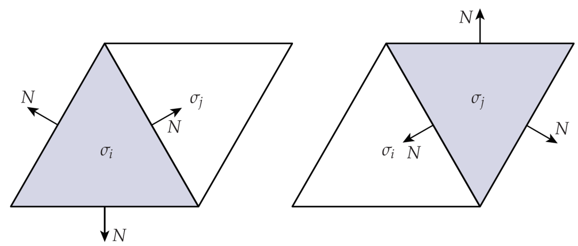

Figure 4: Unit normals across an edge sum to zero.

$$\begin{aligned}\langle \Delta\varphi, f_i \rangle &= \sum_k \langle \Delta\varphi, f_i \rangle_{\sigma_k} \\ &= \sum_k \left(\langle \nabla\varphi, \nabla f_i \rangle_{\sigma_k} + \langle N \cdot \nabla\varphi, f_i \rangle_{\partial\sigma_k} \right).\end{aligned}$$

Now the second term vanishes since each edge of the mesh is traversed twice, and the unit normals cancel each other out (Figure 4).

Writing $\varphi = \sum_j \lambda_j f_j$, we then get a linear system of equations

$$\sum_j \lambda_j \langle \nabla f_i, \nabla f_j \rangle = \langle \rho, f_i \rangle \qquad \text{for all } i,$$

ie. it remains to solve $Ax = b$ for the symmetric matrix $A_{ij} = \langle \nabla f_i, \nabla f_j \rangle$.

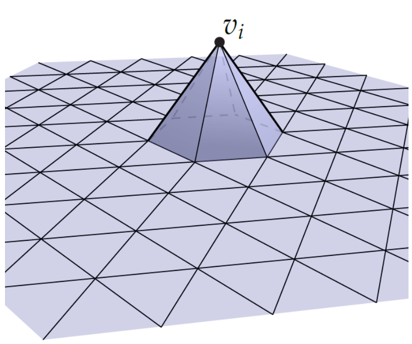

Figure 5: The basis functions f_i.

Given a triangular mesh, we have a convenient choice of basis: the piecewise linear hat functions f_i, which is 1 at the vertex v_i and 0 at all other vertices (Figure 5). Now it is an easy exercise in trigonometry to show that:

- In a triangle with vertex v_i, $\langle f_i, f_i \rangle = (\cot\alpha + \cot\beta)/2$, where α, β are the angles at the other two vertices;
- In a triangle with edge $v_i v_j$, $\langle f_i, f_j \rangle = -\cot\theta/2$, where θ is the angle opposite the edge $v_i v_j$.

Note that with our choice of basis functions, we have $\lambda_j = \varphi(v_j)$. Thus

$$\begin{aligned}\langle \Delta\varphi, f_i \rangle &= \sum_j \lambda_j \langle \nabla f_i, \nabla f_j \rangle \\ &= \frac{1}{2} \sum_{j \sim i} (\cot\alpha_j + \cot\beta_j)(\varphi(v_i) - \varphi(v_j)),\end{aligned}$$

where the angles are as indicated in Figure 6. This is the so-called *cotan formula* for the discrete Laplacian.

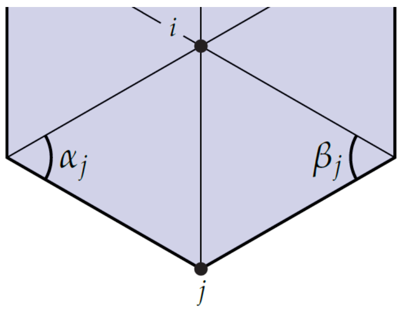

Figure 6: Angles in the cotan formula.

Discrete exterior calculus We now take another approach to discretising the Laplacian, by first looking at how to discretise differential forms. This essentially involves integrating forms over elements of the mesh.

For example, a 1-form α can be represented by its integrals over all edges in the mesh, ie. the numbers $\hat{\alpha}_e := \int_e \alpha$. (Note that we need to first fix an orientation for each edge.) More generally, a k-form that has been integrated over all (oriented) k-dimensional cells of a simplicial mesh is called a *discrete differential* k-*form*.

