When Edge Importance Disagrees: Resistance-Ranked versus Betweenness Pruning on Urban Road Networks

Jianru Shen
University of Montana
Missoula, Montana, USA
js258133@umconnect.umt.edu

Abstract

Sparsifying large road-network graphs by removing a fraction of their edges is a common way to reduce storage and accelerate spatial queries, but the choice of which edges to remove can quietly degrade the network that remains. We compare three edge-importance criteria, uniform random removal, low edge-betweenness removal, and low-effective-resistance removal, on 27 urban road networks drawn from OpenStreetMap. At edge-removal ratios from five to thirty percent we evaluate each criterion along three complementary axes: length-weighted routing utility, the retention of normalized algebraic connectivity on the largest connected component, and the fraction of an arterial backbone, taken from OpenStreetMap road classes, that each criterion destroys. The criteria disagree, and the three axes disagree about effective resistance itself. Deleting the edges with the lowest exact effective resistance preserves routing better than random deletion at every ratio, yet removes the arterial backbone faster than random deletion from five percent onward and holds an early connectivity advantage that fades to statistical parity by twenty-five percent. Betweenness removal is consistently the best of the three on all three axes, at every removal ratio. Deterministic deletion of low-resistance edges, which is distinct from the randomized, reweighted spectral sparsifier, is therefore not a safe proxy for routing importance when compressing real road networks, and a routing-oriented criterion is the more dependable choice.

CCS Concepts

• **Information systems** → **Geographic information systems**; • **Mathematics of computing** → *Graphs and surfaces*; *Spectra of graphs*; • **Theory of computation** → Sparsification and spanners.

Keywords

road networks, graph sparsification, effective resistance, betweenness centrality, OpenStreetMap

ACM Reference Format:
Jianru Shen. 2026. When Edge Importance Disagrees: Resistance-Ranked versus Betweenness Pruning on Urban Road Networks. In *The 34th ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '26), November 03–06, 2026, Riverside, CA, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3841645.3843396

SIGSPATIAL '26, Riverside, CA, USA

ACM ISBN 979-8-4007-2950-8/2026/11
https://doi.org/10.1145/3841645.3843396

1 Introduction

Road networks are naturally represented as graphs whose nodes are intersections and whose edges are road segments weighted by physical length. As these graphs grow to metropolitan scale, many spatial systems reduce their size by sparsification, removing a fraction of the edges to lower storage cost and to accelerate downstream tasks such as routing, indexing, and network analysis. Sparsification is attractive because a road network contains many edges that seem redundant, but whether the reduced graph still behaves like the original depends entirely on which edges are removed, a choice that has received far less scrutiny than the compression itself.

That choice rests on a notion of edge importance, and different notions need not agree. Ranking edges by effective resistance and deleting the lowest first removes the edges whose endpoints are most strongly connected by alternative routes; ranking by edge betweenness and deleting the lowest first instead removes the connections that carry the least shortest-path traffic. The first criterion is electrical, the second path-based, and on the nearly planar, strongly arterial graphs that describe road networks it is not obvious whether an edge judged redundant by one is also unimportant for the other. We test this empirically by comparing these two rankings and uniform random deletion on 27 urban road networks from OpenStreetMap, each sparsified at removal ratios from five to thirty percent and evaluated not by whether the criterion optimizes its own objective, but by whether the reduced network remains useful.

We make three contributions. First, we evaluate these criteria along three complementary axes, routing utility, normalized algebraic connectivity, and an externally defined arterial backbone, two of which do not reuse betweenness in their definitions. Second, we show that effective-resistance scoring is not a safe guide for compressing road networks: the three axes return three different verdicts on it, better than random deletion on routing, worse on the arterial backbone, and an early connectivity advantage that fades under moderate compression. Third, we show that low-betweenness deletion is the dependable default for compressing a road network while keeping it useful for routing, giving practitioners a concrete and easily applied recommendation.

