



HAL
open science

Testing Interestingness Measures in Practice: A Large-Scale Analysis of Buying Patterns

Martin Kirchgessner, Vincent Leroy, Sihem Amer-Yahia, Shashwat Mishra

► **To cite this version:**

Martin Kirchgessner, Vincent Leroy, Sihem Amer-Yahia, Shashwat Mishra. Testing Interestingness Measures in Practice: A Large-Scale Analysis of Buying Patterns. International Conference on Data Science and Advanced Analytics, Oct 2016, Montreal, Canada. hal-01407787

HAL Id: hal-01407787

<https://hal.science/hal-01407787v1>

Submitted on 2 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Testing Interestingness Measures in Practice: A Large-Scale Analysis of Buying Patterns

Martin Kirchgessner, Vincent Leroy,
Sihem Amer-Yahia, Shashwat Mishra
Univ. Grenoble Alpes, CNRS, LIG
F-38000 Grenoble, France
firstname.lastname@imag.fr

Intermarché Alimentaire International
STIME
17 Allée des Mousquetaires
91078 Bondoufle cedex, France
contact@mousquetaires.fr

Abstract—Understanding customer buying patterns is of great interest to the retail industry. Association rule mining is a common technique for extracting correlations such as *people in the South of France buy rosé wine* or *customers who buy paté also buy salted butter and sour bread*. Unfortunately, sifting through a high number of buying patterns is not useful in practice, because of the predominance of popular products in the top rules. As a result, a number of “interestingness” measures (over 30) have been proposed to rank rules. However, there is no agreement on which measures are more appropriate for retail data. Moreover, since pattern mining algorithms output thousands of association rules for each product, the ability for an analyst to rely on ranking measures to identify the most interesting ones is crucial. In this paper, we develop CAPA (Comparative Analysis of PAtterns), a framework that provides analysts with the ability to compare different rule rankings. We report on how we used CAPA to compare 34 interestingness measures applied to patterns extracted from customer receipts of more than 1,800 stores for a period of one year.

I. INTRODUCTION

Ever since databases have been able to store basket data, many techniques have been proposed to extract useful insights for analysts. One of the first, association rule mining [1], also remains one of the most intuitive. Association rules are often used to summarize consumer trends in a set of transactions or as input to a classifier [2]. The problem is the very high number of rules, typically in the order of millions. That is exacerbated by the lack of thorough studies of which of the many interestingness measures for ranking rules [3] is most appropriate for which application domain. We present CAPA, a framework to compare the rankings that result from applying different interestingness measures to association rules in the retail domain. CAPA relies on a flexible architecture and on *j*LCM [4], our parallel and distributed pattern mining algorithm that runs on MapReduce. The use of real datasets and a close collaboration with experienced domain experts from *Intermarché*, one of the largest retailers in France, has led us to select the most relevant measures to rank association rules in the food retail domain.

Our dataset contains 290 million receipts from 1,884 stores in all of France, gathered over one year, 2013. Mining this data results in a huge number of rules. For example, using a minimum support of 1,000 *j*LCM mines 2,746,418 frequent rules of the form *customer segment* \rightarrow *product category*. Out

of these, 15,063 have a confidence of 50% or higher. Table I shows a ranking of the top-10 rules according to 3 different interestingness measures proposed in [3]. If we denote rules as $A \rightarrow B$, *confidence* is akin to precision and is defined as the probability to observe B given that we observed A , i.e., $P(B|A)$. *Piatetsky-Shapiro* [5] combines how A and B occur together with how they would if they were independent, i.e., $P(AB) - P(A)P(B)$. *Pearson’s* χ^2 , measures how unlikely observations of A and B are independent. This example shows that these measures result in different rule rankings.

The first question we ask ourselves is **how different are the rule rankings produced by existing interestingness measures in the retail domain?** We examine the rankings produced by 34 measures [3], [6]. This effort was conducted for 3 mining scenarios designed by experienced analysts from the marketing studies department of *Intermarché*. In the first scenario, *demo_assoc*, the analyst provides a target product category and expects rules of the form *customer segment* \rightarrow *category*, i.e. *customers who belong to the described segment purchase products in the target category*. In the other two scenarios, the analyst provides a target product p and expects rules of the form *set of products* $\rightarrow p$. Such rules are either extracted based on a receipt-centric view, where products are grouped by receipt (*prod_assoc_t*), or based on a customer-centric view, where products are grouped by customer across several receipts (*prod_assoc_c*). Our first finding is that existing interestingness measures can be automatically grouped into 6 families of similar measures, regardless of the mining scenario.

We then conducted a user study with two experienced domain experts from *Intermarché* in order to address the following question: **out of the 6 families of interestingness measures, which ones are meaningful?** Our study lets analysts choose one of 3 mining scenarios along with target products or categories, as well as a (hidden) ranking measure. Their interactions with the resulting list of association rules were observed and their feedback recorded. Overall, ranking rules by decreasing confidence was preferred because of its intuitive nature. Moreover, analysts were willing to trade some confidence (akin to precision) for a significant gain in support (akin to recall). This is achieved by the *Piatetsky-Shapiro’s* measure [5]. However, this measure promotes rules containing very frequent products that are interpreted as noise by analysts.

TABLE I

TOP-5 DEMOGRAPHICS ASSOCIATION RULES, ACCORDING TO DIFFERENT INTERESTINGNESS MEASURES. RULES ARE DENOTED {AGE, GENDER, DEPARTMENT} \rightarrow *product category*. PRODUCT CATEGORIES WERE TRANSLATED TO ENGLISH, FRENCH DEPARTMENTS WERE LEFT UNCHANGED.

| by confidence | | by Piatetsky-Shapiro [5] | | by Pearson's χ^2 | |
|---------------------------------------|---------------------|--------------------------|---------------------------|-----------------------|---------------------------|
| {> 65, F, Aube} | \rightarrow Dairy | {*, *, Nord} | \rightarrow Liquids | {*, *, Somme} | \rightarrow Cut cheese |
| {> 65, F, Aveyron} | \rightarrow Dairy | {*, *, Nord} | \rightarrow Soft drinks | {*, F, Somme} | \rightarrow Cut cheese |
| {> 65, F, Val de Marne} | \rightarrow Dairy | {*, *, Nord} | \rightarrow Beers | {> 65, *, Morbihan} | \rightarrow Fresh milk |
| {> 65, F, Seine S ^t Denis} | \rightarrow Dairy | {*, *, Nord} | \rightarrow Spreads | {> 65, *, Somme} | \rightarrow Cut cheese |
| {> 65, F, Haute Saone} | \rightarrow Dairy | {*, F, Nord} | \rightarrow Soft drinks | {*, *, Finistere} | \rightarrow Canned pork |

This drawback is alleviated by providing the ability to filter out rules containing uninteresting products.

