

HW 7 Solutions up front

Randomized Algorithms

CSE 312 Spring 24
Lecture 27

Announcements

Last two concept checks (27, 28) available tonight.

CC27 (today's content) due Friday morning

CC28 (end-of-quarter-wrap-up) due Monday morning

You can fill that out tonight, don't need to look ahead at lecture content.

We'll have time at the end of tomorrow's lecture for questions. Bring practice problems or general questions if you have any.

What's a randomized algorithm?

A randomized algorithm is an algorithm that uses randomness in the computation.

Well, ok.

Let's get a little more specific.

Two common types of algorithms

Las Vegas Algorithm

Always tells you the right answer

Takes varying amounts of time.

Monte Carlo Algorithm

Usually tells you the right answer. Sometimes the wrong one.

A classic Las Vegas Algorithm

Remember Quick Sort?

Pick a "pivot" element

Move all the elements smaller than the pivot to the left subarray (in no particular order)

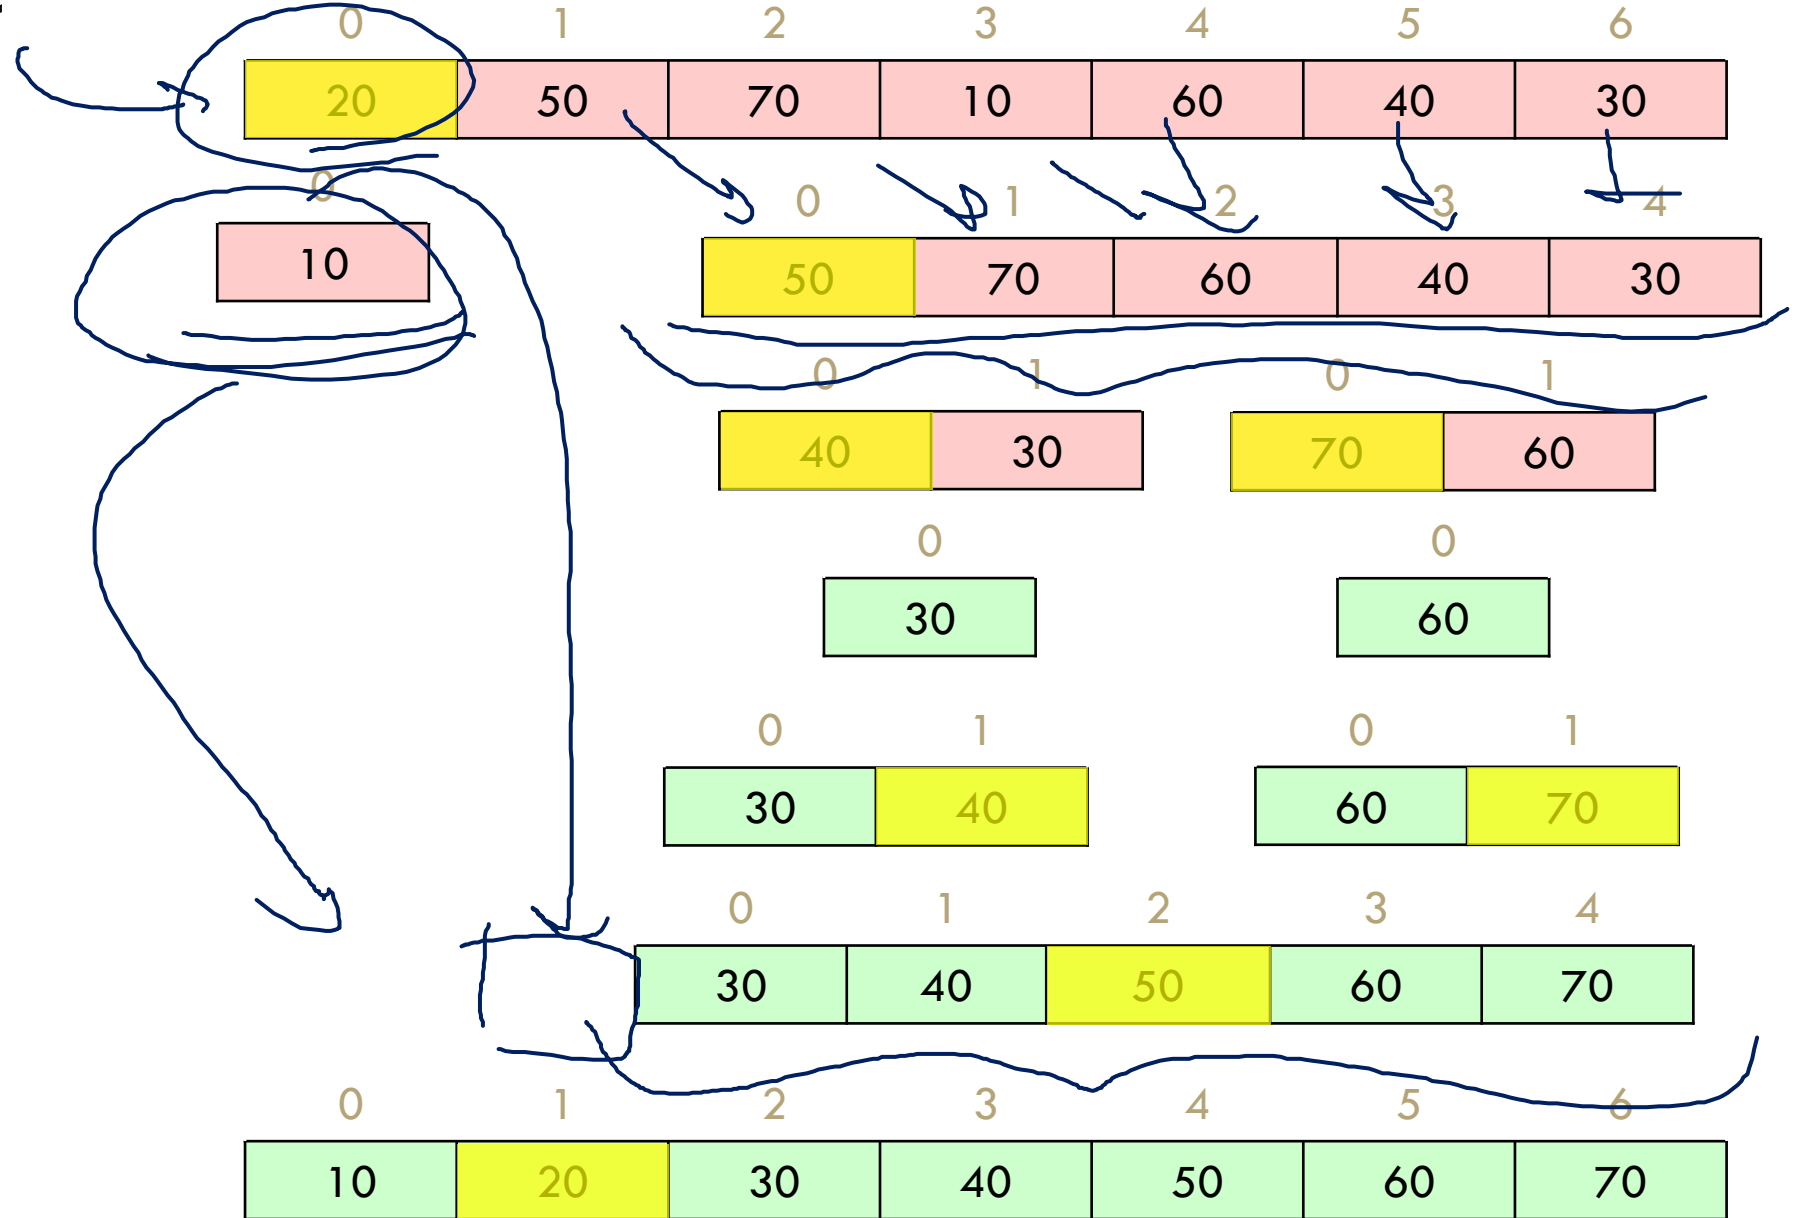
Move all elements greater than the pivot element to the right subarray (in no particular order)

Make two recursive calls

It's sometimes implemented as a Las Vegas Algorithm.

That is, you'll always get the same answer (there's only one sorted array) but the time can vary.

