

A Unifying Framework for Spectrum-Preserving Graph Sparsification and Coarsening: Appendix

A Empirical validation of the approximation in equation (10)

In order to derive our graph reduction algorithm, we assume that the entries of the M_e associated to different edges are approximately entrywise uncorrelated (Main Text Section 3.3). Similar to how the variance of the sum of independent random variables is the sum of their individual variances, this assumption allows us to approximate the expected squared Frobenius error of the final reduced graph $\mathbb{E}[\|\mathbf{L}_G^\dagger - \mathbf{L}_G^\dagger\|_F^2]$ as a sum over the sequence of probabilistic actions to individual edges:

$$\underbrace{\mathbb{E}\left[\left\|\sum \Delta\mathbf{L}^\dagger\right\|_F^2\right]}_{\text{true error}} \approx \underbrace{\sum \mathbb{E}\left[\|\Delta\mathbf{L}^\dagger\|_F^2\right]}_{\text{estimated error}}. \quad (18)$$

In Figure SI 1, we empirically validate this assumption for networks with a variety of structures. In fact, the true error is statistically equal to or less than the estimated error. Thus, the estimated error may be used by StopCriterion in Algorithm 1.

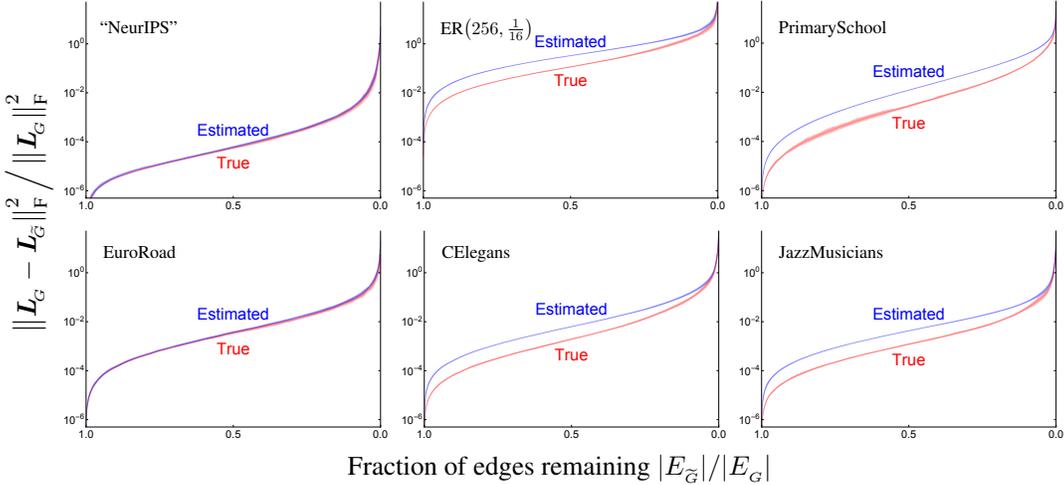


Figure SI 1: **The approximation of uncorrelated changes to \mathbf{L}^\dagger is nearly exact or a conservative estimate.** We test the validity of equation (18) using a variety of datasets: *Top left*: a triangular mesh of the text “NeurIPS” (567 nodes and 1408 edges); *Top middle*: an Erdős–Rényi model (256 nodes and $p = 1/16$); *Top right*: a weighted social network of face-to-face interactions between primary school students, with initial edge weights proportional to the number of interactions between pairs of students (236 nodes and 5899 edges) from [53]; *Bottom left*: a transportation network of European cities and roads between them (1039 nodes and 1305 edges) from [56]; *Bottom middle*: the *C. elegans* posterior nervous system connectome (269 nodes and 2902 edges) from [52]; and *Bottom right*: a collaboration network of Jazz musicians (198 nodes and 2742 edges) from [51]. We applied Algorithm 1, prioritizing edge reduction (allowing for deletion, contraction, and reweighting), and setting $q = 1/16$ and $d = 1/4$. We recorded the estimated error and the true error in \mathbf{L}^\dagger as a function of amount of reduction. Shading denotes one standard deviation about the mean for 32 runs of the algorithm. In general, the estimated error serves as an approximate upper bound of the true error in \mathbf{L}^\dagger (although it is nearly exact for graphs with a geometric quality). The validity of the approximation allows one to use a bound on the estimated error as a StopCriterion in Algorithm 1.

B Derivation of the optimal probabilistic action to an edge

As discussed in Section 3.4, we seek to minimize:

$$\mathcal{C} = \mathbb{E}\left[\|\Delta\mathbf{L}^\dagger\|_F^2\right] - \beta^2\mathbb{E}[r], \quad (19)$$

subject to

$$\mathbb{E}[\Delta \mathbf{L}^\dagger] = \mathbf{0}. \quad (20)$$

When reducing multiple edges, $\mathbb{E}[r]$ is additive and $\mathbb{E}[\|\Delta \mathbf{L}^\dagger\|_F^2]$ is approximately additive (see Appendix Section A). Thus, we partition this minimization into a sequence of subproblems, treating each perturbation to an edge individually.

Recall that

$$\Delta \mathbf{L}^\dagger = \underbrace{f\left(\frac{\Delta w}{w_e}, w_e \Omega_e\right)}_{\text{nonlinear scalar}} \times \underbrace{\mathbf{M}_e}_{\text{constant matrix}}, \quad \text{where} \quad f = -\frac{\frac{\Delta w}{w_e}}{1 + \frac{\Delta w}{w_e} w_e \Omega_e}.$$

We now derive the optimal probability of deleting (p_d), contracting (p_c), or reweighting ($1 - p_d - p_c$) a given edge e , along with the change to its weight (Δw) in the case of the latter.

