Graphs and Networks Lecture 13

Physical Metaphors for Graphs

Daniel A. Spielman October 9, 2013

13.1 Overview

We will examine physical metaphors for graphs. We begin with a spring model and then discuss resistor networks.

13.2 Review: Interpolation on a graph

Say we have a graph $G = (V, E)$, a special set of vertices W, and some labels $l : W \to \mathbb{R}$ whose values we only know at W. We would like guess the labels at the remaining vertices $V - W$. We use the vector $\boldsymbol{x}$ to denote the labels that we guess. For $u \in W$ we require $\boldsymbol{x}(u) = l(u)$, and for $u \in V - W$ we require that

$$\boldsymbol{x}(u) = \frac{1}{d(u)} \sum_{(u,w)\in E} \boldsymbol{x}(w). \tag{13.1}$$

In the previous lecture, we saw that this problem has a solution. Today, we will continue to study this inference problem using two different physical metaphors.

To begin, imagine that every edge is a rubber band (or spring), and that every vertex is a little ring to which all its edge rubber bands are connected. For each vertex $v \in W$, we nail the corresponding ring down onto the real line at $l(v)$. We could let the other rings settle into position, and guess that $\boldsymbol{x}(u)$ is the location of ring u. Hey, why not?

13.3 Hooke's Law

We can use Hooke's law to figure out where all the rings in $V - W$ should wind up. Assume that each rubber band is an ideal spring with spring constant 1. Actually, if you have weights on the edges, you could make these the spring constants. If a rubber band connects vertices u and v, then Hooke's law tells us that the force it exerts at node u is in the direction of v and is proportional to the distance between u and v. Let $\boldsymbol{x}(u)$ be the position of each vertex u. Then the force the rubber band between u and v exerts on u is That is,

$$\boldsymbol{x}(v) - \boldsymbol{x}(u).$$

In a stable configuration, all of the vertices that have not been nailed down must experience a zero net force. That is

$$\sum_{v:(u,v)\in E} \boldsymbol{x}(v) - \boldsymbol{x}(u) = 0 \quad \Longrightarrow \quad \sum_{v:(u,v)\in E} \boldsymbol{x}(v) = d(u)\boldsymbol{x}(u). \quad \Longrightarrow \quad \frac{1}{d(u)} \sum_{v:(u,v)\in E} \boldsymbol{x}(v) = \boldsymbol{x}(u).$$

So, each vertex that is not nailed down is the average of its neighbors, same as the requirements given by equations (13.1). Consequently, our spring-based inference is exactly equivalent to the inference problem we defined last time.

In the weighted case, we would have for each $u \in V - W$

$$\frac{1}{d(u)} \sum_{v:(u,v)\in E} w_{u,v}\boldsymbol{x}(v) = \boldsymbol{x}(u), \tag{13.2}$$

where here we define

$$d(u) = \sum_{v:(u,v)\in E} w_{u,v}$$

to be the weighted degree of node u.

A function $\boldsymbol{x}$ that satisfies these equations for each vertex $u \in V - W$ is said to be *harmonic* on $V - W$.

Recall, last time we showed that a solution to these equations exists when edges have equal weight, by setting up a random walk on the graph. We can use the same approach to show that a solution exists when we have arbitrary positive edge weights, by using a random walk which transitions from its current vertex position u to each neighbor v with probability $w_{u,v}/d(u)$.

13.4 Energy

Physics also tells us that the nodes will settle into a position that locally minimizes the potential energy. The potential energy of an ideal linear spring with constant w when stretched to length l is

$$\frac{1}{2}wl^2.$$

So, the potential energy in a configuration $\boldsymbol{x}$ is given by

$$\mathcal{E}(\boldsymbol{x}) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{(u,v)\in E} w_{u,v}(\boldsymbol{x}(u) - \boldsymbol{x}(v))^2. \tag{13.3}$$

