
An Improved A* Search Algorithm for Road Networks Using New Heuristic Estimation

2021 Summer STEM Institute

Kevin Chen

Abstract

Finding the shortest path between two points in a graph is a fundamental problem that has been well-studied over the past several decades. Shortest path algorithms are commonly applied to modern navigation systems, so our study aims to improve the efficiency of an existing algorithm on large-scale Euclidean networks. The current literature lacks a deep understanding of certain algorithms' performance on these types of networks. Therefore, we incorporate a new heuristic function, called the k -step look-ahead, into the A search algorithm and conduct a computational experiment to evaluate and compare the results on road networks of varying sizes. Our main findings are that this new heuristic yields a significant improvement in runtime particularly for larger networks when compared to standard A*, as well as that a higher value of k is needed to achieve optimal efficiency as network size increases. Future research can build on this work by implementing a program that automatically chooses an optimal k value given an input network. The results of this study can be applied to GPS routing technologies or other navigation devices to speed up the time needed to find the shortest path from an origin to a destination, an essential objective in our daily lives.*

1. Introduction

Finding the shortest path from an origin to a destination is a fundamental, well-studied problem faced by researchers and mathematicians for at least the past century. Even in primitive societies, finding shortest paths was a com-

mon, essential objective. However, extensive research in the shortest path problem began relatively late as compared to other combinatorial optimization problems, likely due to its apparent simplicity (Schrijver, 2012). In recent years, these algorithms have been applied to various modern navigation systems such as GPS routing technologies on road networks.

The problem is typically represented by a graph consisting of nodes and edges, where the goal is to find the shortest path between two given nodes such that the sum of the weights of the traversed edges is minimized. When applying this theoretical model to road networks, nodes often represent intersections and edges correspond to road segments that are weighted by the length of the segment. Perhaps the most well-known shortest path algorithm was presented by Dutch computer scientist Edsger W. Dijkstra in 1959 (Dijkstra & others, 1959), which remains one of the most efficient methods to this day (Zeng & Church, 2009). Dijkstra's algorithm finds the optimal solution to the shortest path problem with nonnegative edge weights by expanding neighboring nodes uniformly in all directions. Another efficient but less-studied technique is the A* (pronounced A-star) algorithm (Hart et al., 1968), which builds off of Dijkstra's by incorporating a heuristic function that takes into consideration the estimated remaining distance from the current node to the destination, thus limiting the search area.

The heuristic function of the A* search algorithm is a critical part of its implementation because it highly influences the algorithm's runtime and solution optimality (Fu et al., 2006). While shortest path problems in general are a well-studied topic, past studies notably lack a detailed analysis on how the heuristic estimating method of A* affects its computational performance on real-life networks of various sizes.

Therefore, the primary objective of this paper is to introduce a new heuristic for A* that, to our knowledge, has not before been implemented and analyzed in the environment of real road networks. This new *look-ahead heuristic* involves the idea of deciding which node in the graph to expand next based on the likelihood that one of the node's

*Research project conducted during the 2021 Summer STEM Institute. The Summer STEM Institute (SSI) is a six-week virtual summer program where high school students learn how to design and conduct data science and other computational research projects. To view other SSI distinguished projects, please visit www.summersteminstitute.org.

children is on the shortest path. We will then evaluate the computational efficiency of this new version of A* and compare it to the standard A* and Dijkstra’s algorithm by testing them on road networks of three different sizes. The results of this experiment can be used in the future to determine which of these algorithms is the most optimal to use on Euclidean networks under different circumstances.

2. Literature Review

While a plethora of past studies have dealt with efforts to compare shortest path algorithms, many of them evaluate their computational performance on randomly generated graphs as opposed to real networks (Cherkassky et al., 1996; Biswas et al., 2005; Magzhan & Jani, 2013). This is problematic when it comes to real-life applications because synthetic graphs lack geographical reality and the most efficient shortest path algorithms on random networks may not practically be ideal when used on empirical systems (Zeng & Church, 2009). Furthermore, these studies exclude the A* algorithm in their comparisons, but this decision is somewhat reasonable because A* functions best in the context of Euclidean space instead of randomized networks. While Zhan and Noon (Zhan & Noon, 1998) examine the performance of multiple shortest path algorithms on several large-scale U.S. road networks, they do not mention A*. Additionally, Chan et al. (Chan et al., 2016) compare the efficiency of six different algorithms including A*, but do so on a relatively small bus route that is not representative of a typical large transportation network.

On the other hand, other previous experiments carry out direct comparisons between A* and Dijkstra’s algorithm on complex road networks, but do so without varying certain features of the algorithms (Rachmawati & Gustin, 2020; Jain, 2021). They generally find the standard A* to discover the shortest path faster in those environments. Zeng and Church (Zeng & Church, 2009) analyze the effects of varying the type of data structure used to maintain the set of candidate nodes (e.g. k -array heaps, double buckets, approximate buckets) in A* and Dijkstra’s. Moreover, in the previously mentioned papers that study the application of shortest path algorithms to road networks, Euclidean distance is the most commonly used heuristic function in A*, as it ideally describes the lower bound of the true cost from a node to the destination in Euclidean space. Some more studies attempt to incorporate an “overdo” heuristic into A* that sacrifices solution optimality for a shorter runtime (Engineer, 2001; Jacob et al., 1999). Finally, others explore a variation of A* known as the ALT algorithm, which uses a different lower-bounding technique based on landmarks and the triangle inequality as opposed to Euclidean bounds (Goldberg & Harrelson, 2005; Goldberg et al., 2006).