The constraint (20) requires that this reweight satisfies

$$\frac{p_d}{1 - w_e \Omega_e} - \frac{p_c}{w_e \Omega_e} + (1 - p_d - p_c) \mathbb{E}[f | \text{reweight}] = 0, \quad (21)$$

where we have used the following limits:

$$\begin{aligned} \text{deletion:} & \quad \frac{\Delta w}{w_e} \rightarrow -1, & f & \rightarrow (1 - w_e \Omega_e)^{-1} \\ \text{contraction:} & \quad \frac{\Delta w}{w_e} \rightarrow +\infty, & f & \rightarrow -(w_e \Omega_e)^{-1}. \end{aligned} \quad (22)$$

Likewise, the cost function (19) for acting on the edge e becomes:

$$\mathcal{C} = \left(\frac{p_d}{(1 - w_e \Omega_e)^2} + \frac{p_c}{(w_e \Omega_e)^2} + (1 - p_d - p_c) \mathbb{E}[f^2 | \text{reweight}] \right) m_e^2 - \beta^2 (r_d p_d + r_c p_c), \quad (23)$$

where r_d and r_c are the number of prioritized items that would be removed by a deletion or contraction, respectively.

For a fixed p_d and p_c , $\mathbb{E}[f | \text{reweight}]$ is fixed by equation (21). As $\frac{\partial^2 f}{\partial \Delta w^2} > 0$ everywhere, the inequality $\mathbb{E}[f^2 | \text{reweight}] \geq \mathbb{E}[f | \text{reweight}]^2$ becomes an equality under minimization of (23).

Thus, if an edge is to be reweighted, it will be changed by the unique Δw satisfying

$$\frac{p_d}{1 - w_e \Omega_e} - \frac{p_c}{w_e \Omega_e} - (1 - p_d - p_c) \frac{\frac{\Delta w}{w_e}}{1 + \frac{\Delta w}{w_e} w_e \Omega_e} = 0. \quad (24)$$

Clearly, the space of allowed solutions lies within the simplex \mathcal{S} : $0 \leq p_d, 0 \leq p_c, p_d + p_c \leq 1$. The additional constraint $-1 \leq \frac{\Delta w}{w_e} \leq \infty$ further implies that $p_c \leq w_e \Omega_e$ and $p_d \leq 1 - w_e \Omega_e$. Hence, we substitute (24) into (23), and minimize it over this domain (given $m_e, w_e \Omega_e, \tau_e$, and β). After some careful elementary calculus, we obtain the solution provided in Figure 1 of the Main Text.

C Lifting the matrices of a contracted graph

Here, we provide a detailed rationale for the definitions given in Section 3.5, namely, the choice of $\mathbf{L}_{\bar{G}}$ and $\mathbf{L}_{\bar{G}}^\dagger$, and how to “lift” these matrices to the original dimension $|V_G| \times |V_G|$ when edges have been contracted.

Recall the following definitions:

$$\mathbf{L}_{\bar{G}} = \mathbf{W}_n^{-1} \mathbf{B}^\top \mathbf{W}_e \mathbf{B}, \quad (25)$$

$$\mathbf{L}_{\bar{G}}^\dagger = (\mathbf{L}_{\bar{G}} + \mathbf{J})^{-1} - \mathbf{J}, \quad (26)$$

$$\mathbf{L}_{\bar{G},l} = \mathbf{C}^\top \mathbf{L}_{\bar{G}} \mathbf{W}_n^{-1} \mathbf{C}, \quad (27)$$

$$\mathbf{L}_{\bar{G},l}^\dagger = \mathbf{C}^\top \mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1} \mathbf{C}, \quad (28)$$

where

$$\mathbf{J} = \frac{1}{\bar{\mathbf{1}}^\top \bar{\mathbf{w}}_n} \bar{\mathbf{1}} \bar{\mathbf{w}}_n^\top, \quad (29)$$

$$\mathbf{C} = \{c_{ij}\} = \begin{cases} 1 & \text{node } j \text{ in supernode } i \\ 0 & \text{otherwise.} \end{cases} \quad (30)$$

The above definitions ensure that the lifted $\mathbf{L}_{\bar{G},l}^\dagger$ of the contracted graph is identical to the $w_e \rightarrow \infty$ limit of the original \mathbf{L}_G^\dagger .

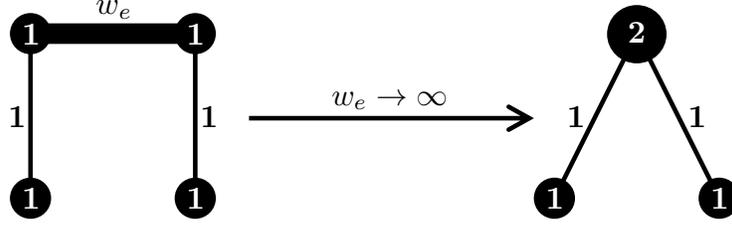


Figure SI 2: **Contracting the center edge of a line graph.** *Left:* Original graph G with large weight w_e on the center edge. *Right:* Reduced graph \tilde{G} obtained by contracting this edge ($w_e \rightarrow \infty$). Note that the weight of the contracted nodes sum to give the weight of the resulting supernode in the reduced graph.

To illustrate the consistency of these definitions, we consider a concrete example: the line graph with 3 edges, where the center edge is to be contracted (Figure SI 2). Let the center edge have weight $w_e \gg 1$, while the other two have a fixed weight of 1.

