

Algebraic graph calculus

We describe a graph-theoretic analogue of vector calculus. The linear operators of vector calculus (gradient, divergence, laplacian) correspond to the matrices naturally associated to graphs (incidence matrix, adjacency matrix). This analogy is useful for formalizing the discretization of some problems in image and surface processing that are often defined in a continuous setting.

1 Reminder of vector calculus

Vector calculus deals with functions and vector fields defined in \mathbf{R}^3 .

1.1 Functions and vector fields

A *function* (or *scalar field*) is a map $u : \mathbf{R}^3 \rightarrow \mathbf{R}$. A *vector field* is a map $\mathbf{v} : \mathbf{R}^3 \rightarrow \mathbf{R}^3$. Vector fields are written in bold.

Let us fix some typical names for the coordinates. The coordinates of a point in \mathbf{R}^3 are written as (x, y, z) . If \mathbf{v} is a vector field, then $\mathbf{v} = (a, b, c)$ where a, b and c are three scalar fields called the components of \mathbf{v} . We denote the partial derivatives of a function using subindices, for example $a_y := \frac{\partial a}{\partial y}$.

1.2 Differential operators

The *gradient* of a function u is a vector field ∇u defined by

$$\nabla u = (u_x, u_y, u_z)$$

The *divergence* of a vector field $\mathbf{u} = (a, b, c)$ is a scalar field $\text{div}(\mathbf{u})$ defined by

$$\text{div}(\mathbf{u}) = a_x + b_y + c_z$$

The *curl* of a vector field $\mathbf{u} = (a, b, c)$ is another vector field $\text{curl}(\mathbf{u})$ defined by

$$\text{curl}(\mathbf{u}) = (c_y - b_z, a_z - c_x, b_x - a_y)$$

Finally, the *laplacian* of a scalar field u is the scalar field Δu defined by

$$\Delta u = u_{xx} + u_{yy} + u_{zz}.$$

Notice that, except for the curl, all these operations can be defined in \mathbf{R}^N . However, the curl is specific to three dimensions. There is a similar operator in two dimensions, which we call also the curl and computes a scalar field $\text{curl}(\mathbf{u})$ from a vector field $\mathbf{u} = (a, b) : \mathbf{R}^2 \rightarrow \mathbf{R}^2$

$$\text{curl}(\mathbf{u}) = b_x - a_y$$

Notice that it is the last component of the 3D curl.

The curl is also defined in dimension 7. Let $\mathbf{u} = (u^1, \dots, u^7)$ be a vector field in \mathbf{R}^7 , then

$$\text{curl}(\mathbf{u}) = \begin{pmatrix} u^2_4 - u^4_2 + u^3_7 - u^7_3 + u^5_6 - u^6_5 \\ u^3_5 - u^5_3 + u^4_1 - u^1_4 + u^6_7 - u^7_6 \\ u^4_6 - u^6_4 + u^5_2 - u^2_5 + u^7_1 - u^1_7 \\ u^5_7 - u^7_5 + u^6_3 - u^3_6 + u^1_2 - u^2_1 \\ u^6_1 - u^1_6 + u^7_4 - u^4_7 + u^2_3 - u^3_2 \\ u^7_2 - u^2_7 + u^1_5 - u^5_1 + u^3_4 - u^4_3 \\ u^1_3 - u^3_1 + u^2_6 - u^6_2 + u^4_5 - u^5_4 \end{pmatrix}$$

where a sub-index i denotes a partial derivative in the i -th dimension of \mathbf{R}^7 . And analogously we can define the 6-dimensional curl by taking the last component (resulting in a scalar field).

1.3 Differential identities and properties

The most important identity is $\Delta u = \text{div}(\text{grad}(u))$, that can be used also as the definition of Δ .

Other identities involving the curl are $\text{curl}(\nabla u) = 0$ and $\text{div}(\text{curl}(\mathbf{u})) = 0$.

The functions u such that $\nabla u = 0$ on \mathbf{R}^3 are the constants.

The vector fields \mathbf{v} such that $\text{curl}(\mathbf{v}) = 0$ are called *conservative*, *irrotational* or *integrable*. They are of the form $\mathbf{v} = \nabla u$ for some function u called the *potential* of \mathbf{v} .