To fill in the aforementioned gaps in the current literature,

our main contribution is to improve A*’s performance on large road networks by integrating a heuristic function that has the potential to decrease runtime while guaranteeing the optimal path.

3. Purpose

As alluded to earlier, the goal of this study is to improve upon the performance of the A* search algorithm on Euclidean road networks by implementing a k -step look-ahead heuristic. Specifically, we break this problem down into the following steps:

1. We aim to evaluate the computational performance of Dijkstra’s algorithm, the standard A* algorithm, and the A* algorithm with a k -step look-ahead heuristic function on three real road networks of different sizes. We compare the efficiency of these algorithms in terms of program runtime.
2. We will then use the results of the computational experiment to determine which of the algorithms is best to use for different sizes of the shortest path problem on Euclidean networks.

4. Shortest Path Algorithms

In the past, scientists and mathematicians have derived various algorithms for solving the shortest path problem. In this paper, we focus on evaluating Dijkstra’s algorithm and the A* algorithm, and we provide an overview of each of these methods in the following sections.

4.1. Notation and Definitions

In this section, we discuss the notation necessary for understanding graphs and shortest path algorithms. Most of the notation introduced here will be used throughout section 4. We then rigorously define the shortest path problem.

Let $G = (V, E)$ be a *directed*, *undirected*, or *mixed graph* consisting of a set of *vertices* (i.e. *nodes*) V and a set of *edges* E . We denote $e_{x,y} \in E$ as the edge from node x to node y , and $c(x, y)$ as the *cost* (i.e. *weight*) of that edge. We also say that a node n' is *adjacent* to node n if there exists an edge $e_{n,n'}$. In the context of shortest path algorithms, we let $d(n)$ be the total cost of the shortest path from the source node to node n . We also call this value the *distance value* of node n . We let the *source node* of the shortest path problem be s and the *target node* be t . We also have $\text{dist}(x, y)$ represent the Euclidean (i.e. straight-line) distance between nodes x and y .

Next, let

$$P = ((n_1, n_2), (n_2, n_3), \dots, (n_{m-1}, n_m))$$

be a *path* of edges connecting node n_1 to n_m where each (n_i, n_{i+1}) is an edge of the path connecting adjacent nodes n_i and n_{i+1} for all $1 \leq i < m$. We say that a path P is *valid* if and only if $y_j = x_{j+1}$ for all $1 \leq j < m - 1$, where (x_j, y_j) and (x_{j+1}, y_{j+1}) are the j th and $(j + 1)$ th elements of P , respectively. We define a path to be *acyclic* if the path does not pass through any same node more than once. We call node n' a *successor* of n if n' comes after n in a given path. In addition, a path between node n and n' is said to have *length* k if the path passes through exactly k edges (note that this is a different quantity from total cost of the path).

The goal of the *single-pair shortest path problem* is to find the path P between s and t , where $n_1 = s$ and $n_m = t$, that minimizes the sum

$$\sum_{i=1}^{m-1} c(n_i, n_{i+1})$$

over all possible values of m . Other common variations of the shortest path problem exist besides the single-pair problem. These include the single-source problem, where the objective is to find the shortest paths from one node to all other nodes, and the all-pairs problem, where a shortest path must be found between all pairs of nodes in the graph (Madkour et al., 2017).

4.2. Dijkstra's Algorithm

First introduced by Dijkstra in 1959 (Dijkstra & others, 1959), this algorithm finds the optimal path on a graph with nonnegative weights from a single source node to a target or set of targets. While the algorithm is more commonly used to solve the single-source shortest path problem, it will be implemented in this study to handle a single pair of nodes for the purpose of comparing its computational efficiency with that of A*.

Dijkstra's method, as outlined in Algorithm 1, utilizes a min-priority queue to store the set of unvisited nodes, leading to more efficient computations compared to other basic data structures (Zeng & Church, 2009). Since our work focuses on finding the shortest path from an origin to a particular destination, the algorithm terminates when the target node t is marked as visited. Asymptotically, the algorithm has a worst-case time complexity of

$$O((|V| + |E|) \log |V|)$$

when a standard binary heap is implemented as the priority queue data structure (Cormen et al., 2009).

4.3. A* Algorithm

The A* search algorithm is a best-first search method (Russell & Norvig, 2010) and has a general procedure similar

Algorithm 1 Dijkstra's Algorithm

Dijkstra(G, s, t):

Input: A graph $G = (V, E)$, start node $s \in V$ and target node $t \in V$.

Output: The shortest path of edges from node s to node t .

