

# Geometry algorithms

Damon Wischik, Computer Laboratory, Cambridge University. Lent Term 2022


## Contents

<b>1</b>	<b>A* algorithm</b>	<b>1</b>
<b>2</b>	<b>Segment intersection</b>	<b>3</b>
<b>3</b>	<b>Jarvis's march</b>	<b>5</b>
<b>4</b>	<b>Graham's scan</b>	<b>8</b>
<b>5</b>	<b>Solutions to exercises</b>	<b>10</b>




## 1. A\* algorithm

Suppose we're searching for the shortest path between two vertices, on a graph which is embedded in space. In other words, suppose that every vertex has spatial coordinates—meaning that we can measure spatial distances and directions between them. The spatial distance won't be the correct graph distance, since it ignores reachability.




spatial distance = 24




graph distance = 26

*Image credit: Red Blob Games<sup>1</sup>*

Nonetheless, can we use spatial knowledge to speed up path-finding? A simple greedy strategy is to prioritize the search in the direction of our destination. We can use a priority queue just as in Dijkstra's algorithm, but prioritize vertices by their spatial distance to the destination, in contrast to Dijkstra's algorithm which prioritizes by graph distance from the origin and thus searches blindly in all directions.




greedy search



blind search

Greedy search does *not* work, in the following sense: if we stop the search as soon as we reach the destination vertex, we might end up with a suboptimal path.

Here's a simple example to illustrate. The solid lines are graph edges with weights, and the dotted lines show spatial distance to the destination vertex  $t$ . The greedy algorithm, starting from  $s$ , will first add  $a$  and  $b$  to the queue, pick  $b$  next because it's nearer the destination, add  $t$  to the queue, then pop  $t$ . So it will stop with the path  $s \rightarrow b \rightarrow t$ , when we really want  $s \rightarrow a \rightarrow b \rightarrow t$ .



<sup>1</sup>There is an excellent tutorial about A\* at Red Blob Games, <http://www.redblobgames.com/pathfinding/a-star/introduction.html>

## PROBLEM STATEMENT AND ALGORITHM

The setting is the same as for Dijkstra's algorithm. Consider a directed graph with edge costs  $\geq 0$ , and suppose we're given a start vertex  $s$  and a destination vertex  $t$ . We seek a shortest path from  $s$  to  $t$ . 'Shortest' has its usual meaning — the cost of a path is the sum of its edge weights, and we seek a path of minimum cost.

The  $A^*$  algorithm requires that we have available an heuristic distance function  $h(v \text{ to } t)$ . (This might for example be the straight-line distance from  $v$  to  $t$ .) It proceeds exactly like Dijkstra's algorithm, except that it uses a different key to sort the priority queue:

$$\text{key}(v) = v.\text{distance} + h(v \text{ to } t)$$

where  $v.\text{distance}$  is its best guess so far for the distance from  $s$  to  $v$ . It terminates as soon as it pops  $t$ .

**Theorem (Correctness).** *Assume that the heuristic distance is always less than or equal to the graph-theoretic distance, i.e. that*

$$h(v \text{ to } t) \leq d(v \text{ to } t) \quad \text{for all } v$$

where  $d(v \text{ to } t)$  is the cost of the shortest path from  $v$  to  $t$ . The  $A^*$  algorithm terminates and, when it does,  $t.\text{distance} = d(s \text{ to } t)$ .

We must use the version of dijkstra as given in lecture notes, which allows vertices to re-enter to explore after they've been popped. (Some textbooks present a different version of Dijkstra's algorithm, which does not allow re-entry.)

**Exercise 1<sup>o</sup>.** Use the  $A^*$  algorithm, by hand, to find a path from  $s$  to  $t$  in the four-node graph above.


**Exercise 2<sup>\*</sup>.** Prove that when  $A^*$  terminates it has found the correct distance to  $t$ . [Hint. It's an induction, but not the same type of induction as used in Dijkstra's algorithm. Look at the proof for case (i) of Bellman's algorithm.]

## 2. Segment intersection

Do two line segments intersect? This is a simple question, and a good starting point for many more interesting questions in computational geometry.




Let's start with a simpler problem. Is the point  $q$  above or below the dotted line? The answer doesn't need anything more than basic school maths: if  $q_y > (p_y/p_x)q_x$  then it's above.



