

Bellman-Ford and All-Pairs Shortest Paths

Lecture 19

April 3, 2025

Part I

SSSP with Negative Weights

Last Time

Want to solve the single-source shortest-path problem: given a vertex s , what is $\text{dist}(s, t)$ for each vertex t ?

- Weighted graphs: length is the sum of the weights of the edges.

Last Time

Want to solve the single-source shortest-path problem: given a vertex s , what is $\text{dist}(s, t)$ for each vertex t ?

- Weighted graphs: length is the sum of the weights of the edges.

Dijkstra's algorithm makes “guesses” based on shortest path seen so far, and at each step “confirms” the smallest guess.

Last Time

Want to solve the single-source shortest-path problem: given a vertex s , what is $\text{dist}(s, t)$ for each vertex t ?

- Weighted graphs: length is the sum of the weights of the edges.

Dijkstra's algorithm makes “guesses” based on shortest path seen so far, and at each step “confirms” the smallest guess.

- Correct *as long as the edge weights are all non-negative.*


Last Time

Want to solve the single-source shortest-path problem: given a vertex s , what is $\text{dist}(s, t)$ for each vertex t ?

- Weighted graphs: length is the sum of the weights of the edges.

Dijkstra's algorithm makes “guesses” based on shortest path seen so far, and at each step “confirms” the smallest guess.

- Correct *as long as the edge weights are all non-negative*.
- May give the wrong answer if there are negative edge weights!



DP: Attempt 1

Why don't we use DP? (It worked for DAGs!)

DP: Attempt 1

Why don't we use DP? (It worked for DAGs!)

Subproblem definition:

Recurrence:

Evaluation Order:

DP: Attempt 2

How do we avoid circular dependencies? (Hint: add a subproblem parameter.)

Subproblem definition:

Recurrence:

Evaluation Order:

DP Pseudocode

Let $LSP(v, \ell)$ be the length of the shortest path from s to v that uses at most ℓ edges.

SSSP-DP(G, s):

Initialize LSP as a $V \times V$ matrix

Set $LSP[s, 0] = 0$ and $LSP[v, 0] = \infty$ for all $v \neq s$

for ℓ from 1 to $V - 1$:

for all vertices v :

$LSP[v, \ell] = LSP[v, \ell - 1]$

for all edges (u, v) :

$LSP[v, \ell] = \min(LSP[v, \ell], LSP[u, \ell - 1] + w(u, v))$

return $LSP[u, V - 1]$ for all $u \in V$

DP Pseudocode

Let $LSP(v, \ell)$ be the length of the shortest path from s to v that uses at most ℓ edges.

SSSP-DP(G, s):

Initialize LSP as a $V \times V$ matrix

Set $LSP[s, 0] = 0$ and $LSP[v, 0] = \infty$ for all $v \neq s$

for ℓ from 1 to $V - 1$:

for all vertices v :

$LSP[v, \ell] = LSP[v, \ell - 1]$


for all edges (u, v) :

$LSP[v, \ell] = \min(LSP[v, \ell], LSP[u, \ell - 1] + w(u, v))$

return $LSP[u, V - 1]$ for all $u \in V$

Efficiency?

DP Example



Simplification

Observation: we don't *really* care that the path we find for $\text{LSP}(u, \ell)$ has at most ℓ edges...

Simplification

Observation: we don't *really* care that the path we find for $LSP(u, \ell)$ has at most ℓ edges...

May as well always use our current best guess at the distance!

Simplification

Observation: we don't *really* care that the path we find for $\text{LSP}(u, \ell)$ has at most ℓ edges...

May as well always use our current best guess at the distance!

```
BellmanFord( $G, s$ ):
```

```
    Set  $s.\text{dist} = 0$  and  $v.\text{dist} = \infty$  for all  $v \neq s$ 
```

```
    for  $\ell$  from 1 to  $V - 1$ :
```

```
        for all edges  $(u, v)$ :
```

```
             $v.\text{dist} = \min(v.\text{dist}, u.\text{dist} + w(u, v))$ 
```

```
    return  $u.\text{dist}$  for all  $u \in V$ 
```

Simplification

Observation: we don't *really* care that the path we find for $\text{LSP}(u, \ell)$ has at most ℓ edges...

May as well always use our current best guess at the distance!

```
BellmanFord( $G, s$ ):  
  Set  $s.\text{dist} = 0$  and  $v.\text{dist} = \infty$  for all  $v \neq s$   
  for  $\ell$  from 1 to  $V - 1$ :  
    for all edges  $(u, v)$ :  
       $v.\text{dist} = \min(v.\text{dist}, u.\text{dist} + w(u, v))$   
  return  $u.\text{dist}$  for all  $u \in V$ 
```

Correctness?