- 1: **unvisitedQueue** := priority queue of unvisited nodes sorted by d value
 - 2: **previous** := an empty map
 - 3: $d(n) := \infty$ for all nodes n in G
 - 4: $d(s) := 0$
 - 5: **while** unvisitedQueue not empty and t not visited **do**
 - 6: Extract node i from unvisitedQueue
 - 7: **for all** j adjacent to i **do**
 - 8: **if** j is visited **then**
 - 9: Continue
 - 10: **dTemp** := $d(i) + c(i, j)$
 - 11: **if** **dTemp** < $d(j)$ **then**
 - 12: $d(j) := \mathbf{dTemp}$
 - 13: **previous**[j] := i
 - 14: Reorder unvisitedQueue
-

to Dijkstra's, but the difference comes with the concept of a heuristic function that estimates the remaining distance from node n to the target node t . This crucial feature causes A* to run faster than Dijkstra's in many past computational experiments (Zeng & Church, 2009; Rachmawati & Gustin, 2020; Engineer, 2001). A* functions by choosing which path to extend based on the node that minimizes the value of

$$f(n) = g(n) + h(n)$$

where n is the next node in the path, $g(n)$ is the current distance value of n , and $h(n)$ is the predetermined heuristic value of n . Using this heuristic, A* is able to concentrate its search in the direction of the target by prioritizing nodes with a lower h value (i.e. nodes that are estimated to be closer to the destination). This main distinction between A* and Dijkstra's can be seen in Algorithm 2.

The most commonly used heuristic estimator when dealing with real networks is Euclidean distance, as this metric always describes the lowest bound of the distance between any two points in space (Sharma & Pal, 2015; Schultes, 2008). Euclidean distance is an example of an *admissible heuristic*, which is defined as a heuristic that never overestimates the true lowest possible cost from a node n to t , thereby guaranteeing that the shortest path to t is not overlooked. Additionally, a heuristic is considered *consistent* (or *monotone*) if it satisfies

$$h(n) \leq c(n, p) + h(p) \text{ and } h(t) = 0,$$

where p is a successor of any node n in the graph. In the context of A*, this means that the cost by which a node

Algorithm 2 A* Search Algorithm

 $A^*(G, s, t, h)$:

Input: A graph $G = (V, E)$, start node $s \in V$, target node $t \in V$, and heuristic function $h(n)$.

Output: The shortest path of edges from node s to node t .

```

1: unvisitedQueue := priority queue of unvisited nodes
   sorted by  $f$  value
2: previous := an empty map
3:  $f(n) := \infty$  for all nodes  $n$  in  $G$ 
4:  $g(n) := \infty$  for all nodes  $n$  in  $G$ 
5:  $f(s) := h(s)$ 
6:  $g(s) := 0$ 
7: while unvisitedQueue not empty and  $t$  not visited do
8:   Extract node  $i$  from unvisitedQueue
9:   for all  $j$  adjacent to  $i$  do
10:    if  $j$  is visited then
11:      Continue
12:     $gTemp := g(i) + c(i, j)$ 
13:    if  $gTemp < g(j)$  then
14:       $g(j) := gTemp$ 
15:       $f(j) := g(j) + h(j)$ 
16:       $previous[j] := i$ 
17:   Reorder unvisitedQueue
    
```

is reached must be the lowest possible once that node is expanded, given that the graph only includes nonnegative edge weights. When the heuristic is simply Euclidean distance, this property of consistency is analogous to the triangle inequality because it essentially states that the straight-line distance between n and t is never greater than the sum of the distance between n and p plus the distance between p and t . A consistent heuristic is always admissible, but the converse is not necessarily true. This can be proven by induction (Dechter & Pearl, 1985).

While the time complexity of A* depends heavily on the search space and the heuristic function used, its worst-case performance on an unbounded search space is $O(b^d)$, where b is the branching factor (i.e. the average number of children per node, or the outdegree) and d is the depth of the shortest path solution (Russell & Norvig, 2010).

4.4. The k -step Look-ahead Heuristic

We aim to improve the computational efficiency of the A* search algorithm by integrating the k -step look-ahead heuristic, in which k is a parameter that we choose. First, we start with a couple of definitions. We define $C(P, t)$ to be a *cut function* that represents the section of path P up until node t , if t is on path P . For instance, if $P = ((a, b), (b, t), (t, c), (c, d))$, then $C(P, t) = ((a, b), (b, t))$. If P does not pass through t , then $C(P, t) = P$. If X is

the set of all visited nodes in G , then let $S_{G, X, n, k}$ be the set of all acyclic paths from n of length k that do not pass through any nodes in X .

The intuitive explanation behind the k -step look-ahead heuristic is that the heuristic scans next possible paths to a certain extent from a given node and selects the estimated cost of the path that seems most promising. Specifically, instead of estimating the remaining cost to get from node n to target node t directly using Euclidean distance, we do so by iteratively finding the cost of each possible path of length k to the unvisited successors of n , added to the Euclidean distance between that successor node and t . We then select the lowest of those calculated values and assign that to $h(n)$. More formally, we have

$$h(n) = \min_{P \in S_{G, X, n, k}} \left[\left(\sum_{(x, y) \in C(P, t)} c(x, y) \right) + \text{dist}(n', t) \right]$$

where n' is the final node in path $C(P, t)$.