The vector fields \mathbf{v} such that $\text{div}(\mathbf{v}) = 0$ are called *divergence-free*, *volume-preserving*, *solenoidal* or *incompressible*. They are of the form $\mathbf{v} = \text{curl}(\mathbf{u})$ for some vector field \mathbf{u} called the *vector potential* of \mathbf{v} .

The scalar fields u such that $\Delta u = 0$ are called *harmonic functions*.

The following identities are immediate applications of the product rule for derivatives:

$$\begin{aligned} \nabla(fg) &= f\nabla g + g\nabla f \\ \text{div}(f\mathbf{g}) &= f\text{div}(\mathbf{g}) + \mathbf{g} \cdot \nabla f \end{aligned}$$

1.4 Integral calculus

The divergence theorem:

$$\int_{\Omega} \text{div}(\mathbf{g}) = \int_{\partial\Omega} \mathbf{g} \cdot \mathbf{ds}$$

Combining the divergence theorem with the product rule we obtain the integration by parts formula.

$$\int_{\partial\Omega} f \mathbf{g} \cdot \mathbf{ds} = \int_{\Omega} f \operatorname{div}(\mathbf{g}) + \int_{\Omega} \mathbf{g} \cdot \nabla f$$

Thus, if at least one of the two functions vanishes on the boundary of Ω

$$0 = \int_{\Omega} f \operatorname{div}(\mathbf{g}) + \int_{\Omega} \mathbf{g} \cdot \nabla f$$

or, in another notation

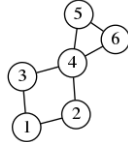
$$\langle f, \operatorname{div}(\mathbf{g}) \rangle = \langle -\nabla f, \mathbf{g} \rangle$$

thus that the operators div and $-\nabla$ are adjoint to each other. Integrating by parts twice we obtain that the operator Δ is self-adjoint.

2 Graphs and their matrices

A *graph* is $G = (V, E)$ where V is a set called the *vertices* of G , and E is a subset of $V \times V$ called the *edges* of G .

We assume always that the set V is finite, and its elements are numbered from 1 to n . Thus, the set E is also finite (the cardinal is at most n^2) and we assume that the elements of E are numbered from 1 to m .

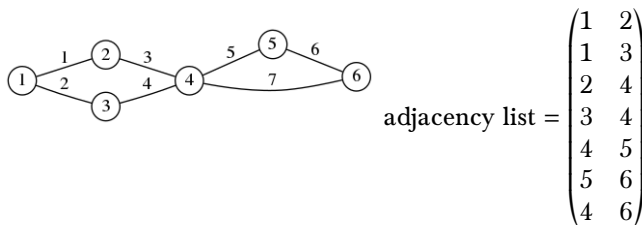


$$V = \{1, 2, 3, 4, 5, 6\}$$

$$E = \{\{1, 2\}, \{1, 3\}, \{2, 4\}, \{3, 4\}, \{4, 5\}, \{5, 6\}, \{4, 6\}\}$$

2.1 The adjacency list

Given a graph of n vertices and m edges, the adjacency list is a matrix of m rows and 2 columns that contains the pairs of vertices connected by each edge. The entries of this matrix are integers on the set $\{1, \dots, n\}$. Thus, if the k -th row is (i, j) , this means that edge k connects vertices i to j .



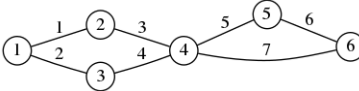
$$\text{adjacency list} = \begin{pmatrix} 1 & 2 \\ 1 & 3 \\ 2 & 4 \\ 3 & 4 \\ 4 & 5 \\ 5 & 6 \\ 4 & 6 \end{pmatrix}$$

The adjacency list is a very efficient representation for sparse graphs (where the number of edges is proportional to the number of vertices). However, it is not very interesting from the algebraic point of view. We will see in the following three other matrices that have a very rich algebraic interpretation.

2.2 The adjacency matrix A

Given a graph of n vertices and m edges, the adjacency matrix is a square matrix $A = a_{ij}$ of size $n \times n$. The entries of A are zeros and ones, with $a_{ij} = 1$ if there is an edge from i to j and $a_{ij} = 0$ otherwise.