Simplification

Observation: we don't *really* care that the path we find for $\text{LSP}(u, \ell)$ has at most ℓ edges...

May as well always use our current best guess at the distance!

```
BellmanFord( $G, s$ ):
```

```
    Set  $s.\text{dist} = 0$  and  $v.\text{dist} = \infty$  for all  $v \neq s$ 
```

```
    for  $\ell$  from 1 to  $V - 1$ :
```

```
        for all edges  $(u, v)$ :
```

```
             $v.\text{dist} = \min(v.\text{dist}, u.\text{dist} + w(u, v))$ 
```


```
    return  $u.\text{dist}$  for all  $u \in V$ 
```

Correctness?

Proof by induction: after iteration i , $v.\text{dist} \leq \text{LSP}(v, i)$ for all v .

Just One Catch

What happens to Bellman-Ford on the following graph?



Negative Cycles

If our graph has a cycle with negative *total* length, our argument that the shortest path must use at most $V - 1$ edges breaks down.

Negative Cycles

If our graph has a cycle with negative *total* length, our argument that the shortest path must use at most $V - 1$ edges breaks down.

This is no accident—if there's a negative cycle, the concept of a shortest path doesn't make sense, since we can always shorten our path by taking the cycle more times!*

Negative Cycles

If our graph has a cycle with negative *total* length, our argument that the shortest path must use at most $V - 1$ edges breaks down.

This is no accident—if there's a negative cycle, the concept of a shortest path doesn't make sense, since we can always shorten our path by taking the cycle more times!*

(Technically, it's actually the concept of a shortest *walk* that breaks down. But note that all our algorithms implicitly are computing the shortest walk and using the fact that the shortest walk *must* be a path as long as there are no negative cycles. In fact, if we really want to find the shortest *path* with no restrictions on the weights at all, this turns out to be one of the canonically hard problems we study in the last part of the course!)

Detecting Negative Cycles

How can we (algorithmically) tell if a graph has a negative cycle?

Detecting Negative Cycles

How can we (algorithmically) tell if a graph has a negative cycle?

Observation: a negative cycle will cause BellmanFord to continue trying to update distances even after $V - 1$ iterations.

Detecting Negative Cycles

How can we (algorithmically) tell if a graph has a negative cycle?

Observation: a negative cycle will cause BellmanFord to continue trying to update distances even after $V - 1$ iterations.

```
BellmanFord( $G, s$ ):
```

```
    Set  $s.dist = 0$  and  $v.dist = \infty$  for all  $v \neq s$ 
```

```
    for  $\ell$  from 1 to  $V - 1$ :
```

```
        for all edges  $(u, v)$ :
```

```
             $v.dist = \min(v.dist, u.dist + w(u, v))$ 
```

```
    for all edges  $(u, v)$ :
```

```
        if  $v.dist > u.dist + w(u, v)$ : return “Negative cycle!”
```

```
    return  $u.dist$  for all  $u \in V$ 
```


Detecting Negative Cycles

How can we (algorithmically) tell if a graph has a negative cycle?

Observation: a negative cycle will cause BellmanFord to continue trying to update distances even after $V - 1$ iterations.

BellmanFord(G, s):

Set $s.\text{dist} = 0$ and $v.\text{dist} = \infty$ for all $v \neq s$

for ℓ from 1 to $V - 1$:

for all edges (u, v) :

$v.\text{dist} = \min(v.\text{dist}, u.\text{dist} + w(u, v))$

for all edges (u, v) :

if $v.\text{dist} > u.\text{dist} + w(u, v)$: return “Negative cycle!”

return $u.\text{dist}$ for all $u \in V$

Exercise: modify the pseudocode to output a negative cycle.

Negative Weights Takeaways

If G has no negative cycles, we can use BellmanFord to solve the SSSP problem in $O(VE)$ time.

- Can also *check* if G has a negative cycle!

If G may have negative cycles, the SSSP problem is (depending on your definitions) either (1) not well-defined, or (2) one of the canonical hard problems we'll consider in the last part of the class.

Part II

All-Pairs Shortest Paths

Recall APSP

Previously, we fixed a single source s and asked for $\text{dist}(s, t)$ for all targets t .

Recall APSP

Previously, we fixed a single source s and asked for $\text{dist}(s, t)$ for all targets t .

More general question: output $\text{dist}(s, t)$ for *all pairs* of s and t .
(Similar to SSSP, we will need to assume no negative cycles.)

Recall APSP

Previously, we fixed a single source s and asked for $\text{dist}(s, t)$ for all targets t .

More general question: output $\text{dist}(s, t)$ for *all pairs* of s and t .

(Similar to SSSP, we will need to assume no negative cycles.)

Naïve approach: run SSSP from each vertex.