Quick Sort

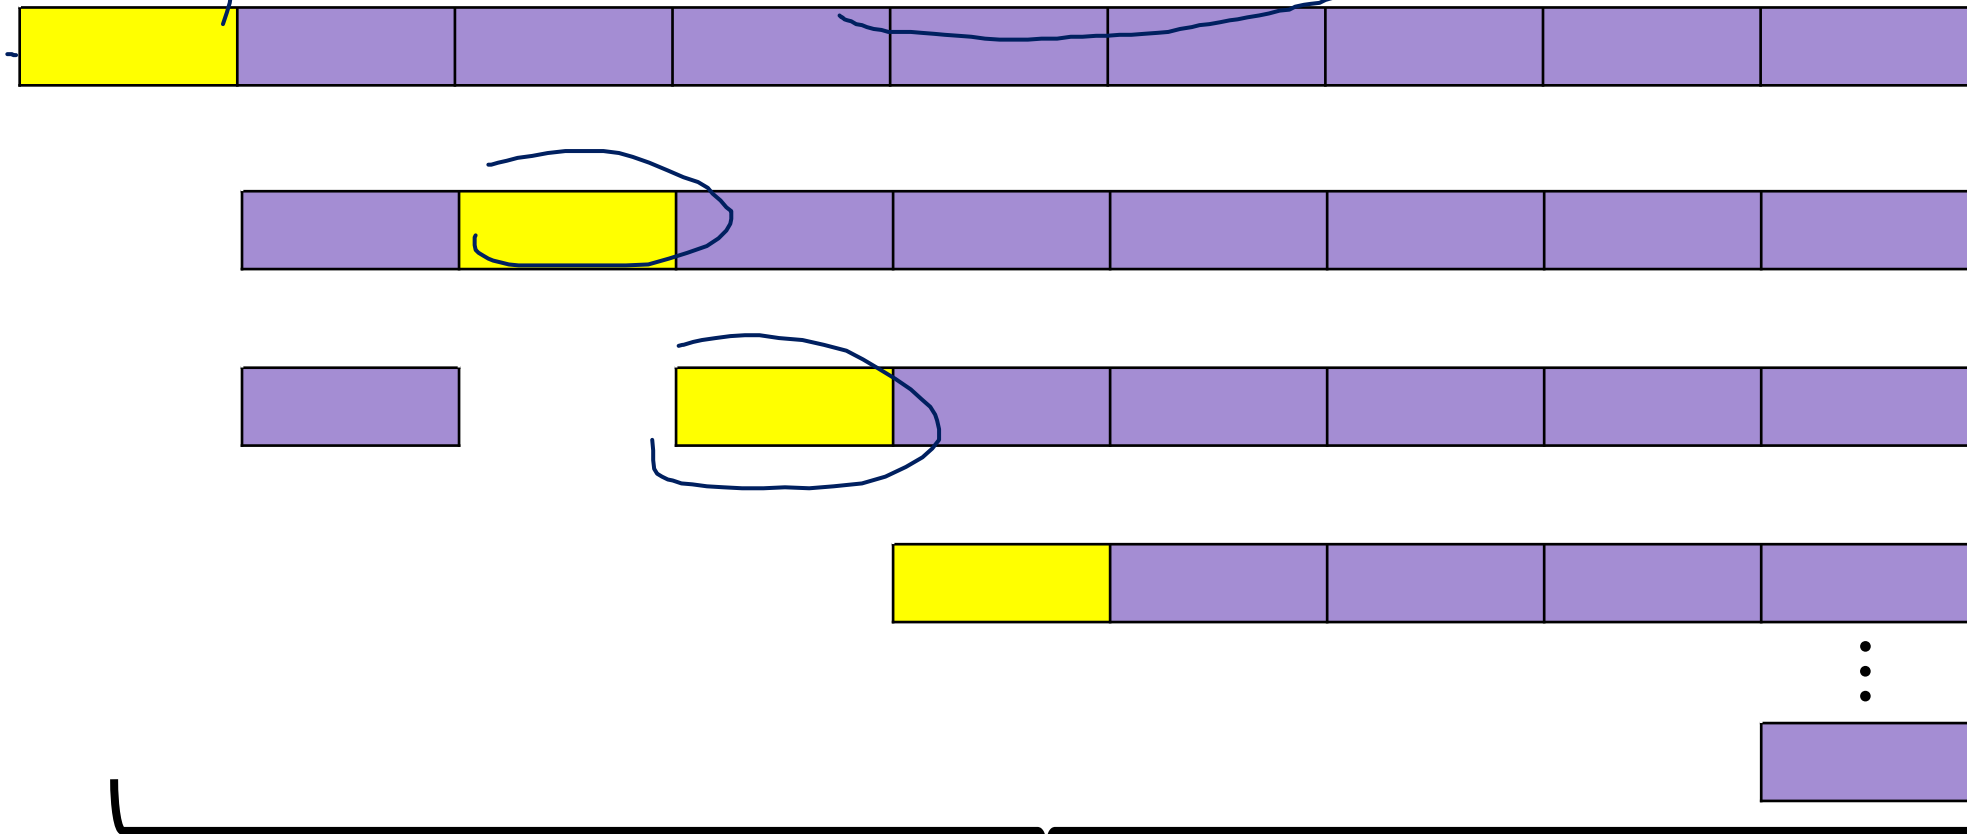


How long does it take?

Well.. it depends on what pivots you pick.

Total time:

$$\sum_{i=1}^n i = O(n^2)$$



About n levels

$O(i)$ work when i elements remaining.

For Simplicity

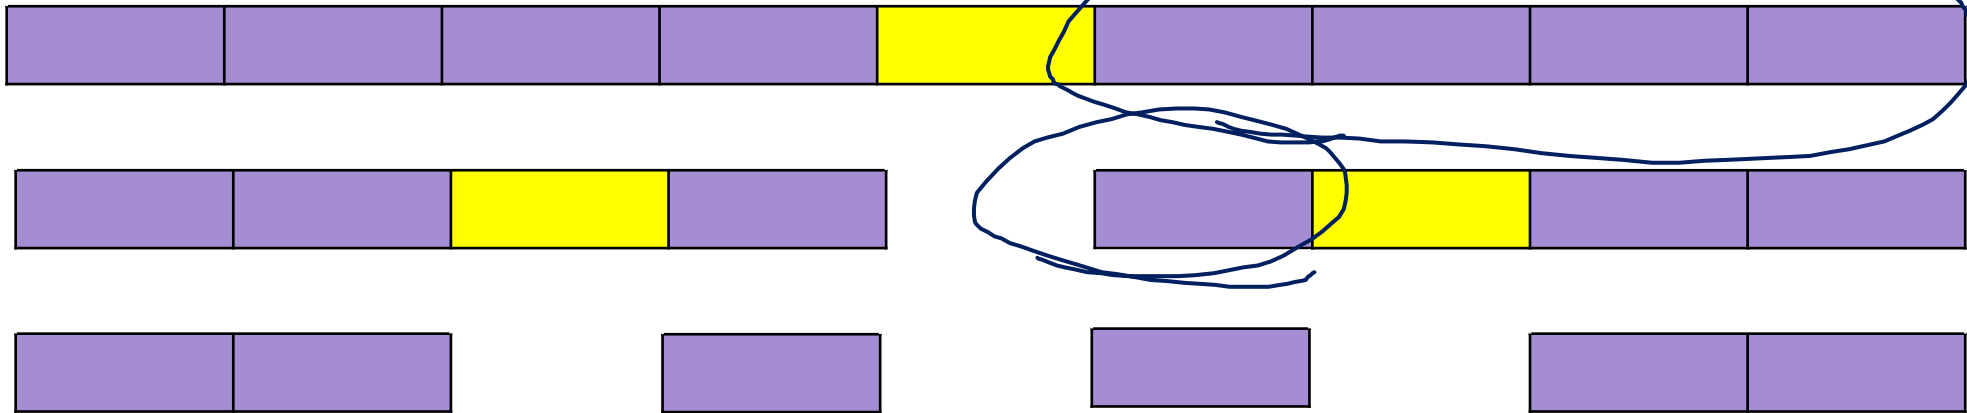
We'll talk about how quicksort is really done at the end.

For now an easier-to-analyze version:

```
if (elements remaining > 1)
    pick a pivot uniformly at random
    split based on pivot
    sortedLeft = QuickSort(left half)
    sortedRight = QuickSort(right half)
    return (sortedLeft pivot sortedRight)
```


What leads to a good time?

Pivots closer to the middle would be better.