For example, to find the estimated cost from n to t using the 1-step look-ahead heuristic, we calculate

$$c(n, n') + \text{dist}(n', t)$$

for every unvisited child node n' of n , and assign $h(n)$ to the smallest resulting value.

We emphasize two details implied from our definition of the k -step look-ahead heuristic that become increasingly important as the value of k gets larger:

1. As previously mentioned, the “look-ahead” path P that our algorithm takes must be acyclic, i.e. it should not revisit any nodes previously passed through earlier in the same look-ahead path. Otherwise, the process could be simulating a sub-optimal path where nodes are revisited, potentially leading to unnecessary calculations. This is particularly true when dealing with undirected graphs where it is possible to travel between nodes in both directions. The list `alreadySeen` in Algorithm 3 keeps track of which nodes have already been looked through in a given look-ahead path.
2. If the algorithm reaches the target node during its look-ahead process, that particular look-ahead path should terminate and the returned value should be the cost of that path from node n to t . This prevents the algorithm from unnecessarily extending to successor nodes beyond the desired target node. This is the reasoning behind finding the cost of path $C(P, t)$ as opposed to the entire cost of P .

Algorithm 3 shows example pseudocode for computing the 2-step look-ahead heuristic value of a node n .

Algorithm 3 2-step Look-Ahead Heuristic

2-step Look-ahead(G, n, t):

Input: A Euclidean graph $G = (V, E)$, node $n \in V$, and target node $t \in V$.

Output: $h(n)$, or the estimated cost of the shortest path from n to t .

```

1: if  $n = t$  then
2:    $h(n) := 0$ 
3: else
4:   alreadySeen := an empty list only containing  $n$ 
5:   possibleVals := an empty list only containing  $\infty$ 
6:   for all  $n_1$  adjacent to  $n$  do
7:     if  $n_1$  has not been visited then
8:       if  $n_1 = t$  then
9:         Add  $c(n, n_1)$  to possibleVals
10:      Continue
11:     Add  $n_1$  to alreadySeen
12:     for all  $n_2$  adjacent to  $n_1$  do
13:       if  $n_2$  is not visited and not in al-
readySeen then
14:         Add  $c(n, n_1) + c(n_1, n_2) +$ 
 $\text{dist}(n_2, t)$  to possibleVals
15:       Remove  $n_1$  from alreadySeen
16:    $h(n) := \min(\text{possibleVals})$ 
17: return  $h(n)$ 

```

The k -step look-ahead is a more accurate heuristic of the shortest path between n and t than direct Euclidean distance because it takes into account some of the edges that the real shortest path from s to t will actually traverse. This suggests that a shortest path will be found quicker with this new heuristic than by using standard A*, particularly for relatively small values of k (for larger k , the number of calculations the heuristic has to make tends to increase exponentially, which can lead to slower runtimes). Furthermore, this estimator is still admissible, as the resulting value for $h(n)$ will never be greater than the cost of the true shortest path from n to t , so the optimal solution is guaranteed on Euclidean networks.

5. Methodology

In this section, we discuss the procedure used to perform our computational experiment on the efficiency of Dijkstra’s algorithm, standard A*, and A* with a k -step look-ahead heuristic for various values of k on sample road networks. To our knowledge, the new heuristic has not before been evaluated on Euclidean networks in the current literature. All tests are run using the Python programming language (K. Hong, 2020) in Google Colaboratory on a personal computer with a 1.3 GHz Intel Core i7 processor and

16 GB of RAM.

5.1. Test Network Details

We implement the shortest path algorithms on three different road networks as described in Table 1: the relatively small Downtown Brooklyn network, the medium-sized Jersey City network, and the larger New York City network (see appendix A.1 for visual layouts). We choose to include datasets of different sizes for the purpose of determining which algorithms perform comparatively well or poorly for various scales of networks. The data is obtained from the OSMnx Street Networks Dataverse, a compilation found on Harvard Dataverse that includes the street networks of every city, town, urban area, county, census tract, and Zillow-defined neighborhood in the United States (Boeing, 2017a; 2020). The networks are created from OpenStreetMap data using the OSMnx software (Boeing, 2017b). In the networks, nodes correspond to street intersections, edges represent the road segments connecting those intersections, and weights correspond to the lengths of those road segments in meters. Additionally, for each network, the data source provides node and edge lists. The node lists contain each node in the network, its assigned ID, and its geographic coordinates, while the edge lists include information about the two nodes that each edge connects, the cost of the edge, and whether the edge is one-way only, among other features. Note that these networks are considered mixed graphs because they contain both directed and undirected edges.

Each road network dataset must also be pre-processed in order to remove groups of nodes that are entirely disconnected from the rest of the network, thus ensuring a fully navigable road system. Lastly, before the final experiment is run, all algorithms are tested on several source-target pairs for all three networks to confirm that they consistently find the same optimal path.

