Spring 2009	**CSC 466:** Knowledge Discovery from Data	Alexander Dekhtyar

Lab 2: Association Rules Mining

Due date: Tuseday, April 21, midnight.

Lab Assignment

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

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 Tuesday, April 14 lab. A (single) group of three students may be formed to accommodate the extra student should we have one.

Dataset

The assignment is based on the `Extended BAKERY` dataset from Labs 0 and 1.

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 and represent **exactly the same receipts** as the files you used in Labs 0 and 1 to create relational database.

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 datasets wiki:

http://wiki.csc.calpoly.edu/datasets/wiki

The wiki page for the EXTENDED BAKERY dataset is:

http://wiki.csc.calpoly.edu/datasets/wiki/ExtendedBakery

Individual .CSV files are available from wiki pages devoted to sub-datasets:

http://wiki.csc.calpoly.edu/datasets/wiki/ExtendedBakery1000
http://wiki.csc.calpoly.edu/datasets/wiki/ExtendedBakery5000
http://wiki.csc.calpoly.edu/datasets/wiki/ExtendedBakery20K
http://wiki.csc.calpoly.edu/datasets/wiki/ExtendedBakery75K

Additionally, the zip file containing all CSV files can be downloaded from

http://wiki.csc.calpoly.edu/datasets/wiki/apriori

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 an ordered by `Goods.Id` list of goods purchased from the bakery.

   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.

   Example:

   ```
   1,0,0,0,0,0,0,0,1,0,0,0,0,0,0,0,1,0,0,0,0,0,0,0
   ```

3. Items Table. Files `XXXXi.csv`. Each line of the file represents a single tuple from the `Items` table. The columns are:

   - Receipt(number), Quantity, Item

   Example:

   ```
   1,3,7
   1,4,15
   1,2,49
   1,5,44
   ```

   (note, that the item IDs may be out of order)

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) 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*.

Deliverables

You should discover **skyline frequent itemsets** and **skyline association rules** in each of the four `EXTENDED BAKERY` datasets.

Submit the following:

- The list of skyline frequent itemsets and the list of skyline association rules you discovered in each of the four datasets.

  For each *skyline frequent itemset* specify:

  1. All *items* in it. Use `Goods.Flavor` and `Goods.Food` attribute values to describe each item.
  2. The support of the itemset.

  For each *skyline association rule* specify:

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

  Additionally, you may include an 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 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 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:

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

Note: The TA's program refers to items by their Ids, not full names. Your code should supply full names.

Submission Instructions

You will use the handin tool to submit your files. Each pair submits exactly one copy of all materials. Put all your files in a single archive (zip or gzipped tar), name it `lab02.zip` or `lab02.tar.gz` and submit as follows:

```
$ handin dekhtyar-grader lab02 lab02.zip
```