CS 4700: Foundations of Artificial Intelligence

Fall 2017

Instructor: Prof. Haym Hirsh

Lecture 5

Piazza course website is up:

piazza.com/cornell/spring2017/cs4700

The researchers found that when apples were scarce, the agents quickly learned to attack one another -- zapping, or “tagging” their opponent with a ray that temporarily immobilized them. When apples were abundant, the agents preferred to co-exist more peacefully.

Rather chillingly, however, the researchers found when they tried this same game with more intelligent agents that drew on larger neural networks -- a kind of machine intelligence designed to mimic how certain parts of the human brain work -- they would “try to tag the other agent more frequently, i.e. behave less cooperatively, no matter how we vary the scarcity of apples,” they wrote in a blog post on DeepMind’s website.

New research from DeepMind, Alphabet Inc.’s London-based artificial intelligence unit could ultimately shed light on this fundamental question.

Leibo said that the agents used in the apple-gathering and Wolfpack experiments had no short-term memory, and as a result could not make any inferences about the intent of the other agent. “Going forward it would be interesting to equip agents with the ability to reason about other agent’s beliefs and goals,” he said.
In the meantime, it might be wise to keep a few spare apples around.

→ Finds environment plays big role in fostering AI cooperation

When our robot overlords arrive, will they decide to kill us or cooperate with us?

New research from DeepMind, Alphabet Inc.’s London-based artificial intelligence unit could ultimately shed light on this fundamental question.

Artificial intelligence changes the way it behaves based on the environment it is in, much like humans do, according to the latest research from DeepMind .

[…] During the work, they found it is possible for AI to act in an "aggressive manner" when it feels it is going to lose out, but agents will work as a team when there is more to be gained.

By **MATT BURGESS**

1 hour ago

Understanding Agent Cooperation

So, depending on the situation, having a greater capacity to implement complex strategies may yield either more or less cooperation. The new framework of sequential social dilemmas allows us to take into account not only the outcome of the interaction (as in the Prisoner's dilemma), but also the difficulty of learning to implement a given strategy.

behaviour of complex multi-agent systems such as the economy, traffic, and environmental challenges.

Self-interested people often work together to achieve great things. Why should this be the case, when it is in their best interest to just care about their own wellbeing and disregard that of others?

Robots

Actors, teachers, therapists – think your job is safe from robots? Think again

Thanks to advances in artificial intelligence, many jobs that weren't considered ripe for automation suddenly are

Dan Tynan in San Francisco

@tynanwrites

Thursday 9 February 2017 03.00 EST

Meet 'Botlr', a towel-delivering assistant that's already being experimented with at Aloft Hotels. Photograph: Botlr

Most popular in US

Wikipedia bans Daily Mail as 'unreliable' source

How can Obama smile at a time like this? I think he knows something | Jonathan Jones

Donald Trump: a man so obnoxious that karma may see him reincarnated as himself...

Sarah Palin touted as US ambassador to Canada? You betcha!

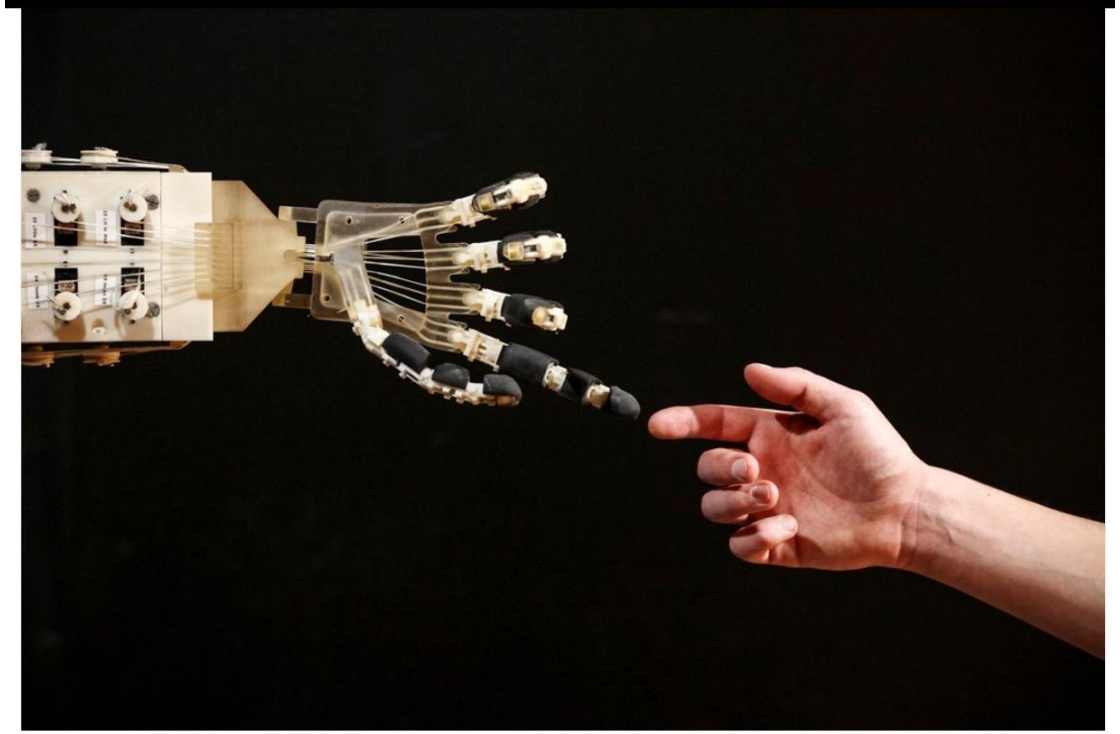

Oli Scarff / Getty

Is AI a Threat to Christianity?

Are you there, God? It's I, robot.

JONATHAN MERRITT | FEB 3, 2017 | TECHNOLOGY

In his relatively short tenure, Pope Francis has been hard at work welcoming spiritual seekers into the Catholic Church. He's refused to judge LGBT people, sought to integrate divorced couples, and extended priests' ability to forgive abortion. But Francis's wide arms have arguably never stretched further than a mass in 2014 when he suggested the church would baptize Martians.

"If—for example—tomorrow an expedition of Martians came ... and one says, 'But I want to be baptized!' What would happen?" Pope Francis asked. "When the Lord shows us the way, who are we to say, 'No, Lord, it is not prudent! No, let's do

Women's Wear Daily

+ Business + Beauty + Men's + Runway + Accessories + Eye + ...

Today's Digital Daily

SUBSCRIBE LOG IN Search WWD'S 2M+ Stories

FASHION FASHION FASHION

"It's 2017, and the new thing is called AI," he said. "And it's more powerful than the Internet, because the Internet didn't think, the Internet didn't know you. So that shift is about to happen again, where there is a new type of way that we connect and get information. It's about to enter society and it knows a lot. It's coming and it's going to be awesome."

Will.i.am and Farfetch's Jose Neves

Topics ranged from AI to the importance of data to Will.i.am's regrets about not having studied computer science.

By Lorelei Marfil on February 7, 2017

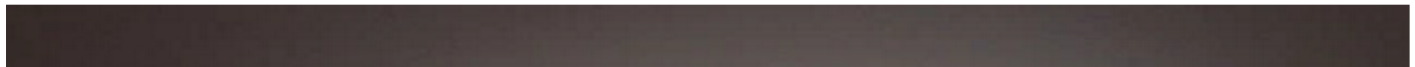

ESSENTIALIST

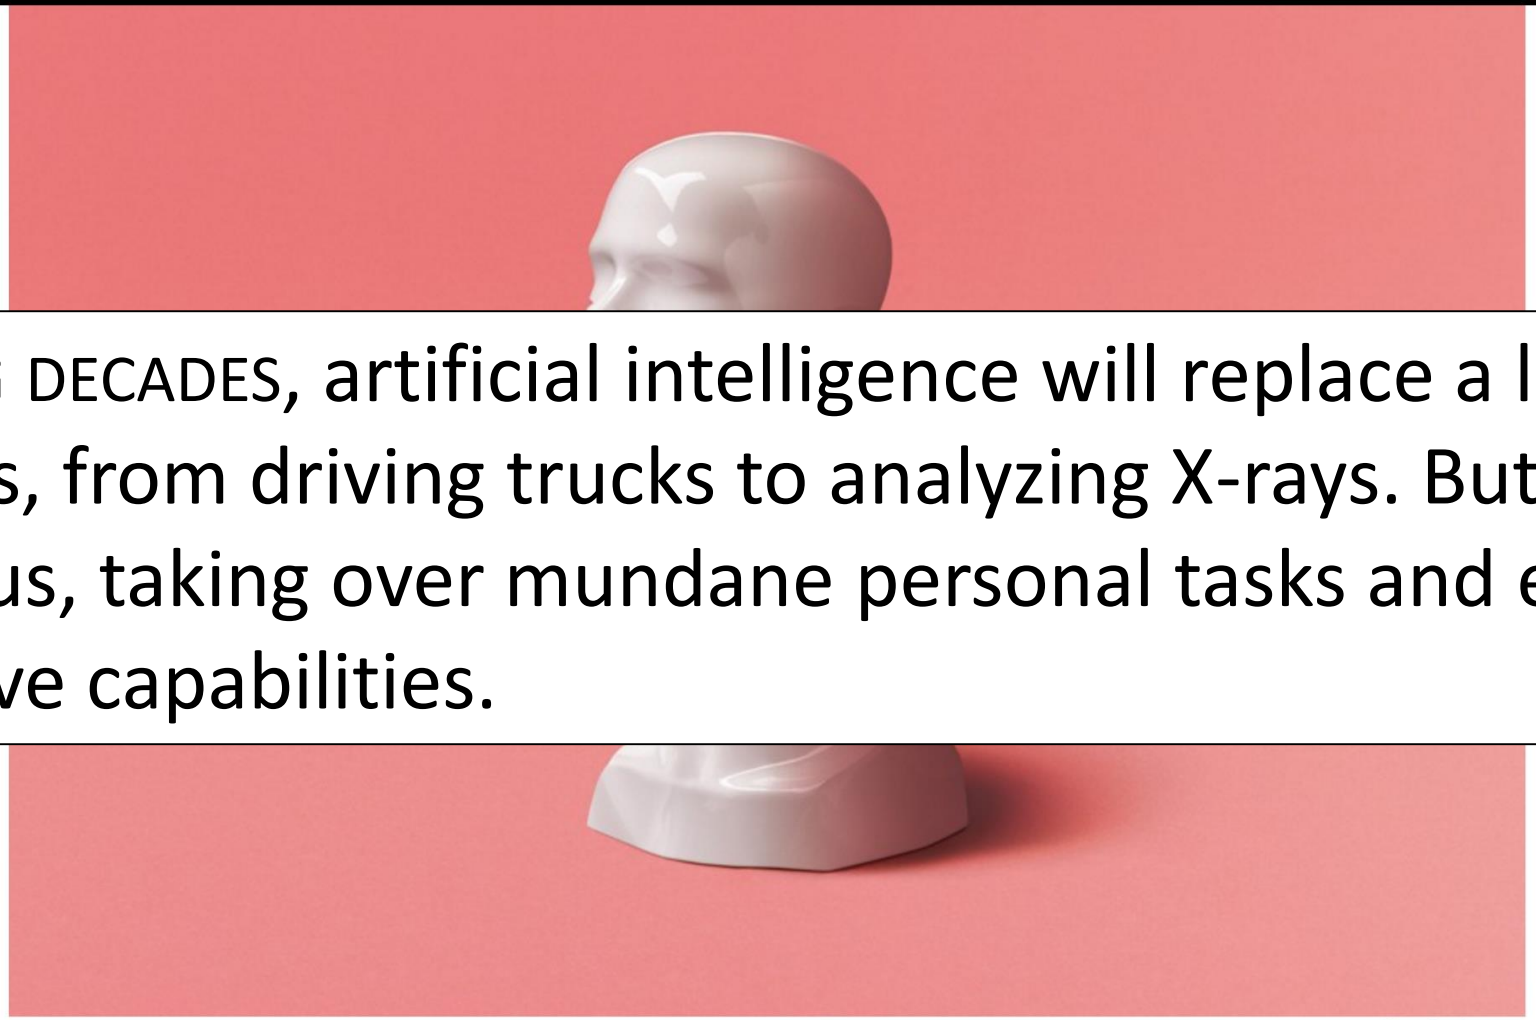

Alvaro Dominguez

THE COMING DECADES, artificial intelligence will replace a lot of human jobs, from driving trucks to analyzing X-rays. But it will also work *with* us, taking over mundane personal tasks and enhancing our cognitive capabilities.

Our Bots, Ourselves

How the descendants of Siri and Alexa could change our daily lives, thoughts, and relationships

MATTHEW HUTSON | MARCH 2017 ISSUE | TECHNOLOGY

Datamation Hangout:

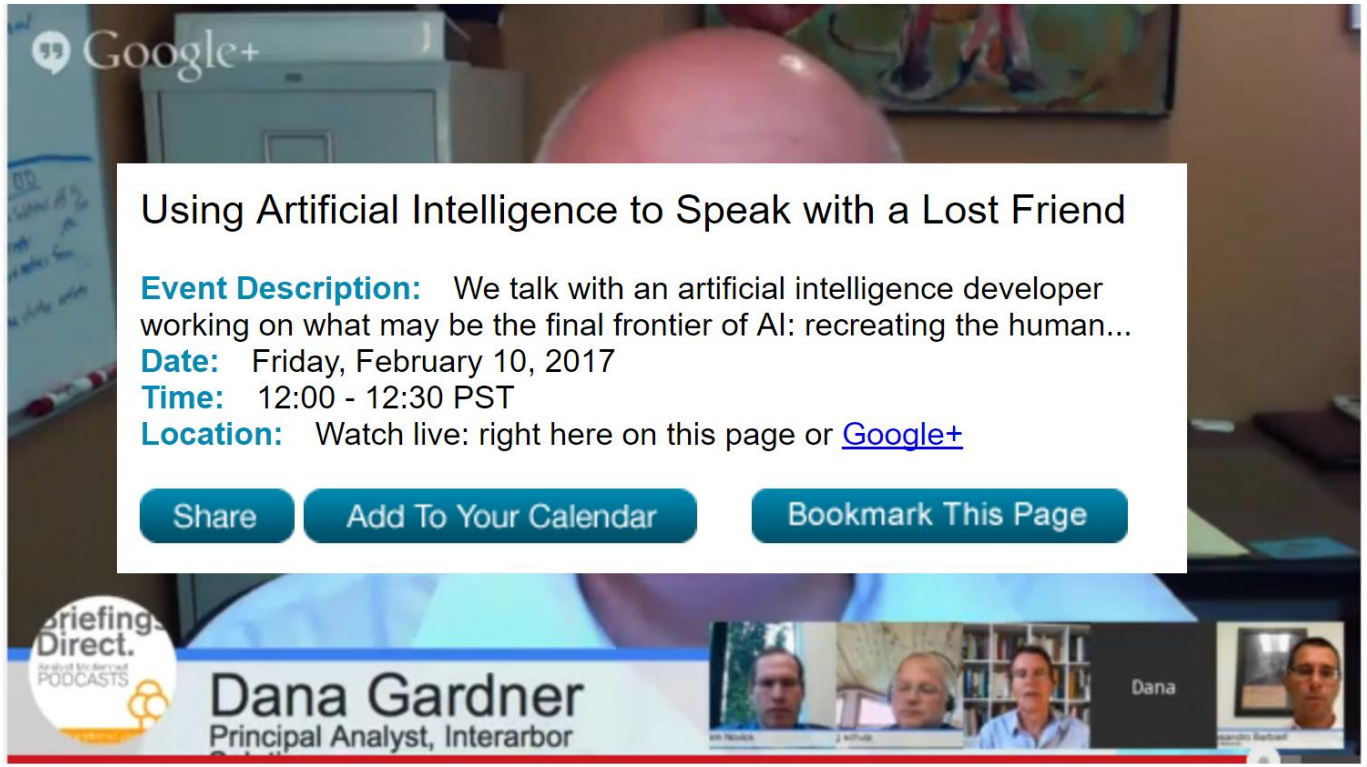

Video streams on this page on Friday, February 10, 12 noon PT.