5.2. Algorithm Implementations

In our computational experiment, twelve algorithms are tested on each network: Dijkstra’s, standard A*, and A* with a k -step look-ahead heuristic where k is each of 1 to 10 inclusive. For each network, a certain number of source-target node pairs are randomly selected from the network (with the exception of New York City), and each algorithm

Table 1. Characteristics of test road networks.

Road network	# of nodes	# of edges
Downtown Brooklyn	149	281
Jersey City	2232	5320
New York City	55136	141085

Table 2. Experimental details for each test road network.

Road Network	# of source-target pairs	# of runs per pair
Downtown Brooklyn	200	30
Jersey City	100	10
New York City	1	1

is run a fixed number of times for each source-target pair. Specifics are outlined in Table 2. We then record the resulting average runtime and average nodes expanded (i.e. nodes visited) for an algorithm’s execution (see section 6). As Table 2 shows, we design our experiment differently between road networks in terms of the number of times each algorithm is run. This is because the overall runtimes of the algorithms significantly increase with the size of the network, and runtimes on the smaller networks must be recorded at a greater degree of precision. Since we run the algorithms on only one distinct shortest path problem in the New York City dataset, we deliberately choose the source and destination nodes to be on opposite sides of the network in order to encapsulate the network’s size and complexity.

To implement a priority queue data structure to efficiently store unvisited nodes, we utilize Python’s *heapq* module, which represents a binary tree and can order a list by priority in linear time. Moreover, we incorporate a calculation of Euclidean distance between two given nodes in meters using the haversine formula since the geographic coordinates of each node is known (Jain, 2021). The formula states that

$$d = 2r \arcsin \left(\sqrt{\sin^2(c_1) + \cos(\phi_1) \cos(\phi_2) \sin^2(c_2)} \right)$$

where $c_1 = \left(\frac{\phi_2 - \phi_1}{2} \right)$ and $c_2 = \left(\frac{\lambda_2 - \lambda_1}{2} \right)$

- (ϕ_1, λ_1) and (ϕ_2, λ_2) are the coordinates of the two points in radians,
- r is the radius of the Earth, and
- d is the great-circle distance between the two points.

Finally, for our implementations of all variations of A*, the final recorded runtimes do not include the program’s graph pre-processing steps of calculating the Euclidean distance between each node in the network and the target node.

6. Results

The results of our computational experiment with shortest path algorithms on road networks are displayed in this section. As previously described, we run the A* algorithm

Table 3. Comparison of Algorithm Performance on Downtown Brooklyn Network.

Algorithm	Avg Runtime (s)	Avg # of Nodes Expanded
Dijkstra’s	0.0055	69.81
Standard A*	0.0030	34.44
A*, $k = 1$	0.0028	30.59
A*, $k = 2$	0.0026	27.83
A*, $k = 3$	0.0029	26.92
A*, $k = 4$	0.0032	25.59
A*, $k = 5$	0.0048	26.91
A*, $k = 6$	0.0057	28.98
A*, $k = 7$	0.0089	29.39
A*, $k = 8$	0.0119	29.34
A*, $k = 9$	0.0225	31.28
A*, $k = 10$	0.0368	30.68

with a k -step look-ahead heuristic on each of the three sample road networks for all k from 1 to 10 inclusive. Additionally, we include Dijkstra’s algorithm to provide context of the algorithms’ performance. As is a common convention (Zeng & Church, 2009; Zhan & Noon, 1998; Chan et al., 2016; Rachmawati & Gustin, 2020), we use runtime to judge an algorithm’s efficiency. For each source-target pair, all algorithms we implement guarantee the same shortest path. See appendix A.2 for additional visualizations of the number of nodes expanded as a function of k .

6.1. Small Sample Network

Results for all twelve algorithms on the Downtown Brooklyn street network are shown in Table 3. In terms of runtime, a 2-step look-ahead heuristic achieves the greatest efficiency by a small margin for this particular network at roughly 0.0026 seconds on average and 10.5% faster than standard A*. Moreover, the heuristics with $k = 1$ and $k = 3$ also perform better than regular A*. Figure 1 shows that as k increases, A*’s average runtime grows exponentially. While $k = 2$ achieves the quickest runtime, A* expands the least number of nodes on average with a 4-step look-ahead. In addition, when $k = 10$, the largest value of k that we test, the runtime is significantly greater than that of Dijkstra’s while still expanding less than half as many nodes.

6.2. Medium Sample Network

As shown in Table 4, the A* algorithm efficiency is able to be improved for 6 different values of k on the medium-sized Jersey City road network. Most notably, the best performance occurs when $k = 4$, running 32.6% faster than regular A* at approximately 0.2719 seconds on average. This same value of k also results in the least average num-

Table 4. Comparison of Algorithm Performance on Jersey City Network.