$V \setminus V$	1	2	3	4	5	6
1	0	1	1	0	0	0
2	1	0	0	1	0	0
3	1	0	0	1	0	0
4	0	1	1	0	1	1
5	0	0	0	1	0	1
6	0	0	0	1	1	0



Notice that this matrix has somewhat less information than the adjacency list, because the ordering of the edges is lost. Thus, there is a unique way to compute the adjacency matrix from the list, but many $m!$ different ways to get the list from the matrix. We can choose an arbitrary canonical ordering of the edges (for example, in lexicographic order).

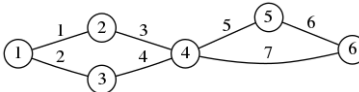
2.3 The Laplacian matrix L

Let A be the adjacency matrix of a graph G . If we sum the values of all the elements of the i -th row, we obtain the number of edges going out of vertex i (called the degree of the vertex). Let us put the vector with all the degrees in the diagonal of a matrix D ; in octave/matlab notation $D = \text{diag}(\text{sum}(A))$. The Laplacian matrix of G is defined as

$$L = A - \text{diag}(\text{sum}(A))$$

In the typical case where A is symmetric with 0 on the diagonal, the matrix L is the same as A with minus the degree of each vertex on the diagonal entries.

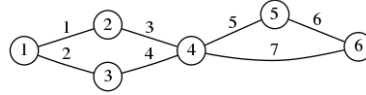
$V \setminus V$	1	2	3	4	5	6
1	-2	1	1	0	0	0
2	1	-2	0	1	0	0
3	1	0	-2	1	0	0
4	0	1	1	-4	1	1
5	0	0	0	1	-2	1
6	0	0	0	1	1	-2



2.4 The incidence matrix B

Given a graph of n vertices and m edges, the incidence matrix is a rectangular matrix $B = b_{ij}$ of m rows and n columns. The entries of B are zeros, ones and minus ones given by the edges of the graph: if the k -th edge goes from i to j , then, on the k -th row there are values -1 and 1 on positions i and j respectively; there are zeros everywhere else.

$E \setminus V$	1	2	3	4	5	6
1	-1	1	0	0	0	0
2	-1	0	1	0	0	0
3	0	-1	0	1	0	0
4	0	0	-1	1	0	0
5	0	0	0	-1	1	0
6	0	0	0	0	-1	1
7	0	0	0	-1	0	1



Notice that the incidence matrix contains the same information as the adjacency list (including the order of the edges).

There is an interesting relationship between the incidence matrix and the Laplacian matrix, that can be checked algebraically:

$$L = -B^T B$$

This identity is the discrete analogue of $\Delta = \text{div grad}$, as we will explain below.

2.5 The unsigned incidence matrix C

The incidence matrix B defined above is signed, on each row there are two non-zero entries whose values are -1 and 1 . Thus the sum of any row is zero. We can write the matrix B as $B = B_1 - B_0$, where the matrices B_0 and B_1 have only zeros and ones, with a single non-zero entry per row.

It will be useful later to consider the *unsigned incidence matrix* C , defined as $C = \frac{1}{2}(B_0 + B_1)$, or equivalently $C = \frac{1}{2}|B|$. The rows of the matrix C sum to one.

The following relations are immediate to verify

$$A = 2C^T C - B^T B/2$$

$$\text{deg} = 2C^T C + B^T B/2$$

where deg is an $n \times n$ diagonal matrix, whose values are the degrees of each vertex.

3 Vector calculus on graphs

Most of the constructions that we have described on the vector calculus reminder above have a direct correspondence in the case of graphs.

3.1 Analogies

The correspondence between vector calculus and graph theory is laid out in the following table. The main idea is that scalar fields correspond to functions defined on vertices, and vector fields correspond to functions defined on edges.