Artificial intelligence, it seems, now accomplishes nearly any task: it drives cars, helps us shop, decides which of our friends to display on Facebook. But Eugenia Kuyda is pushing AI into a still

Tweets by @Datamation

Datamation.com @Datamation
Business Leaders are Bullish on AI - Over the next three years, most executives are planning to invest in AI, c... ow.ly/QWvk509ohgy
13h

Datamation.com @Datamation
3 Steps to Acclimatizing Your Organization to the Cloud - How does an established company move from in-house or... ow.ly/iZNS509nZUe
16h

Datamation.com @Datamation
8.4 Billion 'Things' to Light Up the IoT in 2017 - The number of connected devices feeding data into the Intern... ow.ly/RnR6509mqRV

INNOVATION

Realdoll builds artificially intelligent sex robots with programmable personalities

DIGITAL TRENDS By Luke Dormehl · Published February 06, 2017

File photo: People in costume dance during the 19th annual Halloween costume carnaval in West Hollywood, California October 31, 2006. (REUTERS/Mario Anzuoni)

Sex doll manufacturer Realdoll is dipping its toe (and we don't want to know which other body parts) into the world of artificial intelligence and robotics with a forthcoming robot sex assistant that promises to form a "real bond" with its, erm, users.

The new system is made up of several components, which will roll out over the course of this year and next. It will begin with the Harmony AI app, scheduled for

More from Fox News


Tamron Hall bolts NBC


Authorities chase drug smugglers


Flying cars in 1-3 years? Uber


Mariah Carey wears lingerie to the gym


Is Hollywood becoming mean-


Incredible gecko slips out of its own

Latest in Science

NASA figured out how to make Venus-resistant electronics

9h ago

SpaceX aims to launch ISS resupply mission on February 18th

10h ago

TSA debunks its own airport

Google uses AI to sharpen low-res images

'Blade Runner' style image enhancement is just a neural network away.

Jon Fingas, @jonfingas
02.07.17 in Gadgetry

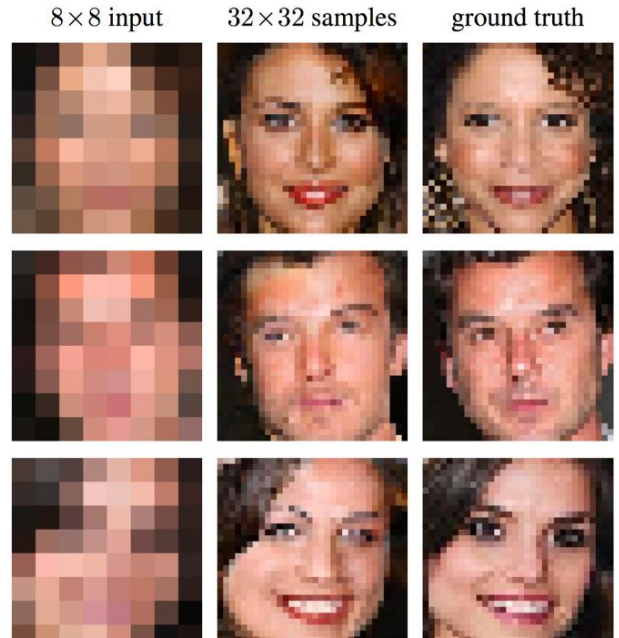

Today

- Uninformed Search (R&N Ch 3)

Thursday, February 16

- Informed Search (R&N Ch 3)