Cut in half is ideal.
 $O(\log n)$ levels.

$$O(n \log n)$$

$O(i)$ work when i elements remaining. -- $O(n)$ per level

Work at each level

How much work do those splits do?

Each call choose a pivot

$O(n)$ total per level

Each element is compared to the pivot and possibly swapped

$O(1)$ per element so $O(n)$ per level.

So as long as we need at most $O(\log n)$ levels, we'll get the $O(n \log n)$ running time we need.

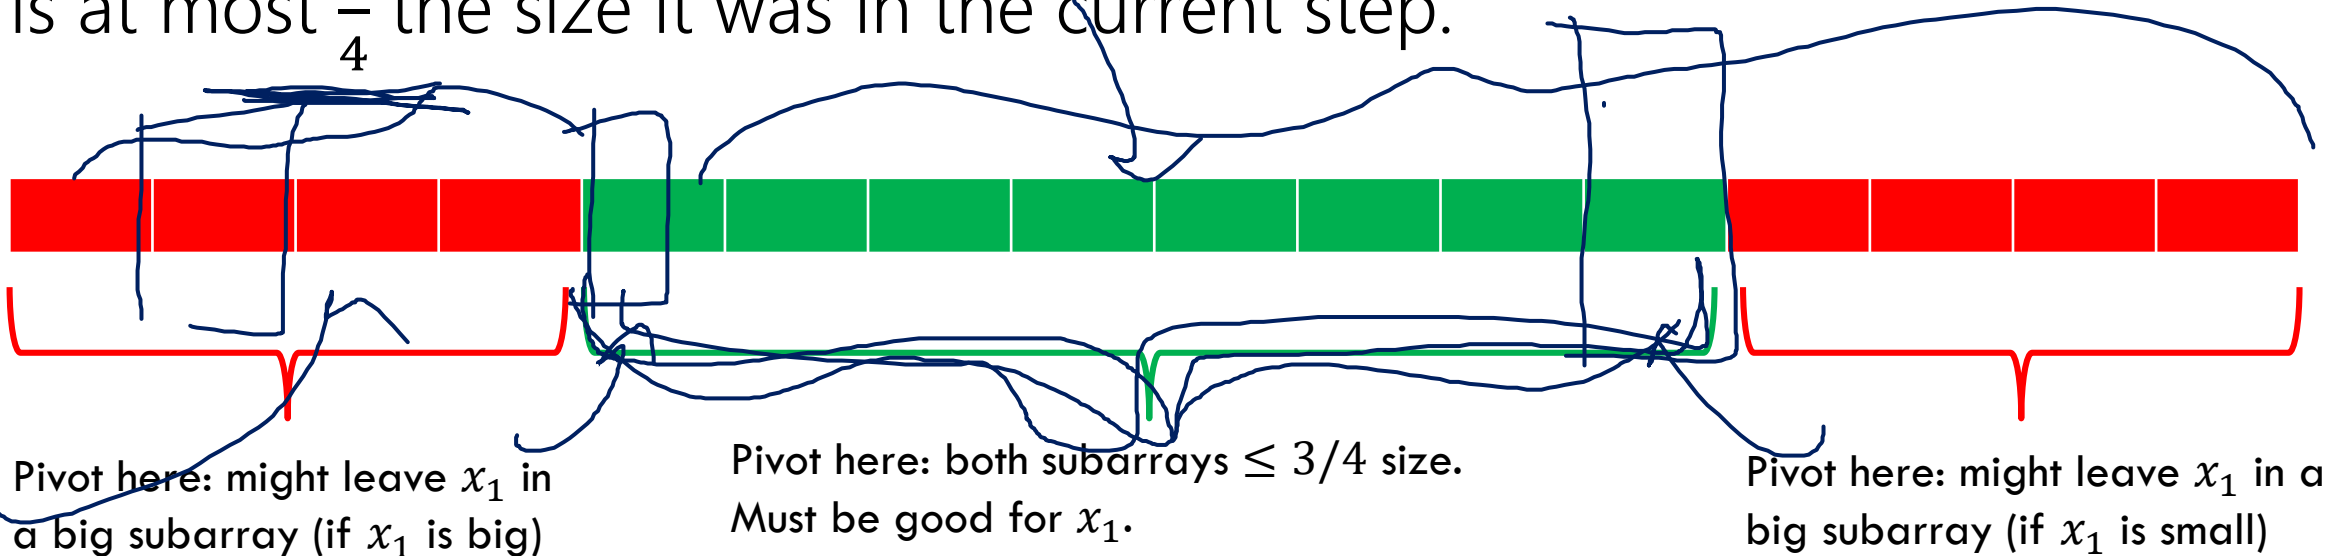
We only get the perfect pivot with probability $\frac{1}{n}$. That's not very likely...maybe we can settle for something more likely.

