Greedy Routing in Urban Road Networks

Miles Aronow, Saray Shai
Computer Science, Wesleyan University, Middletown, CT

Introduction

- Navigation through urban road networks is important for the transportation of people and goods.
- Greedy Routing is a form of navigation from a source to target, where at each iteration, the route proceeds to the neighbor of the current node that is "closest" to the target.
- We investigated two methods of determining the "closest" neighbor as shown below.

Figure 1. Illustration of Greedy Routing

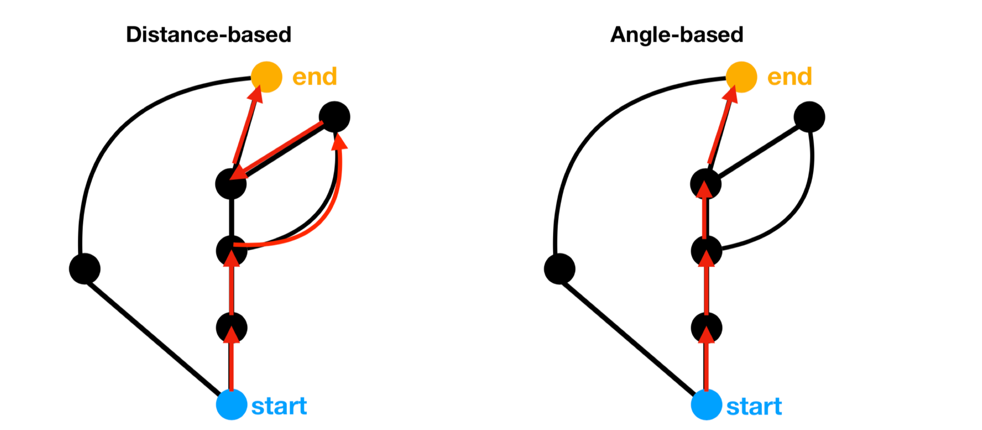

A. Illustration of greedy routing. Distance-based routing (left): proceed to the neighbor that is closest to the target (in Euclidean distance). Angle-based routing (right): proceed to the neighbor that is closest to a straight line between the start and end points.

- Greedy Routing is not guaranteed to successfully reach the target. If at any step, the "closest" neighbor has already been visited, a loop is created and the path fails.
- A path may also be successful, but much longer than the shortest path
- We also studied a modified algorithm where the "closest" neighbor was chosen only from nodes that had not been visited before. This failed if there were no such neighbors
- Greedy routing is useful in establishing the connection between network structure and space since it is expected to perform well exactly when the notion of "closest" in space aligns well with network proximity.
- Understanding the structure of urban transportation networks has applications in reducing congestion and increasing accessibility.

Research Questions

- How successful is greedy routing in urban networks?
- What factors affect the success of greedy routing?
- Where do paths fail and succeed?

Methods

- Data from Open Street Map using the python package osmnx
- Networks consist of approximately 10,000 nodes around the city center
- For each city we randomly sampled source nodes from the network.
- For each source in the sample, we tried greedy routing to every other node in the network and recorded for each source target pair:
  - If GR was successful
  - The length of the greedy path(infinity for failed paths)
  - the length of the shortest path from the source to the target
  - the Euclidean distance(distance of straight line) between source and target
  - The nodes along the greedy path
- We measured Success Rates and GR score defined as the average ratio of shortest path length over greedy path length.
- In this poster only success rate is discussed since the measures were similar.
- A high GR score shows that GR is not only successful but efficient

References

Ruehrup, Stefan (2009). Liu; Chu; Leung (eds.). Theory and Practice of Geographic Routing
Seguin C, van den Heuvel MP, Zalesky A (2018) Navigation of brain networks. Proc Natl Acad Sci
Muscoloni A, Thomas JM, Ciucci S, Bianconi G, Cannistraci CV (2017) Machine learning meets complex networks via coalescent embedding in the hyperbolic space. Nat Commun

Results

The success of greedy routing varies by city. Performance is best in organized grid-like places such as Manhattan and Buenos Aires. Success depends on an array of factors such as a network's 2D-ness, road density, and the existence of one way and dead end roads. By simplifying networks and modifying the algorithm, performance improves in all cities investigated. Furthermore, random rewiring decreases performance suggesting the existing network structure allows for greedy routing to be successful

What factors affect success rate?

