



Algorithms for Data Science

Barna Saha

A new algorithms class!

- Why do we need a new algorithms class?
 - Unprecedented amount of data containing a wealth of information.
 - Example: Twitter receives 6000 tweets per second which amounts to 500 million tweets per day with a storage requirement of ~7000 gigabytes.
 - Traditional algorithms process data in RAM, sequentially and may have high time-complexity
 - Not suitable for processing Twitter data

Characteristics of Big Data

- **VOLUME**
 - Can not store the entire data in the main memory
- **VELOCITY**
 - Data changes frequently. Needs highly efficient processing, often parallel processing.
- **VARIETY & VERACITY**
 - Data coming from many different sources, and often contains noise-adds to the complexity of data processing

This Course

- Develop algorithms to deal with such data
 - Space and Time Efficient
 - Parallel Processing
 - Approximation & Randomization
- Theoretical course with main focus on algorithm analysis
 - Relevant applications will be discussed, and there will be plenty of coding exercises
 - But no software tools will be covered
- Background in basic algorithms (311) and probability (240) are strictly required.

Grading

- Homeworks (3-4) in a group of 2/3
 - Will consist of mathematical problems and/or programming assignments
 - Find your partners early and wisely. Do not come to me with complaints about your partner.
 - 20%
- Midterm
 - 20%
- Final
 - 30%
- In Class Participation
 - In Class Problem Solving & Coding Exercises in Group
 - Every 3 to 4 classes
 - 30%

Office Hours

- Instructor: Thur 11:30 am-12:30pm at CS 232
- Teaching Assistant: Sainyam Galhotra
 - Tue 2-3 pm at EDLAB - LGRT room 223
 - Email: sainyam@cs.umass.edu
- All class related discussions should be done through piazza.
 - Sign up from the course page.
- Course website
 - <https://barnasaha.net/barna/compsci590d-f16/>

Books

- Text Book: We will use reference materials from the following books. **Both can be downloaded for free.**
- **Mining of Massive Datasets**, Jure Leskovec, Anand Rajaraman and Jeff Ullman.
- **Foundations of Data Science**, a book in preparation, by John Hopcroft and Ravi Kannan

An Interesting Problem

- Suppose we see a sequence of items, one at a time.
- We want to keep a single item in memory.
- We want it to be selected at random from the sequence.
- **Easy if we know the number of items “m”**
 - Just draw a random number in between 1 and m
- **What if we do not know m?**

Reservoir Sampling

- ▶ Upon seeing the **first** element—keep it.
 - ▶ The first element is chosen with probability 1.
- ▶ Upon seeing the **second** element—select the second element with probability $\frac{1}{2}$. If the second element is selected discard the first element.
 - ▶ The probability that the second item is sampled = $\frac{1}{2}$
 - ▶ The probability that the first item is sampled = $\frac{1}{2}$
- ▶ Upon seeing the **third** element—select it with probability $\frac{1}{3}$, if selected then discard the element that was previously selected.
 - ▶ The probability that the third item is sampled = $\frac{1}{3}$.
 - ▶ The probability that the second item is sampled = $\frac{2}{3} * \frac{1}{2} = \frac{1}{3}$
 - ▶ The probability that the first item is sampled = $\frac{2}{3} * \frac{1}{2} = \frac{1}{3}$

Reservoir Sampling

- ▶ Upon seeing the **fourth** element—select it with probability $\frac{1}{4}$, if selected then discard the element that was previously selected.
 - ▶ The probability that the fourth item is sampled = $\frac{1}{4}$.
 - ▶ The probability that the third item is sampled = $\frac{3}{4} * \frac{1}{3} = \frac{1}{4}$
 - ▶ The probability that the third item is sampled = $\frac{3}{4} * \frac{1}{3} = \frac{1}{4}$ The probability that the third item is sampled = $\frac{3}{4} * \frac{1}{3} = \frac{1}{4}$
- ▶ Can you generalize the algorithm to any i ?
 - ▶ Upon seeing the i th item—select it with probability $\frac{1}{i}$, if selected, discard the element that was previously selected.
 - ▶ The probability that the i th item is sampled = $\frac{1}{i}$.
 - ▶ The probability that the j th $j < i$ item is sampled = $\frac{i-1}{i} * \frac{1}{i-1} = \frac{1}{i}$

Algorithms Every Data Scientist Should Known: By Josh Wills

<https://blog.cloudera.com/blog/2013/04/hadoop-stratified-random-sampling-algorithm/>

Data scientists, that peculiar mix of software engineer and statistician, are notoriously difficult to interview. One approach that I've used over the years is to pose a problem that requires some mixture of algorithm design and probability theory in order to come up with an answer. Here's an example of this type of question that has been popular in Silicon Valley for a number of years:

Say you have a stream of items of large and unknown length that we can only iterate over once. Create an algorithm that randomly chooses an item from this stream such that each item is equally likely to be selected.

Reservoir Sampling!

algorithms in particular, and talk about how they are used in [Cloudera ML](#), our open-source collection of [data preparation and machine learning algorithms for Hadoop](#).

Applied Reservoir Sampling in Cloudera ML

The first of the algorithms Greg describes is a *distributed* reservoir sampling algorithm. You'll note that for the algorithm we described above to work, all of the elements in the stream must be read sequentially. To create a distributed reservoir sample of size K , we use a MapReduce analogue of the ORDER BY RAND() trick/anti-pattern from SQL: for each element in the set, we generate a random number R between 0 and 1, and keep the K elements that have the largest values of R . This trick is especially useful

Sampling

- A very useful method to obtain appropriate summary of data
- Will learn more in the coming classes
- But needs to be done with care
- Link to video

Next Few Classes

- Probability review before we enter into the more interesting regime!