For the original graph G , the Laplacian and its pseudoinverse are

$$\mathbf{L}_G = \begin{pmatrix} 1 & -1 & 0 & 0 \\ -1 & 1+w_e & -w_e & 0 \\ 0 & -w_e & 1+w_e & -1 \\ -0 & 0 & -1 & 1 \end{pmatrix}, \quad \mathbf{L}_G^\dagger = \frac{1}{8} \begin{pmatrix} 5 + \frac{2}{w_e} & -1 + \frac{2}{w_e} & -1 - \frac{2}{w_e} & -3 - \frac{2}{w_e} \\ -1 + \frac{2}{w_e} & 1 + \frac{2}{w_e} & 1 - \frac{2}{w_e} & -1 - \frac{2}{w_e} \\ -1 - \frac{2}{w_e} & 1 - \frac{2}{w_e} & 1 + \frac{2}{w_e} & -1 + \frac{2}{w_e} \\ -3 - \frac{2}{w_e} & -1 - \frac{2}{w_e} & -1 + \frac{2}{w_e} & 5 + \frac{2}{w_e} \end{pmatrix}.$$

For the contracted graph \tilde{G} , we have

$$\mathbf{W}_n = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \mathbf{J} = \begin{pmatrix} \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Thus, the reduced Laplacian and its pseudoinverse are

$$\mathbf{L}_{\tilde{G}} = \begin{pmatrix} 1 & -1 & 0 \\ -\frac{1}{2} & 1 & -\frac{1}{2} \\ 0 & -1 & 1 \end{pmatrix}, \quad \mathbf{L}_{\tilde{G}}^\dagger = \frac{1}{8} \begin{pmatrix} 5 & -2 & -3 \\ -1 & 2 & -1 \\ -3 & -2 & 5 \end{pmatrix}.$$

When lifted to the original dimensions $|V_G| \times |V_G|$, these become

$$\mathbf{L}_{\tilde{G},l} = \begin{pmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} & 0 \\ -\frac{1}{2} & \frac{1}{2} & \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} & \frac{1}{2} & -\frac{1}{2} \\ 0 & -\frac{1}{2} & -\frac{1}{2} & 1 \end{pmatrix}, \quad \mathbf{L}_{\tilde{G},l}^\dagger = \frac{1}{8} \begin{pmatrix} 5 & -1 & -1 & -3 \\ -1 & 1 & 1 & -1 \\ -1 & 1 & 1 & -1 \\ -3 & -1 & -1 & 5 \end{pmatrix}.$$

Note that the lifted $\mathbf{L}_{\tilde{G},l}^\dagger$ is equal to the $w_e \rightarrow \infty$ limit of the original \mathbf{L}_G^\dagger , as desired. In contrast, the original \mathbf{L}_G diverges, while the lifted $\mathbf{L}_{\tilde{G},l}$ averages the rows and columns of the merged nodes. Moreover, regardless of whether node weights are included in the definitions, using the standard Moore–Penrose pseudoinverse of the reduced Laplacian will yield a lifted pseudoinverse that is not equivalent to the original in the $w_e \rightarrow \infty$ limit.

Additionally, we remark that, while contraction always requires the summing of node weights, it can also lead to the summing of edge weights (when the contracted edge participates in any triangle in the original graph, see Figure SI 3).

D Proof of the relationship between the hyperbolic distance and σ -spectral approximation

In this section, we prove Theorem 1 from Section 5.1

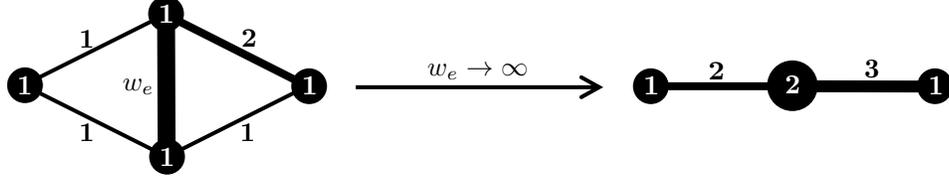


Figure SI 3: **Contracting an edge that participates in triangles.** *Left:* Original graph G containing an edge with large weight w_e that participates in two triangles. *Right:* Reduced graph \tilde{G} obtained by contracting this edge ($w_e \rightarrow \infty$). Note that the two non-contracted edges in each triangle form a single edge in the reduced graph with weight equal to their sum.

Theorem 1. If $d_h(\mathbf{L}_G, \mathbf{L}_{\tilde{G}}) \leq \ln(\sigma)$, then \tilde{G} is a σ -spectral approximation of G .

Proof. Let G be the original graph and \tilde{G} its sparse approximation (no contraction/removing of nodes). Recall the relevant definitions:

\tilde{G} is a σ -spectral approximation of G [6] if

$$\frac{1}{\sigma} \vec{x}^\top \mathbf{L}_G \vec{x} \leq \vec{x}^\top \mathbf{L}_{\tilde{G}} \vec{x} \leq \sigma \vec{x}^\top \mathbf{L}_G \vec{x}, \quad \forall \vec{x} \in \mathbb{R}^{|\mathcal{V}_G|}. \quad (31)$$

We propose to instead measure the hyperbolic distance between the resulting $\mathbf{L}_G \vec{x}$ and $\mathbf{L}_{\tilde{G}} \vec{x}$, namely

$$d_h(\mathbf{L}_G, \mathbf{L}_{\tilde{G}}) \stackrel{\text{def}}{=} \sup_{\vec{x} \perp \mathbf{1}} \left\{ \operatorname{arccosh} \left(1 + \frac{\|(\mathbf{L}_G - \mathbf{L}_{\tilde{G}}) \vec{x}\|_2^2 \|\vec{x}\|_2^2}{2(\vec{x}^\top \mathbf{L}_G \vec{x})(\vec{x}^\top \mathbf{L}_{\tilde{G}} \vec{x})} \right) \right\}, \quad (32)$$

where \mathbf{L}_G and $\mathbf{L}_{\tilde{G}}$ are the Laplacians of G and \tilde{G} , respectively, and \vec{x} is perpendicular to their kernels.