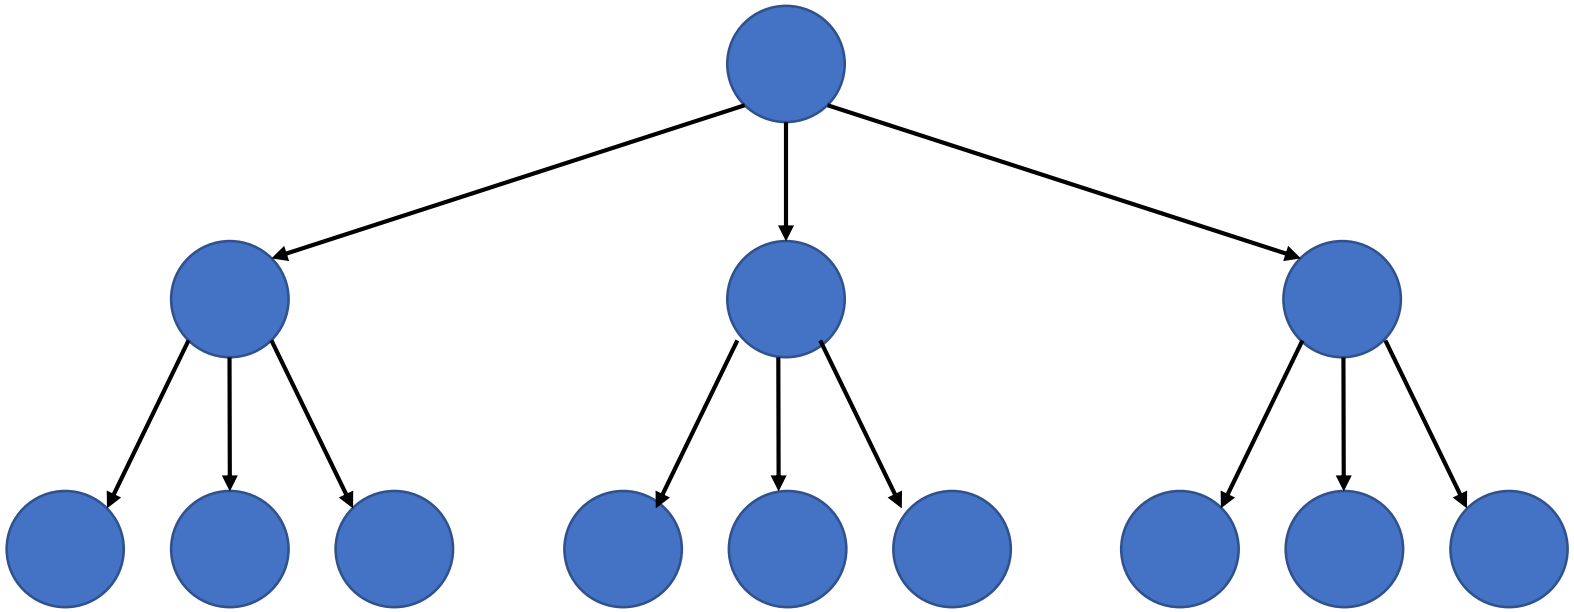

Depth-First Search

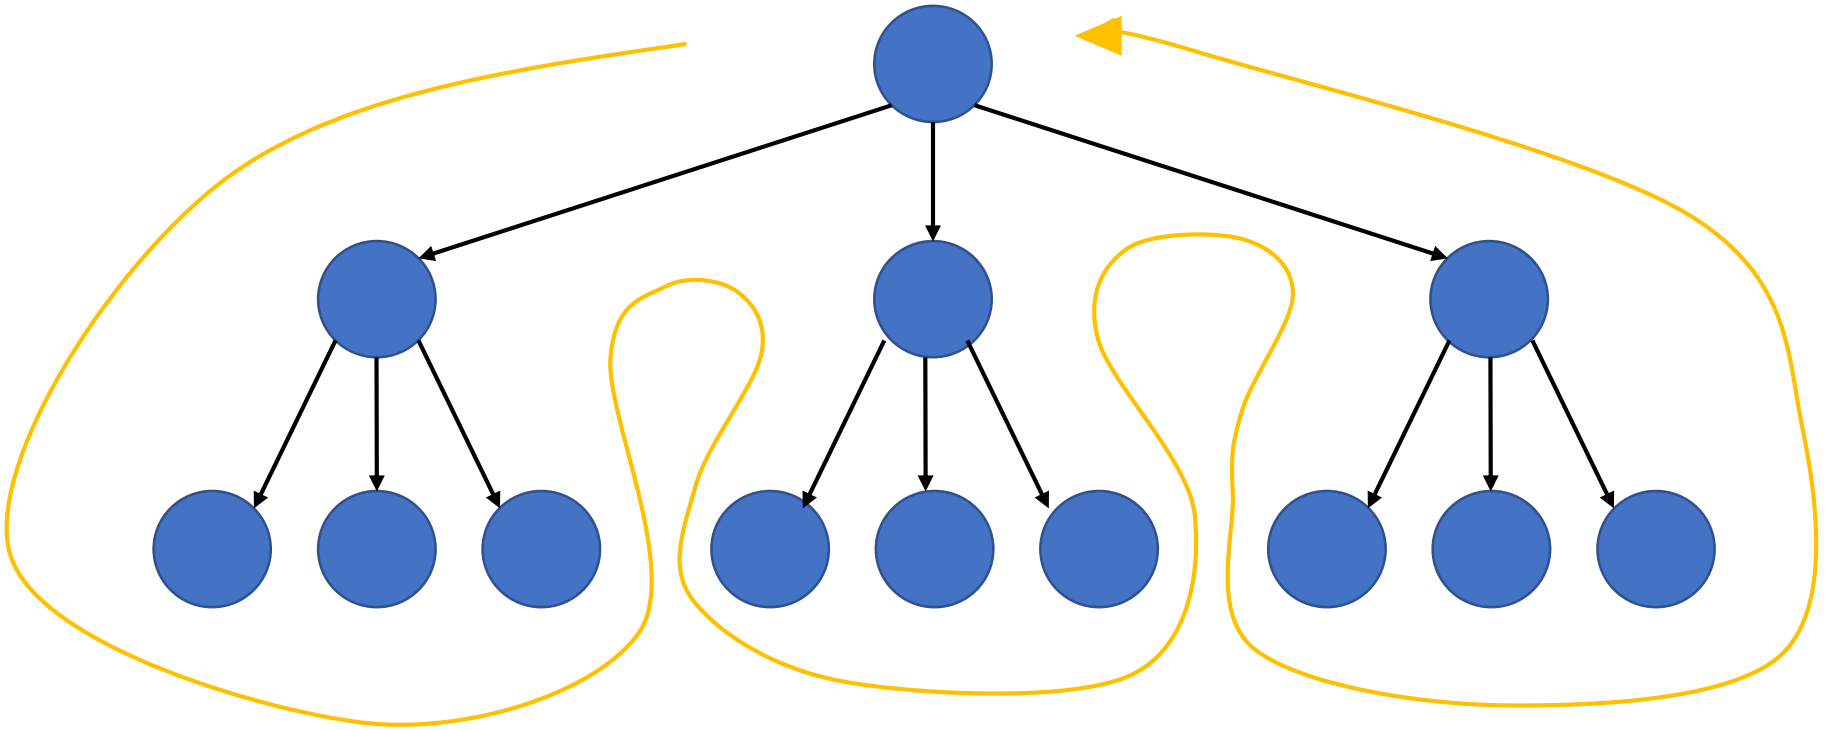

Breadth-First Search

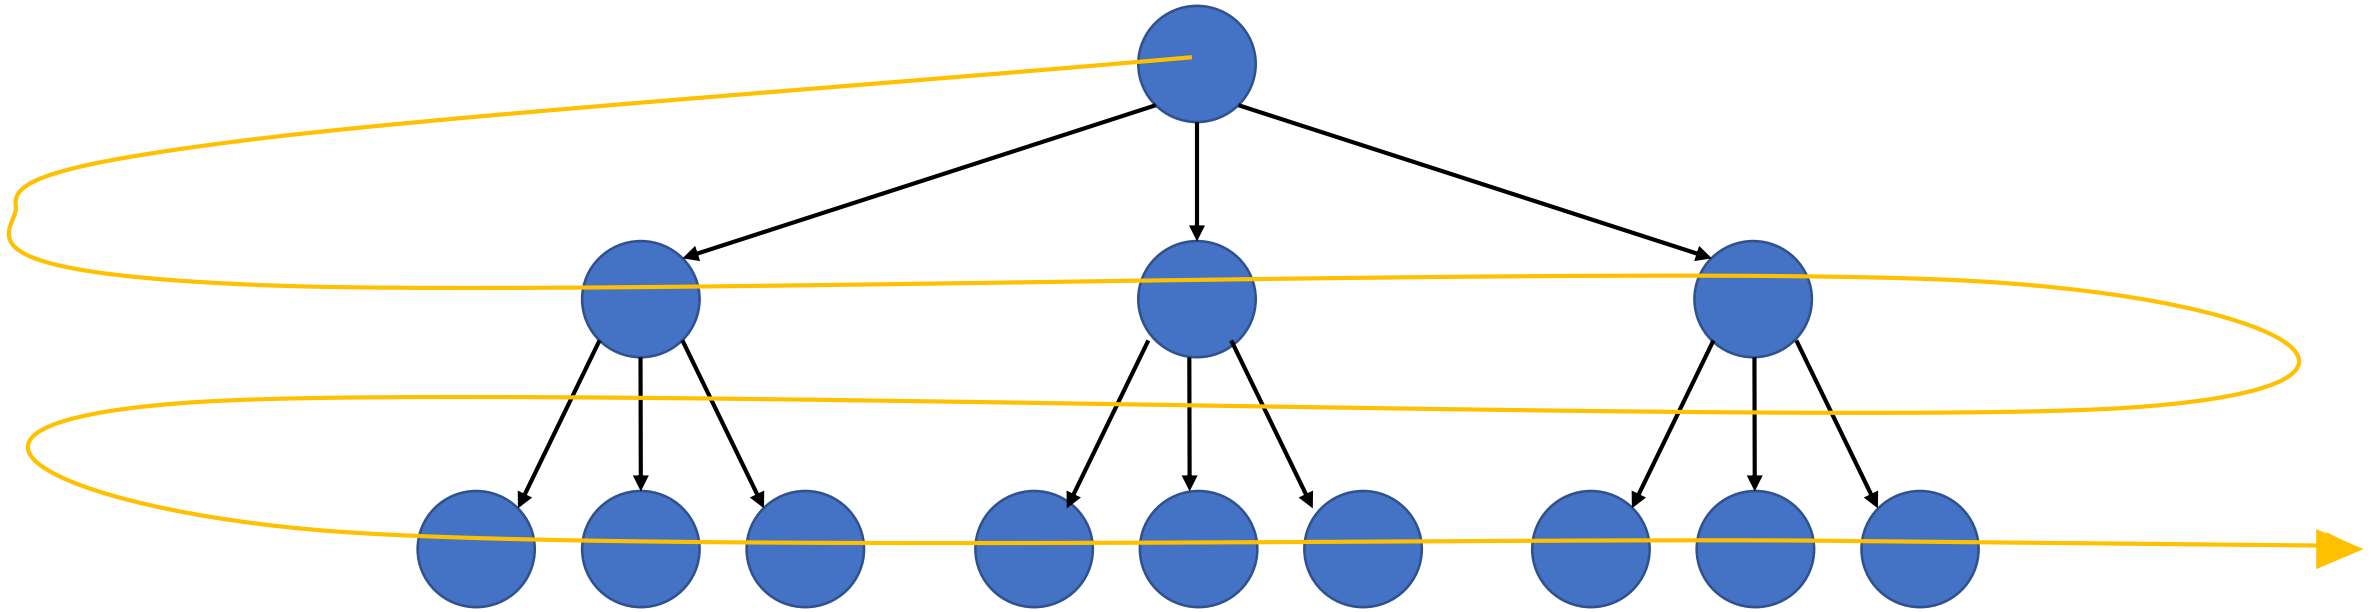

Depth-First Search

```
DFS(s,ops,queue) =
        If goal(s) then return(s);
          Else
                successors ← {}; queue' ← queue;
                For each o ∈ ops that applies to s
                        successors ← successors + apply(o,s);
                queue' ← append(queue',successors);
                If not(empty(queue'))
                        s' ← last(queue);
                        queue' ← remove(s',queue');
                        DFS(s',ops,queue')
                  Else return(FAIL)
Initial call: DFS(initialstate,ops,{})
```

Depth-First Search

```
DFS(s,ops,open) =
        If goal(s) then return(s);
          Else
                successors ← {}; open' ← open;
                For each o ∈ ops that applies to s
                        successors ← successors + apply(o,s);
                open' ← append(open',successors);
                If not(empty(open'))
                        s' ← last(open);
                        queue' ← remove(s',open');
                        DFS(s',ops,open')
                  Else return(FAIL)
Initial call: DFS(initialstate,ops,{})
```

Handling Repeated States

```
DFS(s,ops,open,closed) =
        If goal(s) then return(s);
          Else
                open' ← open;
                If s ∉ closed then
                        closed' ← closed + s;
                        successors ← {};
                        For each o ∈ ops that applies to s
                                successors ← successors + apply(o,s);
                        open' ← append(open',successors);
                   Else closed' ← closed;
                If not(empty(open'))
                        s' ← last(open');
                        open' ← remove(s',open');
                        DFS(s',ops,open',closed')
                Else return(FAIL)
Initial call: DFS(initialstate,ops,{},{})
```

Handling Repeated States

```
DFS(s,ops,open,closed) =
        If goal(s) then return(s);
          Else
                open' ← open;
                If s ∉ closed then
                        closed' ← closed + s;
                        successors ← {};
                        For each o ∈ ops that applies to s
                                successors ← successors + apply(o,s);
                        open' ← append(open',successors);
                  Else closed' ← closed;
                If not(empty(open'))
                        s' ← last(open');
                        open' ← remove(s',open');
                        DFS(s',ops,open',closed')
                Else return(FAIL)
Initial call: DFS(initialstate,ops,{},{})
```

