

CS 327E Milestone 3 due Thursday, 05/03.

This is the first of two milestones that makes use of Apache Beam. The goal of this milestone is to code the Beam pipelines which will be needed to implement your cross-dataset queries.

1. Create a new Jupyter notebook named `milestone3.ipynb`.
2. Identify all the tables in your datamart which contain data that needs to be cleansed. The tables you select should contain either some duplicate records or some fields which have non-conforming data or both. For example, non standard dates, addresses, person names, employer names, locations, genres, etc.
3. For each source table, write a Beam pipeline that normalizes the data from the table and creates a new table with the transformed data. The pipeline should satisfy the following requirements:
  - It should be executed with the Direct Runner
  - It should run a BigQuery query with a `LIMIT` clause over the source table in your datamart such that the number of records being processed is `< 500`
  - It should make a `PCollection` from the BigQuery query results
  - It should implement at least one `DoFn`
  - It should apply at least one `pardo` to the `PCollection`
  - It should write the output to a local file by the name of `output.txt`
  - It should also write the output to a new BigQuery table in your datamart
  - It should be called from your `milestone3.ipynb` notebook.
4. Verify that the BigQuery output tables from the previous step contain a valid primary key. If the output tables are child tables, they should also have a foreign key. Run the SQL statements to verify these constraints from your `milestone3.ipynb` notebook.

### Coding Conventions:

- Each Beam pipeline should be in a file named `<table>_beam.py` where `<table>` is the name of the table being transformed.
- The BigQuery output tables should be named `<table>_Beam` and be stored in your datamart.
- The `DoFn` code should be commented sufficiently to show understanding of the `transform(s)`.

CS 327E Milestone 3 Rubric

**Due Date: 05/03/21**

<p>Implement one or more Beam pipelines that transform the data from your datamart. Sufficiently comment the code to show understanding of the Apache Beam pipeline.</p> <ul style="list-style-type: none"> <li>-100 missing <code>&lt;table&gt;_beam.py</code> from repository</li> <li>-50 code does not implement a DoFn transform</li> <li>-50 code does not pull from or write back to your datamart</li> <li>-20 missing <code>LIMIT</code> clause from input query</li> <li>-40 missing or incorrect <code>pardo</code> call</li> <li>-40 code does not write to output file <code>output.txt</code></li> <li>-50 code does not write to output table <code>&lt;table&gt;_Beam</code></li> <li>-30 code missing comments</li> <li>-40 missing pipeline run call from <code>milestone3.ipynb</code></li> <li>-20 missing or incorrect primary key verification on output table <code>&lt;table&gt;_Beam</code> from <code>milestone3.ipynb</code></li> <li>-20 missing or incorrect foreign key verification on child output table <code>&lt;table&gt;_Beam</code> from <code>milestone3.ipynb</code> (assuming output table is a child table)</li> </ul>	<p>100</p>
<p><code>milestone3.ipynb</code> and <code>&lt;table&gt;_beam.py</code> for each beam pipeline pushed to your group's private repo on GitHub. Your project <b>will not</b> be graded without this submission.</p>	<p><b>Required</b></p>
<p><code>submission.json</code> submitted into Canvas. Your project <b>will not</b> be graded without this submission. The file should have the following schema:</p> <pre>{   "commit-id": "your most recent commit ID from Github",   "project-id": "your project ID from GCP" }</pre> <p>Example:</p> <pre>{   "commit-id": "dab96492ac7d906368ac9c7a17cb0dbd670923d9",   "project-id": "some-project-id" }</pre>	<p><b>Required</b></p>
<p><b>Total Credit:</b></p>	<p><b>100</b></p>