Note that the energy always has to be at least 0. For any $\boldsymbol{x}$ that locally minimizes the energy, the partial derivative of the energy with respect to each variable must be zero. In this case, the variables are $\boldsymbol{x}(u)$ for $u \in V - W$. The partial derivative with respect to $\boldsymbol{x}(u)$ is

$$\frac{1}{2} \sum_{v:(u,v)\in E} w_{u,v} 2(\boldsymbol{x}(u) - \boldsymbol{x}(v)) = \sum_{v:(u,v)\in E} w_{u,v}(\boldsymbol{x}(u) - \boldsymbol{x}(v)).$$

Setting this to zero gives (13.2). Thus, a local minimum of the energy is a solution to the harmonic equations on $V - W$, and conversely, if we have a solution to (13.2), then this gives a local minimum of the potential energy.

We know that equations (13.2) have at least one solution. In a few moments, we will see the solution is in fact unique, provided the graph is connected and W is not empty.

Theorem 13.4.1. *If G is a connected graph with all edges weights positive and $|W| \geq 1$, then the minimizer of* (13.3) *is unique.*

To prove this theorem, we will first show the following lemma:

Lemma 13.4.2. *Consider any graph $G = (V, E)$, $W \subseteq V$, and labels $l : W \to \mathbb{R}$. Suppose $\boldsymbol{x}$ and $\boldsymbol{y}$ are both vectors in $\mathbb{R}^V$ that agree with l on W, and that $\boldsymbol{x}$ is a local minimizer of* (13.3) *over this set. If $\mathcal{E}(\boldsymbol{y}) \leq \mathcal{E}(\boldsymbol{x})$, then for all $(u, v) \in E$*

$$\boldsymbol{x}(u) - \boldsymbol{x}(v) = \boldsymbol{y}(u) - \boldsymbol{y}(v)$$

Proof. For each $(u, v) \in E$, let us define $a_{uv} = \boldsymbol{x}(u) - \boldsymbol{x}(v)$ and $b_{uv} = \boldsymbol{y}(u) - \boldsymbol{y}(v)$. For $0 \leq \epsilon \leq 1$ we also define

$$\boldsymbol{x}_\epsilon \stackrel{\text{def}}{=} (1 - \epsilon)\boldsymbol{x} + \epsilon \boldsymbol{y}.$$

The point $\boldsymbol{x}_\epsilon$ still takes the correct values on nodes in W. If we consider small enough $\epsilon > 0$, then since $\boldsymbol{x}$ is a local minimum, we must have

$$\mathcal{E}(\boldsymbol{x}) \leq \mathcal{E}(\boldsymbol{x}_\epsilon). \tag{13.4}$$

Now, suppose for a contradiction that there exists an edge $(s, t) \in E$ s.t. $a_{st} \neq b_{st}$. Observe the following identity, which holds for all $a, b \in \mathbb{R}$ and $0 \leq \epsilon \leq 1$.

$$(1 - \epsilon)a^2 + \epsilon b^2 - ((1 - \epsilon)a + \epsilon b)^2 = \epsilon(1 - \epsilon)(a - b)^2.$$

Using this identity, we see that

$$\mathcal{E}(x) - \mathcal{E}(\boldsymbol{x}_\epsilon) \geq (1 - \epsilon)\mathcal{E}(x) + \epsilon \mathcal{E}(y) - \mathcal{E}(\boldsymbol{x}_\epsilon) = \sum_{(u,v)\in E} \epsilon(1 - \epsilon)(a_{uv} - b_{uv})^2 \geq \epsilon(1 - \epsilon)(a_{st} - b_{st})^2$$

But if $\epsilon > 0$, this means $\mathcal{E}(\boldsymbol{x}) - \mathcal{E}(\boldsymbol{x}_\epsilon) > 0$, and hence (13.4) is false for all small $\epsilon > 0$. This contradicts our starting assumption that $\boldsymbol{x}$ is a local minimum. Thus $a_{st} = b_{st}$ for all $(u, v) \in E$. □