- Non-negative weights (Dijkstra): $O(V(V + E) \log V)$ time
- Negative weights (Bellman-Ford): $O(V^2 E)$ time

Recall APSP

Previously, we fixed a single source s and asked for $\text{dist}(s, t)$ for all targets t .

More general question: output $\text{dist}(s, t)$ for *all pairs* of s and t .

(Similar to SSSP, we will need to assume no negative cycles.)

Naïve approach: run SSSP from each vertex.

- Non-negative weights (Dijkstra): $O(V(V + E) \log V)$ time
- Negative weights (Bellman-Ford): $O(V^2 E)$ time

Can we do better?

When In Doubt, Try Using DP

How would we write a DP algorithm for this?

Subproblem definition:

Recurrence:

Evaluation Order:

DP Pseudocode

Let $\text{LSP}(u, v, \ell)$ be the length of the shortest path from u to v that uses at most ℓ edges.

APSP-DP(G):

Initialize LSP as a $V \times V \times V$ matrix

Set $\text{LSP}[u, u, 0] = 0$ for all u

Set $\text{LSP}[u, v, 0] = \infty$ for all $u \neq v$

for ℓ from 1 to $V - 1$:

for all vertices u :

for all vertices v :

$\text{LSP}[u, v, \ell] = \text{LSP}[u, v, \ell - 1]$

for all edges (x, v) :

$\text{LSP}[u, v, \ell] = \min \left(\begin{array}{c} \text{LSP}[u, v, \ell], \\ \text{LSP}[u, x, \ell - 1] + w(x, v) \end{array} \right)$

return $\text{LSP}[u, v, V - 1]$ for all $u, v \in V$

DP Pseudocode

Let $\text{LSP}(u, v, \ell)$ be the length of the shortest path from u to v that uses at most ℓ edges.

APSP-DP(G):

Initialize LSP as a $V \times V \times V$ matrix

Set $\text{LSP}[u, u, 0] = 0$ for all u

Set $\text{LSP}[u, v, 0] = \infty$ for all $u \neq v$

for ℓ from 1 to $V - 1$:

for all vertices u :

for all vertices v :

$\text{LSP}[u, v, \ell] = \text{LSP}[u, v, \ell - 1]$

for all edges (x, v) :

$\text{LSP}[u, v, \ell] = \min \left(\begin{array}{l} \text{LSP}[u, v, \ell], \\ \text{LSP}[u, x, \ell - 1] + w(x, v) \end{array} \right)$

return $\text{LSP}[u, v, V - 1]$ for all $u, v \in V$

Efficiency?

Improving Our Approach

Can we come up with a more efficient recurrence?

Subproblem definition: Let $LSP(u, v, \ell)$ be the length of the shortest path from u to v that uses at most ℓ edges.

Recurrence:

Evaluation Order:

Fischer-Meyer Pseudocode

Let $\text{LSP}(u, v, \ell)$ be the length of the shortest path from u to v that uses at most 2^ℓ edges.

FischerMeyer(G):

Initialize LSP as a $V \times V \times (\lceil \log_2(V) \rceil + 1)$ matrix

Set $\text{LSP}[u, v, 0] = w(u, v)$ for all $u, v \in V$

Let $w(u, u) := 0$, and for all $(u, v) \notin E$ let $w(u, v) := \infty$

for ℓ from 1 to $\lceil \log_2(V) \rceil$:

for all vertices u :

for all vertices v :

$\text{LSP}[u, v, \ell] = \min_x (\text{LSP}[u, x, \ell - 1] + \text{LSP}[x, v, \ell - 1])$

return $\text{LSP}[u, v, \lceil \log_2(V) \rceil]$ for all $u, v \in V$

Fischer-Meyer Pseudocode

Let $\text{LSP}(u, v, \ell)$ be the length of the shortest path from u to v that uses at most 2^ℓ edges.

FischerMeyer(G):

Initialize LSP as a $V \times V \times (\lceil \log_2(V) \rceil + 1)$ matrix

Set $\text{LSP}[u, v, 0] = w(u, v)$ for all $u, v \in V$

Let $w(u, u) := 0$, and for all $(u, v) \notin E$ let $w(u, v) := \infty$

for ℓ from 1 to $\lceil \log_2(V) \rceil$:

 for all vertices u :

 for all vertices v :

$\text{LSP}[u, v, \ell] = \min_x (\text{LSP}[u, x, \ell - 1] + \text{LSP}[x, v, \ell - 1])$

return $\text{LSP}[u, v, \lceil \log_2(V) \rceil]$ for all $u, v \in V$

Efficiency?