But it's easy to get tangled up thinking through all the cases (e.g. if  $p_x > 0$  and  $p_y < 0$  do I need to flip the sign?) We can use slightly cleverer maths, namely dot products, to get a cleaner answer:


Let  $p^\top = (-p_y, p_x)$ . If we rotate the vector  $\vec{0p}$  by  $90^\circ$  anticlockwise, we get  $\vec{0p}^\top$ . Now, the sign of  $p^\top \cdot q$ , i.e. of  $-p_y q_x + p_x q_y$ , tells us which side of the dotted line  $q$  is on. The dotted line is called the *extension* of  $\vec{0p}$ .

- $p^\top \cdot q > 0$ :  $q$  is on the left, as you travel along the dotted line in direction  $\vec{0p}$
- $p^\top \cdot q = 0$ :  $q$  is on the line itself
- $p^\top \cdot q < 0$ :  $q$  is on the right



This gives us all the tests we need to decide if two line segments  $\overline{rs}$  and  $\overline{tu}$  intersect:


1. If  $t$  and  $u$  are both on the same side of the extension of  $\overline{rs}$ , i.e. if  $(s-r)^\top \cdot (t-r)$  and  $(s-r)^\top \cdot (u-r)$  have the same sign, then the two line segments don't intersect.
2. Otherwise, if  $r$  and  $s$  are both on the same side of the extension of  $\overline{tu}$ , then the two line segments don't intersect.
3. Otherwise, they do intersect.



Well-written code should test all the boundary cases, e.g. when  $r = s$  or when  $t$  or  $u$  lie on the extension of  $\vec{r}\vec{s}$ . It is however a venial sin to test equality of floating point numbers<sup>2</sup>, because of the vagaries of finite-precision arithmetic, and so the question “How should my segment-intersection code deal with boundary cases?” depends on “What do I know about my dataset and what will my segment-intersection code be used for?”

**Exercise 3.** Two line segments are moving. The first line segment has endpoints  $(-0.2 + 0.1t, -0.1 + 0.1t)$  and  $(0.8 + 0.1t, -0.4 + 0.2t)$ , the second has endpoints  $(-0.36 + 0.52t, 1.1 - 0.3t)$  and  $(2.1 - 0.3t, 0.1 - 0.3t)$ , and  $t \geq 0$ . Do they collide? If so, at what time  $t$ ? Write the most concise pseudocode you can, to solve this problem for arbitrary coefficients.


**Exercise 4.** Consider a horizontal line starting from  $q$  and going infinitely to the right. How many times does it cross an edge? What about from  $r$ ? Devise an algorithm to detect whether a given point is outside a given shape (like  $q$ ) or inside (like  $r$ ), where the shape is specified as a list of non-intersecting polygons and each polygon is specified as a list of points. Discuss the handling of corner cases.




<sup>2</sup>In Python,  $0.1 + 0.1 + 0.1 - 0.3$  gives  $5.55 \times 10^{-17}$ , which is not equal to 0

### 3. Jarvis's march

Given a collection of points, a *convex hull* is what you get if you throw a lasso around all the points, then pull it tight.



Convex hulls are used, among other things, for collision detection. Given a straight line segment on the convex hull of a complex object, all points in that object must lie on the same side of the line. If we can get away with checking the sides of just a few lines, we can drastically speed up checks for collision.




#### MATHEMATICAL DEFINITION

We may as well define convex hull properly, even though for algorithmic discussion that follows we'll be working informally. Formally speaking, given a collection of points  $P = \{p_1, \dots, p_n\}$ , any vector  $q$  is called a *convex combination* if it can be written as

$$q = \alpha_1 p_1 + \dots + \alpha_n p_n \quad \text{where } \alpha_i \geq 0 \text{ for all } i, \text{ and } \sum_{i=1}^n \alpha_i = 1.$$


The *convex hull* of a collection of points is the set of all convex combinations.



The *corner points of the convex hull* are the points  $p \in P$  such that  $p \notin \text{convexhull}(P \setminus \{p\})$ . (These are the points that catch when the lasso is pulled tight—but to turn the informal idea of a lasso pulled tight into precise symbolic logic is rather hard!)

#### THE ALGORITHM

Here is an algorithm to compute the corner points of the convex hull of a collection of points  $P$ . It is due to Jarvis (1973), and was discovered independently by Chand and Kapur (1970).



```

1  let  $q_0$  be the point with lowest  $y$ -coordinate
2  (in case of a tie, pick the one with the largest  $x$ -coordinate)
3
4  draw a horizontal (left→right) line through  $q_0$ 
5  for all other points  $r \in P$ :
6      find the angle  $\theta(r)$  from the horizontal line to  $\overrightarrow{q_0 r}$ , measured  $\odot$ 
7  let  $q_1$  be the point with the smallest angle
8  (in case of a tie, pick the one furthest from  $q_0$ )
9
10  $h = [q_0, q_1]$ 
11 repeatedly:
12     let  $p$  and  $q$  be the last two points added to  $h$  respectively
13     for all other points  $r \in P$ :
14         find the angle  $\theta(r)$  from the extended  $\overrightarrow{pq}$  line to  $\overrightarrow{qr}$ , measured  $\odot$ 
15         pick the point with the smallest angle, and append it to  $h$ 
16         (in case of a tie, pick the one furthest from  $q$ )
17     stop when we return to  $q_0$ 