Do a test before successor generation

Iterative Deepening Search

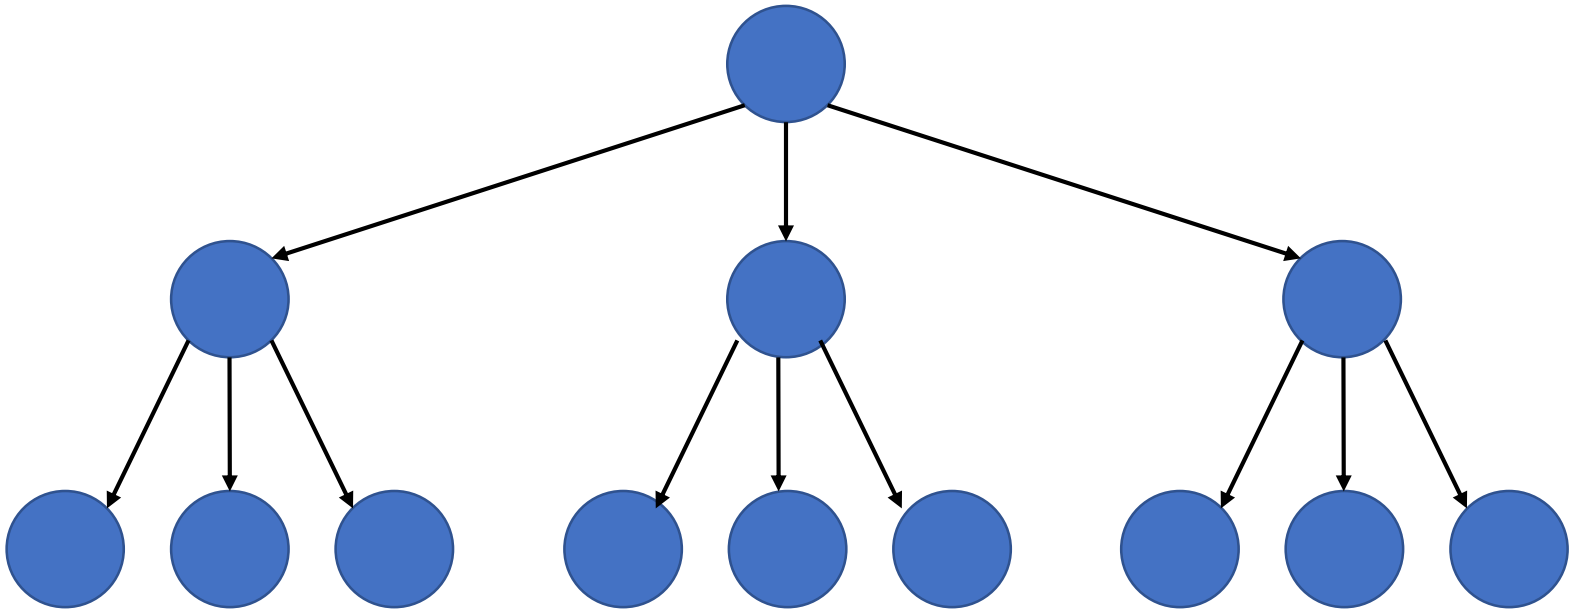

Iterative Deepening Search

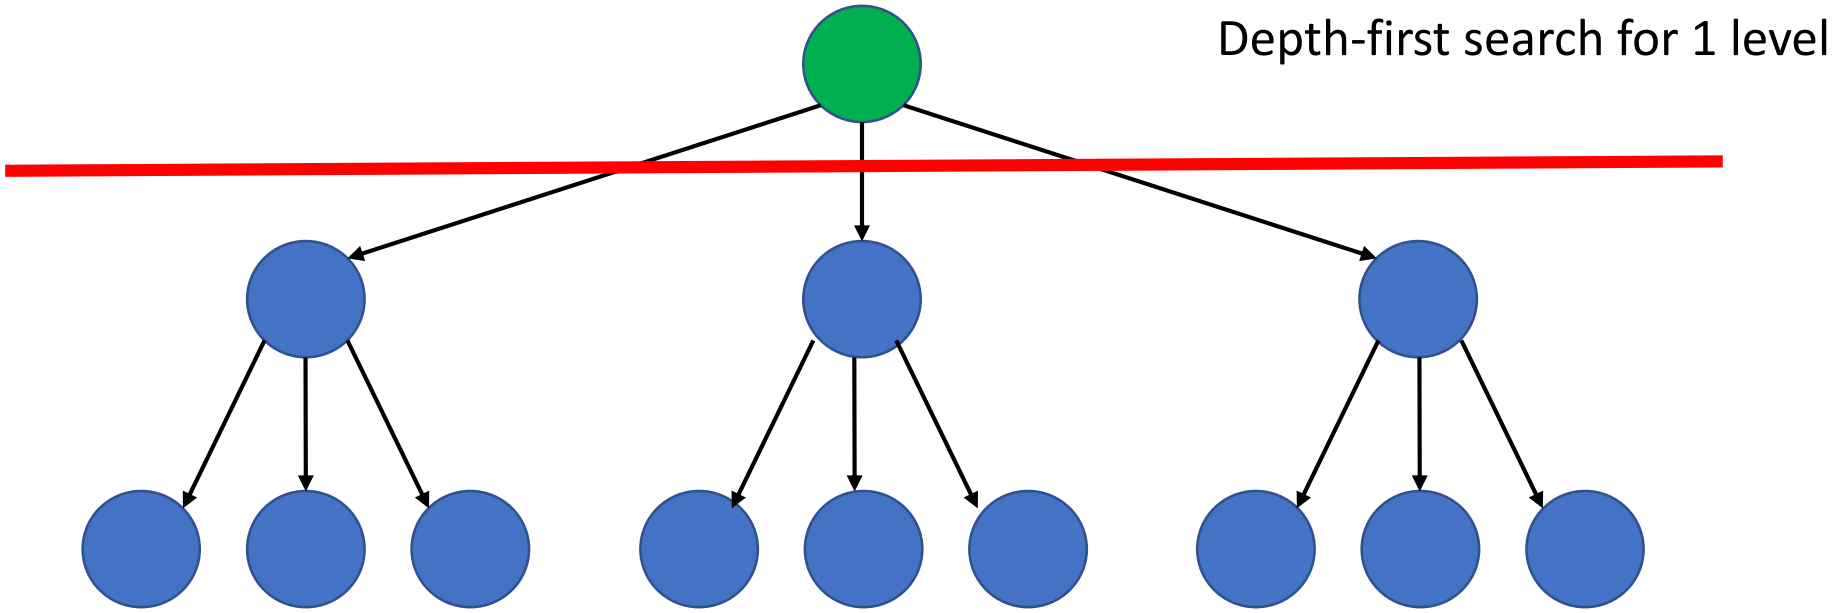

Iterative Deepening Search

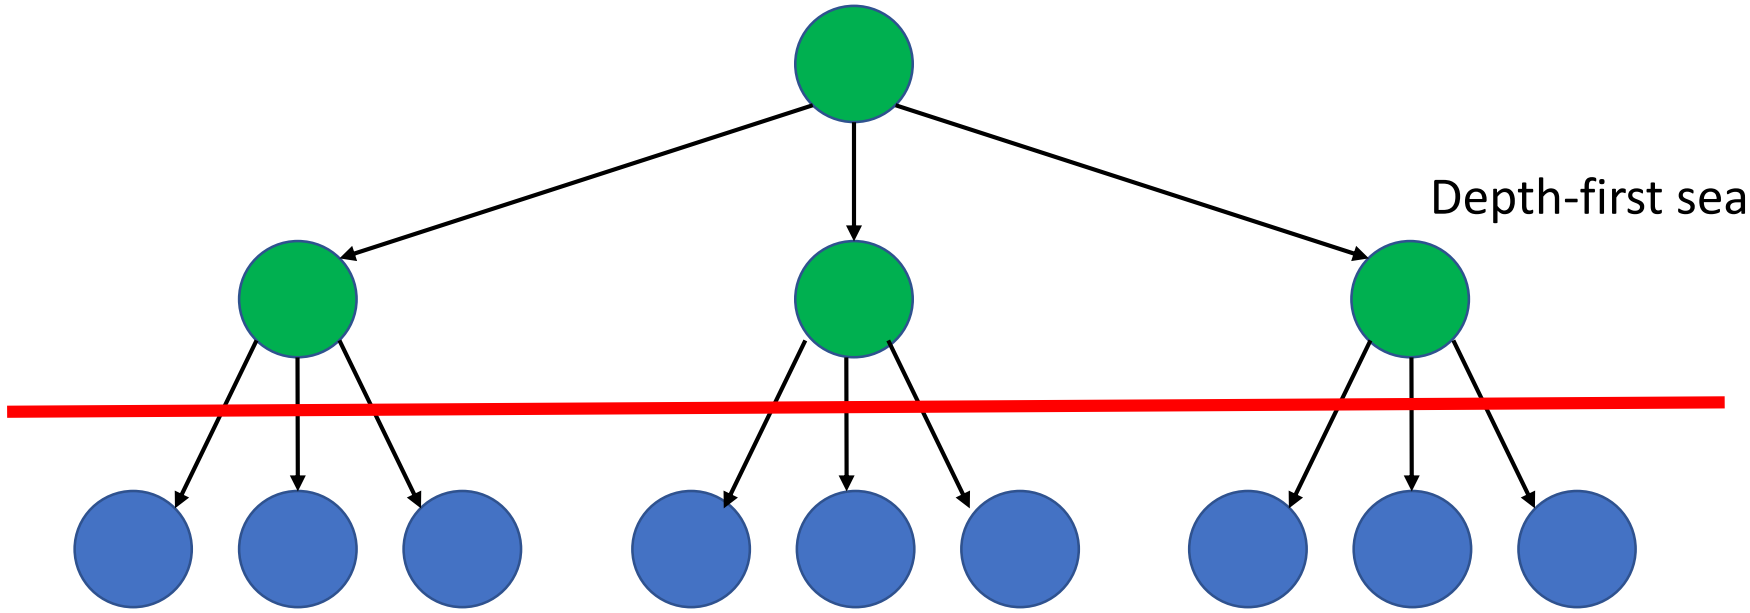

Iterative Deepening Search

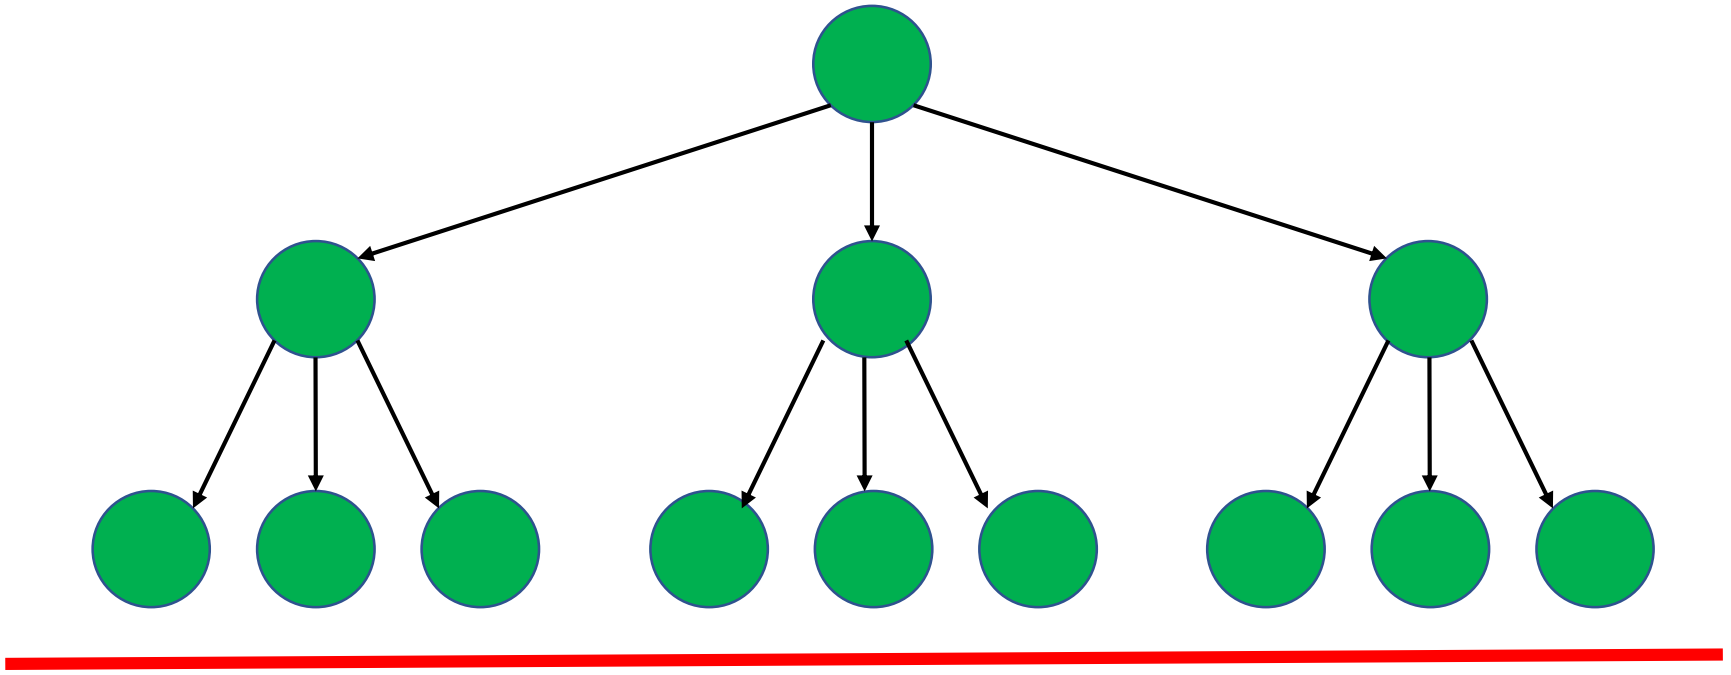

Depth-first search for 3 levels

Iterative Deepening Search

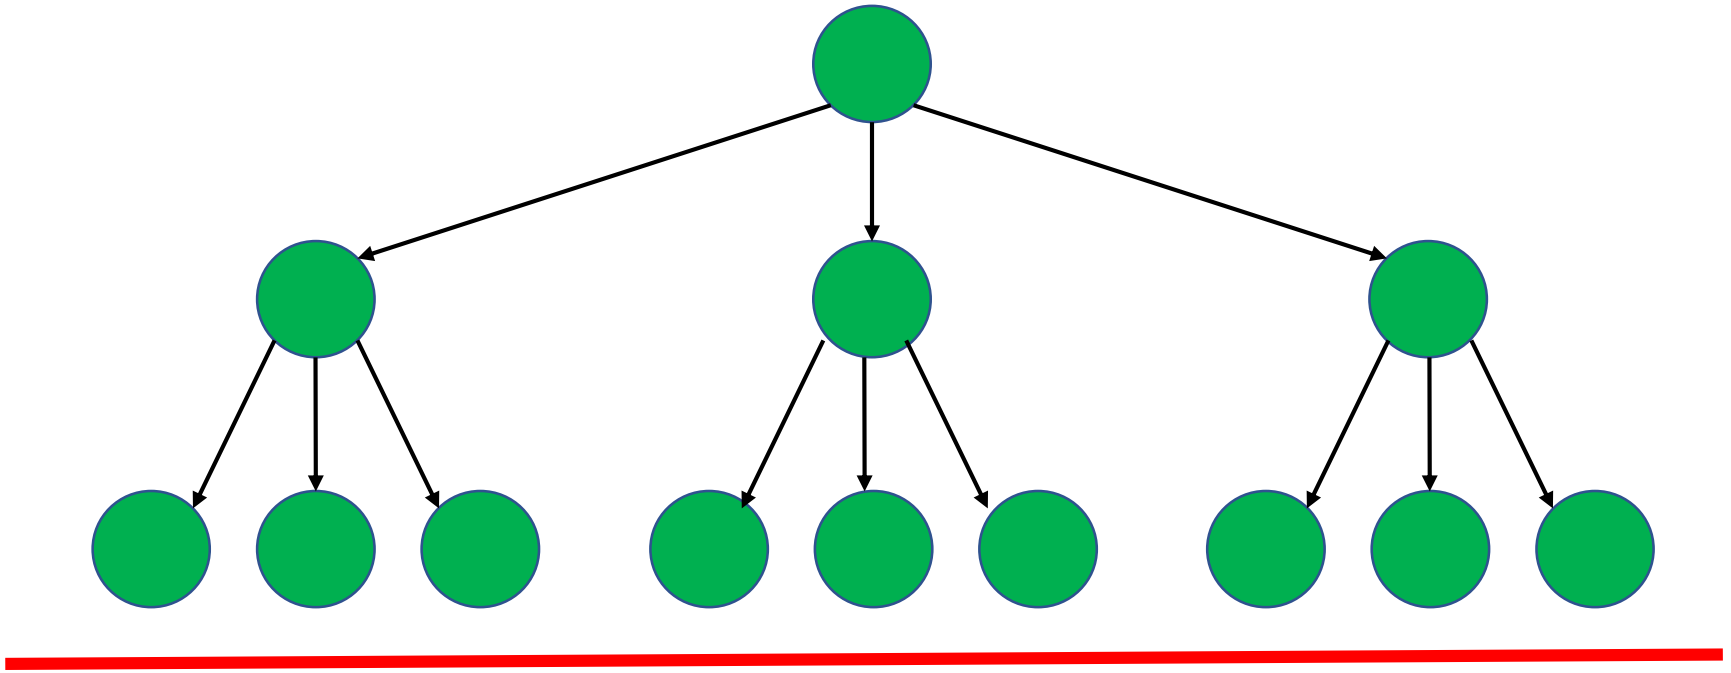

Depth-first search for 3 levels

Etc.

Iterative Deepening Search

Need a depth-first search with a depth bound

Handling Repeated States

```
DFS(s,ops,open,closed) =
        If goal(s) then return(s);
          Else
                open' ← open;
                If s ∉ closed then
                        closed' ← closed + s;
                        successors ← {};
                        For each o ∈ ops that applies to s
                                successors ← successors + apply(o,s);
                        open' ← append(open',successors);
                  Else closed' ← closed;
                If not(empty(open'))
                        s' ← last(open');
                        open' ← remove(s',open');
                        DFS(s',ops,open',closed')
                Else return(FAIL)
Initial call: DFS(initialstate,ops,{},{})
```

Do a test before successor generation

Depth-Bounded Depth-First Search

```
DBDFS(s,ops,open,maxdepth) =
        If goal(s) then return(s);
          Else
                open' ← open;
                if depth(s) < maxdepth
                        successors ← {};
                        For each o ∈ ops that applies to s
                                successors ← successors + apply(o,s);    {depth of each successor is set}
                        open' ← append(open',successors);                {to 1 more than the depth of s}
                If not(empty(open'))
                        s' ← last(open);
                        queue' ← remove(s',open');
                        DBDFS(s',ops,open',maxdepth)
                  Else return(FAIL)
Initial call: DBDFS(initialstate,ops,{},maxdepth)
```

Do a test before successor generation

Iterative Deepening Search

```
IDS(s,ops) =
        i ← 1;
        repeat
                result ← DBDFS(s,ops,{},i);
                 i ← i + 1;
        until result ≠ FAIL;
        return(result)
```