2 Related Work

Effective resistance, which views the graph as an electrical network, is a classical edge-importance quantity [9] that is small when an edge's endpoints are joined by many parallel paths. It is central to spectral sparsification: sampling edges according to their leverage scores $w_e R_e$, the product of an edge's conductance w_e and its effective resistance R_e, and reweighting the survivors yields a sparsifier that provably approximates the Laplacian quadratic form, and hence the spectrum, of the original graph [1, 13]. That guarantee is a property of the randomized, reweighted construction. The deterministic heuristic studied here ranks edges by effective resistance and deletes the least-resistant ones outright, without reweighting; this is not the sparsifier of Spielman and Srivastava [13] and inherits none of its guarantees. A systematic comparison by Hamann et al. [8] evaluates such methods on social networks. Budgeted sparsification that certifies the preservation of persistent connectivity has also been proposed for web graphs [11]. Work on the metric backbone of weighted networks shows that sparsifiers which do not preserve it alter shortest-path structure, and reports effective-resistance sampling among the global methods that do not preserve all shortest paths [12]. These studies use social, biological, web, or synthetic networks and evaluate properties such as diameter, community structure, or spreading dynamics; none asks how common edge-removal criteria jointly affect routing utility and structural preservation on real road networks. Edge betweenness, the fraction of shortest paths through an edge [3, 6], underlies routing-aware reduction; spatial work has reduced road-network size mainly to accelerate queries [7] or for cartographic generalization, not to compare how alternative criteria affect the network that remains.

3 Experimental Setup

3.1 Road Networks

We study 27 urban road networks in the United States, obtained from OpenStreetMap with OSMnx [2]. For each city we download the drivable street network and reduce it to a simple undirected graph whose nodes are intersections and dead ends and whose edges carry the physical length, in meters, of the corresponding road segment. Parallel and bidirectional edges between the same pair of nodes are collapsed to a single edge keeping the shortest length, with the merged edge taking the highest OSM road class among them, and we retain the largest connected component so that every routing query is well defined on the original graph. The networks span sizes from 838 nodes (Burlington, VT) to 48,612 nodes (Phoenix, AZ), with a median of 11,462 nodes and a typical average degree near 2.9. To allow a check of whether the findings depend on urban form, the cities are chosen to cover four broad street-network morphologies: regular grids (for example Denver), historical cores (Savannah), geographically constrained layouts (Missoula), and sprawling networks (Phoenix).

Code and data to reproduce the main results are available at https://github.com/JianruShen/road-pruning-sigspatial.

3.2 Edge-Removal Criteria

Each criterion removes a target fraction of edges from a network, and we evaluate the graph that remains. The three criteria differ only in the order in which edges are removed. All edge scores are computed once on the original graph, and the resulting fixed ranking is used at every removal ratio.

Random removal deletes edges uniformly at random. Because a single draw is noisy, we aggregate every reported quantity over ten random seeds, combining per city before further analysis.

Betweenness removal deletes edges in increasing order of edge betweenness centrality [3, 6], removing the edges that lie on the fewest shortest paths first. Betweenness is computed on the unweighted topology, ignoring road length, so that any advantage of betweenness removal cannot be a tautological consequence of optimizing the length-weighted routing metric. Exact edge betweenness is too costly at these sizes, so we estimate it with Brandes's algorithm from 300 uniformly sampled source nodes per network. On Missoula, small enough to check exactly, the estimated ranking has Spearman correlation 0.96 with the exact one, and the lowest-30-percent edge sets overlap by 88 percent.

Effective-resistance removal deletes edges in increasing order of exact effective resistance, computed on the unweighted topology with unit conductance on every road segment, removing the most electrically redundant edges, those with many parallel alternatives, first. All effective resistances are computed exactly: grounding one node makes the reduced Laplacian nonsingular, and one sparse factorization then yields every resistance from a single solve per edge. The largest network, Phoenix with 67,489 edges, takes roughly 100 seconds on a laptop, and all 27 networks together take under seven minutes.

3.3 Removal Ratios and Evaluation

Each criterion is evaluated at edge-removal ratios of 5, 10, 15, 20, 25, and 30 percent, computed as a fraction of the original edge count. We evaluate each sparsified network along three axes, chosen to provide complementary evidence without directly reusing the pruning scores.

Routing utility. Shortest paths use physical road length as edge weight, so routing reflects geographic distance. Computing all pairs is infeasible on the larger networks, so we estimate routing quantities from a fixed sample of 500 source nodes per city, reused across all removal ratios. Writing $d_G(u, v)$ for the length-weighted shortest-path distance in graph G and S for the sampled sources, the global efficiency [10]

$$E(G) = \frac{1}{|S|\,(n-1)} \sum_{u \in S} \sum_{v \neq u} \frac{1}{d_G(u, v)} \tag{1}$$

averages the inverse shortest-path distance and stays well defined when some pairs become disconnected, contributing zero. We report the relative global-efficiency loss after sparsification, so that lower is better.

