

Runtime and Big O Notation


15-110 – Monday 02/23

Announcements

- **Hw3** was due today
 - How did it go?
- **No Check4** due to spring break. Hw4 is extra-large instead – start early

Announcements

- **Exam1** grades have been released. Median: **83.75/100**. Well done!



Minimum

10.5

Median

83.75

Maximum

100.0

Mean

81.28

Std Dev [?](#)

--

Learning Objectives

- Identify the **worst case** and **best case** inputs of functions
- Compare the **function families** that characterize different functions
- Calculate a specific function or algorithm's efficiency using **Big O notation**

Efficiency = Time = Money

We talk about efficiency a lot in this unit. Why do we care?

Computers are fast, but they can still take time to do complex actions. Faster algorithms can save lives, reduce user frustration, and increase company profits.

A major goal of computer scientists is not just to make algorithms that work, but algorithms that work **efficiently**. To measure this outcome, we need a way to measure how efficient a piece of code is.

Comparing Search Algorithms

Comparing Linear vs. Binary Search

Recall our comparisons of linear search vs. binary search in the previous lectures. How can we compare these algorithms at an **abstract** level?

We could run them on the same input and time them. However, how quickly a program runs varies based on lots of factors (the implementation, the machine, which other programs are running, etc.)

Instead, we'll count the number of **actions** the program takes on a **given input**. This lets us abstract away any noise.

Counting the number of actions

What actions might we count? Some lines of code may compose multiple operations into one line, and some actions may take longer than others to execute on the computer's hardware.

Instead of trying to count every action the computer takes, we could choose some specific action and count how many times the algorithm runs that action based on the **size of the input**.

For example, in linear and binary search we can count the total number of **comparisons to the target** (a basic action) based on the **length of the list** (the size of the input).

Linear vs. Binary Search: Search for 66

```
def linSearch(lst, target):  
    if len(lst) == 0:  
        return False  
    elif lst[0] == target:  
        return True  
    else:  
        return linSearch(lst[1:], target)
```

How many list elements are compared to 66 in this list of 15 elements?
linear search: 9 elements
binary search: 4 elements

```
def biSearch(lst, target):  
    if lst == []:  
        return False  
    else:  
        mid = len(lst) // 2  
        if lst[mid] == target:  
            return True  
        elif target < lst[mid]:  
            return biSearch(lst[:mid], target)  
        else: # lst[mid] < target  
            return biSearch(lst[mid+1:], target)
```

							1 st	4 th	3 rd		2 nd			
12	25	32	37	41	48	58	60	66	73	74	79	83	91	95

Best Case, Worst Case

Best Case and Worst Case

To truly compare the algorithms, it isn't enough to test them on a random example. We want to know how they'll do in the **best case** and in the **worst case**. Those cases are defined based on the properties of the **input** to the function.

Best case: an input of abstract size n that results in the algorithm taking the **least steps possible**.

Worst case: an input of abstract size n that results in the algorithm taking the **most steps possible**.

Best Case and Worst Case – Linear Search

What's the **best case** for linear search?

Answer: a list where the item we search for is in the first position

What's the **worst case** for linear search?

Answer: a list where the item we search for is not in the list.

Best Case and Worst Case – Binary Search

You do: what's the **best case** input and **worst case** input for binary search if we're counting comparisons?

Best Case/Worst Case Actions

How many actions do we perform in the **best case**?

For both linear search and binary search, there's just **one comparison** – when you find the item with the first comparison, you can exit the function immediately.

How many actions in the **worst case**?

In linear search, we have to check **every single element**. If the list has n elements, we do n **comparisons** before knowing the element is not in the list.

What about binary search?

Worst Case Action Count – Binary Search

Each call to binary search compares one item to the target. How many recursive calls (and therefore comparisons) do we make to binary search for different length lists?

List size	Also List size	# Recursive Calls
1	$1 = 2 - 1$	1
$1 * 2 + 1 = 3$	$3 = 4 - 1$	2
$3 * 2 + 1 = 7$	$7 = 8 - 1$	3
$7 * 2 + 1 = 15$	$15 = 16 - 1$	4
$15 * 2 + 1 = 31$	$31 = 32 - 1$	5
	$2^k - 1$	k
	n	$\log_2(n)$


When the input length doubles for linear search, it does **twice** as many comparisons.

When the input length doubles for binary search, it does **just one more comparison**.

It's logarithmic!

Linear Search vs Binary Search

If we map out how many actions linear search and binary search take on the same input sizes, we can see that $\log(n)$ is incredibly fast compared to n . Linear search is exponentially slower in the worst case!



Sidebar: Calculating Efficiency

Our implementation of binary search only looks better than our implementation of linear search because we **only count comparisons**.

Slicing a list also takes additional work, as the computer needs to create a copy of the list. Our recursive implementations of linear and binary search both slice the list on every call.

This is **inefficient** – we're doing more work than we need to. A better approach would be to pass the **reference** of the original list and change the indexes checked instead of changing the list itself.

Function Families

Function Families

When we count the actions taken by algorithms, we don't really care about one-off operations; we care about how the number of actions grows with the **size of the input** (the function parameters).