Iterative Deepening Search

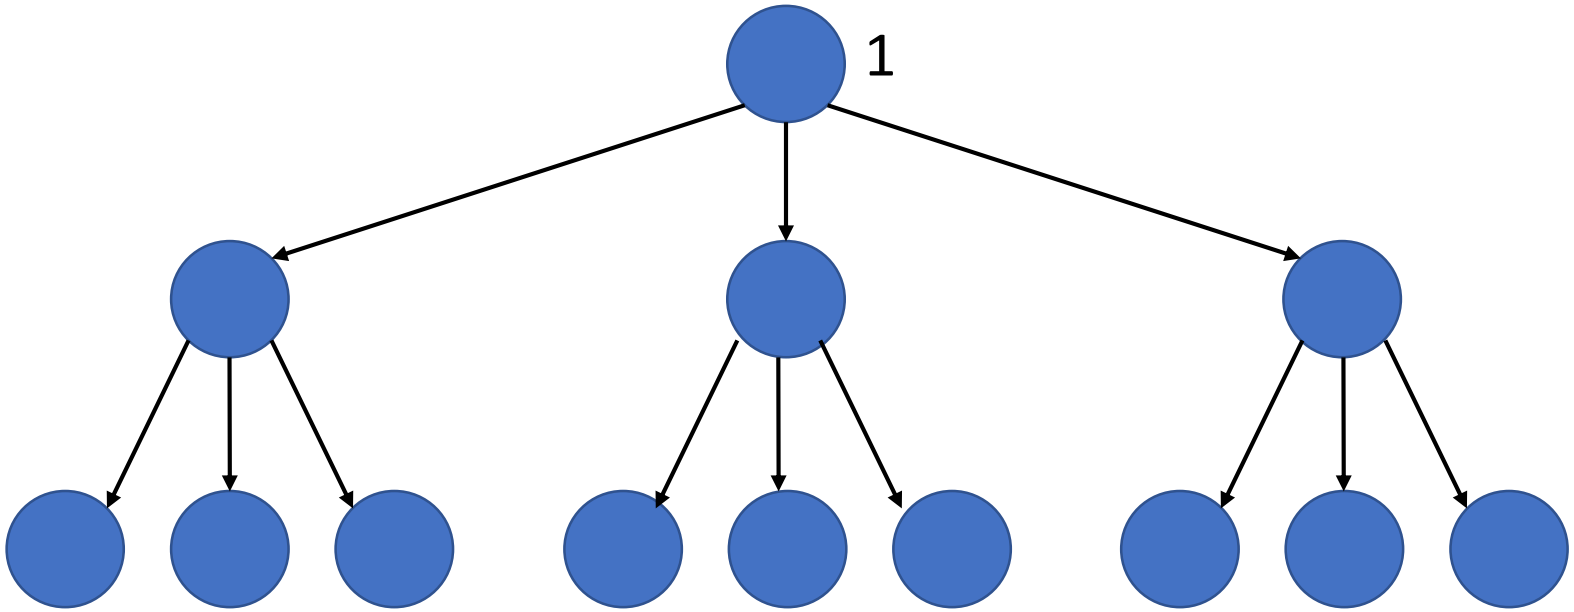

Iterative Deepening Search

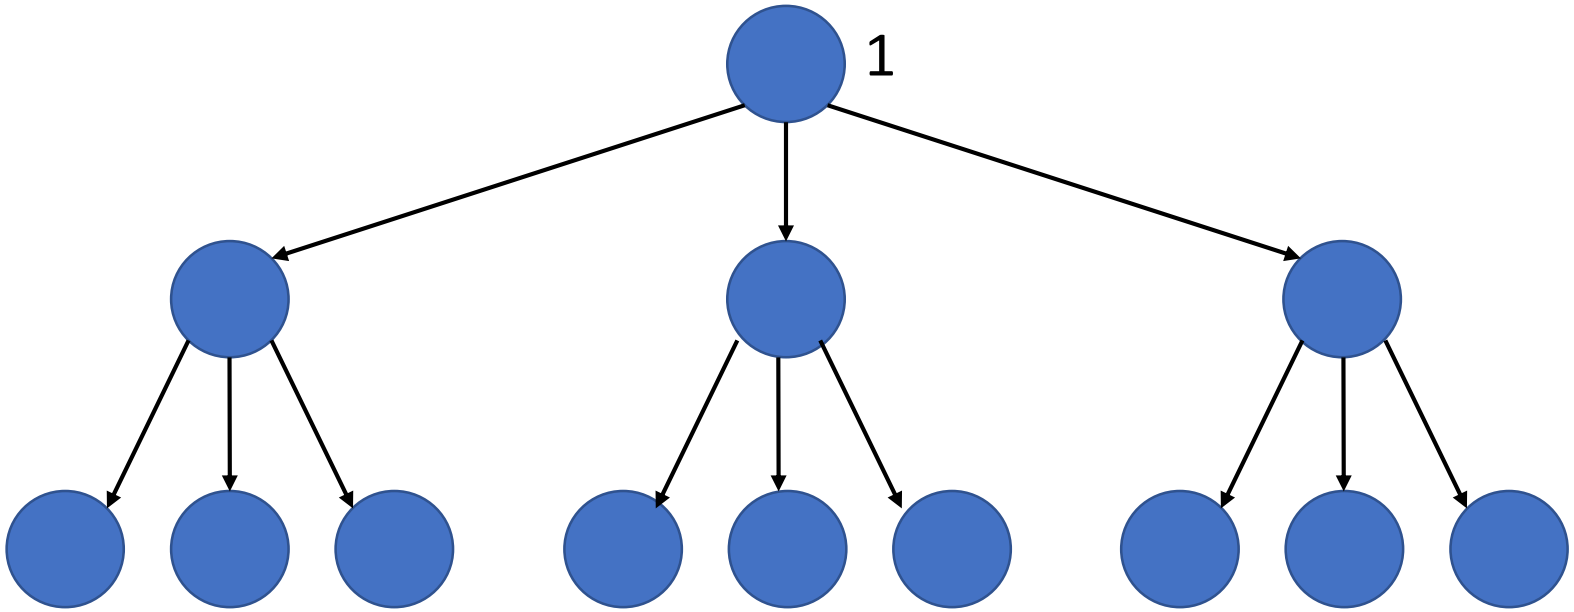

Iterative Deepening Search

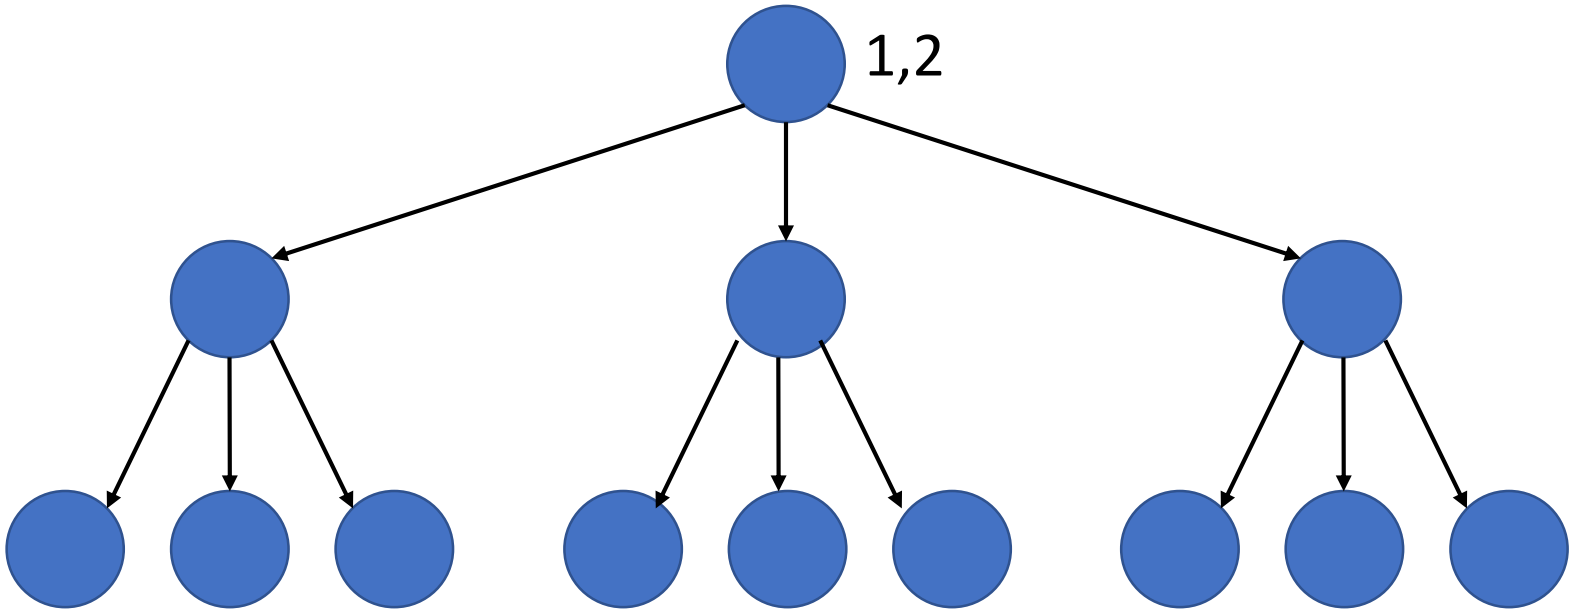

Iterative Deepening Search

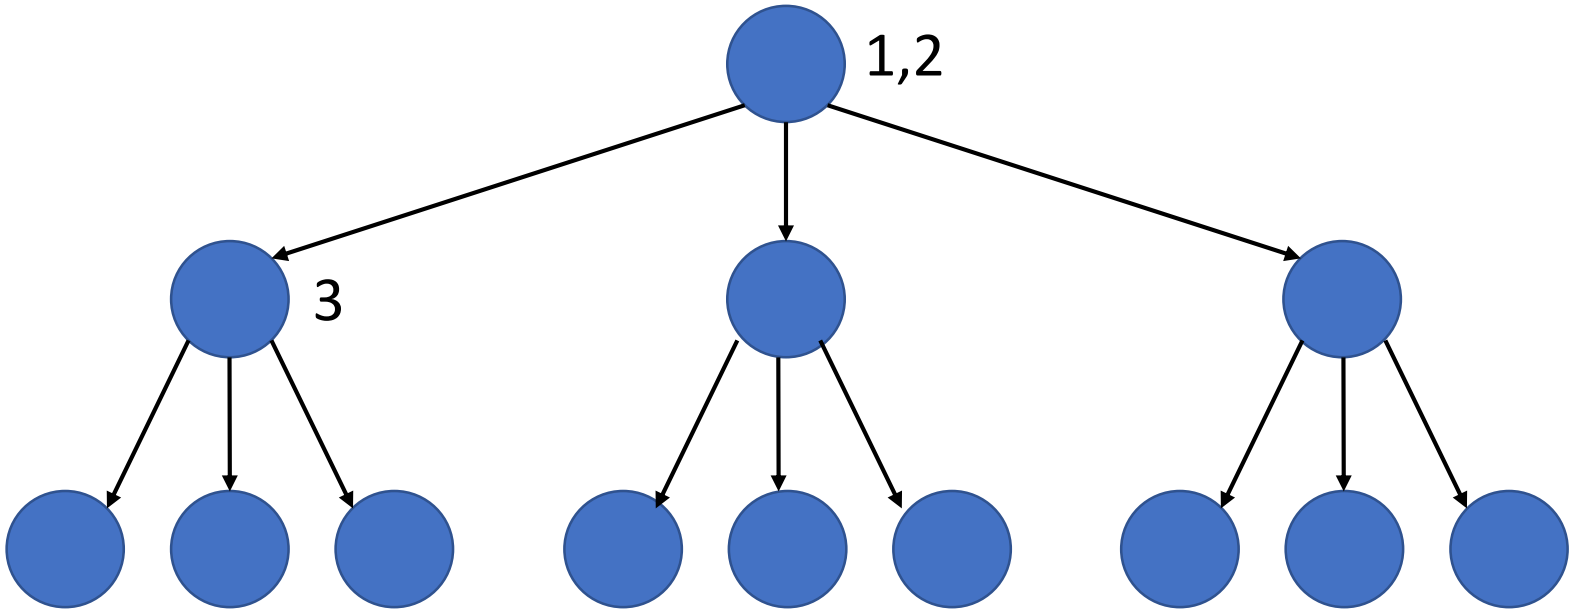

Iterative Deepening Search

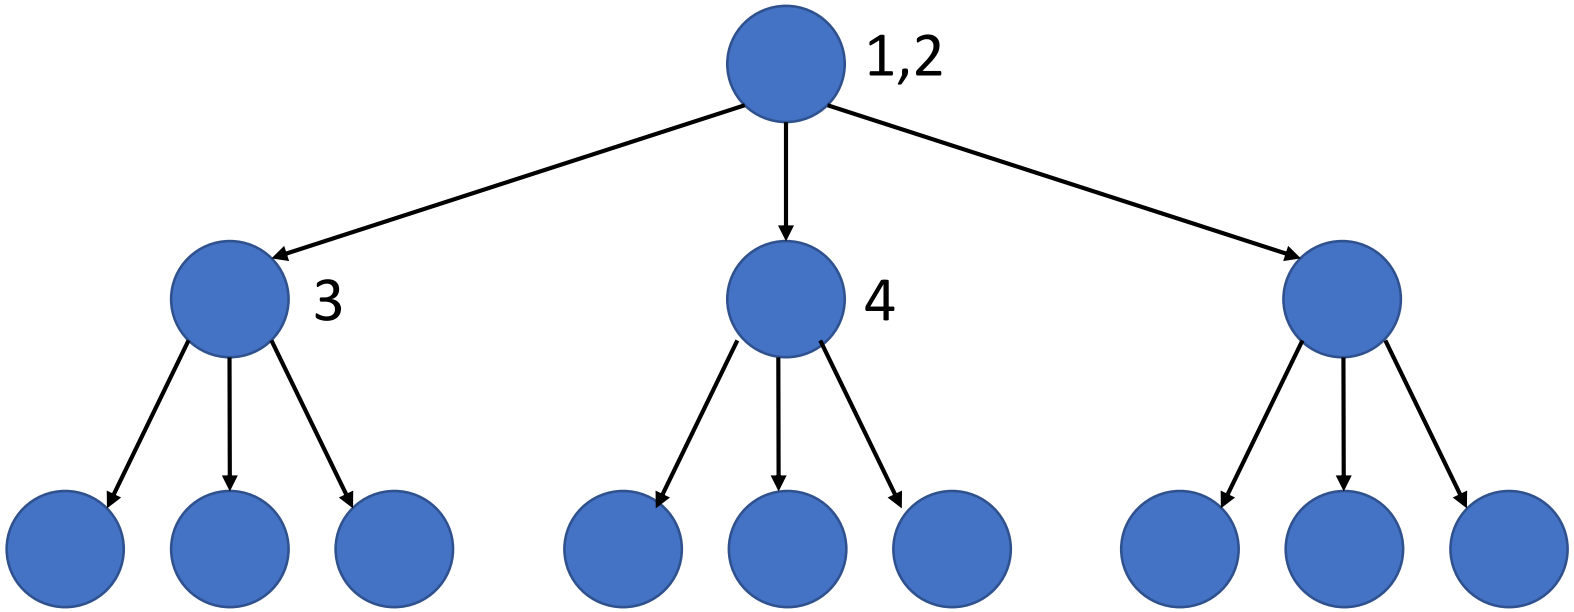

Iterative Deepening Search

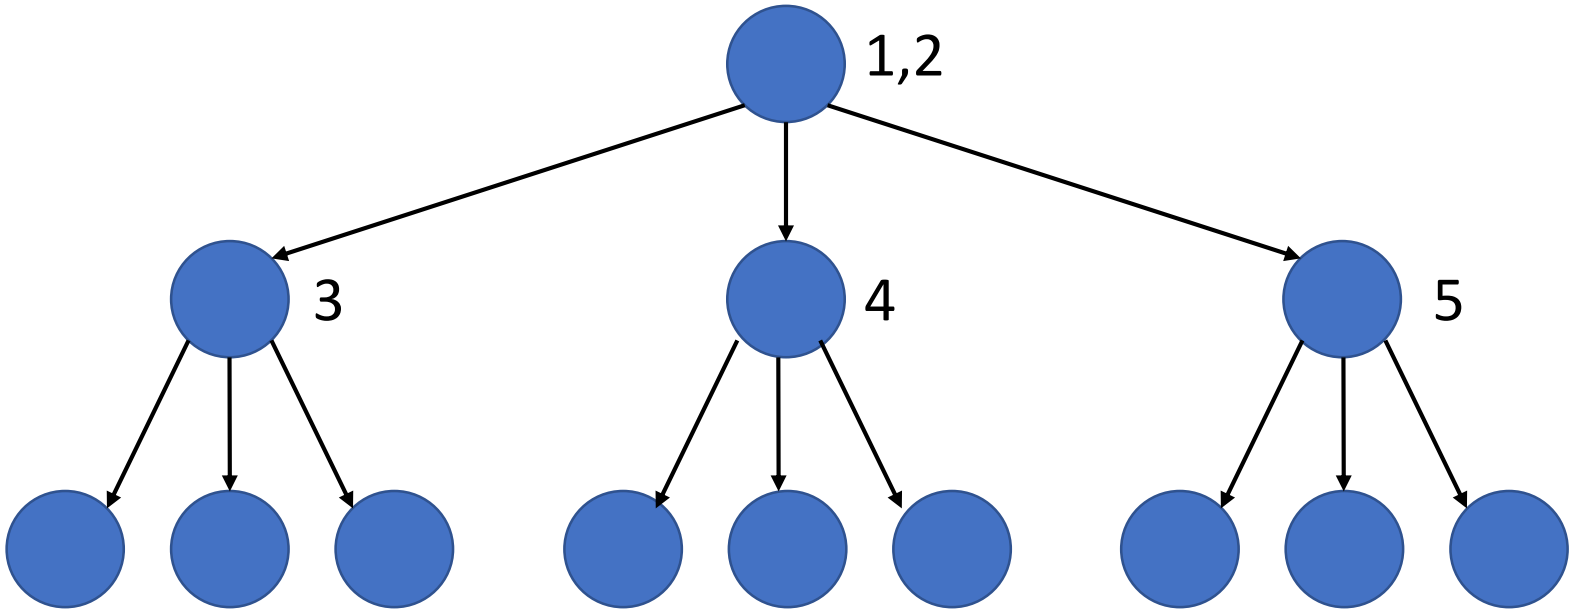

Iterative Deepening Search

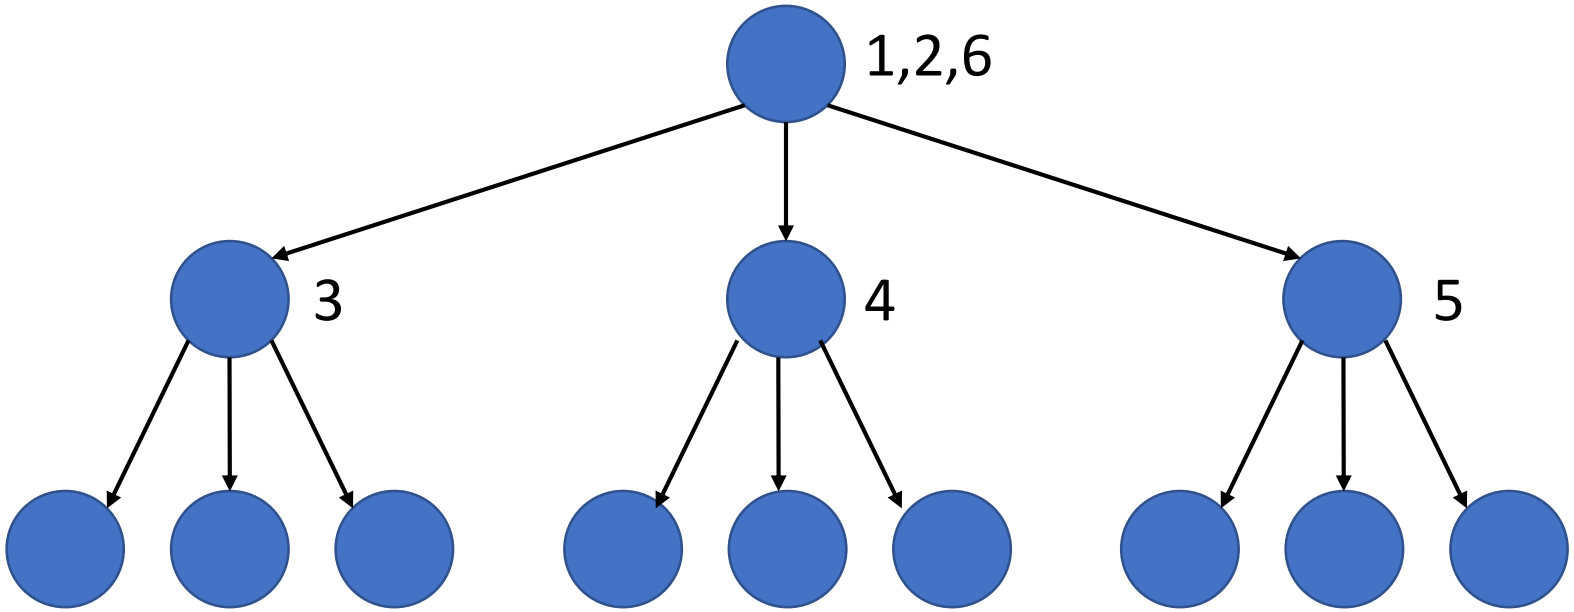

Iterative Deepening Search

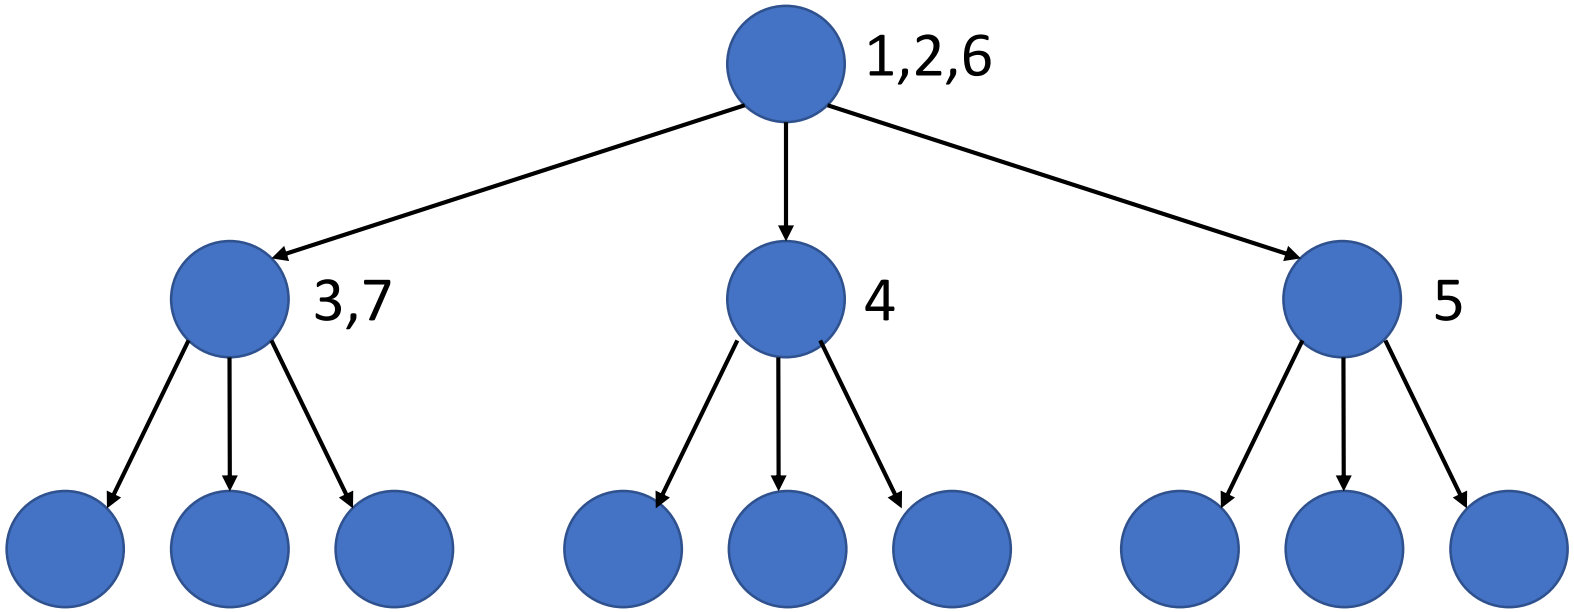

Iterative Deepening Search

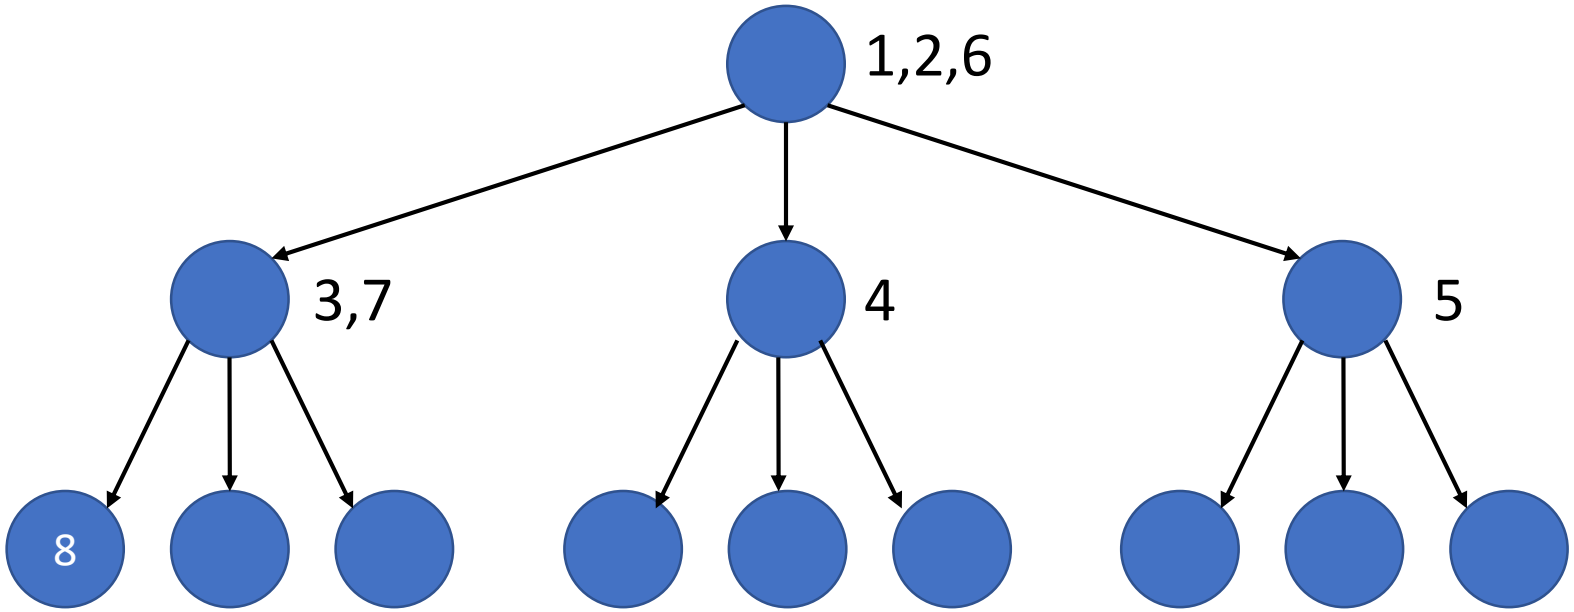

Iterative Deepening Search

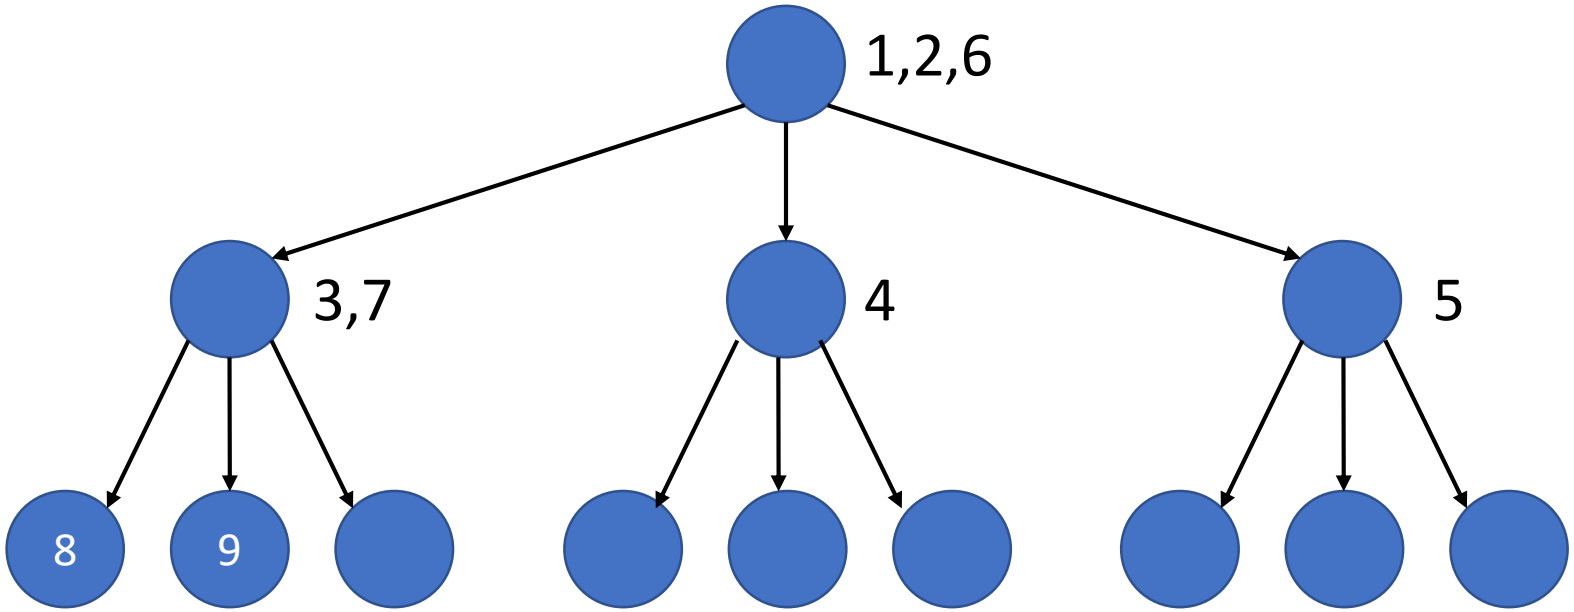

Iterative Deepening Search

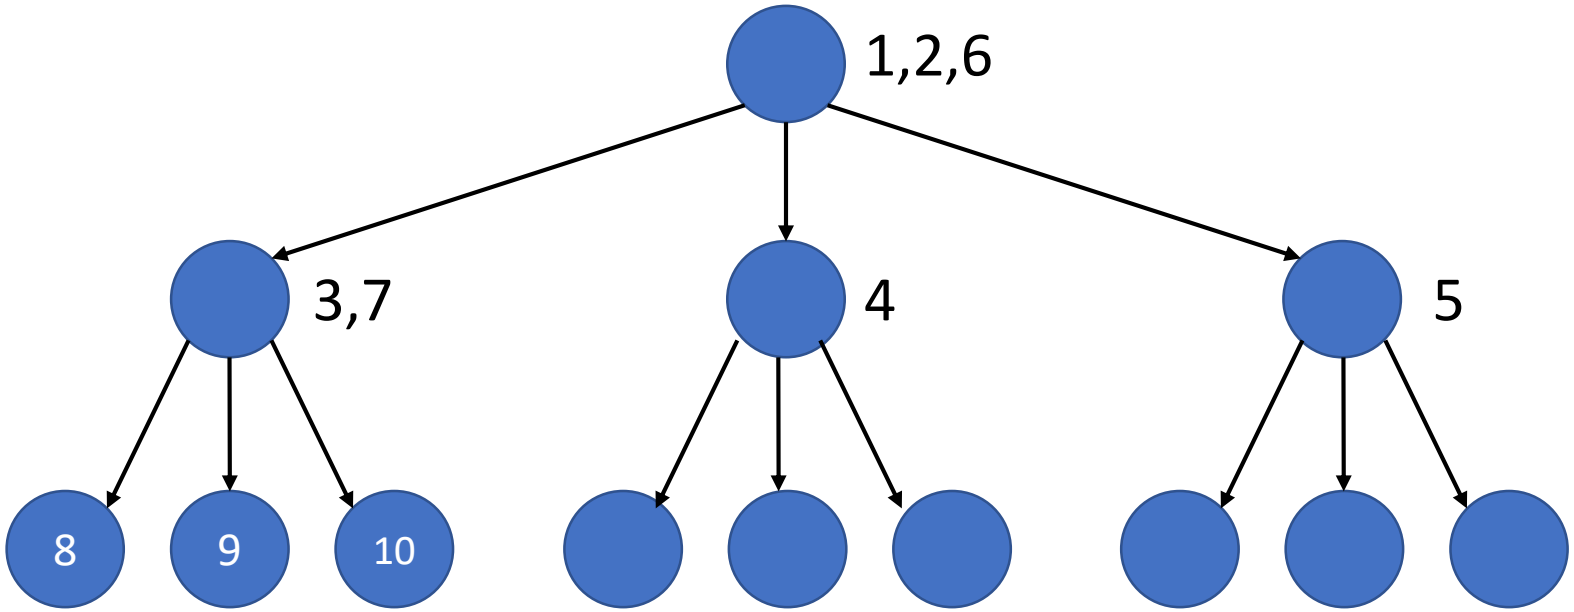

Iterative Deepening Search

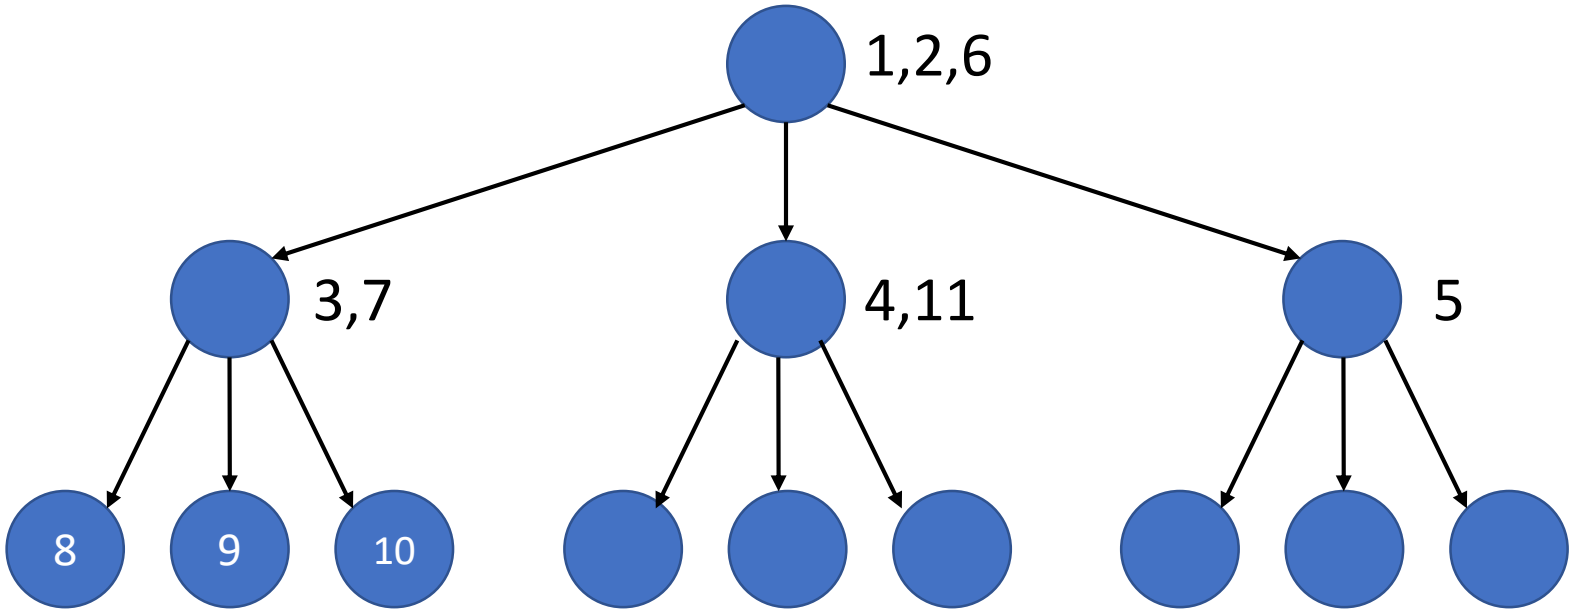

Iterative Deepening Search

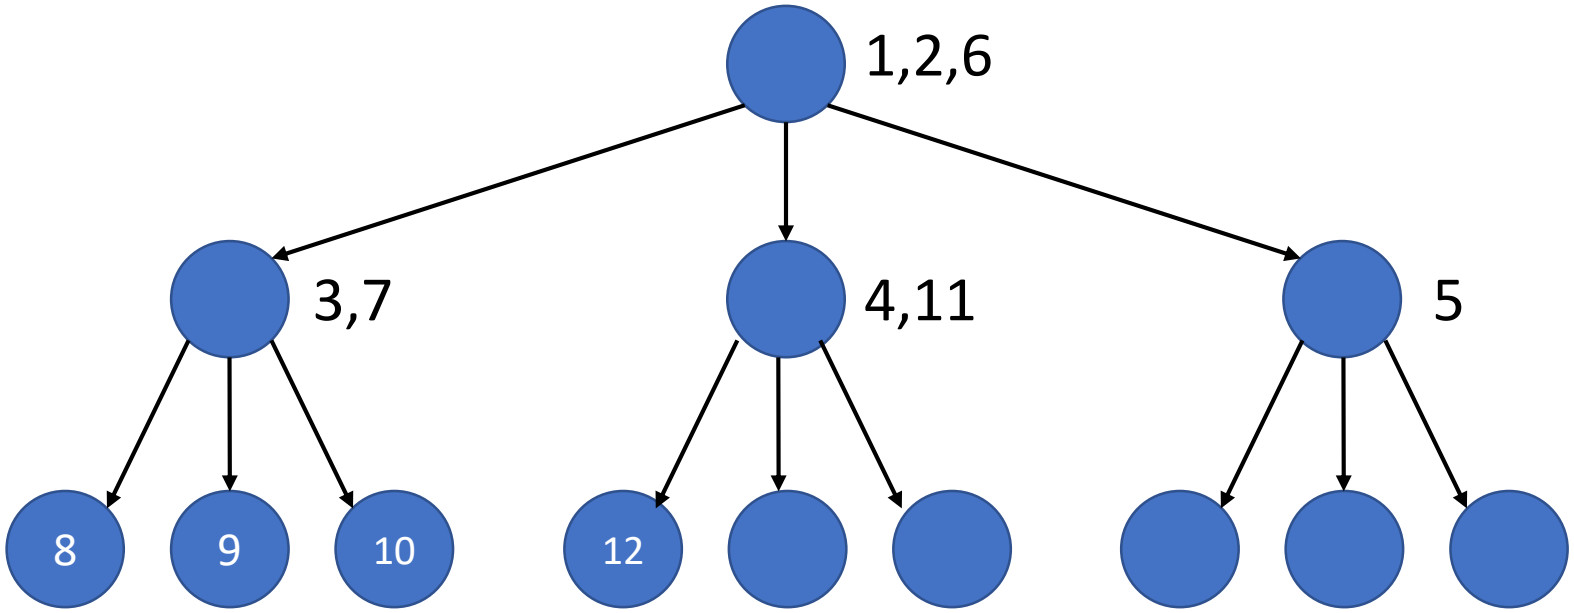

Iterative Deepening Search

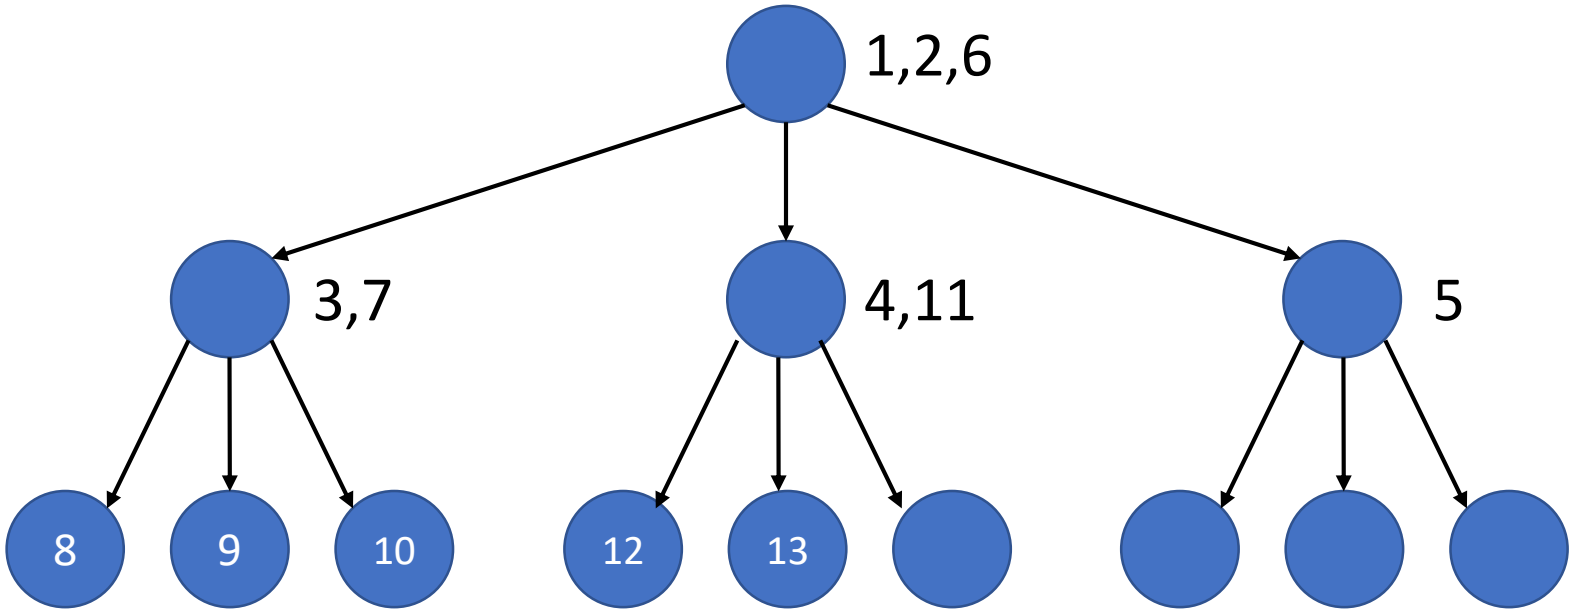

Iterative Deepening Search

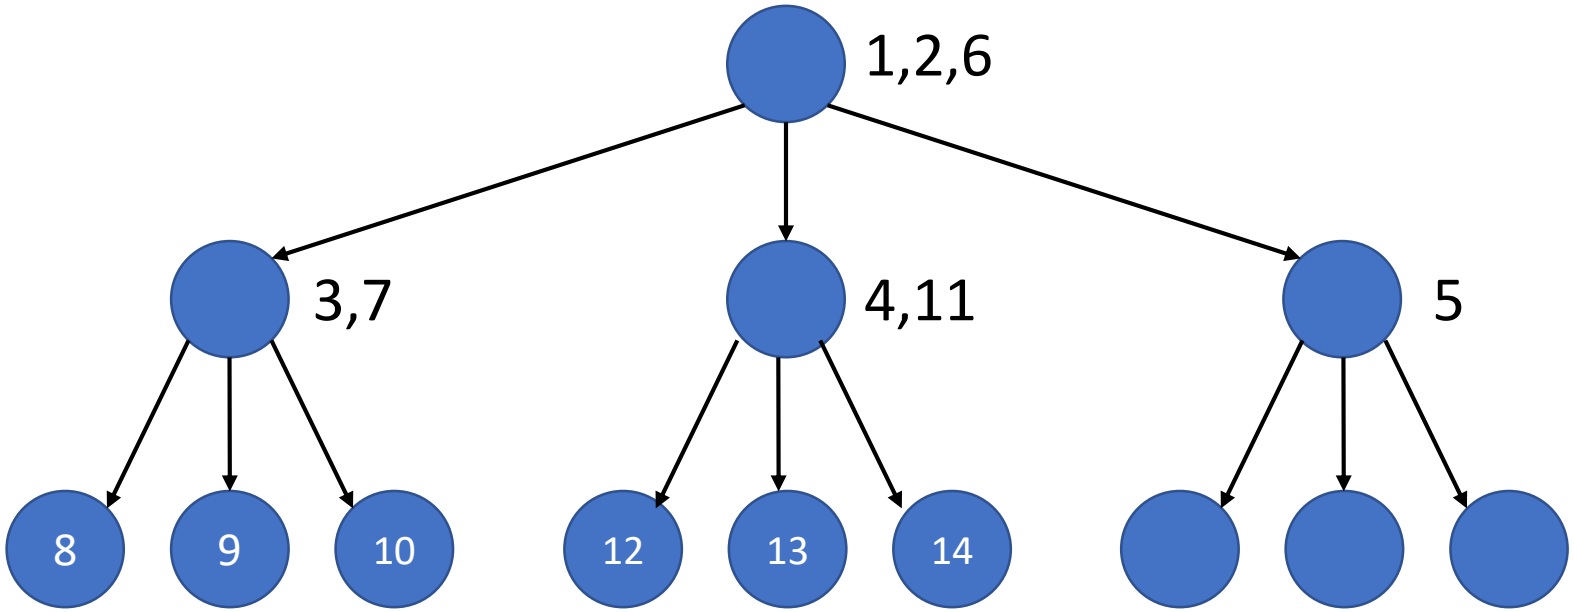

Iterative Deepening Search

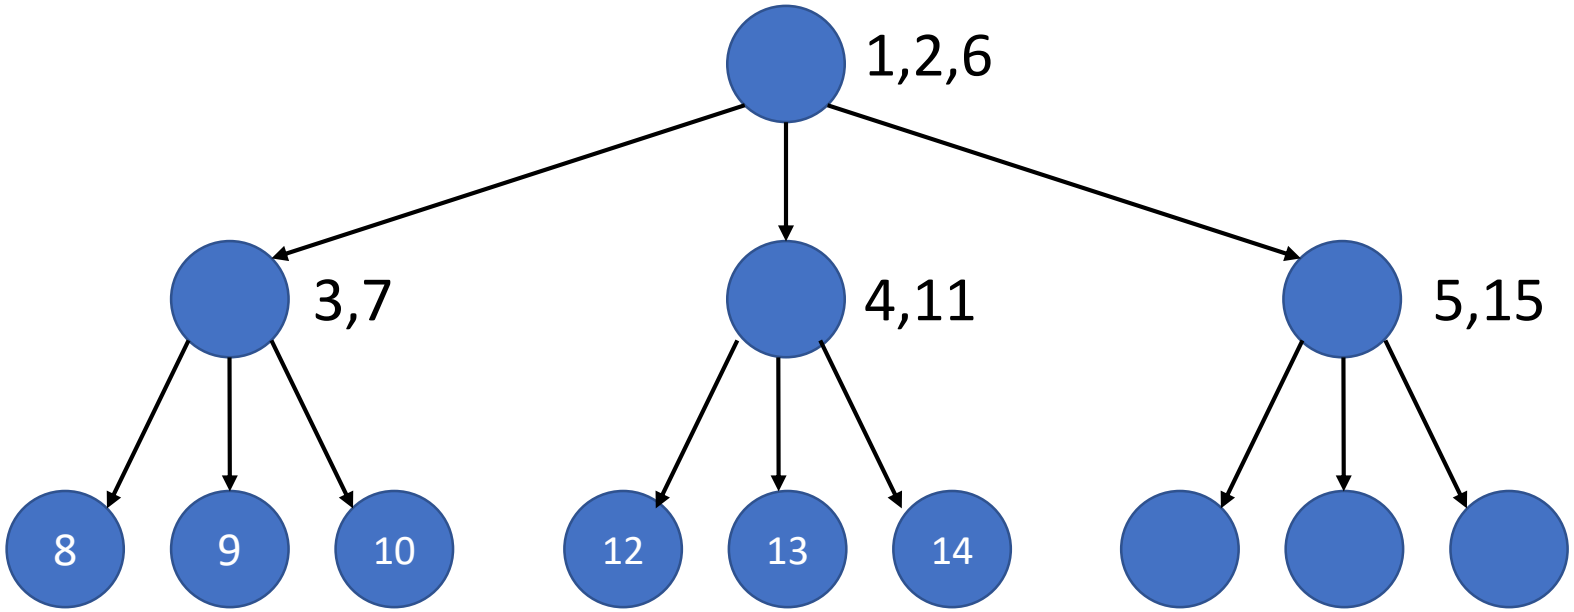

Iterative Deepening Search