Normalized algebraic connectivity. As a topology-only measure of how well the global connectivity structure survives, we use normalized algebraic connectivity, defined as the second-smallest eigenvalue λ_2 of the unweighted normalized Laplacian $\mathcal{L} = I - D^{-1/2} A D^{-1/2}$ [4, 5], where A is the adjacency matrix and D the degree matrix. We compute λ_2 with sparse Lanczos iteration on the largest connected component of the sparsified graph, so that the measure reflects how well the surviving backbone is connected rather than collapsing to zero the moment the graph fragments; fragmentation is assessed separately through the disconnected-pair rate, the fraction of sampled source-target pairs left without a path. We report the retention $\lambda_2(G')/\lambda_2(G)$ relative to the original graph, so that higher is better and values below one indicate lost connectivity.

Arterial backbone. As a third axis, defined without using the evaluated graph scores, we use the road hierarchy that OpenStreetMap assigns to each segment. We define the arterial backbone of a city

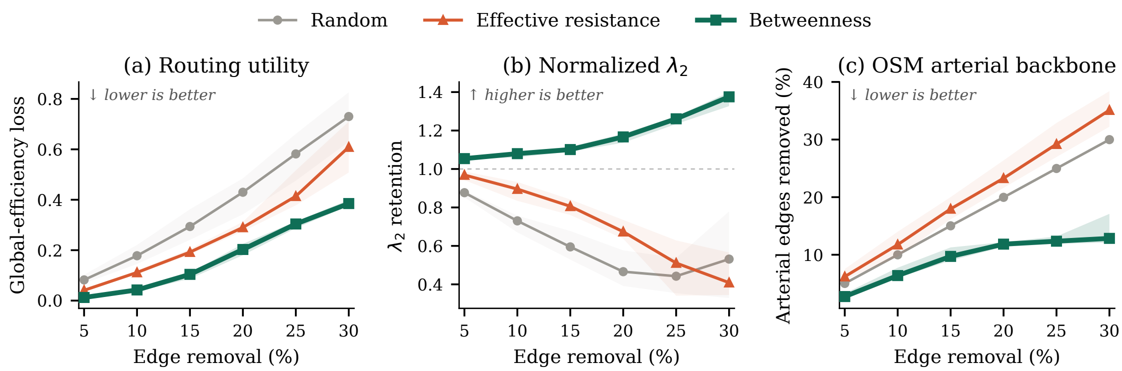

Figure 1: Preservation under increasing edge removal across 27 cities. Lines show city medians, bands bootstrap 95% confidence intervals; lower is better in (a, c), higher in (b). Betweenness removal performs best on all three axes throughout; resistance-ranked removal beats random deletion on routing but strips the arterial backbone faster than random deletion at every ratio. The upturn in (b) near 30% for random removal reflects fragmentation of the giant component, which raises the normalized λ_2 of the smaller surviving core.

as the edges whose OSM class is motorway, trunk, primary, or secondary, including their link ramps, and we measure the fraction of this backbone that each criterion removes. Because the backbone is defined externally by road class rather than by betweenness or any spectral quantity, it gives an external check of whether a criterion preserves higher-order roads intended to support through movement; lower removal is better.

4 Results

Across all 27 cities the three criteria separate cleanly, and the three axes return three different verdicts on resistance-ranked removal (Figure 1, median across cities with a bootstrap 95% confidence interval). We compare criteria with two-sided Wilcoxon signed-rank tests across the 27 cities at each ratio, treating each city as a single matched observation, which avoids the inflated significance that arises when repeated measurements are pooled. The group-median ranking is stable across urban form: betweenness removal is best in every one of the four morphology groups at every ratio on all three axes.

On routing utility (Figure 1a), betweenness removal is best at every ratio and by a widening margin: at 30 percent it loses a median of 0.385 of global efficiency, against 0.609 for resistance-ranked removal and 0.730 for random removal. Resistance-ranked removal beats random deletion in 22 to 25 of 27 cities at every ratio with $p < 10^{-5}$, but it trails betweenness removal everywhere: at 30 percent betweenness removal outperforms resistance-ranked removal in all 27 cities and random removal in all 27 (both $p < 10^{-7}$).