Proof of theorem 13.4.1. We know a local minimizer of (13.3) exists. Denote this by $\boldsymbol{x}$. If $\mathcal{E}(\boldsymbol{y}) \leq \mathcal{E}(\boldsymbol{x})$ for some $\boldsymbol{y}$, then lemma 13.4 tells us $\boldsymbol{x}$ and $\boldsymbol{y}$ have the same difference across all edges. Since W is non-empty, there exists a vertex $u \in W$ where $\boldsymbol{x}$ and $\boldsymbol{y}$ take the same value. We can show that $\boldsymbol{x}$ and $\boldsymbol{y}$ agree on all vertices within distance 1 of w, then on all vertices within distance 2, etc. by induction. Since the graph is connected, we eventually get that $\boldsymbol{y} = \boldsymbol{x}$. Thus all other points have energy strictly greater than $\mathcal{E}(\boldsymbol{x})$. Since this is true for all local minimizers, the local minimizer must be unique, and is also a global minimizer.

□

13.5 Resistor Networks

Given a graph, we can treat each edge as a resistor. If the graph is unweighted, we will assume that the resistor has resistance 1. If an edge e has weight $w(e)$, we will give the corresponding resistor resistance $r(e) = 1/w(e)$. The reason is that when the weight of an edge is very small, the edge is barely there, so it should correspond to very high resistance. Having no edge corresponds to having a resistor of infinite resistance.

The first equation I recall is

$$V = IR,$$

which says that the potential drop across a resistor is equal to the current flowing over the resistor times the resistance. To apply this in a graph, we will define for each edge (a, b) the current flowing *from a to b* to be $\boldsymbol{i}(a, b)$. As this is a directed quantity, we define

$$\boldsymbol{i}(b, a) = -\boldsymbol{i}(a, b).$$

I now let $\boldsymbol{v} \in \mathbb{R}^V$ be the vector of potentials at vertices. Given these potentials (voltages), we can figure out how much current flows on each edge by the formula:

$$\boldsymbol{i}(a, b) = \frac{1}{r_{a,b}} \left(\boldsymbol{v}(a) - \boldsymbol{v}(b)\right) = w_{a,b} \left(\boldsymbol{v}(a) - \boldsymbol{v}(b)\right).$$

I would now like to write this equation in matrix form. The one complication is that each edge comes up twice in $\boldsymbol{i}$. So, to treat $\boldsymbol{i}$ as a vector I will have each edge show up exactly once as (a, b) when $a < b$. I now define the *signed edge-vertex adjacency matrix* of the graph $\boldsymbol{U}$ to be the matrix with rows indexed by edges, columns indexed by vertices, such that

$$\boldsymbol{U}((a, b), c) = \begin{cases} 1 & \text{if } a = c \\ -1 & \text{if } b = c \\ 0 & \text{otherwise.} \end{cases}$$

Define $\boldsymbol{W}$ to be the diagonal matrix with rows and columns indexed by edges and the weights of edges on the diagonals. We then have

$$\boldsymbol{i} = \boldsymbol{W}\,\boldsymbol{U}\,\boldsymbol{v}.$$

Also recall that resistor networks cannot hold current. So, all the flow entering a vertex a from edges in the graph must exit a to an external source. Let $\boldsymbol{i}_{ext} \in \mathbb{R}^V$ denote the external currents, where $\boldsymbol{i}_{ext}(a)$ is the amount of current entering the graph through node a. We then have

$$\boldsymbol{i}_{ext}(a) = \sum_{b:(a,b)\in E} \boldsymbol{i}(a, b).$$

In matrix form, this becomes

$$\boldsymbol{i}_{ext} = \boldsymbol{U}^T \boldsymbol{i} = \boldsymbol{U}^T \boldsymbol{W}\,\boldsymbol{U}\,\boldsymbol{v}. \tag{13.5}$$

The matrix

$$\boldsymbol{L} \stackrel{\text{def}}{=} \boldsymbol{U}^T \boldsymbol{W}\,\boldsymbol{U}$$

will play a very important role in the study of resistor networks. It is called the Laplacian matrix of the graph.