Vector calculus	Graph theory
Base space	Graph vertices V
Tangent space	Graph edges E
$u : \Omega \rightarrow \mathbf{R}$	$u : V \rightarrow \mathbf{R}$
$\mathbf{v} : \Omega \rightarrow \mathbf{R}^3$	$\mathbf{v} : E \rightarrow \mathbf{R}$
Laplacian operator Δ	Laplacian matrix $L \in \mathcal{M}_{n,n}(\mathbf{R})$
gradient operator ∇	incidence matrix $B \in \mathcal{M}_{m,n}(\mathbf{R})$
divergence operator div	matrix $-B^T \in \mathcal{M}_{n,m}(\mathbf{R})$
$\Delta = \text{div grad}$	$L = -B^T B$
scalar field u	$u \in \mathbf{R}^n$
vector field \mathbf{v}	$\mathbf{v} \in \mathbf{R}^m$
vector field ∇u	$Bu \in \mathbf{R}^m$
scalar field Δu	$Lu \in \mathbf{R}^n$
scalar field $\text{div}(\mathbf{v})$	$-B^T \mathbf{v} \in \mathbf{R}^n$
directional derivative $\nabla u(\mathbf{a}) \cdot (\mathbf{b} - \mathbf{a})$	$\nabla u(a, b)$
$\Omega \subseteq \mathbf{R}^3$	$\Omega \subseteq V$
$\partial\Omega \subseteq \mathbf{R}^3$	$\partial\Omega \subseteq E$, defined as $\partial\Omega = E \cap (\Omega \times \Omega^c)$
$\int_{\Omega} \text{div}(\mathbf{v}) = \int_{\partial\Omega} \mathbf{v} \cdot \mathbf{ds}$	$\sum_{a \in \Omega} \text{div}(\mathbf{v})(a) = \sum_{e \in \partial\Omega} \mathbf{v}(e)$
Elliptic PDE $\Delta u = f$	Linear system $Lu = f$
Parabolic PDE $u_t = \Delta u$	First-order Linear ODE System $u_t = Lu$
$\text{div}(D\nabla u)$, $D : \Omega \rightarrow \mathcal{M}_{3,3}(\mathbf{R})$	$-B^T DBu$, $D \in \mathcal{M}_{m,m}$
$g\Delta u$, $g : \Omega \rightarrow \mathbf{R}$	GLu , $G \in \mathcal{M}_{n,n}$
pointwise product uv	Hadamard product $f \odot g$
pointwise product $u\mathbf{v}$	Hadamard product $Cf \odot g$
$\nabla f g = f\nabla g + g\nabla f$	$B(f \circ g) = Cf \odot Bg + Cg \odot Bf$
(nothing)	unsigned incidence matrix $C \in \mathcal{M}_{m,n}(\mathbf{R})$

The curl operator cannot be defined on general graphs, but it can be defined on *planar* graphs, and it satisfies similar identities.

3.2 The graph Laplacian

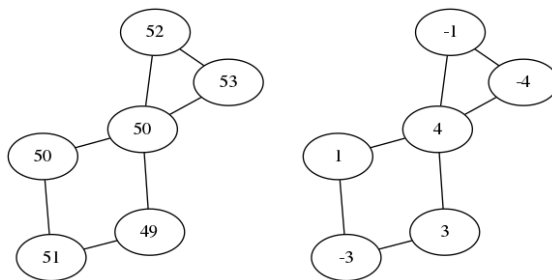
The simplest operator of vector calculus is the Laplacian, transforming scalar fields into scalar fields. It is the simplest because no vector fields are involved, only scalar fields.

Correspondingly, the simplest operator for graphs is also the Laplacian, transforming functions defined on vertices into functions defined on vertices. It is the simplest because no functions defined on edges are involved. Once we have chosen an ordering of the vertices, a scalar field is simply a vector $u \in \mathbf{R}^n$, and the Laplacian operator is defined by a square matrix of size $n \times n$.

Let $G = (V, E)$ be a graph and $u : V \rightarrow \mathbf{R}$ be a scalar field. The **Laplacian** of u is denoted by Δu and is defined as the scalar field $\Delta u : V \rightarrow \mathbf{R}$

$$\Delta u(a) := \sum_{(a,b) \in E} u(b) - u(a)$$

Notice that the sum is performed for a fixed vertex a , and b varies through all the neighbors of a in the graph.