Focus on an element

Let's focus on one element of the array x_1 .

The recursion will stop when every ~~element~~ is all alone in their own subarray.

Call an iteration "good for x_1 " if the array containing x_1 in the next step is at most $\frac{3}{4}$ the size it was in the current step.



Good for x_i

At least half of the potential pivots guarantee x_1 ends up with a good iteration. So we'll use $\mathbb{P}(x_1 \text{ good iteration}) \geq 1/2$

It's actually quite a bit more than half for large arrays – one of the two red subarrays **might** be good for x_1 (just bad for the others in the array) x_1 might be our pivot, in which case it's totally done.

To avoid any tedious special cases for small arrays, just say at least $1/2$.

How many levels?

How many levels do we need to go?

Once x_1 is in a size 1 subarray, it's done. How many iterations does it take?

If we only had good iterations, we'd need

$$\left(\frac{3}{4}\right)^k n \leq 1 \Rightarrow n \leq \left(\frac{4}{3}\right)^k \Rightarrow k \geq \log_{4/3} n.$$

I want (at the end of our process) to say with probability at least $\langle \text{blah} \rangle$ the running time is at most $O(n \log n)$.

What's the probability of getting a lot of good iterations...what's the tool we should use?

pollev.com/robbie

Needed iterations

x_1 is done after $\log_{4/3} n = \frac{\ln(n)}{\ln(\frac{4}{3})} \leq 4 \ln n$ good for x_i iterations.

Let's imagine we do $24 \ln(n)$ iterations. Let X be the number of good for x_i iterations. Let $Y \sim \text{Bin}\left(24 \ln(n), \frac{1}{2}\right)$

$$\mathbb{P}(X \leq 4 \ln n) \leq \mathbb{P}(Y \leq 4 \ln n)$$

Set up for Chernoff

$$\mathbb{P}\left(Y \leq [1 - \delta] \cdot \frac{24 \ln(n)}{2}\right) \leq \exp\left(-\frac{\delta^2 \mu}{2}\right)$$

$$1 - \delta = 1/3 \Rightarrow \delta = 2/3$$

Applying Chernoff

$$\mathbb{P}\left(Y \leq [1 - \delta] \cdot \frac{24 \ln(n)}{2}\right) \leq \exp\left(-\frac{\delta^2 \mu}{2}\right) \leq \exp\left(-\frac{\frac{1}{3^2} \cdot 12 \ln(n)}{2}\right) = \underline{e^{-\frac{8}{3} \cdot \ln(n)}}$$

So, the probability that x_1 is not done after $24 \ln(n)$ iterations is at most $e^{-8 \ln(n)/3} = n^{-8/3}$

~~_____~~

Finishing The bound

So x_i is done with probability at least $1 - n^{-8/3}$

But x_i being done doesn't mean the whole algorithm is done...

This argument so far does apply to any other x_j -- but they aren't independent, so...union bound!

$$\mathbb{P}(\text{algorithm not done}) \leq \sum \mathbb{P}(x_i \text{ done}) = n \mathbb{P}(x_1 \text{ done}) = n \cdot n^{-8/3} = n^{-5/3}$$

$$\mathbb{P}(\text{algorithm done}) > 1 - n^{-5/3}.$$

The Theorem

Quicksort

With probability at least $1 - \frac{1}{n}$, Quicksort runs in time $O(n \cdot \log n)$

This kind of bound (with probability $\rightarrow 1$ as $n \rightarrow \infty$) is called a "high probability bound" we say quicksort needs $O(n \log n)$ time "with high probability"

Better than finding a bound on the expected running time!

Want a different bound?

Want an even better probability? You just have to tweak the constant factors!

Be more careful in defining a “good iteration” or just change $24 \ln(n)$ to $48 \ln(n)$ or $100 \ln(n)$.

It all ends up hidden in the big-O anyway.

