

Victory Lap

CSE 312 Spring 21
Lecture 29

Announcements

One office hour Sunday as you wrap up studying. Link on Ed.

No office hours during the exam.

We'll answer private Ed questions, but only "during-exam" clarifications. E.g. we won't give hints but will clarify wording.

Some confusing wording in real world 3 (around the motivations for fairness) was fixed – more details on Ed.

Thanks to those who pointed out the confusing portions.

None of the questions changed (if you already turned it in you shouldn't have to change anything).

Still have concerns? Robbie is still happy to talk! Also remember concept check 30 is asking for general feedback on all the real worlds.

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)

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.

What Have We Done?

Well let's look back...

Content

Combinatorics (*fancy* counting)

Permutations, combinations, inclusion-exclusion, pigeonhole principle

Formal definitions for Probability

Probability space, events, conditional probability, independence, expectation, variance

Common patterns in probability

Equations and inequalities, “zoo” of common random variables, tail bounds

Continuous Probability

pdf, cdf, sample distributions, central limit theorem, estimating probabilities

Applications

Across CS, but with some focus on ML.

Themes

Precise mathematical communication
Both reading and writing dense statements.

Probability in the “real world”
A mix of CS applications
And some actual “real life” ones.

Refine your intuition
Most people have some base level feeling of what the chances of some event are.
We’re going to train you to have better gut feelings.

Use Your Powers Wisely

We've seen probability can be used in the real world!

But also that it:

Can be counter-intuitive/hard to explain (Bayes Rule/Real World 1)

Probability estimates can depend on the model you're using (Real World 2)

The definition you're using matters (Real World 3)

How (not to) lie with statistics

You now know a lot of the tools that people use to lie with statistics. (See also: [INFO 270](#))

Some patterns to watch out for:

My smoke alarm is going off, please pay for my new house! (analogy from Matt Parker)

Make a model, find that an event that occurred had small probability/fails some statistical test, claim that the **only** explanation is something nefarious occurred.

Better response: could the model be wrong? Is this statistical test appropriate? Once in 100 year events do happen...about once in every hundred years, is this just the one?

How (not to) lie with statistics

See a story about testing?

Remember from Bayes' Rule that you need three numbers to understand a test. (3 of prior, posterior, false positive rate, false negative rate).

Headlines usually give you one number, that often isn't even one of the ones you need for Bayes ("this test is less accurate than a coin flip!").

The article itself, if you're lucky, might give you one or two of the numbers for Bayes – don't forget the prior!

How (not to) lie with statistics

Before being impressed with a number, make sure you understand what it means.

Recent example for Robbie:

I was EXTREMELY excited to see that vaccines have a 90+% efficacy rate.

Then I realized I didn't know what in the world efficacy rate meant.

I was still EXTREMELY excited after I found out what "vaccine efficacy rate" means, but the number is meaningless unless you listen to domain experts on what a good number will be!

How (not to) lie with statistics

We can apply our knowledge to the real world!

But if you're applying in a new domain, get information from domain experts, don't instantly assume because you know Bayes' Rule that you know better than domain experts.

Don't hesitate to use these tools to understand new domains better!

But do keep in mind some things can't be quantified and just because we can use an algorithm doesn't mean we always should.

What to take next?

ML (CSE 446) using probability, linear algebra, and other techniques to extract patterns from data and make predictions.

CSE 421 designing algorithms – very little direct probability, but the combinatorics we did at the beginning will be useful.

We also have a graduate level course in randomized algorithms, but it has a few more prereqs

CSE 447 Natural Language Processing

CSE 490C Cryptography

Other things!

Thank You

This course is always a grind in the way we move from topic to topic.
And you made it through over a year into a pandemic
Over zoom.

Thank you!