Figure 2(a). Stress vs. Success Rate (b). Density(Total Road Length/Area) vs Success Rate

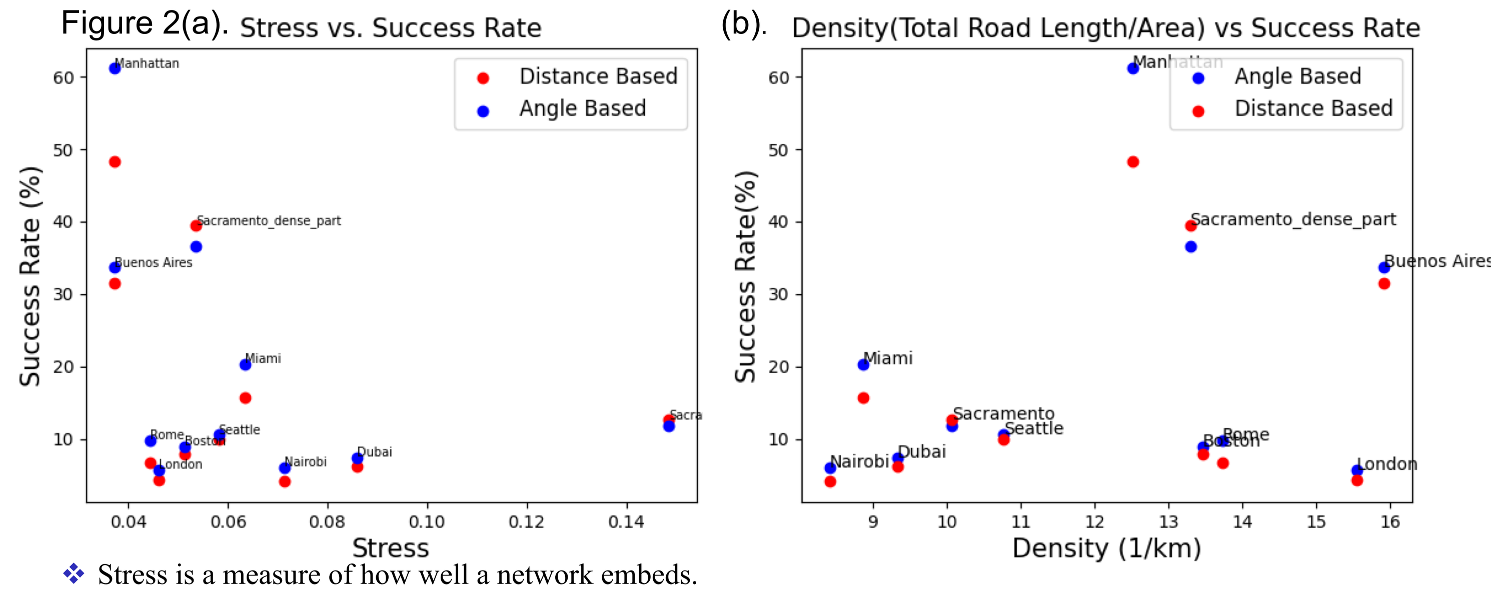

- Stress is a measure of how well a network embeds.
- In this case, the network is embedded in the 2-dimensional Euclidean plane. The lower the stress, the more "2D" a network is.
- We see some correlation between lower stress and higher success rate.
- We hypothesized that greedy routing would be more successful in denser networks.
- There is some correlation, but cities like Rome, Boston, and London demonstrate that cities can be dense and hard to navigate with greedy routing.