We'll call the abstract size of the input for a function **n** . This could be the number of elements in a list or the number of characters in a string.



In math, a **function family** is a set of equations whose outputs all grow at the same rate as their inputs grow. For example, an equation might grow linearly (at the rate of **n**) or quadratically (at the rate of **n^2**).

Similar Algorithms


We saw that linear search executes n comparisons in the worst case, so it belongs to the **linear** function family.

Many other algorithms belong to this family too! For example, a function that counts how many times a character appears in a string and a function that looks for the largest element in an unsorted list would both run in linear time.


On the other hand, binary search belongs to the **logarithmic** function family, as it executes $\log(n)$ comparisons.



Common Function Families




Function Families and Constants



Notice that as n grows, the function family becomes much more important than the constants, and functions with the same function family behave similarly.

Alternate Visualization

Here's another way to think about the function families. Consider what happens when you **double** the size of the input.



Big O Notation

Big O Notation

We often want to express an algorithm's efficiency with a shorthand. Use **Big O notation** to do this.

Find the function family the algorithm belongs to by counting actions in terms of **n** (the size of the input) and identifying the **dominant term** (highest power of n). Ignore constant factors and smaller terms.

# actions	Big O
n	$O(n)$
$3n + 5$	$O(n)$
$20n + n^2 + 1$	$O(n^2)$
$3 \log_2(n)$	$O(\log n)$

Unless specified otherwise, the Big O of an algorithm refers to its runtime in the **worst case** (computer scientists are pessimists).

Caveat: this is a simplified definition. If you take other CS classes, you'll learn more about how Big-O actually works.

Big O Calculation Strategy

How can we count the number of actions in an algorithm?

We can usually get a rough estimate by just counting the number of statements that will run. This is affected by **function calls**, **loops**, **recursion**, and certain **operations** that are influenced by the size of the input.

Sequential Statements are Added

First, if you have statements that run sequentially, the runtime for each statement is **added** to the total runtime.

```
def example(x): # n = x
    x = x + 5 # O(1)
    y = x + 2 # O(1)
    print(x, y) # O(1)
# Total: O(1)
```

Note: $O(1)$ means constant time; the statement is not affected by the size of the input.

Functions Have Their Own Runtimes

What if we call a function? We need to add **that function's runtime** to the overall runtime. Note that it may not be $O(1)$!

```
def printContains(lst, x): # n = len(lst)
    result = linearSearch(lst, x) #  $O(n)$ 
    print(x, "in lst?", result) #  $O(1)$ 
# Total:  $O(n)$ 
```

Loops Multiply Runtimes

Most non-constant runtimes come from some use of loops or recursion. This is because loops let us **repeat actions**, so we have to **multiply** the runtime of the loop body by the number of times the loop repeats.

```
def checkForAll(lst1, lst2): # n = len(lst1) = len(lst2)
    for i in range(len(lst1)): # n repetitions
        printContains(lst2, lst1[i]) # O(n) (from last slide)
# Total: O(n2)
```

Loop Repetitions

Careful – the number of times a loop repeats depends on the loop control variable!

```
def countDigits(n): # n = n
    count = 0 # O(1)
    while n > 0: # loops  $\log_{10}(n)$  times
        n = n // 10 # O(1)
        count = count + 1 # O(1)
    return count # O(1)
# Total:  $O(\log n)$ 
```

Even though this is $\log_{10}(n)$, we don't include the base in the Big O notation because a change of base is just a multiplicative factor.

Conditionals are Sequential

If we multiply loop bodies by the number of repetitions, do we do the same thing to conditionals?

No! Conditionals act more **sequentially** – in the worst case you run the Boolean test, then you run the body. You should add the two together.

```
def valueInBoth(lst1, lst2, value): # n = len(lst1) = len(lst2)
    if linearSearch(lst1, value): # O(n)
        return linearSearch(lst2, value) # O(n)
    return False # O(1)
# Total: O(n)
```

Recursion Multiplies by # Calls

When we use recursion, the amount of work done is the work done by the whole function multiplied by the number of recursive calls that will happen.

```
def factorial(n): # n = n
    if n == 1: # O(1)
        return 1 # O(1)
    else: # O(1)
        return n * factorial(n-1) # n-1 calls
# Total: O(n)
```

Each call does $O(1)$ work but there are n calls total.

Be Careful of Built-in Runtimes!

Functions we define aren't the only ones that can have non-constant runtimes; some of the built-in Python functions and operations have non-constant runtimes too.

```
def countAll(lst): # n = len(lst)
    for i in range(len(lst)): # n repetitions
        count = lst.count(i) # O(n)
        print(i, "occurs", count, "times") # O(1)
# Total: O(n2)
```

We'll let you know on assignments and exams when a built-in method or operation is not constant time. But you are responsible for knowing about `in`, which just implements linear search. `in` applied to a list or string will run in $O(n)$.

Big O Calculation Rules

Type	Rule
Statement or Conditional	Add runtime to following statements (largest wins)
Function Call	Add runtime of the function based on input size
Loop	Multiply number of iterations * work done by body
Recursion	Multiply number of recursive calls * work done by single call

Algorithms Previously Seen...

