

Lab 1: Association Rules Mining

Due date: Monday, October 1, midnight.

Lab Assignment

In this assignment you will analyze collections of market baskets and will determine *frequent itemsets* and *association rules* present in the collections.

Assignment Preparation

This is a pair programming assignment. Each student teams up with a partner. You get to select your partner at the beginning of the Monday, September 24 lab.

Datasets

You are given two different datasets to work with on this assignment: **EXTENDED BAKERY**, **Fantasy Bingo**. The first one is a synthetic dataset simulating actual market baskets (purchases made at a bakery). The second one consists of lists of authors whose books different people have read in a year-long reading challenge called "fantasy bingo".

For each dataset we ask the questions that are translated into the need to find frequent itemsets, and, in some cases - association rules of specific type. We refer to these special frequent itemsets and association rules as Skyline Frequent Itemsets and Skyline Association Rules.

Note: An additional dataset *may* be added to the list at a later time. If this happens, the final deadline to submit your lab will be extended by one lab period.

EXTENDED BAKERY Dataset

The assignment is based on the **Extended BAKERY** dataset. The dataset is a modified version of the **CSC 365 BAKERY** dataset. The **Extended BAKERY** dataset describes the work of a *chain* of bakery shops that sell a variety of pastries and drinks to customers.

The data provided to you for this assignment is the information about purchases made by the bakery chain customers in various locations. The four sub-datasets contain information about 1000, 5000, 20,000 and 75,000 purchases.

For each sub-dataset we provide three files representing the same set of receipts. For simplicity, each file represents the exact purchases: i.e., which items were purchased on which receipt, but **omits other information from the dataset**: the store location, the employee who rang the purchase, the date of the purchase. Additionally, the *quantity* of the purchased item is omitted in two representations of the three listed.

The full description of the dataset is below.

Access to the dataset. All CSV files can be downloaded from the Lab 1 data page

<http://www.csc.calpoly.edu/dekhtyar/466-Fall2018/labs/lab01.html>

The list of **market baskets** for each dataset size is available in **three formats**:

1. Sparse Vector format. Files `XXXX-out1.csv`. Each line of the file has the following format:
 - First column is the receipt Id.
 - Subsequent columns store a list of goods purchased from the bakery ordered by `Goods.Id`.

Example:

1, 7, 15, 44, 49

Receipt 1 contained purchases of a Coffee Eclair, a Blackberry Tart, Bottled Water and a Single Espresso.

2. Full Binary Vector format. Files `XXXX-out2.csv`. Each line of the file has the following format:
 - First column is the receipt id.
 - 50 columns follow, with 0s or 1s as values. A 1 in column $i + 1$ means that a good with `Goods.Id` of i was purchased on the receipt.

- 3 | Aaronovitch, Ben
- 4 | Abbey, Kit
- 5 | Abbey, Lynn
- 6 | Abercrombie, Joe

Note: The value of "**N/A*" is a special value that should be ignored if it is ever discovered in your market baskets (it means that a reader submitting their list of books failed to read one or more books to meet the challenge).

Baskets. The file `bingoBaskets.csv` contains the list of authors for each of the participants in the fantasy bingo. The file has the following format:

`BingoCardId, AuthorId1, AuthorId2, ..., AuthorIdK`

Here, `BingoCardId` is the unique Id of each basket (each reader's submission/bingo card), and `AuthorIdX` is an id of an author (see first column of the `authorlist.psv` file).

Here are a few sample lines:

```
12,3, 68, 91, 166, 183, 224, 237, 259
13,2, 3, 48, 91, 166, 171, 216, 217, 218, 251
```

Please note that all author Ids in each line are sorted by author ID (and all author ids in general are sorted in lexicographical order by the author's last, and first names).

Mining Frequent Itemsets and Association Rules

Your task is to discover the **association rules** that exceed specific given values of *minimum support* and *minimum confidence*.

As discussed in class, mining association rules is a two-step process. On step one, the goal is to discover *frequent itemsets* with support exceeding *minsup*. On step two, the goal is to discover specific *association rules* found within the discovered frequent itemsets.

Algorithms for discovery of both frequent itemsets (**Apriori**) and association rules (**genRules**) will have been discussed in class together with implementation strategies.

Skyline (Maximal) Frequent Itemsets and Association Rules

Association rules mining tends to discover **a lot** of rules in any given dataset. This is due to permutation properties of the rules (e.g., if $A, B \rightarrow C, D$ is

an association rule, then so are $A, B, D \rightarrow C$, $A, B, C \rightarrow D$), and due to the large number of items in a typical dataset.

To make results of your work *observable*, we will be interested only in so-called **skyline** or **maximal** frequent itemsets and association rules.

Definition. A frequent itemset is called a **skyline (maximal)** frequent itemset, if *it is NOT a subset of any other frequent itemset*. An association rule is called a **skyline** association rule if its right side and its left side form a **skyline** frequent itemset.

Informally, **skyline** or **maximal frequent itemsets** are those, that cannot be extended further to form larger frequent itemsets. To constrain the output of your work, you need only to report **skyline** frequent itemsets.

Furthermore, to simplify the process of mining association rules, you shall report **only skyline** association rules in which the right side of the rule contains *a single item*.

Minimum Support and Minimum Confidence

Each of the datasets may require *tuning* the `minsup` and `minconf` parameters - i.e., the minimum support for the frequent itemsets, and the confidence for the association rules. *It is your goal for this assignment to properly tune these parameters and find the best values of `minsup` and `minconf`.*

For the `EXTENDED BAKERY` dataset, all association rules and frequent itemsets were seeded and they are separated from any randomly occurring itemsets by a fairly large gap in support. You need to discover the `minsup` threshold that surfaces these skyline frequent itemsets, as well as the `minconf` value that exposes the actual association rules.