To better understand the Laplacian matrix, let's compute one of its rows. We will do this using the equations we already have:

$$\begin{aligned}\boldsymbol{i}_{ext}(a) &= \sum_{b:(a,b)\in E} \boldsymbol{i}(a,b) \\ &= \sum_{b:(a,b)\in E} \frac{1}{r_{ab}}(\boldsymbol{v}(a) - \boldsymbol{v}(b)) \\ &= \sum_{b:(a,b)\in E} w_{ab}(\boldsymbol{v}(a) - \boldsymbol{v}(b)) \\ &= d(a)\boldsymbol{v}(a) - \sum_{b:(a,b)\in E} w_{ab}\boldsymbol{v}(b).\end{aligned}$$

This gives us the following expression of the entries of $\boldsymbol{L}$:

$$\boldsymbol{L}(a,b) = \begin{cases} d(a) & \text{if } a = b \\ -w_{a,b} & \text{if } (a,b) \in E \\ 0 & \text{otherwise.} \end{cases}$$

In matrix form, we see that $\boldsymbol{L}$ may be expressed as

$$\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{A},$$

where $\boldsymbol{D}$ is the diagonal matrix of weighted degrees and $\boldsymbol{A}$ is the weighted adjacency matrix.

13.6 Energy dissipation

Recall (from physics) that the energy dissipated in a resistor network with currents $\boldsymbol{i}$ is

$$\mathcal{E}(\boldsymbol{i}) \stackrel{\text{def}}{=} \frac{1}{2}\sum_{(a,b)\in E} \boldsymbol{i}(a,b)^2 r_{a,b} = \frac{1}{2}\sum_{(a,b)\in E} \frac{1}{r_{a,b}}(\boldsymbol{v}(a) - \boldsymbol{v}(b))^2 = \frac{1}{2}\sum_{(a,b)\in E} w_{a,b}(\boldsymbol{v}(a) - \boldsymbol{v}(b))^2.$$

This expression should look familiar. Remember for later that it can never be negative.

Let's see that we can express this in terms of the Laplacian as well. Recall that $\boldsymbol{U}\boldsymbol{v}$ is a vector that gives the potential drop accross every edge. So,

$$(\boldsymbol{U}\boldsymbol{v})^T(\boldsymbol{U}\boldsymbol{v})$$

is the sum of the squares of the potential drops accross all the edges. This is almost the expression that we want. We just need to get in the weights. We do this in one of the following ways:

$$\mathcal{E}(\boldsymbol{i}) = (\boldsymbol{W}^{1/2}\boldsymbol{U}\boldsymbol{v})^T(\boldsymbol{W}^{1/2}\boldsymbol{U}\boldsymbol{v}) = (\boldsymbol{U}\boldsymbol{v})^T\boldsymbol{W}(\boldsymbol{U}\boldsymbol{v}) = \boldsymbol{v}^T\boldsymbol{U}\boldsymbol{W}\boldsymbol{U}\boldsymbol{v} = \boldsymbol{v}^T\boldsymbol{L}\boldsymbol{v}.$$

Let me mention some useful spectral properties of the Laplacian.

Theorem 13.6.1. *The Laplacian matrix of a graph is a positive semi-definite graph. If the graph is connected then the nullspace of its Laplacian is spanned by the constant vector.*

Proof. If λ is an eigenvalue $\boldsymbol{L}$ with eigenvector $\boldsymbol{v}$, the for the current $\boldsymbol{i} = \boldsymbol{W}\boldsymbol{U}\boldsymbol{v}$, we have

$$0 \leq \mathcal{E}(\boldsymbol{i}) = \boldsymbol{v}^T \boldsymbol{L}\boldsymbol{v} = \lambda \boldsymbol{v}^T \boldsymbol{v} = \lambda \|\boldsymbol{v}\|^2 .$$

So, $\lambda \geq 0$. It is clear from all of these definitions that $\boldsymbol{L}\mathbf{1} = \mathbf{0}$. The proof that $\mathbf{1}$ actuall spans the nullspace is completely analogous to the proof that the eigenvalue of a walk matrix has multiplicity 1. □