Figure 3. Greedy Routing Success Rate Vs. Euclidean Distance

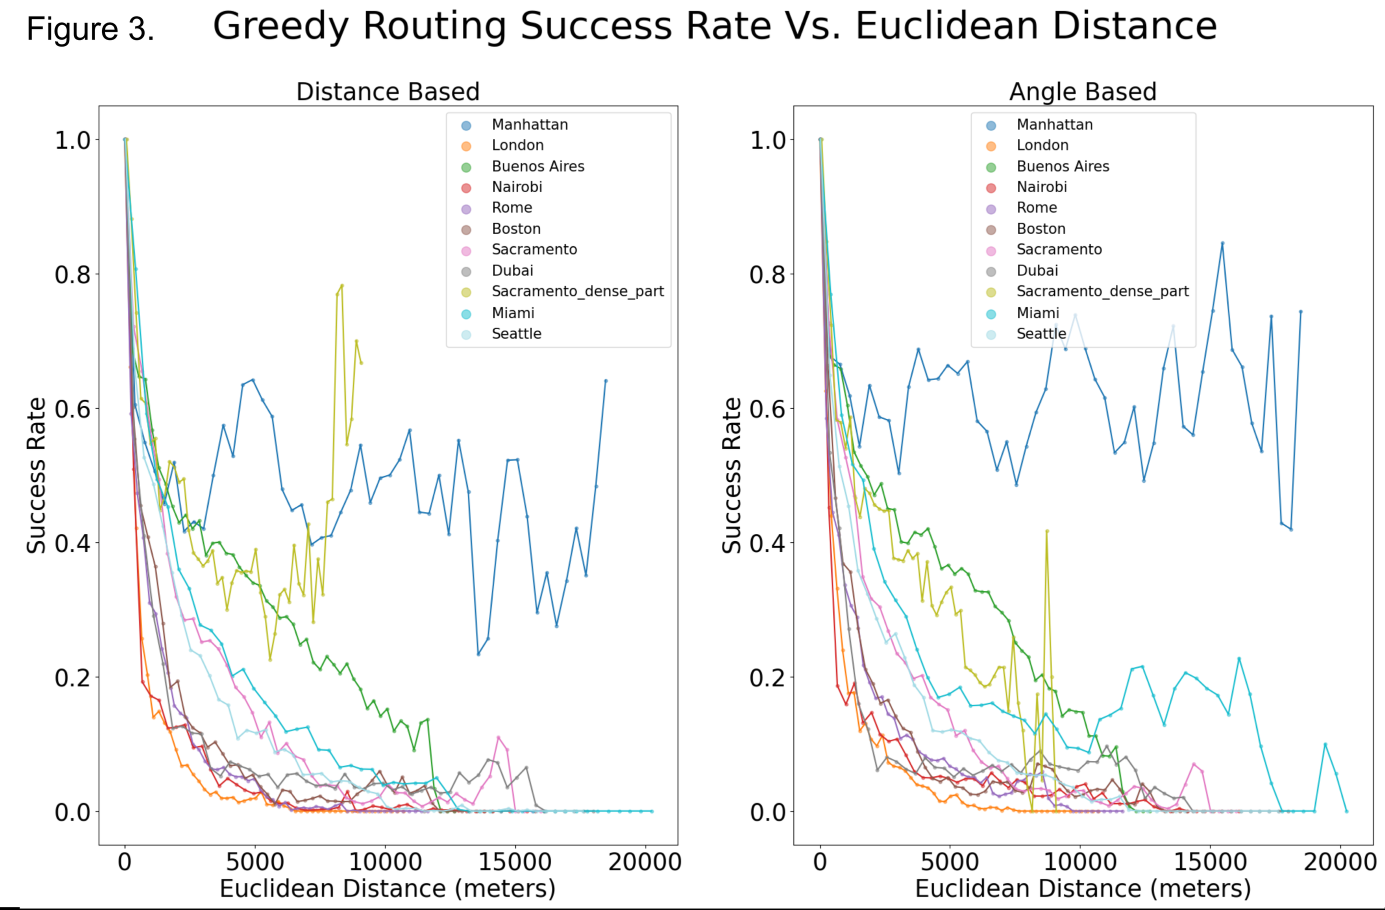

How Successful is Greedy Routing?