Consider the result of a Laplacian acting on such a vector \vec{x} , and decompose the output as a component parallel to \vec{x} with magnitude ℓ_{\parallel} and a component $\vec{\ell}_{\perp}$ perpendicular to \vec{x} :

$$\mathbf{L}_G \vec{x} = \tilde{\ell}_{\parallel} \frac{\vec{x}}{\|\vec{x}\|_2} + \vec{\ell}_{\perp}, \quad \mathbf{L}_{\tilde{G}} \vec{x} = \tilde{\ell}_{\parallel} \frac{\vec{x}}{\|\vec{x}\|_2} + \tilde{\ell}_{\perp}. \quad (33)$$

Hence,

$$\|(\mathbf{L}_G - \mathbf{L}_{\tilde{G}}) \vec{x}\|_2^2 = (\ell_{\parallel} - \tilde{\ell}_{\parallel})^2 + \|\vec{\ell}_{\perp} - \tilde{\ell}_{\perp}\|_2^2, \quad (34)$$

$$\vec{x}^\top \mathbf{L}_G \vec{x} = \ell_{\parallel} \|\vec{x}\|_2, \quad \vec{x}^\top \mathbf{L}_{\tilde{G}} \vec{x} = \tilde{\ell}_{\parallel} \|\vec{x}\|_2. \quad (35)$$

Let $z = \tilde{\ell}_{\parallel} / \ell_{\parallel}$. Substituting (35) into (31), we see that \tilde{G} is a σ -spectral approximation of G if

$$\max \left\{ \frac{\tilde{\ell}_{\parallel}}{\ell_{\parallel}}, \frac{\ell_{\parallel}}{\tilde{\ell}_{\parallel}} \right\} = \max \left\{ z, \frac{1}{z} \right\} \leq \sigma. \quad (36)$$

Now, substituting (34) into (16), we obtain:

$$\begin{aligned} d_{\vec{x}}(\mathbf{L}_G, \mathbf{L}_{\tilde{G}}) &= \operatorname{arccosh} \left(1 + \frac{((\ell_{\parallel} - \tilde{\ell}_{\parallel})^2 + \|\vec{\ell}_{\perp} - \tilde{\ell}_{\perp}\|_2^2) \|\vec{x}\|_2^2}{2\ell_{\parallel} \|\vec{x}\|_2 \tilde{\ell}_{\parallel} \|\vec{x}\|_2} \right) \\ &\geq \operatorname{arccosh} \left(1 + \frac{\tilde{\ell}_{\parallel}^2 - 2\tilde{\ell}_{\parallel} \ell_{\parallel} + \ell_{\parallel}^2}{2\ell_{\parallel} \tilde{\ell}_{\parallel}} \right) \\ &\geq \operatorname{arccosh} \left(\frac{1}{2} \left(z + \frac{1}{z} \right) \right). \end{aligned}$$

Using the identity $\operatorname{arccosh}(x) = \ln(x + \sqrt{x^2 - 1})$,

$$\begin{aligned} d_{\vec{x}}(\mathbf{L}_G, \mathbf{L}_{\tilde{G}}) &\geq \ln \left(\frac{1}{2} \left(z + \frac{1}{z} \right) + \sqrt{\frac{1}{4} \left(z + \frac{1}{z} \right)^2 - 1} \right) \\ &\geq \ln \left(\frac{1}{2} \left(z + \frac{1}{z} \right) + \frac{1}{2} \left| z - \frac{1}{z} \right| \right) \\ &\geq |\ln(z)|. \end{aligned}$$

Thus, if $d_{\vec{x}}(\mathbf{L}_G, \mathbf{L}_{\tilde{G}}) \leq \ln(\sigma) \forall \vec{x} \perp \vec{1}$, then \tilde{G} is a σ -spectral approximation of G , as desired. \square

E Number of edges acted upon per iteration can be $\mathcal{O}(|V|)$

In this section, we study the effect of varying the parameter q , the fraction of sampled edges acted upon, using real-world datasets from different domains (Figure SI 4).

For each iteration of our algorithm, we sample a random independent edge set and act on the fraction q with the lowest β_{*e} (see Main Text Section 4). We find that the resulting error asymptotes around $q \sim 1/16$. We expect that by combining this sampling method with existing algorithmic primitives (eg, [20], see Appendix Section E), our algorithm could achieve a running time of $\tilde{\mathcal{O}}(\langle k \rangle |E|)$, where $\langle k \rangle$ is the average degree (see Main Text Section 4). This would allow it to be used in large-scale applications of graph reduction.