The scalar field u

The scalar field Δu

Just from the definition, we can deduce several properties of the laplacian

1. The sum of all the values of Δu is always zero
2. If $u(a)$ is a local maximum, then $\Delta u(a) < 0$
3. If $u(a)$ is a local minimum, then $\Delta u(a) > 0$
4. If u is constant, then Δu is zero

If we fix an ordering of the vertices, then the scalar fields u and Δu are two vectors in \mathbf{R}^n , and the linear operator $u \mapsto \Delta u$ is given by the matrix $L = A - \text{diag}(\text{sum}(A))$. This follows directly by decomposing the definition of Δ into two sums:

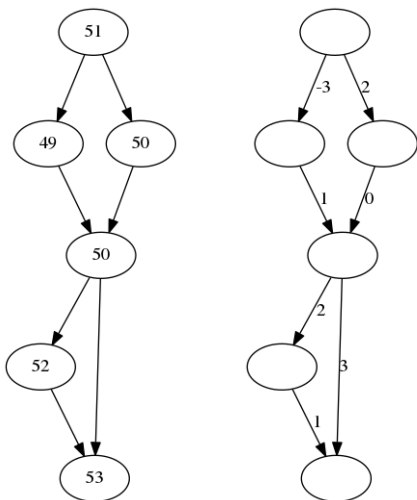
$$\Delta u(a) = \sum_{(a,b) \in E} u(b) - \sum_{(a,b) \in E} u(a) = -u(a)\text{degree}(a) + \sum_{(a,b) \in E} u(b)$$

Notice that the Laplacian has a nice interpretation. If we regard the values of u as a quantity distributed on the vertices of the graph, the condition $\Delta u = 0$ says that the quantity is distributed evenly, or in equilibrium: the amount of quantity at each vertex equals the average amount over its neighbours. In particular, if u is constant then $\Delta u = 0$.

Notice that the matrix L is always singular: a constant vector is an eigenvector of eigenvalue 0. If the graph has k connected components, then we have null vectors that are constant on each connected component, thus the matrix L has rank $n - k$.

3.3 Graph gradient and graph divergence

Recall that scalar fields are functions defined on vertices and vector fields are functions defined on edges. Thus, the gradient transforms a function defined on vertices into a function defined on edges. There is a very natural way of doing that: the value at each edge is obtained as the difference between the values at each side of the edge.



The scalar field u The vector field ∇u

More formally, let $G = (V, E)$ be a graph and $u : V \rightarrow \mathbf{R}$ be a scalar field. The **gradient** of u is the vector field $\nabla u : E \rightarrow \mathbf{R}$ defined by

$$\nabla u(a, b) := u(b) - u(a) \quad \text{for } (a, b) \in E$$

The matrix of this linear map is the incidence matrix B of the graph. Think of the gradient $\nabla u(a, b)$ as the directional derivative of u at point a in the direction of the vector from a to b .

Now let $\mathbf{v} : E \rightarrow \mathbf{R}$ be a vector field. The **divergence** of \mathbf{v} is the scalar field $\text{div}(\mathbf{v}) : V \rightarrow \mathbf{R}$ defined by:

$$\text{div}(\mathbf{v})(a) := \sum_{(a,b) \in E} \mathbf{v}(a, b) - \sum_{(b,a) \in E} \mathbf{v}(b, a) \quad \text{for } a \in V$$

The matrix of this linear map is minus the transposed incidence matrix of the graph $-B^T$.

Notice that the identity $\Delta = \text{div grad}$ is trivial from the definitions above, since both sides are exactly $\sum_{(a,b) \in E} u(b) - u(a)$. Thus, $L = -B^T B$.

3.4 Graph curl

We do not need curls for our application, but let us say some words about them.

These graph-theoretic analogues are easier to understand when we use differential geometry instead of plain vector calculus. In that case, the discrete analogue of k -forms are functions defined over the k -cliques of the graph. Then the exterior derivative is readily built for all values of k , and it contains the gradient, curl and divergence as particular cases. The particularity of 3-dimensional manifolds comes from the fact that in that case 1-forms and 2-forms have the same dimension and can both be interpreted as “vector fields”, thus the curl operator is defined from the exterior derivative $d : \Omega^1 \rightarrow \Omega^2$. In the case of graphs, we cannot in general identify functions defined on edges to functions defined on triangles,

except in one particular case: when the graph is a triangulation. In that case, there is a construction that allows to define the curl of a vector field as a vector field, by traversing the two triangles at each side of an edge. The identity $\text{curl grad} = 0$ is then the sum of 6 values that cancel pairwise, and so on. See the beautiful papers of Oliver Knill for a comprehensive coverage of this.