Figure 4. Success Rate by Method

	Normal Approach		One Way Roads Converted to Two Way		Alternate Algorithm to Avoid Already Visited Nodes		Dead End Roads Removed		All Three Changes	
	Distance	Angle	Distance	Angle	Distance	Angle	Distance	Angle	Distance	Angle
Manhattan	0.48	0.61	0.79	0.8	0.74	0.79	0.49	0.64	0.97	0.97
London	0.04	0.06	0.07	0.11	0.08	0.09	0.08	0.11	0.47	0.53
Buenos Aires	0.31	0.34	0.54	0.50	0.46	0.48	0.41	0.42	0.93	0.93
Nairobi	0.04	0.06	0.04	0.06	0.07	0.09	0.09	0.13	0.32	0.35
Rome	0.07	0.10	0.13	0.15	0.12	0.15	0.11	0.16	0.56	0.63
Boston	0.08	0.09	0.13	0.16	0.11	0.13	0.14	0.15	0.50	0.59
Sacramento	0.13	0.12	0.11	0.13	0.16	0.16	0.19	0.17	0.39	0.41
Dubai	0.06	0.07	0.06	0.05	0.11	0.1	0.08	0.09	0.28	0.30
Sacramento Subsection	0.39	0.37	0.39	0.37	0.48	0.48	0.51	0.51	0.86	0.87
Miami	0.16	0.2	0.21	0.19	0.26	0.30	0.26	0.35	0.71	0.71
Seattle	0.1	0.11	0.12	0.10	0.14	0.15	0.14	0.16	0.35	0.34

- Figure 4 shows the success rates for the various approaches.
- Under the original approach GR was only successful in the very grid-like cities of Manhattan and Buenos Aires.
- Converting one way roads to two way led to large increases for the success rate in the already succeeding cities, but less so in others.
- Similar results were seen with the alternate algorithm and, with the removal of dead end roads.
- Interestingly, combining the three methods led to significant increases in the success rates for all cities.
- Perhaps this supports the claim that greedy routing success depends on an array of factors. Additionally, the angle method(see figure 1) tended to preform the same or slightly better than the distance method.
- In a different experiment we looked at the effect of randomly rewiring edges
- After rewiring 10% of edges success rate fell from 80% to 22% for Manhattan and 54% to 43% for Buenos Aires
- After 100% rewiring success rate fell to less than .1% for both cities.
- This suggest that network structure contributes to the success of greedy routing.

Where do paths fail and succeed?

Figure 5. Examples of failed paths.
Source and target are shown in green. The spot of failure is shown in orange

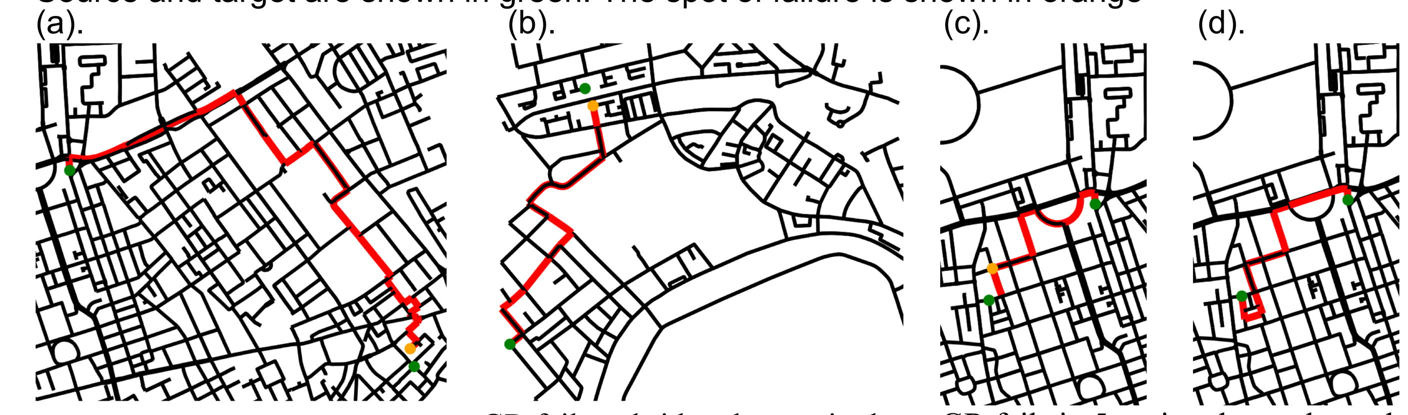

(a) GR navigates almost to the target, but fails at a dead end. (b) GR fails to bridge the gap in the network. (c)–(d) GR fails in 5c using the angle method when it reaches a one way road. Using the distance method in 5d, it maneuvers around.