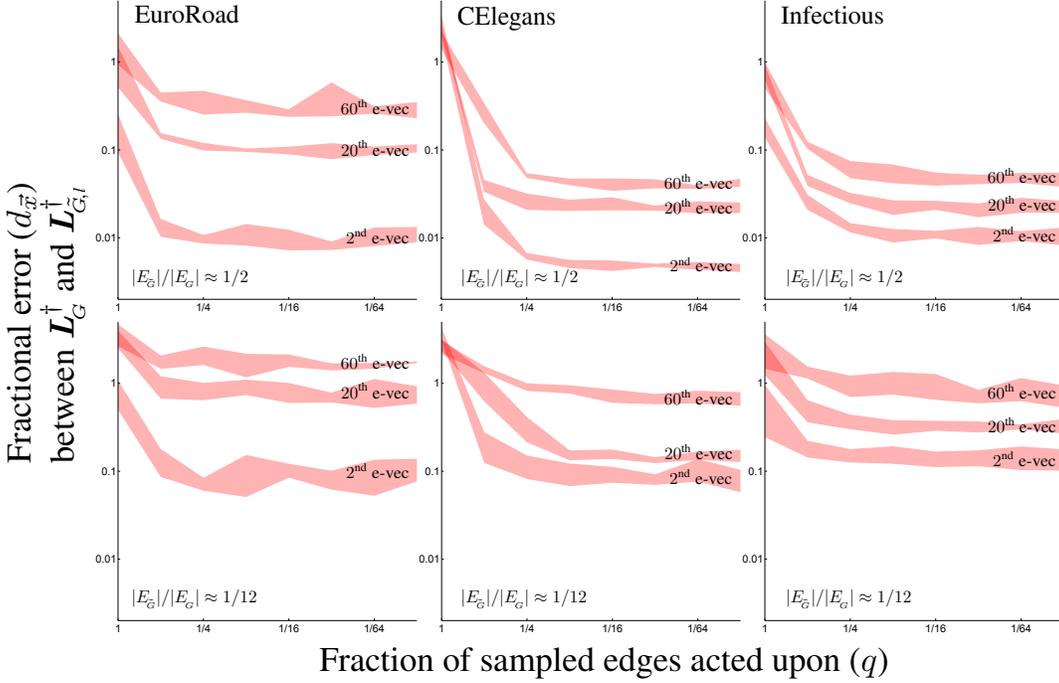


Figure SI 4: **Number of sampled edges acted upon per iteration can be $\mathcal{O}(|V|)$.** We study the effect of varying q , the fraction of the sampled edges that are acted upon per iteration, using three datasets: *Left*: a transportation network of European cities and roads between them (1039 nodes and 1305 edges) from [56]; *Middle*: the *C. elegans* posterior nervous system connectome (269 nodes and 2902 edges) from [52]; and *Right*: a weighted social network of face-to-face interactions during an exhibition on infectious diseases, with initial edge weights proportional to the number of interactions between pairs of people (410 nodes and 2765 edges) from [57]. We prioritize edge reduction (allowing for deletion, contraction, and reweighting). At each iteration, the algorithm randomly samples a maximal independent edge set, and chooses β such that a fraction q of these edges (with the lowest β_{*e}) are acted upon. For each run, we compute the hyperbolic distance $d_{\vec{x}}$ (fractional error) between $\mathbf{L}_G^\dagger \vec{x}$ and $\mathbf{L}_{G,I}^\dagger \vec{x}$, where \vec{x} is one of three eigenvectors of the original Laplacian. *Top* plots display the results when the graph has 1/2 of its original number of edges, and *bottom* plots when it has 1/12. Shading denotes one standard deviation about the mean for 8 runs of the algorithm for a given value of q . Note that a significant fraction ($q \sim 1/16$) of the sampled edges can be reduced each iteration without sacrificing much in terms of accuracy. As, empirically, the size of the independent edge sets are typically $\mathcal{O}(|V|)$, the number of edges acted upon per iteration can likewise be $\mathcal{O}(|V|)$.

F Efficiently computing m_e

As discussed in Main Text Section 4 the main computational bottleneck of our algorithm is computing Ω_e and m_e . For Ω_e , we can draw on the work of [20], which describes a method for efficiently computing ε -approximate values of Ω_e for all edges, requiring $\tilde{\mathcal{O}}(|E| \log |V| / \varepsilon^2)$ time. In this section, we describe an analogous procedure to efficiently compute the m_e .

Recall that the reduced Laplacian is:

$$\mathbf{L}_{\bar{G}} = \mathbf{W}_n^{-1} \mathbf{B}^\top \mathbf{W}_e \mathbf{B},$$

hence, the quantity $\widehat{\mathbf{L}}_{\bar{G}} \stackrel{\text{def}}{=} \mathbf{W}_n^{1/2} \mathbf{L}_{\bar{G}} \mathbf{W}_n^{-1/2}$ is clearly symmetric.

Less obvious is the fact that $\widehat{\mathbf{L}}_{\bar{G}}^\dagger \stackrel{\text{def}}{=} \mathbf{W}_n^{1/2} \mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1/2}$ is also symmetric. This can be seen by noting that $\mathbf{W}_n^{1/2} \mathbf{J} \mathbf{W}_n^{-1/2}$ is symmetric, and using the definition of the inverse (equation (26)):

$$\begin{aligned} \widehat{\mathbf{L}}_{\bar{G}}^\dagger &= \mathbf{W}_n^{1/2} \mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1/2} \\ &= \mathbf{W}_n^{1/2} \left((\mathbf{L}_{\bar{G}} + \mathbf{J})^{-1} - \mathbf{J} \right) \mathbf{W}_n^{-1/2} \\ &= \mathbf{W}_n^{1/2} (\mathbf{L}_{\bar{G}} + \mathbf{J})^{-1} \mathbf{W}_n^{-1/2} - \mathbf{W}_n^{1/2} \mathbf{J} \mathbf{W}_n^{-1/2} \\ &= \left(\mathbf{W}_n^{1/2} \mathbf{L}_{\bar{G}} \mathbf{W}_n^{-1/2} + \mathbf{W}_n^{1/2} \mathbf{J} \mathbf{W}_n^{-1/2} \right)^{-1} - \mathbf{W}_n^{1/2} \mathbf{J} \mathbf{W}_n^{-1/2}. \end{aligned}$$

We also remark that $\widehat{\mathbf{L}}_{\bar{G}}^\dagger$ is indeed the pseudoinverse of $\widehat{\mathbf{L}}_{\bar{G}}$:

$$\widehat{\mathbf{L}}_{\bar{G}}^\dagger \widehat{\mathbf{L}}_{\bar{G}} = \widehat{\mathbf{L}}_{\bar{G}} \widehat{\mathbf{L}}_{\bar{G}}^\dagger = \mathbf{I} - \mathbf{W}_n^{1/2} \mathbf{J} \mathbf{W}_n^{-1/2}$$

The change to the reduced Laplacian $\mathbf{L}_{\bar{G}}$ is given by

$$\Delta \mathbf{L}_{\bar{G}} = \mathbf{W}_n^{-1} \vec{b}_e \Delta w_e \vec{b}_e^\top$$

Thus, by the Woodbury matrix identity, the change to its inverse is

$$\Delta \mathbf{L}_{\bar{G}}^\dagger = f w_e \mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1} \vec{b}_e \vec{b}_e^\top \mathbf{L}_{\bar{G}}^\dagger$$

where f is given by equation (5).