CAPA is made possible with *j*LCM, our distributed pattern mining algorithm that is able to mine millions of patterns in a few minutes [4]. *j*LCM can be constrained to focus on different customer demographics and product taxonomies. Thus, in addition to typical associations between products, it finds associations between customer segments and products and between products and categories.

In summary, this paper presents CAPA, a joint effort between researchers in Academia and business experts at *Intermarché*. The context and goals of the work are provided in Section II. The architecture of CAPA is overviewed in Section III. In Section IV, CAPA is deployed to perform an automatic grouping of measures into 6 groups based on similarities in the raking they produce. These groups are then evaluated by retail experts in Section V leading to insightful findings. The related work is summarized in Section VI. Planned and possible evolutions are finally discussed in Section VII.

II. CONTEXT

A. Dataset

We represent a dataset \mathcal{D} as a set of records of the form $\langle t, c, p \rangle$, where t is a unique receipt identifier, c is a customer, and p is a product purchased by c . When a customer purchases multiple products at the same time, several records with the same receipt identifier t are generated. The set of receipt identifiers is denoted as T . Each receipt identifier is associated with a unique customer, and multiple receipt identifiers can be associated with the same customer. We do not use product price or product cardinality in this work. The complete dataset contains over 290 million unique receipts, spanning 3.5 billion records, generated at a retail chain consisting of 1,884 stores over the whole year of 2013.

The set of customers, \mathcal{C} , contains over 9 million customers. Each customer has demographic attributes. In this study, we focus on 3 attributes: *age*, *gender* and *location*. The attribute *age* takes values in $\{<35, 35-49, 50-65, >65\}$ and the attribute *location* admits French departments as values. Each customer segment is described by a set of user attribute values that are interpreted in the usual conjunctive manner. For example, the segment $\{< 35, Paris\}$ refers to young Parisian customers.

We use $demo(c)$ to refer to the set of attribute values of a customer c . For example, $\{35-49, female, Calvados\}$ represents a 48 year old *female* from the *Calvados* department, whom we will refer to as *Mary*.

The set of products \mathcal{P} contains over 200,000 entries, out of which 55,786 have been sold more than a thousand times. Products are organized in a taxonomy with 19,557 nodes over 4 levels. Products are leaf nodes, and belong to all their ancestor categories. The set of categories a product p belongs to is denoted as $cat(p)$. For example, *chocolate cream* belongs to the categories *Fresh food*, *Dairy*, *Ultra fresh* and *Desserts*.

B. Mining Customer Receipts

1) *Dataset Preparation*: Our analysts are interested in studying two kinds of buying patterns: those representing associations between customer segments and a product category (*young people in the north of France consume sodas*), and those associating a set of products to a single product (*people who purchase pork sausage and mustard also buy Riesling*). In all cases the analyst specifies \mathcal{B} , a set of rule targets.

In the first case, coined `demo_assoc`, \mathcal{B} contains one or more categories. The analyst expects rules of the form *customer segment* \rightarrow *category*, i.e. customers who purchase products in the target category. The second case comes in two variants: `prod_assoc_t`, a receipt-centric view where products are found in the same receipt, and `prod_assoc_c`, a customer-centric view where products are purchased by the same customer over time. In these variants, \mathcal{B} only contains products (as opposed to categories in the first scenario) and the analyst expects rules of the form *set of products* \rightarrow *target product* $p \in \mathcal{B}$. The dataset \mathcal{D} is transformed into a collection of transactions \mathcal{T} that is given as input to the mining process, as summarized in Table II. The set \mathcal{T} is constructed differently for each scenario.

In `demo_assoc`, a transaction is a tuple built for each receipt $\langle t, c, p \rangle$ by associating $demo(c)$ with $cat(p)$. For *Mary*, the record $\langle 234567, Mary, chocolate\ cream \rangle$ is mapped to the transaction $\langle 35-49, female, Calvados, chocolate\ cream, Fresh\ food, Dairy, Ultra\ fresh, Desserts \rangle$. The number of transactions is equal to $|\mathcal{D}|$, and each transaction contains the segments a customer belongs to, and the categories of the product purchased. In `prod_assoc_t`, \mathcal{T} is built by grouping the records in \mathcal{D} by receipt identifier, t . For each t , we generate a transaction as the set of products bought in a single visit to the store $\{p | \langle t, c, p \rangle \in \mathcal{D}\}$. If *Mary* has a store receipt containing the products *cream*, *yogurt*, *cola*, a transaction containing the 3 products is generated. This leads to a total of $|T|$ transactions, where each transaction is a subset of \mathcal{P} . In `prod_assoc_c`, we generate the set of transactions \mathcal{T} by grouping records in \mathcal{D} by customer. For each customer

TABLE II
OUR MINING SCENARIOS AND EXAMPLE ASSOCIATION RULES.

| Target Associations | Input transactions \mathcal{T} | Desired association rules |
|---|--|--|
| demo_assoc: segment \rightarrow category | $\{demo(c) \cup cat(p) \langle t, c, p \rangle \in \mathcal{D}\}$ min support is 1,000 | A segment tends to purchase products in a category. $\{< 35, F, *\} \rightarrow Baby\ food$ |
| prod_assoc_t: product(s) \rightarrow product | $\{\cup_{\langle t, c, j, p_i \rangle \in \mathcal{D}} p_i t \in T\}$ min support is 1,000 | Products purchased <i>simultaneously</i> . $\{vanilla\ cream\} \rightarrow chocolate\ cream$ |
| prod_assoc_c: product(s) \rightarrow product | $\{\cup_{\langle t, j, c, p_i \rangle \in \mathcal{D}} p_i c \in \mathcal{C}\}$ min support is 10,000 | Customers' product associations over time. $\{Pork\ sausage, mustard\} \rightarrow dry\ Riesling$ |

c , we generate a single transaction containing all products ever purchased by c $\{p | \langle t, c, p \rangle \in \mathcal{D}\}$. We obtain $|\mathcal{C}|$ transactions, each of which is a subset of \mathcal{P} . This enables the discovery of patterns occurring over several visits to a store. The number of transactions in `prod_assoc_c` (9,267,961) is less important than in `prod_assoc_t` (290,734,163), but each transaction contains on average 214 products, against 12 for receipts.