Algorithm	Avg Runtime (s)	Avg # of Nodes Expanded
Dijkstra's	1.1554	1192.50
Standard A*	0.4033	304.88
A*, $k = 1$	0.3172	247.29
A*, $k = 2$	0.3196	220.90
A*, $k = 3$	0.2759	203.58
A*, $k = 4$	0.2719	192.81
A*, $k = 5$	0.3088	210.99
A*, $k = 6$	0.3547	220.63
A*, $k = 7$	0.4190	233.78
A*, $k = 8$	0.7356	240.70
A*, $k = 9$	0.9383	254.02
A*, $k = 10$	1.5565	261.78

ber of nodes expanded. We also see a growing exponential behavior in algorithm runtime for values of k greater than 4, depicted in Figure 2. Once k hits 10, A*'s runtime exceeds that of Dijkstra's even though it expands much fewer nodes.

6.3. Large Sample Network

The street system of New York City is by far the largest of the three sample networks we use in this experiment, and all runtimes surpass 10 minutes. We test each algorithm once on a chosen source-target pair in this network, and find that an 8-step look-ahead heuristic achieves optimal performance at just 684.10 seconds, as indicated in Table 5. Using this heuristic results in a 22.2% shorter runtime than

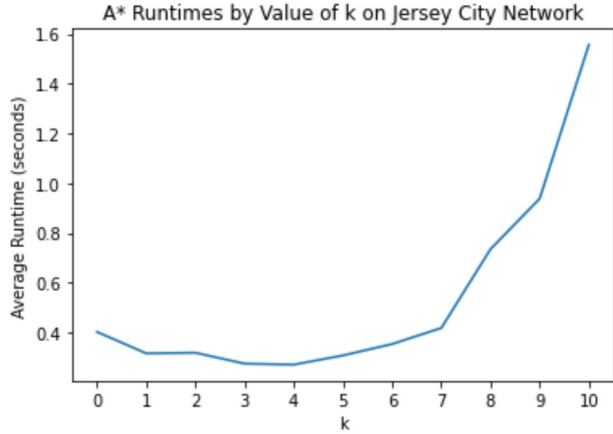


Figure 2. Algorithm runtime in seconds as a function of k in our implementation of A* search algorithm with a k -step look-ahead heuristic when run on the Jersey City network. A* achieves optimal efficiency using a 4-step look-ahead heuristic for this particular network.

Table 5. Comparison of Algorithm Performance on New York City Network.

Algorithm	Runtime (s)	# of Nodes Expanded
Dijkstra's	1389.98	54346
Standard A*	879.43	23370
A*, $k = 1$	829.41	20498
A*, $k = 2$	771.66	19200
A*, $k = 3$	722.54	18323
A*, $k = 4$	760.31	17769
A*, $k = 5$	712.97	17934
A*, $k = 6$	811.00	18959
A*, $k = 7$	761.65	20418
A*, $k = 8$	684.10	20559
A*, $k = 9$	850.23	20264
A*, $k = 10$	1047.57	20348




Figure 1. A visualization of the average algorithm runtime in seconds as a function of k in A* implementation with a k -step look-ahead heuristic when run on the Downtown Brooklyn network. A value of $k = 0$ indicates the standard A* algorithm. Optimal efficiency for this network is achieved with a 2-step look-ahead, but its improvement from standard A* is negligible.

using standard A*. In total, 9 different values of k perform better than regular A*. Figure 3 visualizes this data, showing that the runtimes stay relatively low when k is between 3 and 8, but skyrockets once k exceeds 8. Additionally, A* expands the least number of nodes in the New York City network when using a 4-step look-ahead heuristic. When $k = 10$, A* runs more efficiently than Dijkstra's while also expanding less than half as many nodes.

7. Discussion

Through a computational experiment, our goal has been to determine whether implementing a k -step look-ahead heuristic would improve the efficiency of the A* search algorithm, and to decide which algorithms work best on dif-

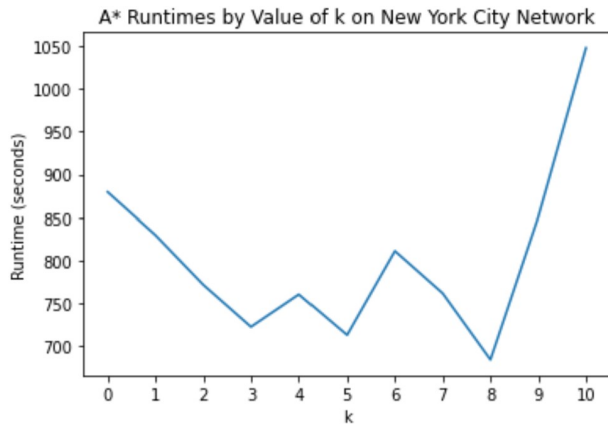


Figure 3. Algorithm runtime in seconds as a function of k for A* with a k -step look-ahead heuristic when run on the New York City network. A* achieves optimal efficiency using an 8-step look-ahead heuristic for this particular network, but also performs relatively well when $k = 3$ and $k = 5$.

ferent sizes of road network. Here, we address this problem by interpreting the results from section 6.

