Class Activity

Activity 1: Hidden Majority Vote

Do you think AI tools make students learn better? (**Yes = 1**, **No = 0**)

1. Flip a fair coin.
2. If **Heads**, **report your true answer** (1 or 0).
3. If **Tails**, **report a random number** (0 or 1) using another coin flip or random generator.
4. Do this process **10 times** and record the values.
5. Enter your 10 reported values into the shared Google Sheet (one row per student).

Discussion

- Why does this protect privacy?
- What happens if someone flips more tails than others?
- Why does the estimate stay unbiased on average?
- How does the estimate change if you increase the number of simulated responses?
- What trade-off do we see between privacy and accuracy?

Activity 2 : Hidden Salary

1. Everyone secretly chooses a "salary" .
2. The first student secretly adds a random number r to their own salary and writes down the total.
3. Pass the running total to the next person.
4. Each new person adds their own salary to the total and passes it on.
5. The last student pass the final total to the first student.
6. The first student then subtracts the random number r to get the true total of all salaries.

Discussion

- Why does this protect privacy?
- What information does each person actually share in this method?
- How does this method scale if the group gets larger or the order changes?
- How is this different from "anonymizing" data after it's collected?
- Do you need to know all the raw data to update the average?
- Extension: Welford's Rule.

Activity 3 : Adaptive Learning Rate

1. You are each a **temperature sensor** in Granville.
2. Each sensor receives an **initial temperature reading** (e.g., between 68°F and 80°F). These readings are noisy, some are too high, some too low.
3. Each sensor is assigned a **learning rate ( )** that determines how much it trusts the network average:
4. In each round, calculates the **network average.**
5. Each sensor updates its own estimate and put it in Google Sheet.

Discussion

- How do small and large values affect the **speed** and **stability** of convergence?
- Did all sensors reach the same final temperature?
- If not, why might small differences remain?
- What happens if one sensor's initial reading is a large outlier? How does the network handle it?
- What other systems use consensus updating?
- How is this update rule similar to **Bayesian updating**?

Final Discussion

Extensions Discussion

How do these build on the same idea? Updating summaries, not raw data?

- **Likelihood ratio:** sequential probability ratio tests