Expectation of Mathematical Level, ECS 256, Fall 2022

ECS 256 has not been taught since Winter 2016, and has been revised since then. You can view a topical outline at *https://heather.cs.ucdavis.edu/ matloff/256/BlurbF22.html*.

I was gratified to see a robust Pass 1 pre-enrollment for the course, for Fall 2022, 50 enrolled and 53 on the waiting list. Since ECS 256 is a Core course, in the Applications section of the CS PhD requirements, it's important to discuss expectations.

The course prerequisites are:

> A calculus-based probability course such as MAT 135, STA 131A or ECS 132; basic linear algebra; programming skills.

Many if not most students have forgotten much of their undergrad coursework in probability and linear algebra. That is fine, and we will review those topics at the start of the course. But what is key is mathematical intuition and facility.

In order to gauge whether your mathematical background is sufficient for this course, see the following examples. Note that these are NOT specifically topics that will arise in the course, but they are examples of the level of mathematical sophistication needed to do well in the course (e.g. for a grade of A- or better). If you can solve at least 2 of the 3 problems, you should do well in this course.

Problem 1

> The field of Recommender Systems involves prediction of ratings. The typical first example is the MovieLens data, consisting of numerical ratings of movies by various users. There are 943 users and 1682 movies, but each user has reviewed only some of the movies. The goal is to fill in the missing ratings. What rating would user i give to movie j, if he/she were to watch the movie and rate it?
>
> Now consider the example of a dataset called House Voting, in the UCI Machine Learning Data Repository set. There are 435 legislators in the US House of Representatives, and there were 17 bills in the dataset. Each representative votes Yes or No on any particular bill, but not all representatives voted on all bills. The goal is to predict how each representative would have voted on the bills they missed.
>
> Question — Explain why the House Voting dataset and the MovieLens dataset are of the same form. What are the analogs in the former of users and movies in the latter? (Note: You are not being asked how to fill in the missing ratings, just to name the correspondences.)

Problem 2

> In calculus, say we have a function f, with first and second derivatives f' and f'', respectively. We find a local minimum of f by setting $f' = 0$ and solving for the root. It is a local minimum if $f'' > 0$.
>
> Question — WHY? Why does $f'' > 0$ tell us that this is a local minimum rather than a local maximum? (Don't look this up on the Web; see if you can reason it out.)

Problem 3

Suppose X is a positive integer-valued random variable. You may recall that its expected value is a weighted sum of certain probabilities:

$$EX = \sum_{i=1}^{\infty} iP(X = i)$$

But one can show that this is equivalent to a sum of other probabilities:

$$EX = \sum_{j=1}^{\infty} P(\text{something involving } j)$$

Question — Using the fact that $i = \sum_{j=1}^{i} 1$, and and reversing the order of summation (start with $\sum_i \sum_j$ but convert to $\sum_j \sum_i$) determine what the "something involving j" is above. You may wish to guide your thinking by considering a small example in which X is 1 or 2, with probability $1 - p$ and p, respectively.

Comments:

- Problem 1 should be quick and easy for grad students, but it mystified many of the undergrads who took my ECS 172 course last quarter.
- Ideally, Problem 2 is something you asked yourself back when you took calculus. Not only would that show that you have math intuition, but even more important, it would show that you care about intuition, rather than just blindly accepting equations.
- Successfully solving Problem 3 shows that you have good facility in manipulating mathematical expressions. The level of this problem is something like the level of some problems we will have for homework in 256.