Linear Search runs in $O(n)$ and **Binary Search** runs in $O(\log n)$, as discussed before.

drawGrid runs in $O(n^2)$ where n is the size of the grid, because we loop over n rows and for each row we loop over n columns.

Towers of Hanoi runs in $O(2^n)$ where n is the number of discs because each function call triggers two more function calls of size $n-1$.

Activity: Calculate the Big O of Code

Activity: predict the Big O runtime of the following piece of code.

```
def sumEvens(lst): # n = len(lst)
    result = 0
    for i in range(len(lst)):
        if lst[i] % 2 == 0:
            result = result + lst[i]
    return result
```

Complex Big O Example

Let's look at a more complex example together:

```
1: def example(lst): # n is len(lst)
2:     result = []
3:     for i in range(0, len(lst), 2):
4:         if lst[i] != lst[i+1]:
5:             average = (lst[i] + lst[i+1]) / 2
6:             if average in lst:
7:                 result.append(average)
8:     return count
```

Runtime: $\text{constant} + n/2 * (\text{constant} + \text{constant} + n + \text{constant}) =$
 $\text{constant} + \text{constant} * n^2 + \text{constant} * n = \mathbf{O(n^2)}$

Line 3 iterates $n/2$ times – we should multiply that by the work done by the loop body.

Line 4 is a conditional with a constant check – add it to the rest of the loop body.

Line 6 is a conditional with a $O(n)$ check – add n to the rest of the body.

Lines 2, 5, 7, and 8 don't depend on the size of the input; they're constant actions.

Big O Helps Us Fix Inefficiency

Why should we go through all this bother? Calculating the Big O of a function that runs slowly might help us identify **how to make the function more efficient**. This is especially important because we usually can't guarantee a small input.

```
def countPairs(lst): # n = len(lst)
    count = 0
    for i in range(len(lst)): # n loops
        for j in range(len(lst)): # n loops
            if i + 1 == j: # adjacent
                if lst[i] == lst[j]:
                    count = count + 1

    return count
# total: O(n^2)
```

```
def countPairs(lst): # n = len(lst)
    count = 0
    for i in range(len(lst)-1): # n-1 loops
        if lst[i] == lst[i+1]: # adjacent
            count = count + 1
    return count
# total: O(n)
```

Additional Learning: High-Speed Trading

Want more examples of how efficiency impacts real life? Check out this podcast episode on high-speed computer trading (where milliseconds make the difference between profit and loss):

<https://radiolab.org/episodes/267124-speed>

Learning Objectives

- Identify the **worst case** and **best case** inputs of functions
- Compare the **function families** that characterize different functions
- Calculate a specific function or algorithm's efficiency using **Big-O notation**

Extra Slides:

Common Big-O Runtimes

O(1) is Constant Time

```
def swap(lst, i, j): # n = len(lst)
    tmp = lst[i]
    lst[i] = lst[j]
    lst[j] = tmp
```

Does the runtime of this algorithm depend on the number of items in the list?

Answer: No.

This algorithm is **constant time** or **O(1)**; its time does not change with the size of the input.

$O(\log n)$ is Logarithmic Time

```
def countDigits(n): # n = n
    count = 0
    while n > 0:
        n = n // 10
        count = count + 1
    return count
```

Every time you increase n by a factor of 10, you run the loop one more time. All the operations in the loop are constant time. Similar to binary search, the algorithm is **logarithmic time**, or $O(\log n)$.

Even though this is $\log_{10}(n)$, we don't include the base in the Big-O notation because a change of base is just a multiplicative factor.

$O(n)$ is Linear Time

```
def countdown(n): # n = n
    for i in range(n, -1, -5):
        print(i)
```

This code will loop $n/5$ times overall. If we double the size of n , how many more times do we go through the loop?

Answer: We double the number of times through the loop. That is **linear time**, or **$O(n)$** , as it is proportional to the size of n . Stepping by 5 doesn't change the function family.

$O(n^2)$ is Quadratic Time

```
def multiplicationTable(n): # n = n
    for i in range(1, n+1):
        for j in range(1, n+1):
            print(i, "*", j, "=", i*j)
```

This seems tricky at first, but note that **every iteration** of the outer loop will do **all the work** of the inner loop.

The inner loop does n total iterations (with $O(1)$ work in its body). This is repeated n times by the outer loop. Therefore, the entire runtime is $O(n^2)$.

$O(2^n)$ is Exponential Time

```
# n = discs
def moveDiscs(start, tmp, end, discs):
    if discs == 1:
        print("Move one disc from",
              start, "to", end)
    else:
        moveDiscs(start, end, tmp, discs - 1)
        moveDiscs(start, tmp, end, 1)
        moveDiscs(tmp, start, end, discs - 1)
```

This is Towers of Hanoi.
Every time we add 1 disc
we double the number of
moves. That's **exponential
time**, or $O(2^n)$.

$$O(2^{n+1}) = O(2^n) + O(2^n)$$