```

## ANALYSIS


At each step of the iteration, we search for the point  $r \in P$  with the smallest angle  $\theta(r)$ , thus the algorithm takes  $O(nH)$  where  $n$  is the number of points in  $P$  and  $H$  is the number of points in the convex hull. (As with Ford-Fulkerson, running time depends on the content of the data, not just the size.)

**Performance note.** The algorithm says “find the point  $r$  with the smallest angle  $\theta(r)$ ”. We could use trigonometry to compute  $\theta$  — but there is a trick to make this faster. (Faster in the sense of ‘games get more frames per second’, but no difference in the big- $O$  sense.) If all the points we’re comparing are on the same side of dotted line, as they are at all steps of Jarvis’s march, then

$$\theta(r_1) < \theta(r_2) \iff r_2 \text{ is on the left of the extended line } \overrightarrow{qr_1}$$

and we’ve seen how to compute this true/false value with just some multiplications and additions.






\* \* \*

Jarvis's march is very much like selection sort: repeatedly find the next item that goes into the next slot. In fact, most convex hull algorithms resemble some sorting algorithm.

**Exercise 5<sup>o</sup>.** Jarvis's march starts by picking a point  $q_0 \in P$  with the lowest  $y$ -coordinate, breaking ties by choosing the larger  $x$ -coordinate. Give an example to show that, without this tie breaking rule,  $q_0$  might not be a corner point.

## 4. Graham's scan

Here is another algorithm for computing the convex hull, due to Ronald Graham (1972). It builds up the convex hull by scanning through all the points in a fan, backtracking when necessary.



```

1 let  $r_0$  be the point with lowest  $y$ -coordinate
2 (in case of a tie, pick the one with the largest  $x$ -coordinate)
3
4 draw a horizontal (left→right) line through  $r_0$ 
5 for all other points  $r$ :
6     find the angle from the horizontal line to  $\overrightarrow{r_0 r}$ , measured  $\odot$ 
7 let  $r_1, \dots, r_{n-1}$  be the sorted list of points, lowest angle to highest
8
9  $h = [r_0, r_1]$ 
10 for each  $r_i$  in the sorted list of points,  $i \geq 2$ :
11     if  $r_i$  isn't on the left of the extension of the final segment of  $h$ :
12         # backtrack
13         repeatedly delete points from the end of  $h$  until  $r_i$  is
14         append  $r_i$  to  $h$ 

```


The diagram on the next page shows how the algorithm proceeds. Each row in the diagram shows it working on a new  $r_i$ , and the side-by-side panels show steps in backtracking.

### ANALYSIS

The initial sort takes time  $O(n \log n)$ , where  $n$  is the number of points in the set. During the scan, each point  $r_i$  is added to the list once, and it can be removed at most once, so the loop is  $O(n)$ .

To save some trigonometrical calculations, the same trick as in Section 3 works.

**Exercise 6.** Graham's algorithm scans points in order, according to a certain angle. Explain carefully what happens when two points have exactly the same angle. Does the code still work? [Hint. The code can crash! Make sure you identify all the cases where the code crashes or gives an incorrect answer.]



## 5. Solutions to exercises

**Exercise 1.** It's worth noting that the dotted line distances do indeed satisfy the requirement for correctness:  $h(a \text{ to } t) = 4 \leq d(a \text{ to } t) = 5$ , and  $h(b \text{ to } t) = 1 \leq d(b \text{ to } t) = 4$ .

```

pop s, s.distance=0
  relax a, b
  toexplore = [a(distance=3,key=7), b(distance=5,key=6)]
pop b, b.distance=5
  relax t
  toexplore = [a(distance=3,key=7), t(distance=9,key=9)]
pop a, a.distance=3
  relax b
  toexplore = [b(distance=4,key=5), t(distance=9,key=9)]
pop b, b.distance=4
  relax t
  toexplore = [t(distance=8,key=8)]
pop t, t.distance=8

```

**Exercise 2.** *This is a delicate proof. There is obviously some sort of induction involved... but the Dijkstra-style argument (“prove by induction that when a vertex is popped it has the correct distance”) doesn’t work, and the example from Exercise 1 demonstrates why. Instead, the induction has to be along the vertices of the correct path, just like the proof of correctness of Bellman’s algorithm.*