On normalized algebraic connectivity (Figure 1b), a topology-only measure complementary to routing, the early advantage of resistance-ranked removal decays as compression grows. Under light removal it preserves connectivity better than random deletion, as its redundancy intuition predicts: at 5 percent its median λ_2 retention is 0.969, above random at 0.877, and it is better in 23 of 27 cities with $p < 10^{-2}$. The advantage shrinks monotonically. Random deletion overtakes it in 4 of 27 cities at 5 percent and 13 of 27 at 30 percent, and from 25 percent onward the paired difference is statistically indistinguishable from zero with $p > 0.6$, although the medians still differ, 0.409 against 0.531 at 30 percent. Betweenness removal, in contrast, retains the normalized λ_2 of its largest component above that of the original graph throughout, rising from 1.05 to 1.38 as it sheds weakly connected peripheral edges; it exceeds resistance-ranked removal in 25 of 27 cities at 30 percent and random in 23 of 27 (both $p < 10^{-3}$). Fragmentation separates random deletion from both ranked criteria: at 30 percent removal its median disconnected-pair rate is 75 percent, against 37 for betweenness and 48 for resistance-ranked removal, and the two ranked criteria are statistically tied at 20 and 25 percent.

On the arterial backbone (Figure 1c), defined entirely outside the graph by road class, the verdict is the harshest and arrives earliest. Resistance-ranked removal deletes a larger share of the backbone than random deletion from the lightest ratio we test: at 5 percent it has already removed a median of 6.2 percent of the backbone against 5.0 for random, and it removes more in 19 to 22 of 27 cities at every ratio with $p < 10^{-2}$. At 30 percent it deletes a median of 35.1 percent of the backbone against 30.0 for random, while betweenness removal deletes only 12.8 percent and preserves the backbone better than resistance-ranked removal in 26 of 27 cities and better than random in all 27 (both $p < 10^{-7}$). One structural fact explains all three verdicts, visible directly on the map (Figure 2): many low-resistance edges lie in densely interconnected arterial regions rather than on peripheral streets, so deleting them thins redundancy inside the arterial core. Routing tolerates this because parallel arterials substitute for one another, which is why resistance-ranked removal still beats random deletion; connectivity tolerates it only briefly, with the advantage draining away as the core thins; the road-class axis does not tolerate it at all. Betweenness removal does the opposite, shedding the peripheral edges that few shortest paths use and leaving the arterial skeleton, and with it routing utility and connectivity, largely intact.

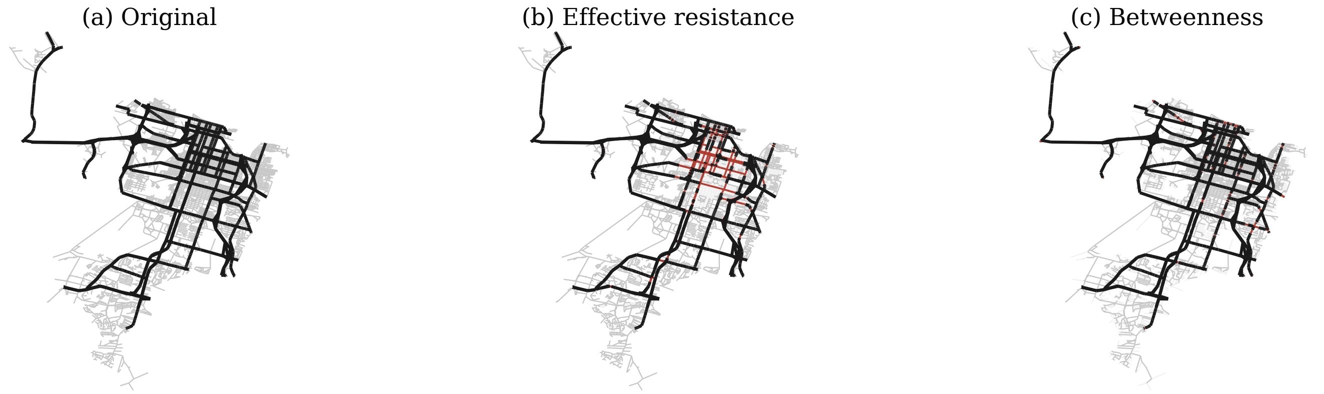

Figure 2: Savannah after 30% edge removal. Effective-resistance removal deletes 32% of the OSM arterial edges and disrupts the through-road skeleton; betweenness removal deletes 12% and largely preserves it. Arterial roads dark, minor roads pale, removed arterial edges dark red, other removed edges nearly invisible.

5 Discussion

Practical recommendation. For practitioners the recommendation is direct but comparative. When the goal is to compress an urban road network while keeping it useful for routing, low-betweenness removal is the dependable default among the criteria tested: across 27 cities it performs best on all three evaluated axes at every compression ratio, and it is the only criterion that largely spares the arterial backbone. It is not lossless: at 30 percent removal even betweenness removal loses 0.385 of global efficiency and disconnects 37 percent of sampled pairs, so aggressive compression needs task-specific validation. Deterministic low-effective-resistance removal is not a reliable alternative. It keeps more routing utility than random deletion but removes arterial edges at a higher rate throughout, and its connectivity advantage is gone by twenty-five percent.