3.5 Graph subsets and their boundaries

It is often necessary to deal with subset of graphs (for example, when we want to interpolate a function which is known only over some vertices). In order to do algebra with them, we model subsets as diagonal operators that contain the indicator function of the subset as the diagonal entries. This model is used for subsets of vertices and subsets of edges.

Notations: Let $X = \{1, \dots, n\}$ (or any finite ordered set) and $Y \subseteq X$. Let a be a vector of length n and A a matrix of size $n \times n$. We use the following, somewhat ambiguous, abuses of notation:

$\text{diag}(A) \in \mathbf{R}^n$: the vector with the elements on the diagonal of A
 $\text{diag}(a) \in \mathbf{R}^{n \times n}$: the diagonal matrix whose diagonal is a . $\mathbf{1}_Y \in \mathbf{R}^n$: the indicator vector of the subset Y $Y = \text{diag}(\mathbf{1}_Y) \in \mathbf{R}^{n \times n}$: the diagonal operator of Y

This last notation is very convenient in image processing, because it represents point-wise multiplication by a binary image as a linear operator (with the same name as the binary image). The diag operator has the same semantics as that of octave/matlab.

Let $G = (V, E)$ be a graph with n vertices and m edges, and let $\Omega \subseteq V$. To avoid introducing new letters, we denote also by $\Omega = \omega_{ij}$ the $n \times n$ matrix that contains the indicator function of this set in its diagonal: $w_{ii} = 1$ if $i \in \Omega$ and $w_{jj} = 0$ otherwise. Notice that the matrix $I - \Omega$ corresponds to the complementary set Ω^c .

We define the **boundary** of a subset of vertices $\Omega \subseteq V$ as the subset of edges $\partial\Omega \subseteq E$ that go from Ω to Ω^c . Notice that $\partial\Omega = E \cap (\Omega \times \Omega^c)$ in set notation. Since $\partial\Omega$ is a subset of edges, it corresponds to a diagonal matrix, also named $\partial\Omega$, of size $m \times m$. In matrix notation we have

$$\partial\Omega = \text{diag}(B\text{diag}(\Omega))$$

where B is the incidence matrix of the graph. We can also write $\mathbf{1}_{\partial\Omega} = B\mathbf{1}_\Omega$.

3.6 Equations on graphs

Now that we have described the differential and boundary operators operator in matrix form, it is immediate to write the discrete analogues of several linear PDE. This is very beautiful because the analytic properties of the corresponding PDE are recovered by elementary linear algebra.

3.6.1. Laplace equation on the whole graph:

$$Lu = 0$$

If the graph is connected, the matrix L has rank $n - 1$ thus its kernel is one-dimensional, corresponding to the constant solutions $u = c$.

3.6.2. Poisson equation on the whole graph, with data $f : V \rightarrow \mathbf{R}$:

$$Lu = f$$

has a unique solution unless f is constant.

3.6.3. Laplace equation on a subset $\Omega \subseteq V$, with Dirichlet boundary conditions $f : \Omega^c \rightarrow \mathbf{R}$:

$$\Omega Lu + (I - \Omega)(u - f) = 0$$

Notice that this is written as an $n \times n$ linear system, but it has a diagonal part corresponding to the values of u outside of Ω . Notice also that the values of f at the vertices that have no neighbors in Ω only appear in the diagonal part. The values of f inside Ω do not appear at all (are cancelled out).