2) *Mining Scenarios*: Given a frequency threshold $\varepsilon \in [1, n]$, an itemset P is said to be *frequent* in a transactions set \mathcal{T} iff $support_{\mathcal{T}}(P) \geq \varepsilon$ where $support_{\mathcal{T}}(P)$ is the number of transactions in \mathcal{T} that contain all items in P . As indicated in Table II, we set the frequency threshold to different values in different scenarios because they differ in the cardinalities of their transactions. Moreover, because marketing actions are decided and applied nation-wide, they are expected to concern at least 1,000 customers, and preferably more than 10,000.

An itemset P is *closed* iff there exists no itemset $P' \supset P$ such that $support_{\mathcal{T}}(P) = support_{\mathcal{T}}(P')$ [7]. The number of closed itemsets can be orders of magnitude less important than the number of itemsets, while providing the same amount of information on \mathcal{T} . Several algorithms, including ours, focus on extracting frequent closed itemsets, increasing performance and avoiding redundancy in results [8], [9].

We consider our mining scenarios described in Section II-B1. Each scenario leads to the construction of a collection of transactions \mathcal{T} , where a transaction is a set of items. Given \mathcal{T} , a frequency threshold ε , we find all closed frequent itemsets, and use them to derive association rules [10]. Each itemset P implies an association rule of the form $A \rightarrow B$ where A, B is a partition of P . A is the antecedent of the rule, and B its consequent. In `demo_assoc`, A is a customer segment and B is a single product category. In `prod_assoc_t` and `prod_assoc_c`, A is a set of products ($A \subseteq \mathcal{P}$) and B is a single product. Analysts generally focus on particular products or product categories. This is why they specify the list of targets \mathcal{B} in each scenario. Table II contains example association rules extracted from our dataset.

C. Interestingness Measures

Large datasets often contain millions of frequent closed itemsets, and each of them may lead to several association rules. The ability to identify valuable rules is therefore of the utmost importance to avoid drowning analysts in useless information. Association rules $A \rightarrow B$ were originally selected using thresholds for support ($support_{\mathcal{T}}(A \cup B)$) and confidence ($\frac{support_{\mathcal{T}}(A \cup B)}{support_{\mathcal{T}}(A)}$) [1]. However using two separate

values, and guessing the right threshold is not natural. Furthermore, support and confidence do not always coincide with the interest of analysts. Hence, a number of interestingness measures that serve different analyses were proposed in the literature [3], [11], [12]. Table III summarizes the measures we use in this work. The first column contains the name of the measure, the second its expression. The last column will be referred to later.

D. Goal

Our goal is to help analysts test and compare the rankings produced by different interestingness measures on rules extracted from \mathcal{D} . An analyst can specify one of 3 mining scenarios, `demo_assoc`, `prod_assoc_t`, and `prod_assoc_c`, and one or several targets (categories in the case of `demo_assoc`, products in the case of the other two), and CAPA generates as many rule rankings as the number of interestingness measures.

III. ARCHITECTURE

Figure 1 contains the main components of CAPA and their interactions. The first module is **acquisition and storage**. Sales records are produced locally at each store, and are loaded daily into a data center ①. Records are stored in a *sales* table, and are augmented with customer segments coming from the *customers* table ②. CAPA's **curation** module is used to build transactions. The analysts selects a mining scenario and a set of input targets ③, which are used to generate the appropriate collection of transactions \mathcal{T} ④. CAPA's **mining** component relies on *jLCM*, an open-source pattern mining library that we developed [4], to compute a set of association rules matching the input targets ⑤. CAPA's **exploitation** component computes the quality of produced rules according to each interestingness measure ⑥, and loads them into a database. Results are presented to the analyst through a web application ⑦. We now describe the details of CAPA.

A. Acquisition and storage

Each of the 1,884 stores locally maintains a log of all customer transactions completed during the day. Whenever a customer checks out, a receipt is generated, indicating the list of products purchased, their price, as well as potential discounts. These receipts are logged under the form of $\langle r, c, p \rangle$ triples and stored in a write-ahead log. Once a day, during the store's closing time, this log is transmitted to the main data center that centralizes all sales records.

TABLE III

INTERESTINGNESS MEASURES OF A RULE $A \rightarrow B$. *, \triangleright INDICATE MEASURES PRODUCING IDENTICAL RULE RANKINGS WHEN B IS FIXED. \diamond , \dagger , \ominus , \otimes INDICATE MEASURES THAT ALWAYS PRODUCE THE SAME RULE RANKING. $|T|$ IS THE NUMBER OF TRANSACTIONS. $P(A) = \text{support}(A)/|T|$.