A Clever Subproblem

Subproblem definition:

Recurrence:

Evaluation Order:

Clever DP Pseudocode

Let $\text{LSP}(u, v, r)$ be the length of the shortest path from u to v that only uses intermediate vertices from $\{1, 2, \dots, r\}$.

APSP-CleverDP(G):

Initialize LSP as a $V \times V \times (V + 1)$ matrix

Set $\text{LSP}[u, v, 0] = w(u, v)$ for all $u, v \in V$

Let $w(u, u) := 0$, and for all $(u, v) \notin E$ let $w(u, v) := \infty$

for r from 1 to V :

 for all vertices u :

 for all vertices v :

$$\text{LSP}[u, v, r] = \min \left(\begin{array}{c} \text{LSP}[u, v, r - 1], \\ \text{LSP}[u, r, r - 1] + \text{LSP}[r, v, r - 1] \end{array} \right)$$

return $\text{LSP}[u, v, V]$ for all $u, v \in V$

Clever DP Pseudocode

Let $\text{LSP}(u, v, r)$ be the length of the shortest path from u to v that only uses intermediate vertices from $\{1, 2, \dots, r\}$.

APSP-CleverDP(G):

Initialize LSP as a $V \times V \times (V + 1)$ matrix

Set $\text{LSP}[u, v, 0] = w(u, v)$ for all $u, v \in V$

Let $w(u, u) := 0$, and for all $(u, v) \notin E$ let $w(u, v) := \infty$

for r from 1 to V :

 for all vertices u :

 for all vertices v :

$$\text{LSP}[u, v, r] = \min \left(\begin{array}{c} \text{LSP}[u, v, r - 1], \\ \text{LSP}[u, r, r - 1] + \text{LSP}[r, v, r - 1] \end{array} \right)$$

return $\text{LSP}[u, v, V]$ for all $u, v \in V$

Efficiency?

Simplification

Similar to Bellman-Ford, we don't *really* care that the path we find for $\text{LSP}(u, v, r)$ doesn't use vertices larger than r —we may as well use whatever best path we've found so far!

Simplification

Similar to Bellman-Ford, we don't *really* care that the path we find for $\text{LSP}(u, v, r)$ doesn't use vertices larger than r —we may as well use whatever best path we've found so far!

FloydWarshall(G):

Set $\text{dist}[u, v] = w(u, v)$ for all $u, v \in V$

Let $w(u, u) := 0$, and for all $(u, v) \notin E$ let $w(u, v) := \infty$

for all vertices r :

for all vertices u :

for all vertices v :

$\text{dist}[u, v] = \min \left(\begin{array}{c} \text{dist}[u, v], \\ \text{dist}[u, r] + \text{dist}[r, v] \end{array} \right)$

return $\text{dist}[u, v]$ for all $u, v \in V$

Simplification

Similar to Bellman-Ford, we don't *really* care that the path we find for $\text{LSP}(u, v, r)$ doesn't use vertices larger than r —we may as well use whatever best path we've found so far!

FloydWarshall(G):

Set $\text{dist}[u, v] = w(u, v)$ for all $u, v \in V$

Let $w(u, u) := 0$, and for all $(u, v) \notin E$ let $w(u, v) := \infty$

for all vertices r :

for all vertices u :

for all vertices v :

$\text{dist}[u, v] = \min \left(\begin{array}{c} \text{dist}[u, v], \\ \text{dist}[u, r] + \text{dist}[r, v] \end{array} \right)$

return $\text{dist}[u, v]$ for all $u, v \in V$

Correctness?

Simplification

Similar to Bellman-Ford, we don't *really* care that the path we find for $\text{LSP}(u, v, r)$ doesn't use vertices larger than r —we may as well use whatever best path we've found so far!

FloydWarshall(G):

Set $\text{dist}[u, v] = w(u, v)$ for all $u, v \in V$

Let $w(u, u) := 0$, and for all $(u, v) \notin E$ let $w(u, v) := \infty$

for all vertices r :

for all vertices u :

for all vertices v :

$\text{dist}[u, v] = \min \left(\begin{array}{c} \text{dist}[u, v], \\ \text{dist}[u, r] + \text{dist}[r, v] \end{array} \right)$

return $\text{dist}[u, v]$ for all $u, v \in V$

Correctness? Proof by induction: after iteration i of outermost loop, $\text{dist}[u, v] \leq \text{LSP}(u, v, i)$ for all $u, v \in V$.

Problems and Algorithms From Today

- Shortest paths from s with no negative cycles
 - Bellman-Ford, $O(VE)$
 - Can also check for negative cycles
- All-pairs shortest paths with no negative cycles
 - Negative weights: Floyd-Warshall, $O(V^3)$
 - Non-negative weights: Dijkstra from each vertex, $O(V(V + E) \log V)$ (Floyd-Warshall is more efficient if $E \approx V^2$)