For the `Fantasy Bingo` dataset, you need to explore the data in order to come up with reasonable ranges for `minsup` and `minconf`. You can start by finding the support for each singleton itemset in each dataset, and building the frequency histogram. This should give you an idea what type of support (in terms of absolute or relative values) can be considered "relevant" in each case. From there, proceed to use to code you develop as a research tool to find the best parameter values.

Code

You can use any language you want. Python is a good choice due to powerful tools for parsing and manipulation of data. Java is a good choice because it will allow you to have a well-designed implementation of all functionality.

You are to write **the entirety of the code** from scratch without borrowing from any of the existing machine learning packages that may be available in your programming language of choice.

If you are using Python, you are allowed to use NumPy for data manip-

ulation and parsing. You can use `Pandas` data frame manipulation functionality. In Java or other programming languages, you are allowed to use comparable packages/libraries/APIs.

Deliverables

You shall discover **skyline frequent itemsets** and **skyline association rules** in each of the four `EXTENDED BAKERY` datasets. Additionally, discover **skyline frequent itemsets** and **skyline association rules** in the `Fantasy Bingo` dataset.

Submit the following:

- A report containing showing discussing your work. The report shall contain the following information:
 1. Description of your procedures to find the optimal `minSup` and `minConf` parameters for your datasets.
 2. Table of `minSup` and `minConf` parameters for each of the four `EXTENDED BAKERY` datasets. Include in this table, the `minSup` and `minConf` values selected and the total number of skyline association rules induced.
 3. A list of skyline association rules for the largest `EXTENDED BAKERY` dataset.
 4. Information about the optimal `minSup` and `minConf` parameters for the `Fantasy Bingo` dataset.
 5. List of skyline frequent itemsets for the `Fantasy Bingo` dataset.
 6. List of skyline association rules for the `Fantasy Bingo` dataset.

For each *skyline frequent itemset* specify:

1. All *items* in it. Use `Goods.Flavor` and `Goods.Food` attribute values to describe each item in the `EXTENDED BAKERY` dataset. Report author names for the `Fantasy Bingo` dataset.
2. The support of the itemset.

Notice, that `Goods.Flavor` and `Goods.Food` attributes (as well as the author names and transcription factor names) are NOT present in the input market baskets (all the formats described above contain only numeric indexes). It is the job of your software to report these attributes given the good ids.

For each *skyline association rule* specify:

1. All items on the left side of the rule (`Goods.Food+Goods.Flavor`, or `Author Name`).
2. The item on the right side of the rule (`Goods.Food+Goods.Flavor`, or `Author Name`).

3. The support of the rule.
 4. The confidence of the rule.
- Any software you have written to discover association rules.

In general, a program for association rules discovery should take as input the following parameters:

1. **Filename.** Name of the CSV file containing the dataset. Your program can use any of the formats made available to you.
2. **minSup.** The minimum support number for frequent itemset and association rule discovery.
3. **minConf.** The minimum confidence number for association rule discovery.
4. **additional filenames.** Names of any additional files needed to produce output.

Optionally, the name of the output file may be passed to your program as well.

Additionally, you may include any optional flags that specify whether:

- all rules/frequent itemsets or skyline rules/frequent itemsets should be returned. (the default behavior is to print skylines).
- only rules, only frequent itemsets or both rules and frequent itemsets shall be printed.

Generally speaking, you may elect to implement your software in any way you like: e.g., you can split reading/parsing data, frequent itemset search and association rules discovery into three separate pieces of code if this is more convenient for you.

- A **README** file which contains the following information (at least):
 - Names of all students in the pair/team.
 - Specification of which type(s) of input format your program(s) take(s).
 - Instructions on how your code should be run. This is especially important if you implemented association rules mining as a sequence of separate programs.

Note: Frequent itemset/association rule lists that you submit can be the output of your program(s), as long as your program output follows the guidelines specified above.

Note: Each **EXTENDED BAKERY** dataset incorporates a specific set of association rules (and frequent itemsets), that stand out. You may have to try your program with a number of **minConf** and **minSup** parameters until you discover all of them, but overall, the **separation between the frequent itemsets/association rules** and all other itemsets/candidate rules is **very robust**.

Training Dataset

To help you calibrate your discovery process, we are providing one more dataset for you. The dataset contains 1000 market baskets and has the following **association rules** seeded in it:

Lemon Cake → Single Espresso
Blackberry Tart → Apple Danish
Napoleon Cake → Gongolais Cookie
Apple Tart and Berry Tart → Blueberry Tart

All these rules have support of at least 10% and a confidence of at least 90%. Note that other rules (e.g., Berry Tart and Blueberry Tart → Apple Tart will also exist in the dataset and would need to be reported).

The training dataset can be downloaded from the course web page. The `example.zip` file contains the following four files inside the `example` directory:

<code>out1.csv</code>	market baskets in sparse vector format
<code>out2.csv</code>	market baskets in full binary vector format
<code>lab2-example-output</code>	output of the TA's rule mining program
<code>Rules2.xml</code>	an XML file specifying the rules found in the dataset

Submission Instructions

Report submission. While I prefer hardcopies of the reports, reports in soft copy only will also be accepted. Reports shall be word-processed, with the results of running your program included where necessary.

Submit your report as a PDF file. Submit your `README` file. Create a zip or a tar.gz archive of your code (any any other files you need to make your programs run). Submit these three files separately.

Submission. Use the `handin` tool to submit your deliverables. Each pair submits exactly one copy of all materials. Name your report `Lab01-report.pdf`. Name your code archive `lab01.zip` Use the following command to submit

```
$ handin dekhtyar lab01 Lab01-report.pdf README lab01.zip
```