Lifting this change back to the original dimension via equation (28) gives

$$\Delta \mathbf{L}_{\bar{G},i}^\dagger = f w_e \mathbf{C}^\top \mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1} \vec{b}_e \vec{b}_e^\top \mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1} \mathbf{C}$$

In particular, as $\mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1}$ is symmetric, $\Delta \mathbf{L}_{\bar{G},i}^\dagger$ is also symmetric, thus we can write the Frobenius norm as

$$\left\| \Delta \mathbf{L}_{\bar{G},i}^\dagger \right\|_F = f w_e \vec{b}_e^\top \mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1} \mathbf{C} \mathbf{C}^\top \mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1} \vec{b}_e \quad (37)$$

$$= f m_e \quad (38)$$

Note that the definition of m_e provided in Section 3.3 of the Main Text (equation (9)) applies to the case of unit node weights, and the general expression is given by

$$m_e = w_e \vec{b}_e^\top \mathbf{L}_{\bar{G}}^\dagger \mathbf{L}_{\bar{G}}^\dagger \mathbf{W}_n^{-1} \vec{b}_e, \quad (39)$$

where we have used $\mathbf{C} \mathbf{C}^\top = \mathbf{W}_n$.

Thus, we can express m_e in terms of $\widehat{\mathbf{L}}_{\bar{G}}^\dagger$:

$$\begin{aligned} m_e &= w_e \vec{b}_e^\top \mathbf{W}_n^{-1/2} \widehat{\mathbf{L}}_{\bar{G}}^\dagger \widehat{\mathbf{L}}_{\bar{G}}^\dagger \mathbf{W}_n^{-1/2} \vec{b}_e \\ &= w_e \left\| \widehat{\mathbf{L}}_{\bar{G}}^\dagger \mathbf{W}_n^{-1/2} \vec{b}_e \right\|_2^2. \end{aligned}$$

We can now use the Johnson–Lindenstrauss lemma to build a structure from which one can efficiently compute approximations of m_e . Let \mathbf{Q} be a random projection matrix of size $k \times n$, where $k = \mathcal{O}(\log n / \varepsilon^2)$, then one can compute ε -approximations of m_e as follows:

$$m_e \approx w_e \left\| \mathbf{Q} \widehat{\mathbf{L}}_{\bar{G}}^\dagger \mathbf{W}_n^{-1/2} \vec{b}_e \right\|_2^2.$$

Let $\mathbf{Z} = \mathbf{Q} \widehat{\mathbf{L}}_{\bar{G}}^\dagger$, and denote the i^{th} rows of \mathbf{Q} and \mathbf{Z} by \vec{q}_i and \vec{z}_i , respectively. Then, one can make k calls to an efficient algebraic multigrid solver (we used the *pyamg* package [58]) to obtain approximate solutions to $\widehat{\mathbf{L}}_{\bar{G}}^\dagger \vec{z}_i = \vec{q}_i$ for the k rows of \mathbf{Z} . An approximation to the m_e of any edge can now be computed by taking the difference between the columns of $\mathbf{Z} \mathbf{W}_n^{-1/2}$ corresponding to the two nodes joined by this edge, and taking the squared 2-norm of the result.

E.1 Constructing the projection matrix

Care must be taken in constructing the projection matrix Q . In particular, its rows must be orthogonal to the null space of $\widehat{\mathbf{L}}_G$, namely $\vec{w}_n^{1/2}$. In addition, the columns must be nearly unit length. To this end, we initialize Q as a random matrix with entries $\{1/\sqrt{k}, -1/\sqrt{k}\}$ with equal probability and iterate the following steps:

1. For each column, scale its values such that it has unit length
2. For each row, subtract its weighted mean $\vec{q}_i^\top \vec{w}_n^{1/2} / \vec{1}^\top \vec{w}_n^{1/2}$

We iterate this procedure until the columns have nearly unit lengths, to within a factor sufficiently smaller than ε .

As a proof of concept, in Figure [SI 5](#) we show the approximate m_e as a function of their exact values.

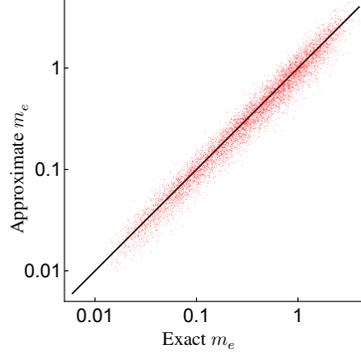


Figure SI 5: **Efficient approximation of m_e** . As a proof of concept, we compare the approximation of m_e (computed using the procedure described in this section) with their exact values. Here, we consider a 64×64 torus graph (4096 nodes and 8192 edges), where the edge weights are randomly distributed as $\exp(U(-2, 2))$, where $U(a, b)$ is the uniform distribution. To calculate the approximate m_e , we project from 4096 to 33 dimensions, resulting in approximations that are typically within a factor of 1.27 of the exact value.

G Perturbations to eigenvalues of the Laplacian pseudoinverse

We first provide the lowest order change in the eigenvalues of \mathbf{L}_G^\dagger . Then, we show how it relates to the Frobenius norm of the perturbation, explicitly relating it to our graph reduction algorithm.

Consider an inverse Laplacian \mathbf{L}^\dagger , which has an eigenvector \vec{x} (without loss of generality, assume $\|\vec{x}\|_2 = 1$) with associated eigenvalue λ . If we perturb \mathbf{L}^\dagger by $\varepsilon \Delta \mathbf{L}^\dagger$, we can solve for the first-order corrections to this “eigenpair” as follows:

$$\begin{aligned} (\mathbf{L}^\dagger + \varepsilon \Delta \mathbf{L}^\dagger)(\vec{x} + \varepsilon \Delta \vec{x}) &= (\lambda + \varepsilon \Delta \lambda)(\vec{x} + \varepsilon \Delta \vec{x}) \\ (\mathbf{L}^\dagger - \lambda) \Delta \vec{x} &= (\Delta \lambda - \Delta \mathbf{L}^\dagger) \vec{x} + \mathcal{O}(\varepsilon), \end{aligned}$$

where we have used $\mathbf{L}^\dagger \vec{x} = \lambda \vec{x}$.