In this context, the *discrete exterior derivative* $\hat{d}$ is almost trivial to define, by Stokes' theorem. Since $\int_\Omega d\alpha = \int_{\partial\Omega} \alpha$, we see that summing a discrete differential k-form (corresponding to α) over all boundary components of a $(k+1)$-cell gives the integral of a $(k+1)$-form (namely, $d\alpha$) over that cell—in other words, a discrete differential $(k+1)$-form corresponding to $d\alpha$.

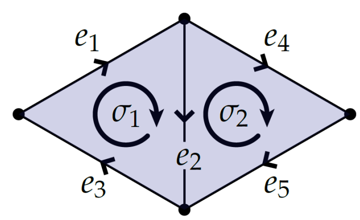

Figure 7: The discrete exterior derivative $\hat{d}$.

To illustrate, suppose we have a discrete differential 1-form $\hat{\alpha}$ on the simplicial mesh shown in Figure 7. Then the discrete differential 2-form $\hat{d}\hat{\alpha}$ is given by

$$
\begin{aligned}
(\hat{d}\hat{\alpha})_1 &= \hat{\alpha}_1 + \hat{\alpha}_2 + \hat{\alpha}_3, \\
(\hat{d}\hat{\alpha})_2 &= \hat{\alpha}_4 + \hat{\alpha}_5 - \hat{\alpha}_2.
\end{aligned}
$$

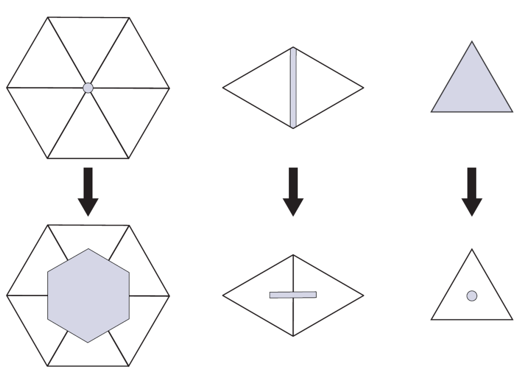

Figure 8: Elements of a dual mesh.

Note that in general, any set of numbers assigned to the oriented k-cells of a simplicial mesh can be considered as a discrete differential k-form, ie. not every discrete differential form has to come from a continuous one.

There is also a discrete analog of the Hodge star, known as the *diagonal Hodge star* $\hat{\star}$, which takes k-cells in the original (primal) mesh to $(n-k)$-cells in a *dual mesh* (Figure 8). Since we want the "density" of the dual form to be the same as that of the primal form, we scale the values of the form by the volume ratio:

$$\hat{*}\hat{\alpha}_i := \frac{|\sigma_i^*|}{|\sigma_i|}\hat{\alpha}_i,$$

where $|\sigma_i|, |\sigma_i^*|$ denote the volumes of the primal and dual cells respectively.

Returning to the Poisson equation, we will transfer the Laplace-Beltrami operator to discrete differential forms by writing $\Delta\varphi = *d*d\varphi = \rho$. This process is illustrated in Figure 9.

Starting with the 0-form φ, we get a discrete 0-form $\hat{\varphi}$, which assigns a number $\varphi_i := \varphi(v_i)$ to every vertex. Then $\hat{d}\hat{\varphi}$ assigns the number $(\hat{d}\hat{\varphi})_{ij} = \varphi_j - \varphi_i$ to every directed edge $v_i v_j$. Taking the diagonal Hodge star gives

$$(\hat{*}\hat{d}\hat{\varphi})_{ij} = \frac{|e_{ij}^*|}{|e_{ij}|}(\varphi_j - \varphi_i),$$

and the next exterior derivative sums along the boundary of the dual cell to yield

$$(\hat{d}\hat{*}\hat{d}\hat{\varphi})_i = \sum_{j\sim i}\frac{|e_{ij}^*|}{|e_{ij}|}(\varphi_j - \varphi_i).$$

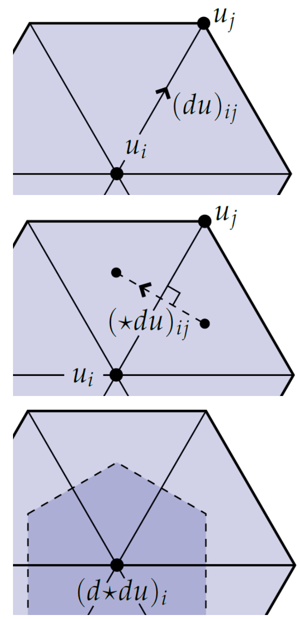

Figure 9: The discrete Laplacian. (Note that u in this figure is φ in the text.)