At the instant  $t$  is first popped, suppose  $t.\text{distance} \neq d(s \text{ to } t)$ . Consider a true shortest path  $s$  to  $t$ . There are two cases for the status of the vertices on this path at the instant in question: (A) all the vertices apart from  $t$  have been popped with the correct distance, (B) there is a sequence of vertices  $s, \dots, u$  which have been popped with the correct distance, then some vertex  $v$  that hasn’t (i.e. either  $v$  not yet popped, or popped with the wrong distance), then zero or more extra vertices before  $t$ . We’ll analyse case (B). Case (A) is similar but simpler.

In case (B), the path is

$$\underbrace{s \rightarrow \dots u}_{\text{popped, correct dist.}} \rightarrow \underbrace{v}_{\text{not}} \rightarrow \dots t$$

When  $u$  was popped with the correct distance, it relaxed  $u \rightarrow v$ , hence at that time

$$\begin{aligned}
 \text{key}(v) &= v.\text{distance} + h(v \text{ to } t) \\
 &\leq u.\text{distance} + c(u \rightarrow v) + h(v \text{ to } t) \quad \text{since } u \rightarrow v \text{ just relaxed} \\
 &= d(s \text{ to } u) + c(u \rightarrow v) + h(v \text{ to } t) \quad \text{by choice of } u \\
 &= d(s \text{ to } v) + h(v \text{ to } t) \quad \text{since } u \rightarrow v \text{ is on a shortest path to } t \\
 &\leq d(s \text{ to } v) + d(v \text{ to } t) \quad \text{by assumption on } h \\
 &= d(s \text{ to } t) \quad \text{since } v \text{ is on a shortest path}
 \end{aligned}$$

Some time after  $u$  was popped,  $t$  gets popped; and when  $t$  was popped it had the wrong distance, hence

$$\text{key}(v) = t.\text{distance} > d(s \text{ to } t).$$

We just showed that, when  $u$  was popped earlier in the execution,  $\text{key}(v) \leq d(s \text{ to } t)$ ; and since keys can only ever decrease it must be that  $v$  gets popped after  $u$  and before  $t$ . But since  $u$  was popped with the correct distance (by construction), and  $u$  relaxes the edge  $u \rightarrow v$ , and that edge is part of a shortest path  $s \rightsquigarrow t$ , then  $v$  must have the correct distance. This contradicts our choice of  $v$ , therefore case (B) is impossible.

A similar but simpler argument shows that case (A) is impossible. Thus, our premise (that when  $t$  was first popped,  $t.\text{distance}$  is incorrect) must be false.

**Exercise 3.** First, define two utility classes: `Point` which is just a tuple with two fields

x and y, and Polynomial which supports evaluation, multiplication, and root-finding. Let  $\text{Polynomial}(a_0, a_1, \dots, a_n)$  denote the polynomial  $a_0 + a_1x + \dots + a_nx^n$ .

```

1 # Transpose a point
2 def T(p): return Point(-p.y, p.x)
3
4 # Define the four coordinates.
5 p = Point(x=Polynomial(-0.2, 0.1), y=Polynomial(-0.1, 0.1))
6 q = Point(x=Polynomial(0.8, 0.1), y=Polynomial(-0.4, 0.2))
7 r = Point(x=Polynomial(-0.36, 0.52), y=Polynomial(1.1, -0.3))
8 s = Point(x=Polynomial(2.1, -0.3), y=Polynomial(0.1, -0.3))
9
10 # Because of the cunning Polynomial class, these are all Polynomials too:
11 rside = T(q-p)*(r-p)
12 sside = T(q-p)*(s-p)
13 pside = T(s-r)*(p-r)
14 qside = T(s-r)*(q-r)
15
16 # At a given instant in time t, do the two line segments intersect?
17 # This is a concise implementation of the method from lectures.
18 def intersect(t):
19     return ((rside(t)>0) != (sside(t)>0)) and ((pside(t)>0) != (qside(t)>0))
20
21 # Find all the instants in time at which one of the "side" tests switches value.
22 events = rside.roots() + sside.roots() + pside.roots() + qside.roots()
23 events = sorted(t for t in events if t>0)
24
25 # LEFT AS AN EXERCISE:
26 # dealing properly with (i) empty events list, (ii) time before first event,
27 # (iii) time after last event.
28
29 # For each interval (u,v), do the two lines intersect at time (u+v)/2?
30 # If they do then return u
31 collisions = [u for u,v in zip(events[:-1],events[1:]) if intersect((u+v)/2)]
32
33 collisions # answer: [2.352]