Taking the inner product with \vec{x} gives

$$\begin{aligned} \vec{x}^\top (\mathbf{L}^\dagger - \lambda) \Delta \vec{x} &= \vec{x}^\top (\Delta \lambda - \Delta \mathbf{L}^\dagger) \vec{x} \\ \Delta \vec{x}^\top (\mathbf{L}^\dagger - \lambda) \vec{x} &= \Delta \lambda \vec{x}^\top \vec{x} - \vec{x}^\top \Delta \mathbf{L}^\dagger \vec{x} \\ 0 &= \Delta \lambda - \vec{x}^\top \Delta \mathbf{L}^\dagger \vec{x}, \end{aligned}$$

where we have used the symmetry of \mathbf{L}^\dagger .

This provides the first-order correction to the eigenvalues of $\mathbf{L}^\dagger + \varepsilon \Delta \mathbf{L}^\dagger$:

$$\Delta \lambda = \vec{x}^\top \Delta \mathbf{L}^\dagger \vec{x}. \tag{40}$$

The correction in [\(40\)](#) is controlled by the operator norm of $\Delta \mathbf{L}^\dagger$,

$$\Delta \lambda = \vec{x}^\top \Delta \mathbf{L}^\dagger \vec{x} \leq \sup_{\|\vec{x}\|_2=1} \|\Delta \mathbf{L}^\dagger \vec{x}\|_2 = \|\Delta \mathbf{L}^\dagger\|_{\text{op}}.$$

Thus, bounding the first-order correction to the eigenvalues,

$$|\Delta\lambda| \leq \|\Delta\mathbf{L}^\dagger\|_{\text{op}}. \quad (41)$$

As the operator norm is bounded by the Frobenius norm (by the Cauchy–Schwarz inequality), the estimated error (ie, $\sum \mathbb{E} \|\Delta\mathbf{L}^\dagger\|_{\text{F}}^2$, equation (18)) provides a conservative bound for the change in the eigenvalues of the resulting reduced graph.

Moreover, as the bound is the same for all eigenvalues of the perturbed \mathbf{L}^\dagger , the *relative* error is more tightly bounded for its largest eigenvalues (those associated with large-scale structure).

H Comparison of graph reduction methods using typical similarity measures

Our proposed hyperbolic distance is not usually used as a measure of similarity. Hence, in this section, we show that other more commonly used measures yield similar results when comparing graph reduction algorithms.

H.1 Sparsification

Figure SI 6 compares our algorithm (prioritizing edge reduction, and excluding the possibility of contraction) with the spectral sparsification algorithm of [20] using a stochastic block model (SBM) with four distinct communities. We choose a highly associative SBM due to the clear separation between the eigenvectors associated with global structure (ie, the communities) and the bulk of the spectrum. Note that these algorithms have different objectives (preserving \mathbf{L}^\dagger and \mathbf{L} , respectively), and both accomplish their desired goal.

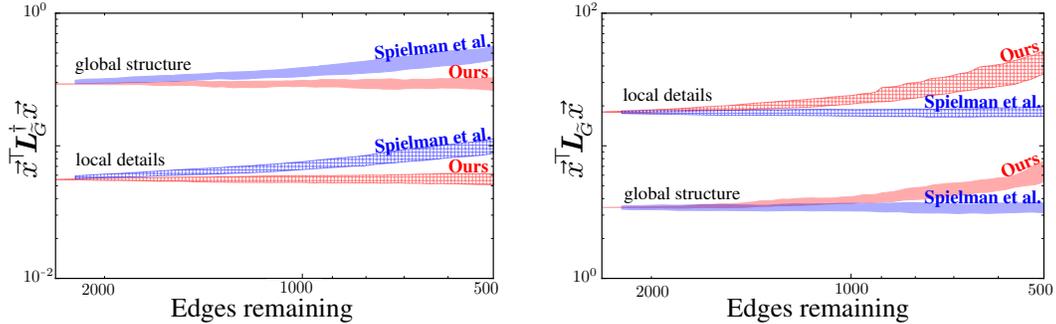


Figure SI 6: **Our sparsification algorithm preferentially preserves global structure.** We compare our algorithm without contraction (in red) with that of Spielman & Srivastava [20] (in blue) using a symmetric stochastic block model (256 nodes, 4 communities, and intra- and inter-community connection probabilities of 2^{-2} and 2^{-6} , respectively). We ran both algorithms 16 times on the same initial graph. For each eigenvector of the original Laplacian, we compute the mean and standard deviation of its quadratic forms (with \mathbf{L}_G^\dagger and with \mathbf{L}_G) as a function of edges remaining. We divide the eigenvectors into two groups: the 3 nontrivial eigenvectors (“global structure”) and the remaining eigenvectors (“local details”), and compute the average mean and average standard deviation for each group. Shading denotes one (average) standard deviation about the (average) mean. *Left:* Laplacian pseudoinverse quadratic form. *Right:* Standard Laplacian quadratic form. Note that the upward bias of the “reciprocal” quadratic form is expected for both algorithms (as $\mathbb{E}[X] \leq 1/\mathbb{E}[1/X]$ for any random variable $X > 0$).

H.2 Coarsening

Figure SI 7 replicates the results of Figure 3 but uses the Laplacian pseudoinverse quadratic form to measure the reduction quality instead of our proposed hyperbolic distance.

Figure SI 8 compares our method with that of Loukas [55], using the average relative error of the k lowest non-trivial eigenvalues of the Laplacian (ie, $\frac{1}{k} \sum_{i=2}^{k+1} |\tilde{\lambda}_i - \lambda_i|/\lambda_i$) to measure the reduction quality.