(3) Similarly to b, GR can't cross a gap in the network. (f) GR doesn't even get close. (g)–(h) GR fails in the first case at a dead end using the distance method. Using the angle method it succeeds.

Figure 6. Visualization of successful targets for a single source

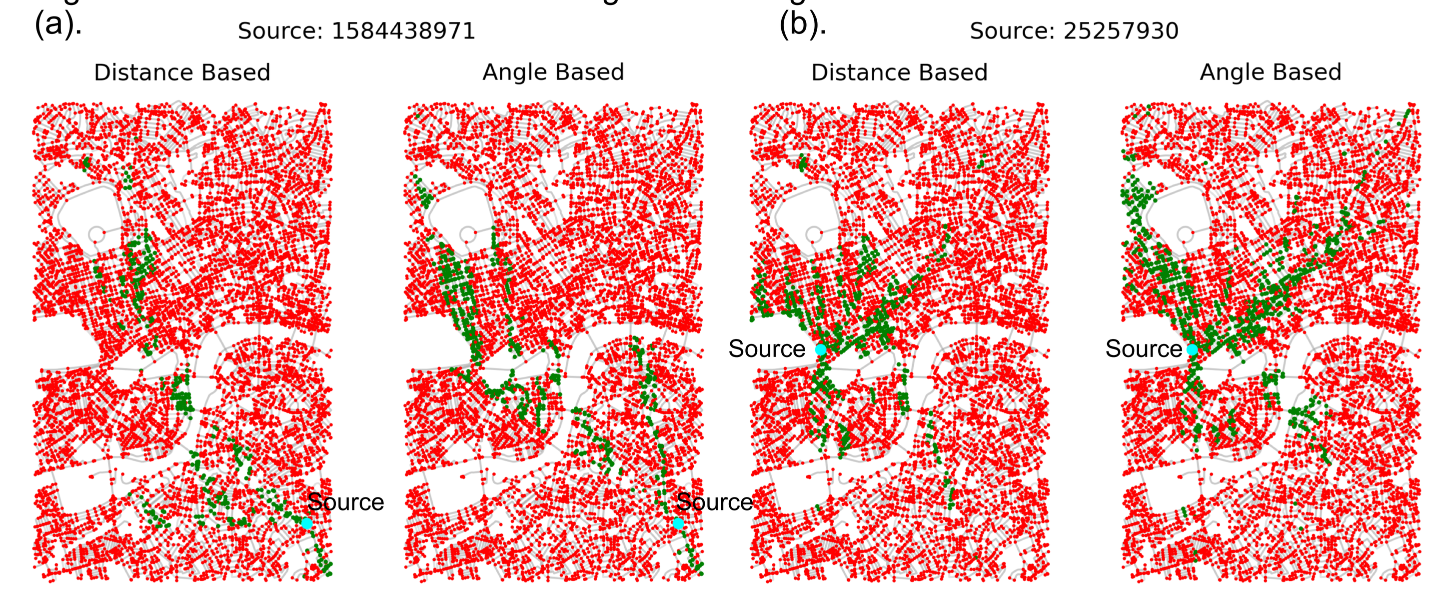

- Greedy routing was run from the source, shown in blue, to all other nodes in the London network.
- The nodes for which greedy routing was successful are shown in green, the nodes for which it failed are shown in red.
- Obstacles, such as the river and park can be seen to cause failures often, but not in all cases.
- Success is correlated with distance as previously seen in Figure 3

