Graphs and Networks Lecture 16

Dynamic and Nonlinear Networks

Daniel A. Spielman October 22, 2013

16.1 Disclaimer

These notes are not necessarily an accurate representation of what happened in class. They are a combination of what I intended to say with what I think I said. They have not been carefully edited.

16.2 Overview

In this lecture we will consider two generalizations of resistor networks: resistor networks with non-linear resistors and networks whose resistances change over time. While they were introduced over 50 years ago, non-linear resistor networks seem to have been recently rediscovered in the Machine Learning community. We will discuss how they can be used to improve the technique we learned in Lecture 13 for semi-supervised learning.

The material on time-varying networks that I will present comes from Cameron Musco's senior thesis from 2012.

16.3 Non-Linear Networks

A non-linear resistor network, as defined by Duffin [Duf47], is a like an ordinary resistor network but the resistances depend on the potential differences across them. In fact, it might be easier not to talk about resistances, and just say that the amount of flow across an edge increases as the potential difference across the edge does. For every resistor e, there is a function

$$\phi_e(v)$$

that gives the flow over resistor e when there is a potential difference of v between its terminals.

We will restrict our attention to functions ϕ that are

a. continuous,

b. monotone increasing,

c. symmetric, by which I mean $\phi_e(-v) = -\phi_e(v)$.

Note that condition c implies that $\phi_e(0) = 0$. For an ordinary resistor of resistance r, we have

$$\phi_e(v) = v/r.$$

However, we can and will consider more interesting functions.

If the graph is connected and we fix the voltages at some of the vertices, then there exists a setting of voltages at the other vertices that results in a flow satisfying flow-in equals flow-out at all non-boundary vertices. Moreover, this flow is unique.

We will prove this in the next section through the use of a generalization of energy dissipation.

16.4 Energy

We define the energy dissipation of an edge that has a potential difference of v to be

$$\Phi_e(v) \stackrel{\text{def}}{=} \int_0^v \phi_e(t)dt.$$

We will show that the setting of the voltages that minimizes the total energy provides the flow I claimed exists.

In the case of linear resistors, where $\phi_e(v) = v/r$,

$$\Phi_e(v) = \frac{1}{2}\frac{v^2}{r},$$

which is exactly the energy function we introduced in Lecture 13.

The conditions on ϕ_e imply that

d. Φ_e is strictly convex[1],

e. $\Phi_e(0) = 0$, and

f. $\Phi_e(-x) = \Phi_e(x)$.

We remark that a function that is strictly convex has a unique minimum, and that a sum of strictly convex functions is strictly convex.

Theorem 16.4.1. *Let $G = (V, E)$ be a non-linear resistor network with functions f_e satisfying conditions a, b and c for every $e \in E$. For every set $S \subseteq V$ and fixed voltages w_a for $a \in S$, there exists a setting of voltages v_a for $a \notin S$ that result in a flow of current that satisfies the flow-in equals flow-out conditions at every $a \notin S$. Moreover, these voltages are unique.*

[1] That is, for all $x \neq y$ and all $0 < \lambda < 1$, $\Phi_e(\lambda x + (1-\lambda)y) \leq \lambda\Phi_e(x) + (1-\lambda)\Phi_e(y)$.

Proof. For a vector of voltages v, define

$$\Phi(v) = \sum_{(a,b)\in E} \Phi_{(a,b)}(v_a - v_b).$$

As each of the functions $\Phi_{(a,b)}$ are strictly convex, Φ is as well. So, Φ has a minimum subject to the fixed voltages. At this minimum point, we know that for every $a \notin S$

$$\begin{aligned} 0 &= \frac{\partial \Phi(v)}{\partial v_a} \\ &= \sum_{b:(a,b)\in E} \frac{\partial \Phi_{(a,b)}(v_a - v_b)}{\partial v_a} \\ &= \sum_{b:(a,b)\in E} \phi_{(a,b)}(v_a - v_b). \end{aligned}$$

We may now set

$$f_{(a,b)} = \phi_{(a,b)}(v_a - v_b).$$

This is a valid flow because for every vertex $a \notin S$ the sum of the flows out of v_a, taken with appropriate signs, is zero.