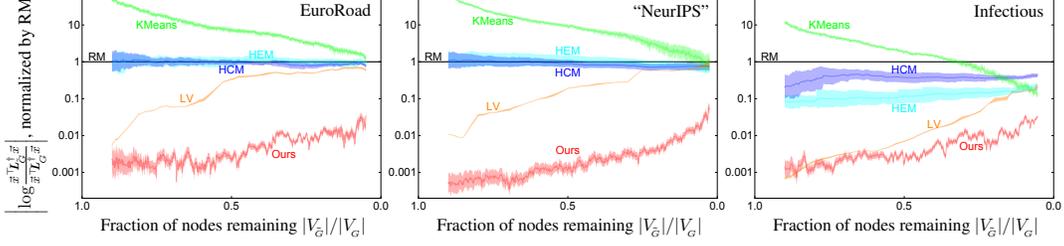


Figure SI 7: **Our coarsening algorithm performs even better when using the quadratic form with L^\dagger .** Here we replicate the experiments in Figure 3. However, instead of using our proposed hyperbolic distance, we consider the logarithm of the fractional change in the Laplacian pseudoinverse quadratic form for \vec{x} the lowest non-trivial eigenvector of the original Laplacian: $|\log(\vec{x}^T L_G^\dagger \vec{x} / \vec{x}^T L_G^\dagger \vec{x})|$. As before, for each algorithm, we plot the mean of this quantity normalized by that obtained by random matching (RM). Shading denotes one standard deviation about the mean for 16 runs of the algorithms. The results are remarkably similar to those obtained using our proposed hyperbolic distance (Figure 3). The most notable deviation is that our algorithm appears to perform *better* when compared using this quadratic form.

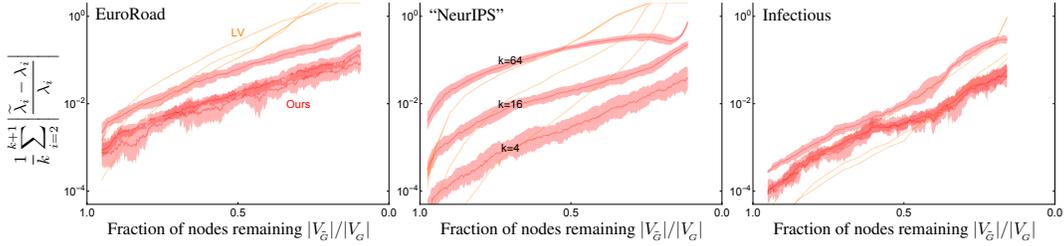


Figure SI 8: **Our algorithm preferentially preserves the lower portion of the Laplacian spectrum.** We compare our coarsening algorithm (**Ours**) with that of Loukas [55] (**LV**) using the same three datasets as in Figure 3. We use the relative error in the k lowest non-trivial eigenvalues of the Laplacian: $\frac{1}{k} \sum_{i=2}^{k+1} |\tilde{\lambda}_i - \lambda_i| / \lambda_i$, a measure of spectral similarity considered in [55]. Shading denotes one standard deviation about the mean for 8 runs of the algorithms. Note that our algorithm performs considerably better when applied to graphs with a geometric quality.

I Applications to graph visualization

Data visualization is an important (and aesthetically pleasing) application of graph reduction. As such, we generated videos of our algorithm reducing several real-world datasets. Figure SI 9 displays several stages of our algorithm applied to a temporal social network. A video of this reduction can be found [here](#); an application to an airport network (a case with both geometric and scale-free aspects) can be found [here](#); an application to the European road network can be found [here](#), and a reduction of a “hierarchical meta-graph” can be found [here](#)¹.

¹Explicit urls for the non-hyperlinked:
youtube.com/watch?v=qqLJc1VUML8; youtube.com/watch?v=tXUr6BRaEI;
youtube.com/watch?v=UVhT0y4Uae0; and youtube.com/watch?v=i3u4kkxMK40.

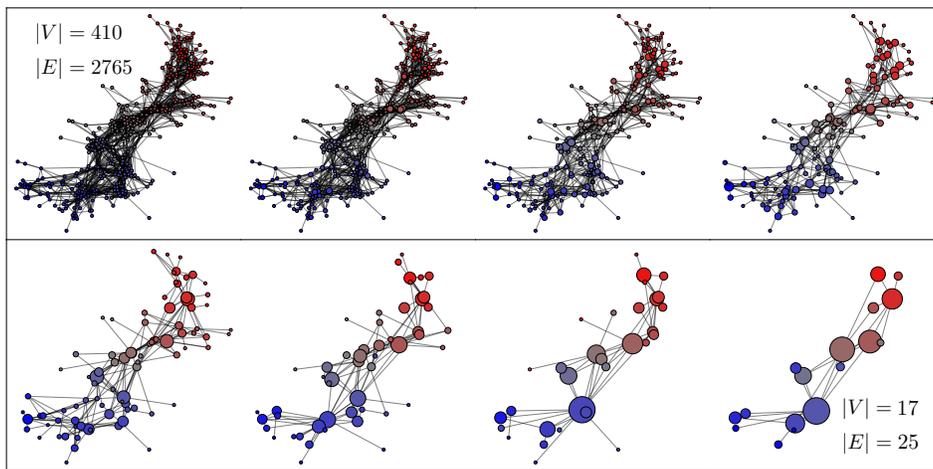


Figure SI 9: **Visualization of our graph reduction algorithm preserving global structure.** We applied our algorithm (prioritizing edge reduction, and allowing for deletion, contraction, and reweighting) to a weighted social network of face-to-face interactions during an exhibition on infectious diseases, with initial edge weights proportional to the number of interactions between pairs of people (410 nodes and 2765 edges) from [57]. Node color indicates the lowest nontrivial eigenvector of the reduced Laplacian, which in this case is aligned with the temporal direction. This graph displays a notable amount of hierarchical clustering (owing to its social nature), which is reflected in the reduced graphs. Eg, our algorithm begins by collapsing small, tightly-knit clusters of several people into one “supernode”, corresponding to groups of people who visited the exhibition together. A video of this reduction can be found [here](#).