

This is active learning for active learning.

Let's start with some definitions.

Active learning is a machine learning method where you actively select data points to learn a function using as little data as possible.

On the other hand, active learning is a pedagogy technique where instructors actively engage students in their lessons rather than passively presenting content.

And if you travel in the right circles you are constantly hearing about active learning. It's all the rage.

So if you also care about active learning, putting them together just makes sense.

Our question will be: how do we utilize active learning techniques so that students can learn active learning better?

Let me start with the case for active learning

Active learning is a proven technique in ML, it's very effective when you need to learn a function but want to use relatively little data.

But despite its usefulness, it's very hard to understand.

There's lots of math

And for reasons I still haven't figured out, there are bandits everywhere. Not sure why no one has captured them yet.

Regardless. With the mathematical and conceptual difficulties students will struggle to learn active learning.

Enter active learning

Active learning is ubiquitous as a pedagogical tool. It has totally revolutionized education from the old sage-on-a-stage just-tell-the-class-what-you-know model.

And you might be intimidated, can I do active learning myself? Am I cool enough to pull off such an advanced pedagogical technique?

Well I've got good news. Active learning is *easy*. I'll show you how by giving you unreflected-on-ideas for how to use active learning.

Based on current research you'll succeed – I have never seen an education paper say active learning was not effective in any context – I think we can be confident that you too can help student actively learn.

Let's see a case study. What we'll do is take a homework problem and turn it into an in-lecture multiple choice poll.

I grabbed this question from Kevin's CSE 546 homework 0,

Let X_1 to X_n be n iid uniform $0,1$ random variables. Let Y be their maximum, you want to find the $E[Y]$.

Some of you may be experiencing flashbacks, but

Let me remind you of the full calculation

[calculation omitted]

Now for us, nothing intimidating here...but for a student seeing the problem for the first time, it would take least 10 minutes to answer. Don't worry about that, just give students the question and about 90 seconds to choose an option on a poll.

here's poll results I think you'd see.

Now this may look like the results of random guessing. But look closer. There are 101 responses and 26 got the answer right.

Now, You'd expect. (hang on, take the log of the likelihood set equal to 0...)

Yes, you'd expect a mere 25.25 correct responses. That's a whole 0.75 more students learning than if you didn't ask the question.

And if somehow fewer than 25% of students get the answer right, just call it a trick question, explain the right answer and poll again. I think you'll find 100% get it right then. And wow, what more could you hope for in showing the effectiveness of a pedagogical intervention.

That's not all; I'm sure you can find other options, but a few more off the top of my head.

Wearing masks will really help them inhabit the persona of a bandit for these purposes. Actually also make sure to tie one hand behind their back

Research project proposals is a great space for active learning – think pair share, where you work by yourself for 2 minutes then tell those around you what you did for the next minute is a great way of generating commentary on research proposals. At the very least will be better than the current NeurIPS process.

I think you'll find that the effect of active learning on the effectiveness of learning looks like this. With more active learning, the outcomes for learners go from bad up to good.

Thank you.