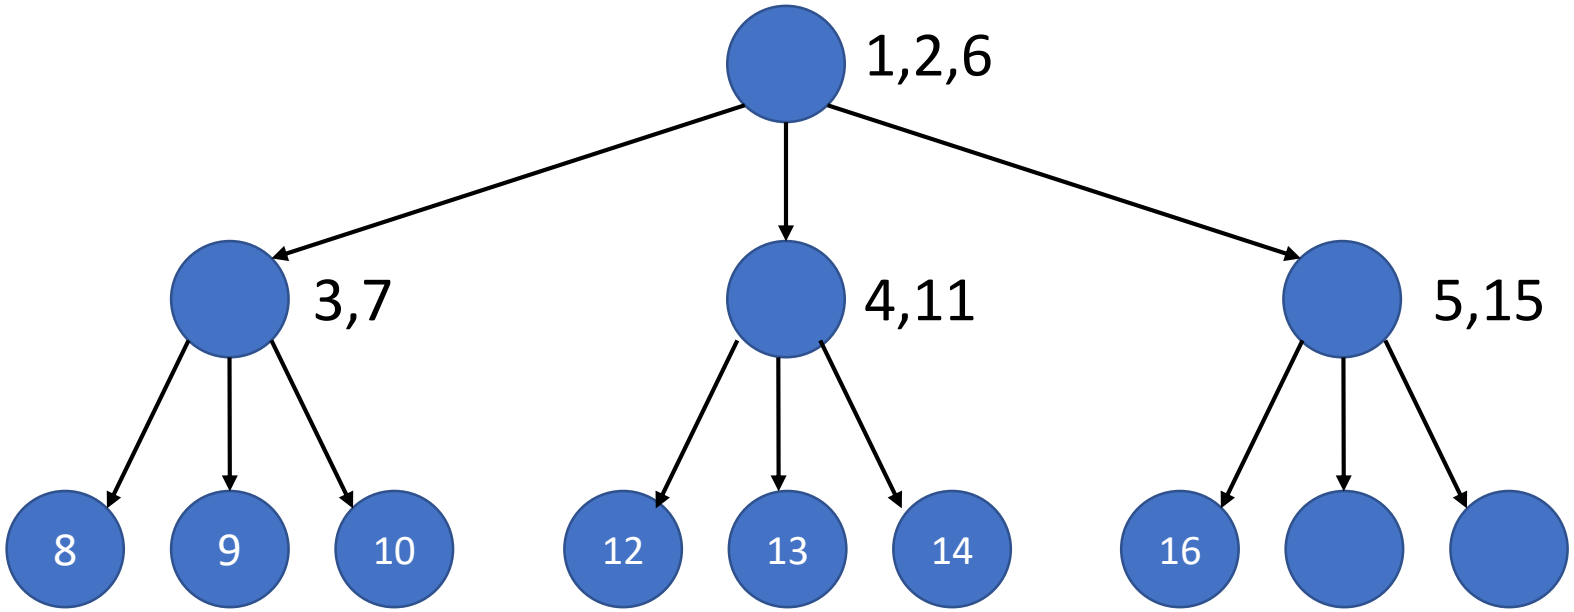

Iterative Deepening Search

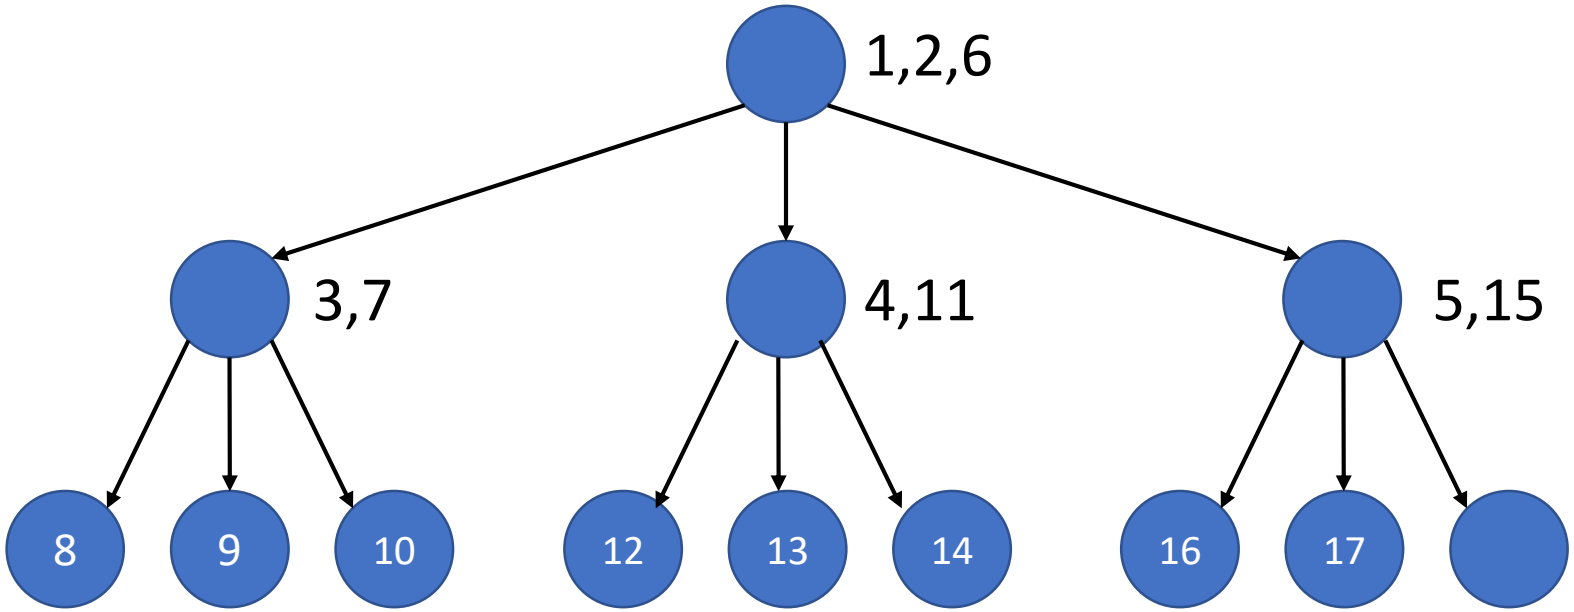

Iterative Deepening Search

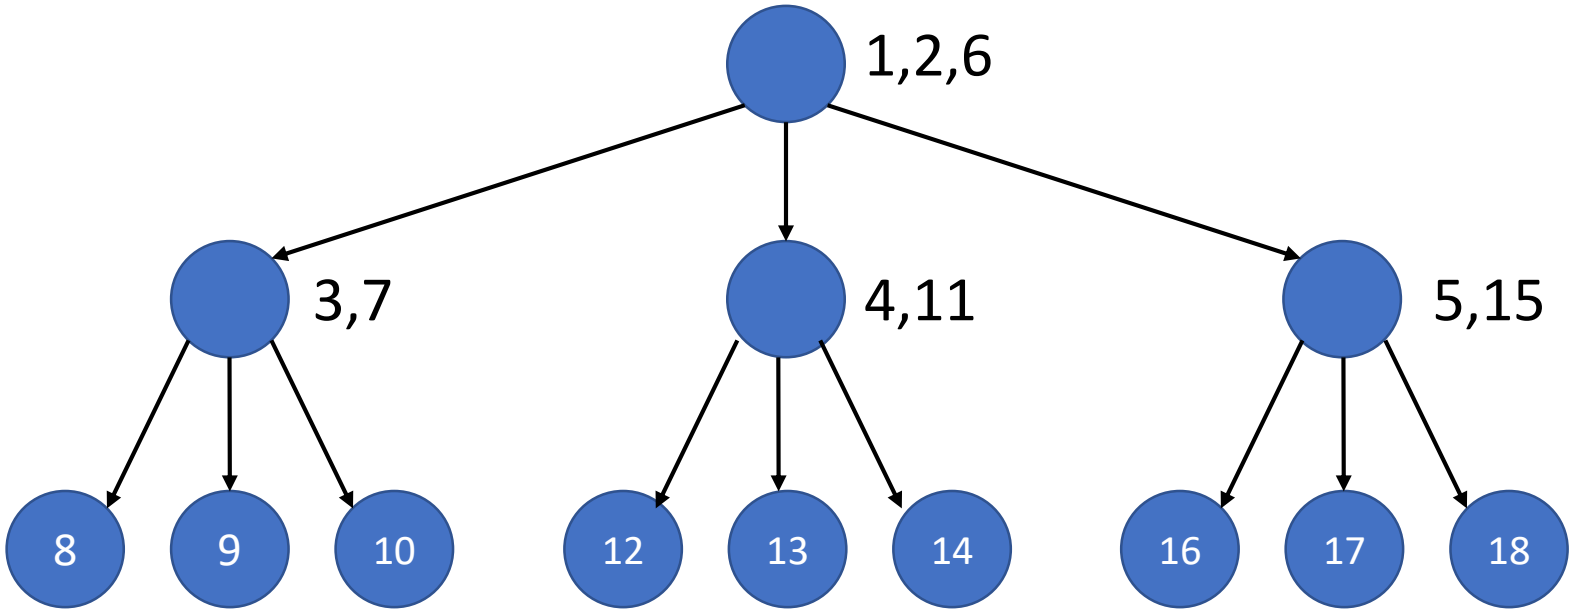

Analysis: Iterative Deepening

Criterion	DFS	BFS	Iterative Deepening
Complete?	No	Yes	
Optimal?	No	Yes	
Time?	∞	$O(b^d)$	
Space?	O(db)	$O(b^d)$	

Analysis: Iterative Deepening

Criterion	DFS	BFS	Iterative Deepening
Complete?	No	Yes	Yes
Optimal?	No	Yes	
Time?	∞	$O(b^d)$	
Space?	O(db)	$O(b^d)$	

Analysis: Iterative Deepening

Criterion	DFS	BFS	Iterative Deepening
Complete?	No	Yes	Yes
Optimal?	No	Yes	Yes
Time?	∞	$O(b^d)$	
Space?	O(db)	$O(b^d)$	

b=2	# States Visited		
Level	Best-First	Iterative Deepening	Ratio
0	1	1	1.000000
1	3	4	1.333333
2	7	11	1.571429
3	15	26	1.733333
4	31	57	1.838710
5	63	120	1.904762
6	127	247	1.944882
7	255	502	1.968627
8	511	1013	1.982387
9	1023	2036	1.990225
10	2047	4083	1.994626
11	4095	8178	1.997070
12	8191	16369	1.998413
13	16383	32752	1.999145
14	32767	65519	1.999542
15	65535	131054	1.999756
16	131071	262125	1.999870
17	262143	524268	1.999931
18	524287	1048555	1.999964
19	1048575	2097130	1.999981