Conversely, for any setting of voltages that results in a flow that has no loss or gain at any $a \notin S$, we can reverse the above equalities to show that the partial derivatives of $\Phi(v)$ are zero. As $\Phi(v)$ is strictly convex, this can only happen at the unique minimum of $\Phi(v)$. $\square$

16.5 Uses in Semi-Supervised Learning

In Lecture 13, I suggested an approach to estimating a function f on the vertices of a graph given its values at a set $S \subseteq V$:

$$\min_{x : f(a) = x(a) \text{ for } a \in S} \sum_{(a,b)\in E} (x(a) - x(b))^2.$$

Moreover, we saw that we can minimize such a function by solving a system of linear equations.

Unfortunately, there are situtations in which this approach does not work very well. In general, this should not be surprising: sometimes the problem is just unsolvable. But, there are cases in which it would be reasonable to solve the learning problem in which this approach fails.

Better results are sometimes obtained by modifying the penalty function. For example, Bridle and Zhu [BZ13] (and, essentially, Herbster and Guy [HL09]) suggest

$$\min_{x : f(a) = x(a) \text{ for } a \in S} \sum_{(a,b)\in E} |x(a) - x(b)|^p,$$

for $1 < p < 2$.

While a well-selected p will often improve accuracy, the drawback of this approach is that we cannot perform the minimization nearly as quickly as we can when $p = 2$.

16.6 Dual Energy

We can establish a corresponding, although different, energy for the flows. Let ψ be the inverse of ϕ. We then define the flow-energy of an edge that carries a flow of f to be

$$\Psi(f) \stackrel{\text{def}}{=} \int_0^f \psi(t)dt.$$

If we minimize the sum of the flow-energies over the space of flows, we again recover the unique valid flow in the network. (The function Φ is implicit in the work of Duffin. The dual Ψ comes from Millar [Mil51]).

In the classical case, Φ and Ψ are the same. While they are not the same here, *their sum is.* We will later prove that when $v = \psi(f)$,

$$\Psi(f) + \Phi(v) = fv.$$

In fact, one can show that for all f and v,

$$\Psi(f) + \Phi(v) \geq fv,$$

with equality only when $v = \psi(f)$.

Theorem 16.6.1. *Under the conditions of Theorem 16.4.1, let f_{ext} be the vector of external flows resulting from the induced voltages. Let f be the flow on the edges that is compatible with f_{ext} and that minimizes*

$$\Psi(f) \stackrel{\text{def}}{=} \sum_{(a,b)\in E} \Psi_{(a,b)}(f_{(a,b)}).$$

Then, f is the flow induced by the voltages shown to exist in Theorem 16.4.1.

Sketch. We first show that f is a potential flow. That is, that there exist voltages v so that for every edge (a,b), $f_{(a,b)} = \phi_{(a,b)}(v_a - v_b)$. The theorem then follows by the uniqueness established in Theorem 16.4.1.

To prove that f is a potential flow, we consider the potential difference that the flow "wants" to induce on each edge, $\psi(f_{(a,b)})$. There exist vertex potentials that agree with these desired potential differences if an only if for every pair of vertices and for every pair of paths between them, the sum of the desired potential differences along the edges in the paths is the same. To see this, arbitrarily fix the potential of one vertex, such as s. We may then set the potential of any other vertex a by summing the desired potential differences along the edges in any path from s.