A final Hodge star will divide the above expression by the area of the dual cell $|C_i|$, but it is customary to move that to the RHS to yield

$$\sum_{j\sim i}\frac{|e_{ij}^*|}{|e_{ij}|}(\varphi_j - \varphi_i) = |C_i|\rho_i,$$

with $\rho_i := \rho(v_i)$. As before, this is a system of linear equations where the matrix is symmetric.

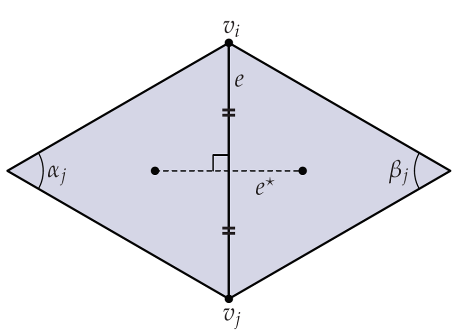

Figure 10: Choosing circumcentres as dual vertices.

The lengths $|e_{ij}^*|$ depend on the choice of dual mesh. One popular option is to choose the dual vertices at the circumcentres of the primal faces (Figure 10), even though this runs into issues with very obtuse triangles. In this case, the dual edges are the perpendicular bisectors of the primal edges, and a short calculation gives

$$\frac{|e_{ij}^*|}{|e_{ij}|} = \frac{1}{2}(\cot\alpha_j + \cot\beta_j).$$

Plugging this back into the previous equation, we recover the cotan formula for the Laplacian.

Remark: Laplacian linear systems We note that under some mild conditions, the linear systems that are involved in the discrete Poisson problem are not only symmetric, but *Laplacian matrices*—named for their connection to the Laplacian operator—with nonnegative entries on the diagonal and nonpositive entries off the diagonal, and with entries summing to 0 in every row. Moreover, these matrices are *sparse*, since most of its entries are 0.

Besides in geometric processing, Laplacian linear systems have applications in other areas such as graph drawing and learning functions on graphs. Thus it was a major theoretical breakthrough when in 2004, Spielman and Teng announced a nearly-linear time algorithm to solve sparse Laplacian systems. The field has steadily progressed since then, and now the best known algorithm can solve a Laplacian system faster than we can sort its nonzero entries!

Conclusion

Mathematics finds itself becoming increasingly relevant to other areas of scientific research, especially in this age of computation. Conversely, in the search for a theoretical framework underlying applied problems, new areas of mathematical research such as numerical analysis, numerical linear algebra, information theory, and many others have spawned and grown.

The field of discrete differential geometry is a beautiful example of this trend. Beyond the broad scope of its applications, perhaps equally important are the new perspectives it brings to a classical subject. It is at such cross-junctions of ideas that I feel that mathematics is always fresh, always young, and always alive.

References

K. Crane, M. Wardetzky, "A Glimpse into Discrete Differential Geometry." *Notices Amer. Math. Soc.*, **64**(10), 2011, 1153–1159.

K. Crane. *Discrete Differential Geometry: An Applied Introduction*, Course notes, Carnegie Mellon University, 18 Oct 2017. Retrieved from `www.cs.cmu.edu/~kmcrane/Projects/DDG/paper.pdf`.

J. Solomon, K. Crane, E. Vouga. "Laplace-Beltrami: The Swiss Army Knife of Geometry Processing." Symposium on Geometry Processing, Cardiff, 7 Jul 2014. Tutorial. Retrieved from `www.cs.cmu.edu/~kmcrane/Projects/Other/SwissArmyLaplacian.pdf`.