20	2097151	4194281	1.999990
n	$2^{n+1}-1$	$\frac{2^{n+2}-n-3}{2^{n+1}-1}$	$\rightarrow 2$

b=10	# States Visited		
Level	Best-First	Iterative Deepening	Ratio
0	1	1	1.000000
1	11	12	1.090909
2	111	123	1.108108
3	1111	1234	1.110711
4	11111	12345	1.111061
5	111111	123456	1.111105
6	1111111	1234567	1.111110
7	11111111	12345678	1.111111
8	1.11E+08	123456789	1.111111
9	1.11E+09	1.235E+09	1.111111
10	1.11E+10	1.235E+10	1.111111
11	1.11E+11	1.235E+11	1.111111
12	1.11E+12	1.235E+12	1.111111
13	1.11E+13	1.235E+13	1.111111
14	1.11E+14	1.235E+14	1.111111
15	1.11E+15	1.235E+15	1.111111
16	1.11E+16	1.235E+16	1.111111
17	1.11E+17	1.235E+17	1.111111
18	1.11E+18	1.235E+18	1.111111
19	1.11E+19	1.235E+19	1.111111
20	1.11E+20	1.235E+20	1.111111
n	$\frac{10^{n+1}-1}{9}$	$\frac{\frac{10^{n+2}-1}{9}-(n+2)}{10^{n+1}-1}$	$\rightarrow \frac{10}{9}$

Analysis: Iterative Deepening

$$1 + (1+b) + (1+b+b^2) + \cdots + (1+b+\cdots b^d)$$

Analysis: Iterative Deepening

$$1 + (1+b) + (1+b+b^2) + \cdots + (1+b+\cdots b^d)$$
$$= \sum_{i=0}^{0} b^i + \sum_{i=0}^{1} b^i + \sum_{i=0}^{2} b^i + \cdots + \sum_{i=0}^{d} b^i$$

Analysis: Iterative Deepening