| Measure | Formula | Group and description | |
|---------------------------------------|---|-----------------------|--|
| One-Way Support | $P(B A) \times \log_2 \frac{P(AB)}{P(A)P(B)}$ | G_1 | Highest confidence Very low recall Favors frequent targets |
| Relative Risk | $P(B A)/P(B \neg A)$ | | |
| Odd Multiplier | $(P(AB)P(\neg B))/(P(B)P(A\neg B))$ | | |
| Zhang | $\frac{P(AB) - P(A)P(B)}{\max(P(AB)P(\neg B), P(B)P(A\neg B))}$ | | |
| Yule's Q \diamond | $\frac{P(AB)P(\neg A\neg B) - P(A\neg B)P(B\neg A)}{P(AB)P(\neg A\neg B) + P(A\neg B)P(B\neg A)}$ | | |
| Yule's Y \diamond | $\frac{\sqrt{P(AB)P(\neg A\neg B)} - \sqrt{P(A\neg B)P(B\neg A)}}{\sqrt{P(AB)P(\neg A\neg B)} + \sqrt{P(A\neg B)P(B\neg A)}}$ | | |
| Odds Ratio \diamond | $(P(AB)P(\neg A\neg B))/(P(A\neg B)P(B\neg A))$ | | |
| Information Gain $*\ominus$ | $\log(P(AB)/(P(A)P(B)))$ | | |
| Lift $*\ominus$ | $P(AB)/(P(A)P(B))$ | | |
| Added Value $*$ | $P(B A) - P(B)$ | | |
| Certainty Factor $*$ | $(P(B A) - P(B))/(1 - P(B))$ | | |
| Confidence / Precision $*\otimes$ | $P(B A)$ | | |
| Laplace Correction $*\otimes$ | $(\text{support}(AB) + 1)/(\text{support}(A) + 2)$ | | |
| Loevinger \dagger | $1 - P(A)P(\neg B)/P(A\neg B)$ | | |
| Conviction \dagger | $P(A)P(\neg B)/P(A\neg B)$ | | |
| Example and Counter-example Rate | $1 - P(A\neg B)/P(AB)$ | | |
| Sebag-Schoenauer | $P(AB)/P(A\neg B)$ | | |
| Leverage | $P(B A) - P(A)P(B)$ | | |
| Least Contradiction | $(P(AB) - P(A\neg B))/P(B)$ | G_2 | Very high confidence Very low recall |
| Accuracy | $P(AB) + P(\neg A\neg B)$ | | |
| Pearson's $\chi^2 \triangleright$ | $ T \times \left(\frac{(P(AB) - P(A)P(B))^2}{P(A)P(B)} + \frac{(P(\neg AB) - P(\neg A)P(B))^2}{P(\neg A)P(B)} \right) + T \times \left(\frac{(P(A\neg B) - P(A)P(\neg B))^2}{P(A)P(\neg B)} + \frac{(P(\neg A\neg B) - P(\neg A)P(\neg B))^2}{P(\neg A)P(\neg B)} \right)$ | G_3 | High confidence Low recall Low sensitivity (to target frequency) |
| Gini Index \triangleright | $P(A) \times (P(B A)^2 + P(\neg B A)^2) + P(\neg A) \times (P(B \neg A)^2 + P(\neg B \neg A)^2) - P(B)^2 - P(\neg B)^2$ | | |
| J-measure | $P(AB)\log\left(\frac{P(B A)}{P(B)}\right) + P(A\neg B)\log\left(\frac{P(\neg B A)}{P(\neg B)}\right)$ | | |
| Φ Linear Correlation Coefficient | $(P(AB) - P(A)P(B))/\sqrt{P(A)P(B)P(\neg A)P(\neg B)}$ | | |
| Two-Way Support Variation | $P(AB) \times \log_2 \frac{P(AB)}{P(A)P(B)} + P(A\neg B) \times \log_2 \frac{P(A\neg B)}{P(A)P(\neg B)} + P(\neg AB) \times \log_2 \frac{P(\neg AB)}{P(\neg A)P(B)} + P(\neg A\neg B) \times \log_2 \frac{P(\neg A\neg B)}{P(\neg A)P(\neg B)}$ | | |
| Fisher's exact test | $\frac{\binom{ T \times P(B)}{ T \times P(AB)} \binom{ T \times P(\neg B)}{ T \times P(A\neg B)}}{\binom{ T }{ T \times P(A)} \binom{ T }{ T \times P(B)}}$ | | |
| Jaccard | $P(AB)/(P(A) + P(B) - P(AB))$ | G_4 | Average confidence Average recall Low sensitivity |
| Cosine | $P(AB)/\sqrt{P(A)P(B)}$ | | |
| Two-Way Support | $P(AB) \times \log_2 \frac{P(AB)}{P(A)P(B)}$ | G_5 | Low confidence High recall |
| Piatetsky-Shapiro | $P(AB) - P(A)P(B)$ | | |
| Kloggen | $\sqrt{P(AB)\max(P(B A) - P(B), P(A B) - P(A))}$ | G_6 | Lowest confidence Highest recall Favors rare targets |
| Specificity | $P(\neg B \neg A)$ | | |
| Recall | $P(A B)$ | | |
| Collective Strength | $\frac{P(AB) + P(\neg B \neg A)}{P(A)P(B) + P(\neg A)P(\neg B)} \times \frac{1 - P(A)P(B) - P(\neg A)P(\neg B)}{1 - P(AB) - P(\neg B \neg A)}$ | | |

We rely on YARN [13] to administer the cluster storing sales records. Data is stored in an HBase database [14], and processed using the MapReduce framework [15]. Sales records are stored in the *sales* table. To avoid redundancy and ease data processing, records are grouped by receipt before being stored in *sales*. Thus, each receipt is a line in the table, and the content of the receipt is stored in the *meta* column family. We leverage HBase's flexibility on columns by recording each product identifier as a column qualifier, with information such as the cardinality and the unit-price as a value. The row key of each receipt is defined as *storeId-day-customerId-receiptId*. The *sales* table is configured to be sorted by row key. This allows operations such as selecting the sales records of a

given store to be efficiently performed in a single scan, while selecting a specific time period can also be done by combining 1,884 ranges (one per store identifier). Given that customer purchases may vary between geographical areas [16] and over time, these operations are frequently used by analysts. This data layout is optimized to perform these selections without incurring unnecessary reads. That enables storage of large amounts of data without increasing the cost of analyzing a fixed number of records. Sales logs transferred from the stores are first stored on the distributed file system HDFS, and then loaded into HBase using MapReduce, as a daily batch job.

Each customer constitutes an entry in the *customers* table, which records the segments she belongs to. After loading the

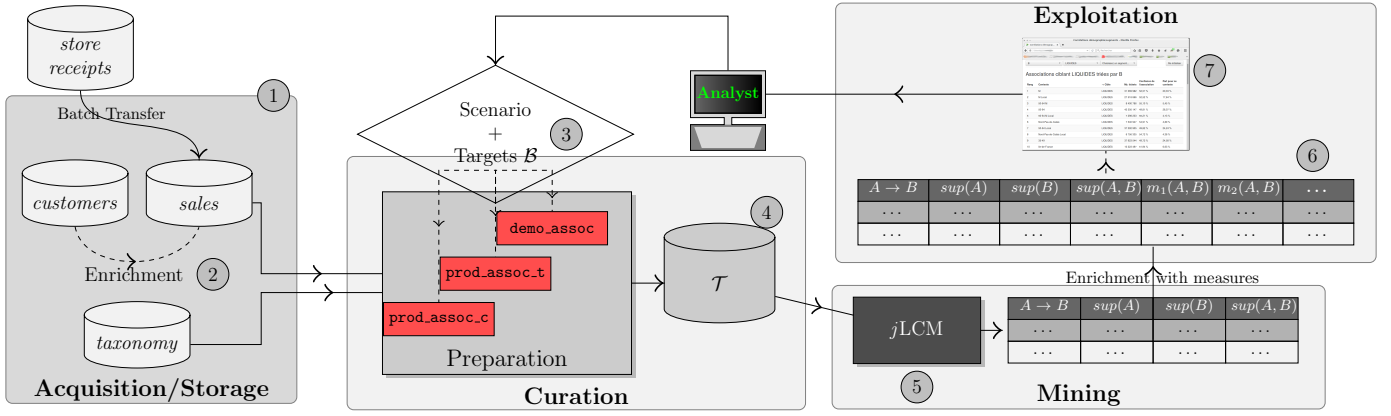


Fig. 1. Overview of the architecture

sales records into the database, we enrich the *sales* table using another MapReduce job. For each new record, the receipt is augmented with the user segments by querying the *customers* table and copying these segments to the *meta* column family in *sales*. Hence, each sales record is assigned a snapshot of the user information at the time the receipt was generated.