Equivalenty, the desired potential differences are realizable if and only if the sum of these desired potential differences is zero around every cycle. To show that this is the case, we use the minimality of the flow. Because $\Psi(f)$ is strictly convex, small changes to the optimum have a negligible effect on its value (that is, the first derivative is zero). So, pushing an ϵ amount of flow around any cycle will not change the value of $\Psi(f)$. That is, the sum of the derivatives around any cycle will be zero. As

$$\frac{\partial}{\partial f}\Psi_e(f) = \psi_e(f),$$

this means that the sum of the desired potential differences around every cycle is zero. $\square$

Theorem 16.6.2. *If* $f = \phi(v)$, *then*

$$\Phi(v) + \Psi(f) = vf.$$

Proof. One can prove this theorem through "integration by parts". But, I prefer a picture. In the following two figures, the curve is the plot of ϕ. In the first figure, the shaded region is the integral of ϕ between 0 and v (2 in this case). In the second figure, the shaded region is the integral of ψ between 0 and $\phi(v)$ (just turn the picture on its side). It is clear that these are complementary parts of the rectangle between the axes and the point $(v, \phi(v))$.

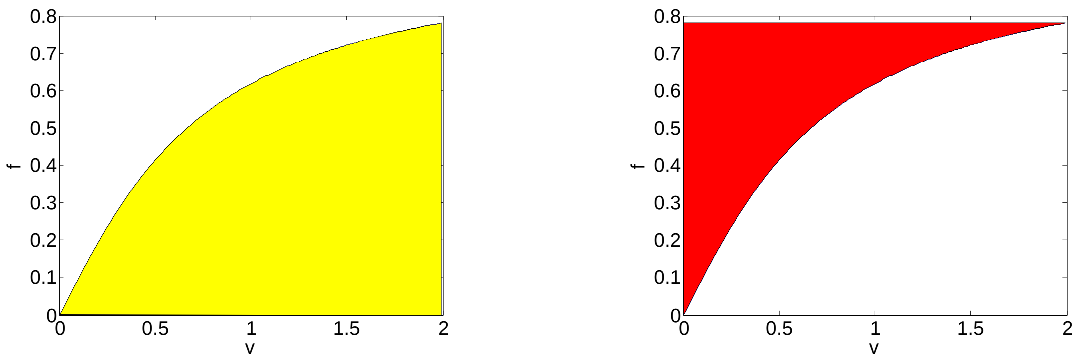

(a) $\Phi(v)$ (b) $\Psi(f)$ when $f = \phi(v)$

$\square$

The bottom line is that almost all of the classical theory can be carried over to nonlinear networks.

16.7 Thermistor Networks

We now turn our attention to networks of resistors whose resistance changes over time. We consider a natural model in which edges get "worn out": as they carry more flow their resistance increases. One physical model that does this is a thermistor. A thermistor is a resistor whose resistance increases with its temperature. These are used in thermostats.

Remember the "energy dissipation" of a resistor? The energy dissipates as heat. So, the temperature of resistor increases as its resistance times the square of the flow through it. To prevent the temperatures of the resistors from going to infinity, we will assume that there is an ambient temperature T_A, and that they tend to the ambient temperature. I will denote by T_e the temperature of resistor e, and I will assume that there is a constant α_e for each resistor so that its resistance

$$r_e = \alpha_e T_e. \tag{16.1}$$

We do not allow temperatures to be negative.

Now, assume that we would like to either flow a current between two vertices s and t, or that we have fixed the potentials of s and t. Given the temperature of every resistor at some moment, we can compute all their resistances, and then compute the resulting electrical flow as we did in Lecture 13. Let f_e be the resulting flow on resistor e. The temperature of e will increase by $r_e f_e^2$, and it will also increase in proportion to the difference between its present temperature and the ambient temperature.

This gives us the following differential equation for the change in the temperature of a resistor:

$$\frac{\partial T_e}{\partial t} = r_e f_e^2 - (T_e - T_A). \tag{16.2}$$

Ok, there should probably be some constant multiplying the $(T_e - T_A)$ term. But, since I haven't specified the units of temperature we can just assume that the constant is 1.

By substituting in (16.1) we can eliminate the references to resistance. We thus obtain

$$\frac{\partial T_e}{\partial t} = \alpha_e T_e f_e^2 - (T_e - T_A).$$

There are now two natural questions to ask: does the system converge, and if so, what does it converge to? If we choose to impose a current flow between s and t, the system does not need to converge. For example, consider just one resistor e between vertices s and t with $\alpha_e = 2$. We then find

$$\frac{\partial T_e}{\partial t} = \alpha_e T_e f_e^2 - (T_e - T_A) = 2T_e - (T_e - T_A) = T_e + T_A.$$

So, the temperature of the resistor will go to infinity.