```

Here is the full-blown idiomatic Python code to implement Point and Polynomial.

```

1 import collections
2 import math
3
4 class Point(collections.namedtuple('Point', ['x','y'])):
5     def __sub__(self, other): return Point(x=self.x-other.x, y=self.y-other.y)
6     def __neg__(self): return Point(x=-self.x, y=-self.y)
7     def __str__(self): return 'Point({}, {})'.format(str(self.x), str(self.y))
8     def __mul__(self, other): return self.x*other.x + self.y*other.y
9     def __call__(self, x): return Point(x=self.x(x), y=self.y(x))
10
11 class Polynomial:
12     def __init__(self, a0=0, *ais):
13         self._coef = []
14         for ai in reversed(ais):
15             if ai==0.0 and len(self._coef)==0: continue
16             self._coef.insert(0, ai)
17         self._coef.insert(0, a0)
18     def __len__(self):
19         return len(self._coef)
20     def __getitem__(self, i):
21         return self._coef[i] if i<len(self._coef) else 0
22     def __neg__(self):
23         return Polynomial(*[-ai for ai in self._coef])
24     def __add__(self, other):
25         return Polynomial(*[self[i]+other[i] for i in range(max(len(self),len(other))])])
26     def __sub__(self, other):
27         return Polynomial(*[self[i]-other[i] for i in range(max(len(self),len(other))])])
28     def __mul__(self, other):
29         f,g = self,other


```

```

30     return Polynomial(*[sum(f[j]*g[i-j] for j in range(i+1))
31                        for i in range(len(f)+len(g)-1)])
32 def __call__(self, x):
33     return sum(a*math.pow(x,i) for i,a in enumerate(self._coef))
34 def __str__(self):
35     terms = [{"a} {e}"].format(a=a, e=" if i==0 else "x"+" if i==1 else "^"+str(i))
36             for i,a in enumerate(self._coef)]
37     return " + ".join(terms)
38 def roots(self):
39     if len(self) > 3:
40         raise Exception("No solution for cubics or higher")
41     p = self
42     if p[1] == 0 and p[2] == 0:
43         return []
44     elif p[2] == 0:
45         return [-p[0]/p[1]]
46     elif p[0] == 0:
47         return [0, -p[1]/p[2]]
48     elif p[1]**2 < 4*p[0]*p[2]:
49         return []
50     else:
51         x = math.sqrt(p[1]**2 - 4*p[0]*p[2])
52         return [(-p[1]+x)/(2.0*p[2]), (-p[1]-x)/(2.0*p[2])]

```

**Exercise 4.** Algorithm: count the total number of segments that cross the infinite horizontal line. Odd  $\Rightarrow$  inside, even  $\Rightarrow$  outside.




To test if a segment  $\overline{p_1 p_2}$  crosses a half-infinite horizontal line: we can actually just apply the test from lecture notes without thinking, and plug in the coordinate  $(+\infty, q.y)$  for the other end of the horizontal line ‘segment’. This can be justified by considering  $(x, q.y)$  as  $x \rightarrow \infty$ , and asking what side it’s on as  $x \rightarrow \infty$ . Alternatively, if we don’t like  $+\infty$  in code, here is what the tests expand to. First, for simplicity, subtract  $q$  from each coordinate so that we’re considering the half-infinite line from  $(0, 0)$ . Then,

```

1 if p1.y and p2.y have the same sign:
2     doesn't cross
3 else if p1.y == p2.y:
4     doesn't cross
5 else if p2.y - p1.y and -p1.x(p2.y - p1.y) - p1.y(p1.x - p2.x) have the same sign:
6     doesn't cross
7 else:
8     does cross

```


What about corner cases? In the picture, the line from  $p$  actually goes through the edge of the polygon, so the number of crossings is (mathematically) not well-defined. What should the algorithm do in such cases? What about the literal corner case, where the line goes precisely through a corner? (Ha ha hah.) I am not aware of any algorithm that deals properly with these issues, as well as coping with finite-precision arithmetic.



## Exercise 5.



## Exercise 6.



- (a) deletes  $r_2$  as soon as it considers  $r_3$ ; this is the desired behaviour
- (b) deletes  $r_3$  when it considers the next point, which is the desired behaviour (except that if there is no next point,  $r_3$  will be included in the output, which is incorrect)
- (c) like (a)
- (d) the algorithm crashes in this case: it deletes  $r_1$  then has no “final segment” for comparison
- (e) like (d) but even worse because we don't actually want to delete  $r_1$
- (f) like (d)