B. Curation

As described in Section II-B1, mining customer receipts begins with the construction of a transactions dataset \mathcal{T} following the requirements of the analyst. This operation is performed using MapReduce jobs executed on the *sales* table. In the case of *demo_assoc*, a single *map* operation is sufficient. The product taxonomy is loaded in memory by all mappers through the distributed cache, and, given a row, for each product registered in the *products* column family, a transaction containing its categories is generated. Customer segments are directly available in the *meta* column family thanks to the enrichment phase and are added to the transactions. As described in Section III-A, records are already grouped by receipt when stored in *sales*, thus no further processing is necessary for an analysis in *prod_assoc.t*. Each line of *sales* generates one transaction containing the set of products. In *prod_assoc.c*, the products bought by a given customer are grouped using a reduce operation with the customer identifier as a key to generate a transaction. In all cases, at the end of this phase the dataset \mathcal{T} is stored on HDFS as a text file, with one line per transaction.

C. Mining

1) *Extracting itemsets using jLCM*: Generating association rules, presented in Section II-B2, firstly requires the extraction of frequent itemsets from \mathcal{T} . We rely on *jLCM*, available as an open-source library [4]. *jLCM* is integrated in a MapReduce job. The itemset extraction job scans the input \mathcal{T} once in the *map* function, and builds for each target item b in \mathcal{B} , a filtered dataset limited to transactions containing b : $\mathcal{T}_{\{b\}} = \{E \in \mathcal{T}, b \in E\}$. This is done using the target items from \mathcal{B} as intermediate keys for the *reduce* function. For each target, *reduce* executes *jLCM* on its filtered dataset. *jLCM* is

a recursive algorithm that enumerates itemsets and computes their frequency. Closed itemsets are returned along with their support, with the exception of singletons that cannot produce association rules. In *demo_assoc*, itemsets should contain a single category only, so all categories except the target one are removed from transactions.

This technique allows us quickly obtain itemsets that satisfy our constraint, i.e. all itemsets contain one of the targets from \mathcal{B} . The job's run-time is dominated by the scan of \mathcal{T} in the *map* phase, which can be accelerated by the addition of worker nodes. On average, each *reduce* task completes in 10 seconds.

2) *Evaluating relevant association rules*: Analysts aim at uncovering interesting association rules expressed as $A \rightarrow B$. Evaluating the interestingness of an association rule requires computing the support of itemsets A , B and $A \cup B$ in \mathcal{T} . The standard method for mining association rules consists in finding all frequent itemsets in the dataset, and then generating the rules. Given that our analysts have specified a restricted set of targets \mathcal{B} , this approach would be wasteful. This motivates our distribution of the itemsets extraction, presented in the previous sub-section. Our itemsets extraction job gives the support of B and $A \cup B$ for all association rules we are interested in (i.e. all B satisfy $B = \{e\}, e \in \mathcal{B}$). This job also materializes, as a prefix tree in a side-output file, the set \mathcal{A} of all antecedent itemsets, whose support needs to be evaluated.

A second MapReduce job completes the evaluation of association rules. Each *map* operation reads a transaction of \mathcal{T} and counts the support all association rules' antecedents. The *reduce* phase uses itemsets in \mathcal{A} as intermediate keys and sums partial counts to obtain the total support. This two-step approach avoids the computation of many itemsets that never appear as a rule antecedent.

D. Exploitation

The quality measures selected require at most $P(A)$, $P(B)$ and $P(A \cup B)$ to be computed, because, given $|\mathcal{T}|$, other probabilities like $P(B|A)$ or $P(A \neg B)$ can be derived from them. Therefore, we denormalize the results of the mining phase in order to store those 3 probabilities with each A and B . The supports of all rules' antecedents (providing $P(A)$)

are centralized and joined to the results of j LCM (which provides $P(AB)$ and $P(B)$). After this denormalization, each row represents an association rule and has enough information to compute its score. This table is then augmented with 34 columns, one for each measure implemented in CAPA, and listed in Table III. Because large numbers are involved, for *Fisher's exact test* we actually use the logarithm of the binomial coefficients, which are computed as logarithms of the gamma function. This makes the calculation feasible, but requires long iterations so we do it in parallel again (this is easy to implement thanks to the denormalization). The complete table is stored in a relational database.

The final component of CAPA is a web application allowing the analyst to explore this augmented table. In any scenario, the analyst picks a measure and selects a target product or category, or a set of target products or categories. Association rules are then returned in a table and sorted according to the selected measure. A rule like *yogurt* \rightarrow *cheese* is displayed with 3 values: *support* (number of customers who bought both cheese and yogurt), *confidence* (fraction of yogurt buyers who also bought cheese), *recall* (fraction of cheese buyers who also bought yogurt). During the user study these figures help analysts quickly judge the volume of sales for each rule.

IV. EMPIRICAL EVALUATION

We present an empirical evaluation of the 34 measures for association rules introduced in Section II-C. Recall that our goal, stated in Section II-D, is to assist the analyst in selecting measures. Our evaluation consists in comparing rankings produced by these measures on retail data to discover which measures differ significantly in practice. We then use that similarity to classify ranking measures into *groups*. We annotate these groups based on the properties common to the group. We discuss key insights obtained from experimentation on each group. The goal of this evaluation is to automatically detect similarities between interestingness measures and reduce the number of candidate measures to present to analysts in the user study (Section V).

We first present in Section IV-A methods used to compare ranked list. Then, we compare the resulting rankings in Section IV-B). We conclude the empirical evaluation with the selection of representative measures in Section IV-C.

A. Ranking similarity measures

In this section, we discuss some methods for comparison of ranked lists. The first three methods are taken from the literature. We then introduce *NDCC*, a new parameter-free ranking similarity designed to emphasize differences at the top of the ranking.