For this reason, I prefer to just fix the voltages of certain vertices. Under these conditions, we can prove that the system will converge. While I do not have time to prove this, we can examine what it will converge to.

If the system converges, that is if the voltages at the nodes converge along with the potential drops and flows across edges, then

$$0 = \frac{\partial T_e}{\partial t} = \alpha_e T_e f_e^2 - (T_e - T_A).$$

To turn this into a relationship between f_e and v_e, we apply the identity $f_e r_e = v_e$, which becomes $f_e \alpha_e T_e = v_e$, to obtain

$$0 = v_e f_e - T_e + T_A.$$

To eliminate the last occurence of T_e, we then multiply by f_e and apply the same identity to produce

$$0 = v_e f_e^2 - v_e/\alpha_e + f_e T_A.$$

The solutions of this equation in f_e are given by

$$f_e = \pm\sqrt{\frac{1}{\alpha_e} + \left(\frac{T_A}{2v_e}\right)^2} - \frac{T_A}{2v_e}.$$

The correct choice of sign is the one that gives this the same sign as v_e:

$$f_e = \frac{1}{2v_e}\left(\sqrt{\frac{(2v_e)^2}{\alpha_e} + T_A^2} - T_A\right). \tag{16.3}$$

When v_e is small this approaches zero, so we define it to be zero when v_e is zero. As v_e becomes large this expression approaches $\alpha_e^{-1/2}$. Similarly, when v_e becomes very negative, this approaches $-\alpha_e^{-1/2}$. If we now define

$$\phi_e(v_e) = \frac{1}{2v_e}\left(\sqrt{\frac{(2v_e)^2}{\alpha_e} + T_A^2} - T_A\right),$$

we see that this function satisfies properties a, b and c. Theorem 16.4.1 then tells us that a stable solution exists.

16.8 Low Temperatures

We now observe that when the ambient temperature is low, a thermistor network produces a minimum s-t cut in a graph. The weights of the edges in the graph are related to α_e. For simplicity, we will just examine the case when all $\alpha_e = 1$. If we take the limit as T_A approaches zero, then the behavior of ϕ_e is

$$\phi_e(v_e) = \begin{cases} 0 & \text{if } v_e = 0 \\ 1 & \text{if } v_e > 0 \\ -1 & \text{if } v_e < 0. \end{cases}$$

We will obtain similar behavior for small T_A: if there is a non-negligible potential drop across an edge, then the flow on that edge will be near 1. So, every edge will either have a flow near 1 or a negligible potential drop. When an edge has a flow near 1, its energy will be near 1. On the other hand, the energy of edges with negligible potential drop will be near 0.

So, in the limit of small temperatures, the energy minimization problem becomes

$$\min_{v:v(s)=0,v(t)=1} \sum_{(a,b)\in E} |v(a) - v(b)|.$$

One can show that the minimum is achieved when all of the voltages are 0 or 1, in which case the energy is the number of edges going between voltage 0 and 1. That is, the minimum is achieved by a minimum s-t cut.

References

[BZ13] Nick Bridle and Xiaojin Zhu. p-voltages: Laplacian regularization for semi-supervised learning on high-dimensional data. In *Eleventh Workshop on Mining and Learning with Graphs (MLG2013)*, 2013.

[Duf47] R. J. Duffin. Nonlinear networks. IIa. *Bull. Amer. Math. Soc*, 53:963–971, 1947.

[HL09] Mark Herbster and Guy Lever. Predicting the labelling of a graph via minimum p-seminorm interpolation. In *Proceedings of the 2009 Conference on Learning Theory (COLT)*, 2009.

[Mil51] William Millar. Cxvi. some general theorems for non-linear systems possessing resistance. *Philosophical Magazine*, 42(333):1150–1160, 1951.