

CS 378 – Big Data Programming

Lecture 12

User Sessions from Logs

Review

- Assignment 5 – Avro Objects
- We'll look at implementation details of:
 - Mapper
 - Combiner
 - Should we use one? Can we use one?
 - Reducer
 - Avro generated Java code

Other Issues

- Running MRUnit tests with Avro objects
 - Codehaus jackson version consistency
 - jackson-mapper-asl
 - jackson-core-asl
 - Avro serialization
- “shaded” JAR file – all dependencies included
 - Except Hadoop JAR – Why?
- Avro field definitions
 - Unions, defaults, ...

Review - Design Pattern

- Structured to hierarchical design pattern
- Data sources linked by some foreign key
- Data is structured and row based
 - For example, from databases
- Data is semi-structured and event based
 - Web logs

Sessionizing Web Logs

- Create user sessions from web logs
- Represents all the actions by a user
- Allows later analysis to “replay” the user actions
- Collect measures and metrics about user behavior
 - Pages viewed, time on page, clicks
 - Path through the site, entry to the site (from a search engine?)

Sessionizing Web Logs

- To start (this or any “big data” application)
- We need to understand the data
 - Fields, values
 - Data size
- We need to define our goal
 - What do we want to end up with

Web Logs

- Let's look at some data
- Logs saved in database
 - Log entries already have structure
 - Tab separated values
 - Easily parsed (lots of work has been done for us)

Web Logs


- Our goal is to aggregate user actions into sessions, so we can better understand
 - User behavior
 - The impact changes have on user behavior
- So what should a session look like?

User Session

- Data about the session as a whole
- List of events (pages viewed, actions taken)
 - Ordered in time
- In our logs, what data is session-wide
- What data is impression/action specific

MapReduce in Hadoop

Figure 2.4, Hadoop - The Definitive Guide



Assignment 6

- Define an Avro object for user session
 - One user session for each unique userID
 - Session will include an array of events
 - Events ordered by timestamp
- Identify data associated with the session as a whole
- Identify data associated with individual events
- Include all the fields in the log entries
- Create enums where requested

Assignment 6

- Run WordCount on `dataSet6.tsv` – see what's in it
 - Modify WordCount to output values for each field:
 - `fieldname:value`
 - Ignore these fields (they have lots of values):
 - `event_timestamp, image_count, initial_price, mileage, referrer, user_id, vin`
- `event_type`
 - Break this into two fields in your schema:
 - Type (enum): use the first word of this field value
 - `change, click, contact, edit, share, show, submit, visit`
 - Subtype (enum): use remainder of the string
 - Example: `ContactForm`

Assignment 6

Recommendations

- Use `dataSet6Small.tsv` for development and testing
- Get your app working with just a few fields populated
 - Session with no events
 - Add events, but just a few fields first
 - Extend the schema
 - Populate the new field(s) in your schema
- Write some unit tests as you go