We are given of a set of association rules \mathcal{R} to rank. We interpret each measure, m , as a function that receives a rule and generates a score, $m : \mathcal{R} \rightarrow \mathbb{R}$. We use $L_{\mathcal{R}}^m$ to denote an ordered list composed of rules in \mathcal{R} , sorted by decreasing score. Thus, $L_{\mathcal{R}}^m = \langle r_1, r_2, \dots \rangle$ s.t. $\forall i > i' m(r_i) < m(r_{i'})$. We generate multiple lists, one for each measure m , from the same set \mathcal{R} . $L_{\mathcal{R}}^m$ denotes a ranked list of association rules

according to measure m where the rank of rule r is given as $rank(r, L_{\mathcal{R}}^m) = |\{r' | r' \in \mathcal{R}, m(r') \geq m(r)\}|$. To assess the dissimilarity between two measures, m and m' , we compute the dissimilarity between their ranked lists, $L_{\mathcal{R}}^m$ and $L_{\mathcal{R}}^{m'}$. We use r^m as a shorthand notation for $rank(r, L_{\mathcal{R}}^m)$.

1) *Spearman's rank correlation coefficient*: Given two ranked lists $L_{\mathcal{R}}^m$ and $L_{\mathcal{R}}^{m'}$, *Spearman's rank correlation* [17] computes a linear correlation coefficient that varies between 1 (identical lists) and -1 (opposite rankings) as shown below.

$$Spearman(L_{\mathcal{R}}^m, L_{\mathcal{R}}^{m'}) = 1 - \frac{6 \sum_{r \in \mathcal{R}} (r^m - r^{m'})^2}{|\mathcal{R}|(|\mathcal{R}|^2 - 1)}$$

This coefficient depends only on the difference in ranks of the element (rule) in the two lists, and not on the ranks themselves. Hence, the penalization is the same for differences occurring at the beginning or at the end of the lists.

2) *Kendall's τ rank correlation coefficient*: *Kendall's τ rank correlation coefficient* [18] is based on the idea of agreement among element (rule) pairs. A rule pair is said to be *concordant* if their order is the same in $L_{\mathcal{R}}^m$ and $L_{\mathcal{R}}^{m'}$, and *discordant* otherwise. τ computes the difference between the number of concordant and discordant pairs and divides by the total number of pairs as shown below.

$$\tau(L_{\mathcal{R}}^m, L_{\mathcal{R}}^{m'}) = \frac{|C| - |D|}{\frac{1}{2}|\mathcal{R}|(|\mathcal{R}| - 1)}$$

$$C = \{(r_i, r_j) | r_i, r_j \in \mathcal{R} \wedge i < j \wedge \text{sgn}(r_i^m - r_j^m) = \text{sgn}(r_i^{m'} - r_j^{m'})\}$$

$$D = \{(r_i, r_j) | r_i, r_j \in \mathcal{R} \wedge i < j \wedge \text{sgn}(r_i^m - r_j^m) \neq \text{sgn}(r_i^{m'} - r_j^{m'})\}$$

Similar to *Spearman's*, τ varies between 1 and -1 , and penalizes uniformly across all positions.

3) *Overlap@k*: *Overlap@k* is another method for ranked lists comparison widely used in Information Retrieval. It is based on the premise that in long ranked lists, the analyst is only expected to look at the top few results that are highly ranked. While *Spearman* and τ account for all elements uniformly, *Overlap@k* compares two rankings by computing the overlap between their top- k elements only.

$$Overlap@k(L_{\mathcal{R}}^m, L_{\mathcal{R}}^{m'}) = \frac{|\{r \in \mathcal{R} \mid r^m \leq k \wedge r^{m'} \leq k\}|}{k}$$

4) *Normalized Discounted Correlation Coefficient*: *Overlap@k*, *Spearman's* and τ sit at two different extremes. The former is conservative in that it takes into consideration only the top k elements of the list whereas the latter two take too liberal an approach by penalizing all parts of the lists uniformly. In practice, we aim for a good tradeoff between these extremes.

To bridge this gap, we propose a new ranking correlation measure coined *Normalized Discounted Correlation Coefficient* or *NDCC*. *NDCC* draws inspiration from *NDCG*, *Normalized Discounted Cumulative Gain* [19], a ranking measure

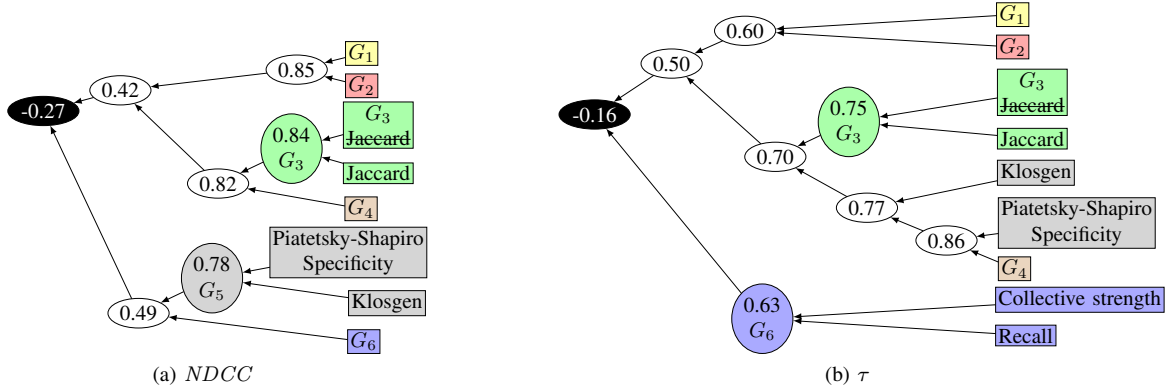


Fig. 2. Hierarchical clustering of interestingness measures for a single target

commonly used in Information Retrieval. The core idea in $NDCG$ is to reward a ranked list $L_{\mathcal{R}}^m$ for placing an element r of relevance rel_r by $\frac{rel_r}{\log r^m}$.

The logarithmic part acts as a smoothing discount rate representing the fact that as the rank increases, the analyst is less likely to observe r . In our setting, there is no ground truth to properly assess rel_r . Instead, we use the ranking assigned by m' as a relevance measure for r , with an identical logarithmic discount. When summing over all of \mathcal{R} , we obtain DCC , which presents the advantage of being a symmetric correlation measure between two rankings $L_{\mathcal{R}}^m$ and $L_{\mathcal{R}}^{m'}$.

$$DCC(L_{\mathcal{R}}^m, L_{\mathcal{R}}^{m'}) = \sum_{r \in \mathcal{R}} \frac{1}{\log(1+r^m) \log(1+r^{m'})}$$

We compute $NDCC$ by normalizing DCC between 1 (identical rankings) and -1 (reversed rankings).

$$NDCC(L_{\mathcal{R}}^m, L_{\mathcal{R}}^{m'}) = \frac{dcc - avg}{max - avg}$$

$$\begin{aligned} \text{where } dcc &= DCC(L_{\mathcal{R}}^m, L_{\mathcal{R}}^{m'}), \text{ max} = DCC(L_{\mathcal{R}}^{m'}, L_{\mathcal{R}}^{m'}) \\ \text{min} &= DCC(L_{\mathcal{R}}^*, L_{\mathcal{R}}^{m'}), \text{ } L_{\mathcal{R}}^* = rev(L_{\mathcal{R}}^m) \\ avg &= (\text{max} + \text{min})/2 \end{aligned}$$

5) *Ranking comparison by example:* We illustrate the difference between all ranking correlation measures with an example in Table IV. This shows correlation of a ranking L^1 with 3 others, according to each measure. $NDCC$ does indeed penalize differences at higher ranks, and is more tolerant at lower ranks.