That’s the power of concentration – the constant coefficient affects the exponent of the probability.

Common Quicksort Implementations


A common strategy in practice is the “median of three” rule.

Choose three elements (either at random or from specific spots). Take the median of those for your pivot

Guarantees you don't have the worst possible pivot.

Only a small constant number of extra steps beyond the fixed pivot (find the median of three numbers is just a few comparisons).

Another strategy: find the true median (very fancy, very impractical: take 421)



Monte Carlo Algorithms

Just some intuition



Algorithms with some probability of failure

There are also algorithms that sometimes give us the wrong answer.
(Monte Carlo Algorithms)

Wait why would we accept a probability of failure?

Suppose your algorithm **succeeds** with probability only $1/n$.

But given two runs of the algorithm, you can tell which is better.

E.g. "find the biggest <blah>" – whichever is bigger is the better one.

How many independent runs of the algorithm do we need to get the right answer with high probability?

Small Probability of Failure

How many independent runs of the algorithm do we need to get the right answer with high probability?

Probability of failure

$$\left(1 - \frac{1}{n}\right)^{k \cdot n} \leq e^{-k}$$

Choose $k \approx \ln(n)$, and we get high probability of success.

So $n \cdot \ln(n)$ (for example) independent runs gives you the right answer with high probability.

Even with very small chance of success, a moderately larger number of iterations gives high probability of success. Not a guarantee, but close enough to a guarantee for most purposes.



Conditional Expectation Practice

Practice with conditional expectations

Consider of the following process:

Flip a fair coin, if it's heads, pick up a 4-sided die; if it's tails, pick up a 6-sided die (both fair)

Roll that die independently 3 times. Let X_1, X_2, X_3 be the results of the three rolls.

What is $\mathbb{E}[X_2]$? $\mathbb{E}[X_2|X_1 = 5]$? $\mathbb{E}[X_2|X_3 = 1]$?

Using conditional expectations

Let F be the event “the four sided die was chosen”

$$\begin{aligned}\mathbb{E}[X_2] &= \mathbb{P}(F)\mathbb{E}[X_2|F] + \mathbb{P}(\bar{F})\mathbb{E}[X_2|\bar{F}] \\ &= \frac{1}{2} \cdot 2.5 + \frac{1}{2} \cdot 3.5 = 3\end{aligned}$$

$\mathbb{E}[X_2|X_1 = 5]$ event $X_1 = 5$ tells us we're using the 6-sided die.

$$\mathbb{E}[X_2|X_1 = 5] = 3.5$$

$\mathbb{E}[X_2|X_3 = 1]$ We aren't sure which die we got, but...is it still 50/50?

Setup

Let E be the event " $X_3 = 1$ "

$$\mathbb{P}(E) = \frac{1}{2} \cdot \frac{1}{6} + \frac{1}{2} \cdot \frac{1}{4} = \frac{5}{24}$$

$$\mathbb{P}(F|E) = \frac{\mathbb{P}(E|F) \cdot \mathbb{P}(F)}{\mathbb{P}(E)}$$

$$= \frac{\frac{1}{4} \cdot \frac{1}{2}}{5/24} = \frac{3}{5}$$

$$\mathbb{P}(\bar{F}|E) = \frac{\mathbb{P}(E|\bar{F}) \cdot \mathbb{P}(\bar{F})}{\mathbb{P}(E)} = \frac{\frac{1}{6} \cdot \frac{1}{2}}{5/24} = \frac{2}{5} \text{ (we could also get this with LTP, but it's good confirmation)}$$

Analysis

$$\mathbb{E}[X_2|X_3 = 1] = \mathbb{P}(F|X_3 = 1)\mathbb{E}[X_2|X_3 = 1 \cap F] + \mathbb{P}(\bar{F}|X_3 = 1)\mathbb{E}[X_2|X_3 = 1 \cap \bar{F}]$$

Wait what?

This is the LTE, applied in the space where we've conditioned on $X_3 = 1$.

Everything is conditioned on $X_3 = 1$. Beyond that conditioning, it's LTE.

$$= \frac{3}{5} \cdot 2.5 + \frac{2}{5} \cdot 3.5 = 2.9.$$

A little lower than the unconditioned expectation. Because seeing a 1 has made it ever so slightly more probable that we're using the 4-sided die.