First, the results agree with our hypothesis that the k -step look-ahead heuristic would lead to shorter runtime than using direct Euclidean distance as the estimator. For the Downtown Brooklyn network, although we are able to slightly improve efficiency for a few values of k , the differences in runtime are negligible ten-thousandths of a second, which is unlikely to be impactful in the real world. On the other hand, using the new heuristic on the Jersey City and New York City networks shortens runtime by at most 0.13 and 195.33 seconds, respectively. In addition, on the Downtown Brooklyn network, 3 values of k improve the efficiency of A*, whereas this number is 6 for the Jersey City network and 9 for New York City. The optimal efficiency occurs when $k = 2$ for Downtown Brooklyn, $k = 4$ for Jersey City, and $k = 8$ for New York City. This suggests not only that the k -step look-ahead heuristic is more effective for larger networks, but that the single optimal value of k increases as the network size grows. The above conclusions seem reasonable because in a larger graph, there is a greater need to expand an algorithm's search in the direction of the target node, which is accomplished by greater values of k to create a more accurate heuristic.

A reason that one shortest path algorithm may run faster than another may be that it visits and expands less nodes, but this metric does not tell the whole story in all cases. For example, in each of the two smaller networks (i.e. Downtown Brooklyn and Jersey City), A* with a 10-step look-ahead heuristic has a considerably longer runtime than Dijkstra's algorithm, but the number of nodes that A* expands in these cases is significantly lower than that of Dijkstra's.

This is explained by the sheer number of steps that A* has to take when calculating the cost of each look-ahead path. As mentioned earlier, the number of steps needed to calculate a k -step look-ahead heuristic value grows exponentially as k increases due to the many possible look-ahead paths that diverge at each node.

Meanwhile, the situation is different with the New York City network. In this case, both the runtime and number of nodes expanded are less for A* with a 10-step look-ahead than for Dijkstra's. This further emphasizes A*'s suitability for large-scale networks, even for relatively high values of k .

A more surprising result of this experiment emerges when looking at data for the number of nodes expanded. The original speculation is that a greater value of k leads to a more accurate heuristic, thus expanding fewer "incorrect" nodes that are not on the shortest path. However, this does not seem to be the case. In particular, for all three networks, the minimal number of nodes expanded during A* search occurs when $k = 4$. Finding a justification to explain this phenomenon can be one direction of future work. Another common theme between all three test networks is that all variations of A* expand significantly fewer nodes than Dijkstra's does, likely due to Dijkstra's tendency to blindly expand nodes in all directions until the destination is found, as it is not guided by a heuristic. For instance, in the New York City network, Dijkstra's algorithm expands 54346 out of 55136 nodes in the entire graph, while all A* implementations go about finding the shortest path more intelligently. Note that the reason Dijkstra's expands nearly all the nodes in the network for this specific case is that we deliberately choose the source and target nodes to be on opposite ends of the network.

While many useful insights are gained from our interpretations, these are preliminary results that may need more testing to solidify. For example, our single run on the New York City dataset for each algorithm may not be fully representative of all possibilities on the network; nevertheless, our analysis in this section is likely safe to generalize and apply to the real world. To further strengthen the validity of our conclusions, future experiments should run these algorithms on many more source-target pairs on several different road networks.

Overall, the results of this study show that incorporating a k -step look-ahead heuristic into the A* search algorithm yields greater efficiency in solving the shortest path problem on transportation networks and can be particularly effective on larger, more complex road systems.

8. Conclusion and Future Work

The shortest path problem is a fundamental combinatorial optimization problem that has been intensely studied for the past several decades. In the current literature, few studies have compared various shortest path algorithms' performance on large, complex road networks (Zeng & Church, 2009; Zhan & Noon, 1998; Engineer, 2001). The main contribution of our work is to implement the A* search algorithm with a k -step look-ahead heuristic to solve the single-pair shortest path problem on various road networks for the first time. We implement Dijkstra's algorithm, A*, and A* with a k -step look-ahead heuristic for all k from 1 to 10 inclusive. We then compare these algorithms' performance on three real road networks of different sizes: the small Downtown Brooklyn network, the medium-sized Jersey City network, and the larger New York City network. From the results of our experiment, we conclude that using the look-ahead heuristic significantly improves the efficiency of A*, especially on larger networks. We also find that the greater the network size, the higher k needs to be in order to achieve optimal efficiency on the given network. Limitations of our experimental methodology include the lack of a diverse range of source-target pairs to fully represent all possible shortest path problems on the New York City network. The results of this study can be incorporated into cutting-edge navigation systems on transportation networks, such as GPS routing technologies, as finding the shortest possible path between two locations is a common objective of our everyday lives.

Some possibilities of related future research directions include incorporating different data structures (i.e. double buckets, approximate buckets) for the priority queue in the implementation of A* with k -step look-ahead, as opposed to the simple binary heap structure we use in this study. Additionally, the average degree of a graph (i.e. the average number of edges that are incident to each node in the graph) may be another factor that affects the efficiency of using a k -step look-ahead heuristic, and should be analyzed in the context of real road networks. Finally, future studies may want to explore the possibility of implementing a program that automatically chooses the most optimal value of k to use in A*. The efficiency of this algorithm can then be evaluated on various Euclidean networks.

9. Acknowledgements

I would like to thank my mentor Jesse Stern for guiding me through the research process, giving me helpful feedback, and making this project possible. I would also like to thank all the staff at the Summer STEM Institute for their outstanding bootcamp program and valuable learning opportunities.