B. Rankings comparison

We perform a comparative analysis of ranking measures, on our 3 mining scenarios summarized in Table II. We generate association rules $A \rightarrow B$ where B is a single product among a set of 64 previously studied by analysts. Overall we obtain 1,651,024 association rules, and we compute one rule ranking per product and per interestingness measure. Our first observation is that the results we obtain for all scenarios lead to the same conclusions. Therefore, we only report numbers for `prod_assoc_c`.

TABLE IV
EXAMPLE RANKINGS AND CORRELATIONS

| | Ranking | Content | | |
|-------|-----------------|----------------------|------------------|-------------|
| | L^1 | r_1, r_2, r_3, r_4 | | |
| | L^2 | r_2, r_1, r_3, r_4 | | |
| | L^3 | r_1, r_2, r_4, r_3 | | |
| | L^4 | r_2, r_3, r_1, r_4 | | |
| | <i>Spearman</i> | τ | <i>Overlap@2</i> | <i>NDCC</i> |
| L^2 | 0.80 | 0.67 | 1 | 0.20 |
| L^3 | 0.80 | 0.67 | 1 | 0.97 |
| L^4 | 0.40 | 0.33 | 0.5 | -0.18 |

While all measures are computed differently, we notice that some of them always return the same ranking for association rules of a given target. We identify them in Table III using symbols. Other notable similarities include *Sebag-Schoenauer* and *lift* (89% of rankings are equal), as well as *Loevinger* and *lift* (87%). This difference between the number of interestingness measures considered (34) and the number of different rankings obtained (25) can easily be explained analytically in the case of a fixed target. Indeed, for a given ranking, $P(B)$ is constant, which eliminates some of the differences between interestingness measures. In addition, some measures only have subtle differences which only appear when selecting extreme values for $P(A)$, $P(B)$ and $P(AB)$, which do not occur in practice in our retail dataset.

1) *Comparative analysis:* We now evaluate similarity between interestingness measures that do not return the same rankings. We compute a 34×34 correlation matrix of all rankings according to each correlation measure described in Section IV-A, and average them over the 64 target products. This gives us a ranking similarity between all pairs of measures. We then rely on hierarchical clustering with average linkage [20] to obtain a dendrogram of interestingness measures and analyze their similarities. The dendrograms for $NDCC$ and τ are presented in Figure 2. For better readability, we merge sub-trees when correlation is above 0.9. To describe the results more easily, we partition interestingness measures into 6 groups, as indicated in the third column in Table III.

G_1 is by far the largest group: in addition to 4 measures that

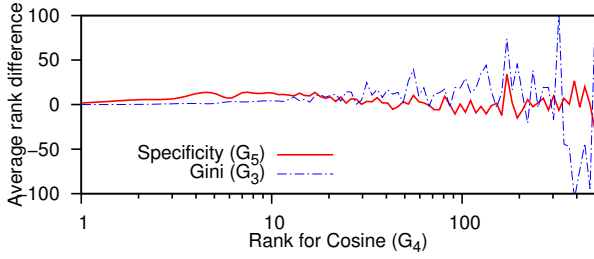


Fig. 3. Rank correlations




Fig. 4. Avg. recall/confidence of the top-20 results of interestingness measures

always generate the same rankings, 14 other measures output similar results. A second group, G_2 , comprising 2 measures, is quite similar to G_1 according to $NDCC$. τ also discovers this similarity, but considers it lower, which shows that it is mostly caused by high ranks. $Jaccard$ is a slight outlier in G_3 according to $NDCC$. Indeed, when focusing on the first 20 elements ($Overlap@20$), only an average of 71% are shared between $Jaccard$ and the rest of G_3 . This situation also occurs between $Klossgen$ and the rest of G_5 . Interestingly, we observe that, according to $NDCC$, G_5 is closest to G_6 and is negatively correlated with the other groups. However, according to τ , G_5 is very similar to G_4 and is negatively correlated with G_6 . This difference between ranking measures illustrates the importance of accounting for rank positions. When the top of the ranking is considered more important, some similarities emerge.

We illustrate this behavior in Figure 3 by displaying correlation between rankings obtained with different interestingness measures. This experiment clearly shows that overall, $cosine$ (G_4) is closer to $specificity$ (G_5) than $Gini$ (G_3), as the rank difference observed in the results is overall smaller. However, when focusing on the top-10 results of $cosine$, $Gini$ assigns closer ranks than $specificity$. This explains the difference in clustering between $NDCC/overlap$ and $\tau/Spearman$.

2) *Annotating groups*: While using hierarchical clustering on interestingness measures allows the discovery of families of measures, and their relative similarity, it does not fully explain which types of results are favored by each of them. We propose to compare their outputs according to the two most basic and intuitive interestingness measures employed in data mining: $recall$ and $confidence$. $recall$ represents the proportion of target items that can be retrieved by a rule, that is, $P(A|B)$. Its counterpart, $confidence$, represents how often the consequent is present when the antecedent is, that is, $P(B|A)$. We present, in Figure 4, the average $recall$ and $confidence$ of the top-20

rules ranked according to each interestingness measure. G_1 contains $confidence$, so it is expected to score the highest on this dimension. G_2 is extremely close to G_1 , but obtains slightly lower $confidence$ and $recall$. We then have, in order of increasing $recall$ and decreasing $confidence$ G_3 , G_4 and G_5 . Finally, G_6 , which contains $recall$, obtains the highest $recall$ but the lowest $confidence$. Figure 4 also shows that executing a Euclidean distance-based clustering, such as k -means, with $recall/confidence$ coordinates would lead to groups similar to the ones obtained with hierarchical clustering. Hence, this analysis is consistent with the hierarchical grouping and the correlation with $NDCC$.