$$\sum_{i=0}^{n} b^i = \frac{b^{n+1} - 1}{b - 1}$$

Analysis: Iterative Deepening

$$1 \quad + \quad (1+b) \; + (1+b+b^2) + \cdots + (1+b+ \cdots b^d)$$

$$= \sum_{i=0}^{0} b^i + \sum_{i=0}^{1} b^i + \; \sum_{i=0}^{2} b^i \quad + \cdots + \sum_{i=0}^{d} b^i$$

$$= \frac{b^1 - 1}{b - 1} + \frac{b^2 - 1}{b - 1} + \cdots + \frac{b^{d+1} - 1}{b - 1}$$

Analysis: Iterative Deepening

$$1 + (1+b) + (1+b+b^2) + \cdots + (1+b+\cdots b^d)$$

$$= \sum_{i=0}^{0} b^i + \sum_{i=0}^{1} b^i + \sum_{i=0}^{2} b^i \ + \cdots + \sum_{i=0}^{d} b^i$$

$$= \frac{b^1 - 1}{b - 1} + \frac{b^2 - 1}{b - 1} + \cdots + \frac{b^{d+1} - 1}{b - 1}$$

$$= \frac{\sum_{i=1}^{d+1}(b^i - 1)}{b - 1} = \frac{\sum_{i=1}^{d+1} b^i - \sum_{i=1}^{d+1} 1}{b - 1} = \frac{\frac{b^{d+2} - 1}{b - 1} - 1 - (d+1)}{b - 1} = O\left(\frac{b^{d+2}}{b^2}\right) = O(b^d)$$

Analysis: Iterative Deepening

Criterion	DFS	BFS	Iterative Deepening
Complete?	No	Yes	Yes
Optimal?	No	Yes	Yes
Time?	∞	$O(b^d)$	$O(b^d)$
Space?	O(bd)	$O(b^d)$	

Analysis: Iterative Deepening

Criterion	DFS	BFS	Iterative Deepening
Complete?	No	Yes	Yes
Optimal?	No	Yes	Yes
Time?	∞	$O(b^d)$	$O(b^d)$
Space?	O(bd)	$O(b^d)$	O(bd)