



Graphs in Heat

Keenan Crane

February 8, 2016



In this note we show how the *heat method* [1] can be used to approximate a solution to the single- or multiple-source shortest path problem. The reader should be forewarned that this technique is primarily of theoretical interest: standard schemes like Dijkstra’s algorithm are almost certainly more useful in practice. Still, it is interesting to see that the heat method can be applied to a problem for which it was not originally designed. Since the method is formulated in terms of a Riemannian manifold (M, g) , the basic idea is to “thicken” the graph and then apply the original formulation, approximating the constituent PDEs using piecewise linear basis functions. Empirically, the final solution approximates the graph distance to within some small factor of the graph diameter. One could potentially improve accuracy by changing the discretization—though again, the practical motivation for doing so is not clear.



Let $G = (V, E)$ be any graph, and let $\ell : E \rightarrow \mathbb{R}$ be an assignment of lengths to edges. If we imagine that the graph is isometrically embedded in Euclidean \mathbb{R}^n as a collection of linear segments, then we can “thicken” the graph by taking the Minkowski sum with a ball of radius $\epsilon > 0$; the expressions we ultimately obtain will not depend on the particular choice of n or ϵ . If the thickened graph M inherits the Riemannian metric g induced by the Euclidean embedding, then we can associate to each edge $ij \in E$ the subset U_{ij} of points in M closest to ij , *i.e.*, the geodesic Voronoi cell. We further associate to these cells functions $\phi_{ij} : M \rightarrow \mathbb{R}$ that vanish outside U_{ij} , and inside U_{ij} interpolate the values 1 at i , and 0 at j . In analogy with Lagrange or “hat” functions used for discretization on a simplicial domain, the corresponding basis functions at vertices are given by $\phi_i|_{U_{ij}} := \phi_{ij}$ for each edge ij incident on i , and zero otherwise. We will use these basis functions to discretize the three steps of the heat method, namely

- I. Solve $(\text{id} - t\Delta)u = u_0$
- II. Evaluate $X \leftarrow \nabla u / |\nabla u|$
- III. Solve $\Delta\phi = \nabla \cdot X$

To discretize the Laplace-Beltrami operator Δ we integrate over the Voronoi regions U_i associated with vertices to obtain an expression for the stiffness matrix $L \in \mathbb{R}^{V \times V}$, namely

$$(Lu)_i := \int_{U_i} \Delta u \, d\mu = \int_{U_i} \nabla \cdot \nabla u \, d\mu = \int_{\partial U_i} N \cdot \nabla u \, d\nu = \sum_{ij \in E} V_\epsilon^{n-1} \frac{u_j - u_i}{\ell_{ij}},$$

where μ and ν are the usual volume measures on \mathbb{R}^n and \mathbb{R}^{n-1} , respectively, N is the outward unit normal at the region boundary, V_ϵ^n is the volume of an n -ball of diameter ϵ , and ℓ_{ij} is the length (or weight) of edge ij . Note that here we have assumed zero-Neumann boundary conditions, corresponding to zero outflow of heat through the walls of the thickened graph. The corresponding lumped mass matrix has entries

$$M_{ii} := \sum_{ij \in E} V_\epsilon^{n-1} \ell_{ij} / 2,$$

where we have ignored the contribution of the small region in the immediate vicinity of the vertex (since it becomes negligible as $\epsilon \rightarrow 0$). The first step of the heat method is therefore discretized as

$$(M - tL)u = u_0,$$

where $u_0 \in \mathbb{R}^{|V|}$ is a sum of Kronecker deltas encoding the source vertices. (As in the simplicial case, a value of 1 already represents an integrated Dirac delta, accounting for the absence of the mass matrix on the right-hand side.) Note that the constant V_ϵ^{n-1} merely scales the solution, which does not affect subsequent calculations. The normalized gradient of u is given along each edge by

$$X_{ij} := \frac{u_j - u_i}{|u_j - u_i|} = \begin{cases} +1, & u_j > u_i, \\ -1, & u_j < u_i \\ 0, & \text{otherwise.} \end{cases}$$

In other words, in the case of graphs, the first two steps of the heat method merely determine a relative ordering of vertices (with respect to distance). Hence, the discrete divergence $b \in \mathbb{R}^{|V|}$ will just count the number of edges with “inflow” (positive gradient) minus the number with “outflow” (negative gradient):

$$b_i := \int_{U_i} \nabla \cdot X \, d\mu = \int_{\partial U_i} n \cdot X \, d\nu = \sum_{ij \in E} n \cdot X_{ij} = \sum_{ij \in E} \text{sgn}(u_j - u_i).$$

The final Poisson equation is then discretized via

$$L\phi = b;$$

since constant functions are in the kernel of L (*i.e.*, a function is determined by its derivative only up to a constant), one shifts the solution $\phi \in \mathbb{R}^{|V|}$ such that it equals zero at the source point.

An example solution is shown above on a random graph with 100 vertices and 200 edges where weights are determined by embedded edge lengths; the black dot marks the source point. Colors plot the distance relative to the (true) maximum distance d_{\max} with respect to the source point. Though it is difficult to say anything about accuracy in general, in this example the mean deviation from the true distance is about 5% relative to d_{\max} . For trees (not surprisingly) we obtain the exact solution.

References

- [1] Keenan Crane, Clarisse Weischedel, and Max Wardetzky. Geodesics in Heat: A New Approach to Computing Distance Based on Heat Flow. *ACM Transactions on Graphics*, 32, 2013.