References

- Biswas, P., Mishra, P., and Mahanti, N. Computational Efficiency of Optimized Shortest Path Algorithms. *Int. J. Comput. Sci. Appl.*, 2(2):22–37, 2005. Publisher: Cite-seer.
- Boeing, G. U.S. Street Network Shapefiles, Node/Edge Lists, and GraphML Files, 2017a.
- Boeing, G. Osmnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65:126–139, 2017b.
- Boeing, G. A multi-scale analysis of 27,000 urban street networks: Every us city, town, urbanized area, and zillow neighborhood. *Environment and Planning B: Urban Analytics and City Science*, 47(4):590–608, 2020.
- Chan, S. Y. M., Adnan, N., Sukri, S. S., and Wan Zainon, W. M. N. An experiment on the performance of shortest path algorithm. 2016.
- Cherkassky, B. V., Goldberg, A. V., and Radzik, T. Shortest paths algorithms: Theory and experimental evaluation. *Mathematical programming*, 73(2):129–174, 1996. Publisher: Springer.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. *Introduction to algorithms*. MIT press, 3rd edition, 2009.
- Dechter, R. and Pearl, J. Generalized best-first search strategies and the optimality of A. *Journal of the ACM (JACM)*, 32(3):505–536, 1985. Publisher: ACM New York, NY, USA.
- Dijkstra, E. W. and others. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1): 269–271, 1959.
- Engineer, F. Fast shortest path algorithms for large road networks. In *Proceedings of 36th annual ORSNZ conference*, 2001.
- Fu, L., Sun, D., and Rilett, L. R. Heuristic shortest path algorithms for transportation applications: State of the art. *Computers & Operations Research*, 33(11):3324–3343, 2006. Publisher: Elsevier.
- Goldberg, A. V. and Harrelson, C. Computing the shortest path: A search meets graph theory. In *SODA*, volume 5, pp. 156–165. Citeseer, 2005.
- Goldberg, A. V., Kaplan, H., and Werneck, R. F. Reach for A*: Shortest Path Algorithms with Preprocessing. In *The shortest path problem*, pp. 93–139. Citeseer, 2006.

- Hart, P. E., Nilsson, N. J., and Raphael, B. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968. Publisher: IEEE.
- Jacob, R., Marathe, M., and Nagel, K. A computational study of routing algorithms for realistic transportation networks. *Journal of Experimental Algorithmics (JEA)*, 4:6–es, 1999. Publisher: ACM New York, NY, USA.
- Jain, J. Shortest Path for Emergency Services: A Comparative Analysis. *International Journal for Research in Applied Science and Engineering Technology*, 9(VI): 426–430, June 2021. ISSN 23219653. doi: 10.22214/ijraset.2021.34974.
- K. Hong. Dijkstra’s shortest path algorithm, 2020.
- Madkour, A., Aref, W. G., Rehman, F. U., Rahman, M. A., and Basalamah, S. A survey of shortest-path algorithms. *arXiv preprint arXiv:1705.02044*, 2017.
- Magzhan, K. and Jani, H. M. A review and evaluations of shortest path algorithms. *International journal of scientific & technology research*, 2(6):99–104, 2013.
- Rachmawati, D. and Gustin, L. Analysis of Dijkstra’s Algorithm and A* Algorithm in Shortest Path Problem. In *Journal of Physics: Conference Series*, volume 1566, pp. 012061. IOP Publishing, 2020. Issue: 1.
- Russell, S. and Norvig, P. *Artificial Intelligence: A Modern Approach*. Pearson Education, 3rd edition, 2010.
- Schrijver, A. On the history of the shortest path problem. *Documenta Mathematica*, 17(1):155–167, 2012.
- Schultes, D. Route Planning in Road Networks. In *Ausgezeichnete Informatikdissertationen*, pp. 22–29, 2008.
- Sharma, S. K. and Pal, B. Shortest path searching for road network using a* algorithm. *International Journal of Computer Science and Mobile Computing*, 4(7): 513–522, 2015.
- Zeng, W. and Church, R. L. Finding shortest paths on real road networks: the case for A. *International journal of geographical information science*, 23(4):531–543, 2009. Publisher: Taylor & Francis.
- Zhan, F. B. and Noon, C. E. Shortest path algorithms: an evaluation using real road networks. *Transportation science*, 32(1):65–73, 1998. Publisher: INFORMS.

A. Appendix

A.1. Sample Road Networks




Figure 4. Downtown Brooklyn network (149 nodes, 281 edges).



Figure 5. New York City network (55136 nodes, 141085 edges).



Figure 6. Jersey City network (2232 nodes, 5320 edges).

A.2. Nodes Expanded vs. k Graphs




Figure 7. Average number of nodes expanded as a function of k on the Downtown Brooklyn network. Minimal nodes are expanded at $k = 4$.




Figure 8. Average number of nodes expanded as a function of k on the Jersey City network. Minimal nodes are expanded at $k = 4$.




Figure 9. Average number of nodes expanded as a function of k on the New York City network. Minimal nodes are expanded at $k = 4$.