Figure 7. Visualizations of Failures

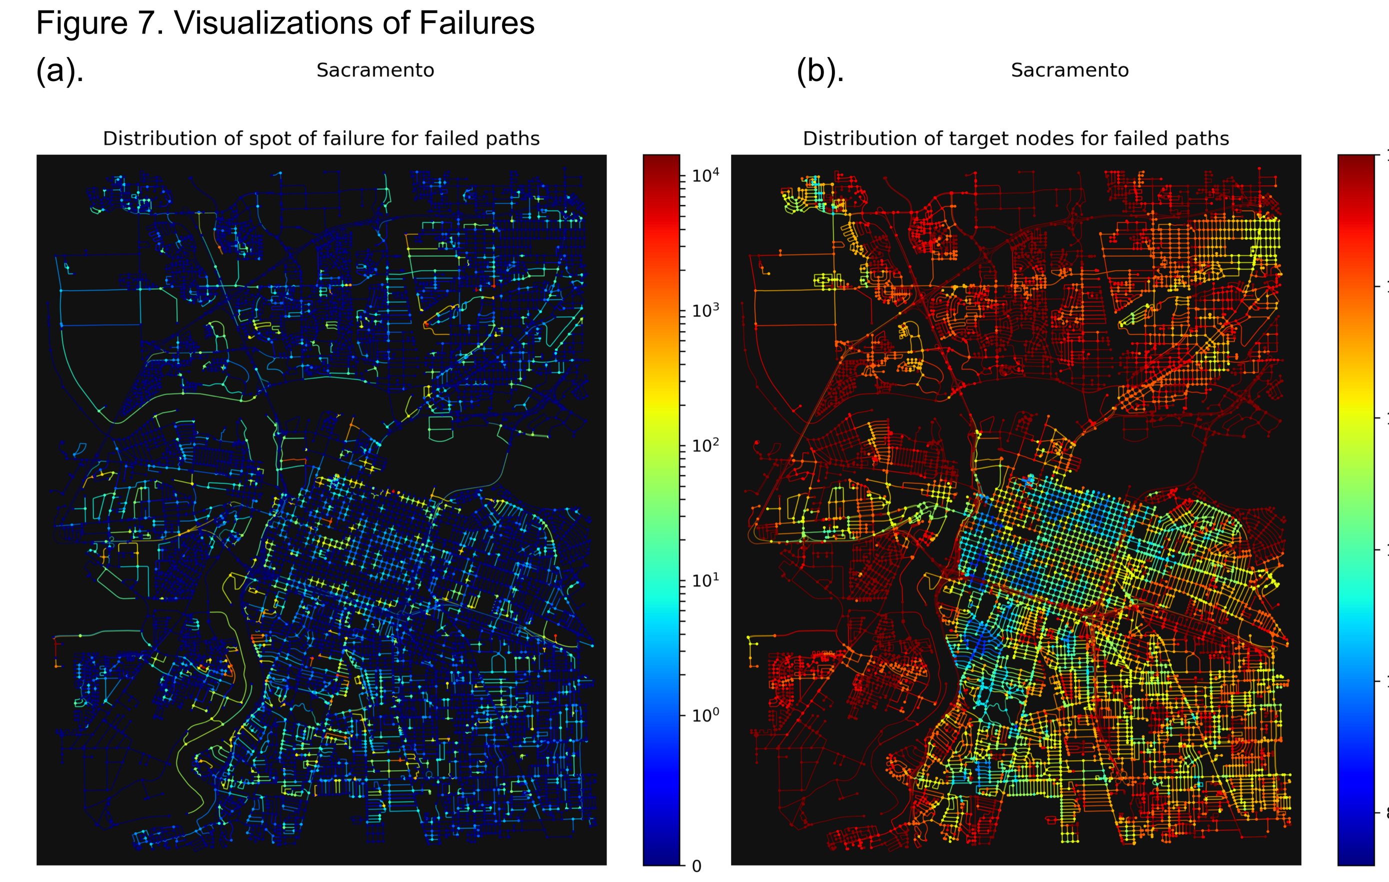

- Recall Greedy routing fails when the "closest" neighbor has already been visited.
- In Figure 7a, nodes are colored corresponding to the number of times the node was a spot of failure across the entire sample.
- Edges are colored corresponding to the average of the two nodes it connects.
- Failures often occur along obstacles such as the river.
- In Figure 7b nodes are colored corresponding to the number of times the node was the target of a failed path.
- You can observe that routing was far more successful to the southeast portion of the city.
- We hypothesized this may be due to density, grid structure, and lack of obstacles.