13.7 Fixing Potentials

We usually don't think of injecting current into a circuit. Rather, we usually attach nodes of a circuit to the terminals of a battery, which induce fixed potentials. Typically, we just attach two vertices to the terminals. Let's call them s and t, and assume that s has been attached to a terminal of potential 1 and that t has been attached to a terminal of potential 0. That is, $\boldsymbol{v}(s) = 1$ and $\boldsymbol{v}(t) = 0$. Then, the potentials all the remaining vertices can be determined by observing that for $x \notin \{s, t\}$, $\boldsymbol{i}_{ext}(x) = 0$ and applying (13.2) to show

$$d(x)\boldsymbol{v}(x) = \sum_{y:(x,y)\in E} \boldsymbol{v}(y) w(x, y). \tag{13.6}$$

Thus, we get another system of linear equations, with boundary conditions fixing the individual variables at s and t. We could also use linear algebra to prove that these equations have a solution. I'll give you a break and use probability instead. We will first show that the equations have a solution, and then prove that it is unique.

13.8 Solving for voltages

We now observe that we can always solve equation (13.5) for $\boldsymbol{v}$, provided that

$$\mathbf{1}^T \boldsymbol{i}_{ext} = 0 \tag{13.7}$$

and the underlying graph is connected. Typically, equations like

$$\boldsymbol{L}\boldsymbol{v} = \boldsymbol{i}_{ext} \tag{13.8}$$

are solvable, except when the matrix is degenerate. Unfortunately, $\boldsymbol{L}$ does have determinant zero, and so is not necessarily invertible.

But, we know a lot about $\boldsymbol{L}$: its nullspace is spanned by $\mathbf{1}$. Equation (13.7) tells us that $\boldsymbol{i}_{ext}$ is orthogonal to the nullspace of $\boldsymbol{L}$, and so, because $\boldsymbol{L}$ is symmetric, $\boldsymbol{i}_{ext}$ is in the range of $\boldsymbol{L}$. So, we can solve (13.5) by just inverting $\boldsymbol{L}$ on its range–which is the space orthogonal to its nullspace. This will provide us with a solution $\boldsymbol{v}$ that is also orthogonal to the nullspace.

Here's a more explicit way of constructing the solution. Let $0 = \gamma_1 < \gamma_2 \leq \cdots \leq \gamma_n$ be the eigenvalues of $\boldsymbol{L}$ and let $\mathbf{1} = \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n$ be a corresponding orthonormal basis of eigenvectors. We have

$$\boldsymbol{i}_{ext} = \sum_{i=1}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right).$$

As

$$\left(\boldsymbol{u}_1^T \boldsymbol{i}_{ext}\right) = 0,$$

we can simplify this to

$$\boldsymbol{i}_{ext} = \sum_{i=2}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right).$$

So, we can just set

$$\boldsymbol{v} = \sum_{i=2}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right)/\gamma_i.$$

This will be a solution to (13.8) because

$$\begin{aligned}
\boldsymbol{L}\boldsymbol{v} &= \boldsymbol{L}\sum_{i=2}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right)/\gamma_i \\
&= \sum_{i=2}^{n} \gamma_i \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right)/\gamma_i \\
&= \sum_{i=2}^{n} \boldsymbol{u}_i \left(\boldsymbol{u}_i^T \boldsymbol{i}_{ext}\right) \\
&= \boldsymbol{i}_{ext}.
\end{aligned}$$

The bottom line is that the conventional wisdom that "one can only solve a system of linear equations if the matrix is invertible" is just too pessimistic. The more optimistic approach that we have taken here is to solve the system by multiplying by the pseudo-inverse instead of the inverse. Since you may find it useful later in life, let me say that the pseudo-inverse of a symmetric matrix is just the inverse on the range of the matrix. For a matrix with eigenvalues $\gamma_1, \ldots, \gamma_n$ and corresponding eigenvectors $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n$, it is given by

$$\boldsymbol{L}^{+} \stackrel{\text{def}}{=} \sum_{i:\gamma_i \neq 0} \gamma_i^{-1} \boldsymbol{u}_i \boldsymbol{u}_i^T.$$