3.6.4. Laplace equation on a subset $\Omega \subseteq V$, with Neumann boundary conditions $g : \partial\Omega \rightarrow \mathbf{R}$:

$$\Omega Lu + (\partial\Omega)(\nabla u - g) = 0$$

Or equivalently, by developing the boundary and gradient operators,

$$[\Omega L + \text{diag}(B \text{diag}(\Omega))B] u = \text{diag}(B \text{diag}(\Omega))g$$

or, in an alternative notation

$$(\text{diag}(\mathbf{1}_\Omega)L + \text{diag}(B\mathbf{1}_\Omega))B u = \text{diag}(B\mathbf{1}_\Omega)g$$

3.6.5. Heat equation on the whole graph with initial condition $u_0 : V \rightarrow \mathbf{R}$:

$$\begin{cases} u_t & = Lu \\ u(0) & = u_0 \end{cases}$$

This is a system of n first-order linear differential equations with constant coefficients. It has a closed-form solution using the matrix exponential $u = e^{tL}u_0$.

3.6.6. Heat equation with source term $f : V \rightarrow \mathbf{R}$ and initial condition $u_0 : V \rightarrow \mathbf{R}$

$$\begin{cases} u_t & = Lu + f \\ u(0) & = u_0 \end{cases}$$

It has likewise a closed-form solution $u = e^{tL}(u_0 - L^{-1}f) - L^{-1}f$. Notice that $L^{-1}f$ only makes sense when f is not a constant vector.

3.6.7. Other combinations are possible, and easy to deduce from the simpler cases: Poisson and Heat equation on subsets with various boundary conditions, etc.

3.7 Riemannian graph geometry

The **isotropic** case of “anisotropic” diffusion in image processing is modelled by terms of the form $g\Delta u$, where g is a positive-real valued function on Ω . In the case of graphs, the function g corresponds to a scalar field $g : V \rightarrow \mathbf{R}$, which we associate to a diagonal $n \times n$ matrix \tilde{g} with the values of g . Then these terms become $\tilde{g}Lu$ in the corresponding discrete model.

Truly **anisotropic** diffusion comes from terms of the form $\text{div}(D\nabla u)$, where the diffusivity D is a field of positive-definite symmetric matrices defined over Ω . In the case of graphs, we have a matrix \tilde{D} , which is also diagonal, but now of size $m \times m$. Then these terms become $\text{div}(D\nabla u)$ in the discrete model. Or, in matrix form, $B^T DBu$.

3.8 Algebraic graph integral calculus

Integral calculus can be generalized readily to graphs. Integrals are replaced by sums over a finite domain, and the various identities of integral calculus (e.g., the divergence theorem) become immediate matrix identities.

Let $G = (V, E)$ be a graph with $V = \{1, \dots, n\}$ and $E = \{1, \dots, m\}$

Let $\Omega \subseteq V$ and let $f : V \rightarrow \mathbf{R}$ be a scalar field. The **integral** of f over Ω is defined as

$$\int_{\Omega} f = \sum_{p \in \Omega} f(p)$$

in matrix notation we have $\int_{\Omega} f := \text{sum}(\Omega f)$. Notice that here f is a vector of length n , Ω is an $n \times n$ matrix, and we are computing the sum of all the components of the vector Ωf to obtain a single number. Notice that f must be defined everywhere, but only the values inside Ω are used; thus, we could have defined f only inside Ω .

An **interface** inside a graph is defined as a set of edges $S \subseteq E$. Given a vector field $\mathbf{v} : E \rightarrow \mathbf{R}$ we define the **flow** of \mathbf{v} through S

$$\int_S \mathbf{v} \cdot \mathbf{ds} := \sum_{e \in S} \mathbf{v}(e)$$

or, in matrix notation, $\int_S \mathbf{v} \cdot \mathbf{ds} = \text{sum}(\tilde{S}\mathbf{v})$ where \tilde{S} is the diagonal matrix containing the indicator function of S . An interesting particular case happens when S is the boundary of some region Ω . We have seen above that the matrix \tilde{S} is then equal to $\text{diag}(B\text{diag}(\Omega))$. This observation leads to the **graph divergence theorem** that says that

$$\int_{\partial\Omega} \mathbf{v} \cdot \mathbf{ds} = \int_{\Omega} \text{div}(\mathbf{v})$$

or, in matrix notation,

$$\mathbf{1}_{\Omega} \cdot (B^T \mathbf{v}) = (B\mathbf{1}_{\Omega}) \cdot \mathbf{v}$$

which is exactly the same thing, written differently.