Interpretation. The broader lesson is that edge importance is not a single quantity: the same criterion can beat random deletion on one axis and lose to it on another. A score capturing one kind of importance, here the electrical redundancy of effective resistance, can mislead for a different purpose, here preserving the arterial structure that carries through traffic; in these networks, resistance-based redundancy systematically assigns low importance to edges of high road class. An edge score should be validated against the task one actually cares about rather than trusted because it is principled in general.

Limitations. Several limitations bound these conclusions. First, we evaluate the deterministic heuristic of deleting low-effective-resistance edges, which we study as a practical heuristic rather than the randomized, reweighted spectral sparsifier; our results speak to that heuristic as an edge-importance criterion and not to the guarantees of the sparsifier, and how a fully randomized construction would behave on road networks is a separate question. Second, effective resistance uses unit conductance; inverse-length or travel-time conductance may rank long arterial segments differently. Third, road-class pruning, degree-two contraction, spanners, and other centrality variants remain untested baselines. Fourth, routing utility is naturally aligned with a shortest-path criterion, but the other two axes do not reuse betweenness and show the effect is not an artifact of how routing is scored. Fifth, we study drivable United States networks; other network types and regions remain open. A structural account of when electrical redundancy and road-class importance diverge, and of the city-to-city variation, is the natural next step.

References

[1] Joshua Batson, Daniel A. Spielman, Nikhil Srivastava, and Shang-Hua Teng. 2013. Spectral Sparsification of Graphs: Theory and Algorithms. *Commun. ACM* 56, 8 (2013), 87–94. doi:10.1145/2492007.2492029

[2] Geoff Boeing. 2017. OSMnx: New Methods for Acquiring, Constructing, Analyzing, and Visualizing Complex Street Networks. *Computers, Environment and Urban Systems* 65 (2017), 126–139. doi:10.1016/j.compenvurbsys.2017.05.004

[3] Ulrik Brandes. 2001. A Faster Algorithm for Betweenness Centrality. *Journal of Mathematical Sociology* 25, 2 (2001), 163–177. doi:10.1080/0022250X.2001.9990249

[4] Fan R. K. Chung. 1997. *Spectral Graph Theory.* American Mathematical Society.

[5] Miroslav Fiedler. 1973. Algebraic Connectivity of Graphs. *Czechoslovak Mathematical Journal* 23 (1973), 298–305.

[6] Linton C. Freeman. 1977. A Set of Measures of Centrality Based on Betweenness. *Sociometry* 40, 1 (1977), 35–41. doi:10.2307/3033543

[7] Robert Geisberger, Peter Sanders, Dominik Schultes, and Daniel Delling. 2008. Contraction Hierarchies: Faster and Simpler Hierarchical Routing in Road Networks. In *Proceedings of the 7th International Workshop on Experimental Algorithms (WEA)*. Springer, 319–333. doi:10.1007/978-3-540-68552-4_24

[8] Michael Hamann, Gerd Lindner, Henning Meyerhenke, Christian L. Staudt, and Dorothea Wagner. 2016. Structure-Preserving Sparsification Methods for Social Networks. *Social Network Analysis and Mining* 6, 1 (2016), 22. doi:10.1007/s13278-016-0332-2

[9] Douglas J. Klein and Milan Randić. 1993. Resistance Distance. *Journal of Mathematical Chemistry* 12, 1 (1993), 81–95. doi:10.1007/BF01164627

[10] Vito Latora and Massimo Marchiori. 2001. Efficient Behavior of Small-World Networks. *Physical Review Letters* 87, 19 (2001), 198701. doi:10.1103/PhysRevLett.87.198701

[11] Jianru Shen. 2026. TopoBudget: Persistent-Connectivity-Preserving Web Graph Sparsification for Reusable Community Analytics. Accepted at WISE 2026. arXiv:2608.07903 [cs.SI]

[12] David Soriano-Paños, Felipe Xavier Costa, and Luis M. Rocha. 2025. Quantifying Edge Relevance for Epidemic Spreading via the Semi-Metric Topology of Complex Networks. *Journal of Physics: Complexity* 6, 3 (2025), 035005. doi:10.1088/2632-072X/adf2ed

[13] Daniel A. Spielman and Nikhil Srivastava. 2011. Graph Sparsification by Effective Resistances. *SIAM J. Comput.* 40, 6 (2011), 1913–1926. doi:10.1137/080734029