While we believe that $NDCC$ better reflects the interpretation of analysts browsing rules, it is important to note that the grouping of interestingness measures created through this evaluation is stable across all 4 correlation measures and for all 3 scenarios. Correlation between different families of measures may vary, but measures within a single family always have a high similarity. Thus, we conjecture that the obtained results are true in the general case of food retailers and we can rely on these groups to reduce the number of options presented to analysts.

C. Selecting representative measures

We summarize the findings of the comparative evaluation in the last column of Table III. We identify 6 families of measures that behave similarly. Each family offers a different trade-off in terms of confidence and recall, and thus ranks association rules differently. We select the quality measure that most represents each family of measures (i.e. with highest average similarity) in order to confront the results of this analysis with the opinion of domain experts in our user study. Taking a general data mining perspective leads us to considering G_3 and G_4 as the most promising families for finding interesting association rules. Indeed, it is important to achieve a good trade-off between $recall$ and $confidence$ in order to find reliable association rules that can be applied in a significant number of cases. Hence, $F1$ score, that combines $recall$ and $confidence$, would prefer G_3 and G_4 to others.

V. USER STUDY

We now report the results of a user study with domain experts from *Intermarché*. The goal of this study is to assess the ability of interestingness measures to rank association rules according to the needs of an analyst. As explained in Section IV, we identified 6 families of measures, and selected a representative of each group for the user study (their names are in bold in Table III). We rely on the expertise of our industrial partner to determine, for each analysis scenario, which family produces the most interesting results. This experiment involved 2 experienced analysts from the marketing department of *Intermarché*. We setup *CAPA* and let analysts select targets multiple times in order to populate the web application's database with association rules (Section III-D). We let our analysts interact with *CAPA* without any time restriction, and collect their feedback in a free text form.

Each analyst firstly has to pick a mining scenario among `demo_assoc`, `prod_assoc_t`, or `prod_assoc_c`. Then she picks a target category or a target product in the taxonomy. In `prod_assoc_t` and `prod_assoc_c`, she also has the option to filter out rules whose antecedent products are not from the same category as the target. Finally, she chooses one of our 6 ranking measures to sort association rules. Neither the name of the measure nor its computed values for association rules are revealed, because we wanted analysts to evaluate rankings without knowing how they were produced.

Resulting association rules are ranked according to a selected measure. Each rule is displayed with its support, confidence and recall, such that analysts can evaluate it at a glance. For each scenario, our analysts are asked which representative measure highlights the most interesting results (as detailed below, in all cases a few of them were chosen).

A. Scrolling behavior

Once the analyst selects a target, *all* matching rules are returned. The initial motivation of this choice was to determine how many results are worth displaying and are actually examined by the analysts. According to the follow-up interview with the analysts, they carefully considered the first ten results, and screened up to a hundred more. Interestingly, analysts mentioned that they also scrolled down to the bottom of the list in order to see which customer segments are not akin to buying the selected category. For example, when browsing demographic association rules, they expected to find $\{50-64\} \rightarrow \text{pet food}$ among top results, but also expected $\{<35, \text{Paris}\} \rightarrow \text{pet food}$ among bottom results. This confirms that all rules should remain accessible. This also indicates that while interestingness measures favor strong associations, it could also be useful to highlight *anti*-rules.

B. Feedback on ranking measures

We let marketing experts explore all 3 scenarios and express their preference towards groups of measures.

In the `demo_assoc` case, G_1 and G_3 were both highly appreciated. G_1 favors rules such as $\{< 35, M, \text{Oise}\} \rightarrow \text{Flat and Carbonated drinks}$. These rules are very specific and thus have a very high confidence (31,58 % in this particular case). However, this comes at the cost of recall (0,08 %). Experts involved in this study value *confidence* much more than *recall*, as their priority is finding rules that they consider reliable. A low support is not necessarily an issue, and can lead to the discovery of surprising niche rules that can be exploited nonetheless. As discussed in Section IV-B2, G_3 offers a more balanced trade-off between confidence and recall, and prioritizes rules such as $\{< 35, *, *\} \rightarrow \text{Baby food}$ (confidence 8,57 %, recall 37,61%). These rules are interesting because they capture a large fraction of the sales of a given category, but are less reliable and generally less surprising. G_2 and G_4 were considered as less interesting than G_1 and G_3 respectively. Their results offer similar trade-offs, but with lower confidence each time. G_5 and G_6 were considered unusable because of their very low confidence.

When experimenting with `prod_assoc`, we observed a slightly different behavior. By default, the analysts favored G_1 and G_2 because of the confidence of their results. Then, we offered the analysts the possibility of filtering the rules to only keep the ones in which the antecedent contains products from the same category as the target. This led to analysts favoring G_3 and G_5 . This difference is caused by an important criterion: the ability of a measure to filter out very popular products. For example, the rule $\{\text{van. cream, emmental}\} \rightarrow \text{choc. cream}$ usually appears just above its shorter version $\{\text{van. cream}\} \rightarrow \text{choc. cream}$, because the first one has a confidence of 32% and the second 31%. However, experts prefer the second one, because *emmental* (cheese) is among the heavy hitters in stores. Its addition to the rule is hence considered insignificant. This “noise” generally increases with *recall*. Hence, when no filtering is available, G_1 is selected, but analysts prefer the *recall* and *confidence* trade-off provided by G_3 and G_5 . Again, G_4 suffered from its proximity to G_3 with lower confidence, while G_6 ’s confidence was too low.

In all cases, analysts mentioned G_6 as uninteresting because it selects rules of low *confidence*. In general, sorting by decreasing *lift* (which is close to sorting by decreasing *confidence*) is the preferred choice. Combined with the minimum support threshold used in the mining phase, this ranking promotes rules that are considered reliable. However, in the case where analysts are given the ability to filter out noisy products (very frequent ones), they prefer the ranking produced by *Piatetsky-Shapiro*’s measure [5]. That could be explained by the fact this measure provides a good compromise between *confidence* and *support*. The noisy products that this measure may introduce can be filtered out by analysts.

VI. RELATED WORK

To the best of our knowledge, CAPA targets datasets which are orders of magnitude bigger (and sparser) than those tested in existing work on ranking association rules. This paper is also the first to complement an algorithmic comparative analysis with a user study involving domain experts.

The definition of quality of association rules is a well-studied topic in statistics and data mining. In their survey [3], Geng *et al.* review 38 measures for association and classification rules. They also discuss 4 sets of properties like symmetry or monotony, and how each of them highlights different meanings of “rule quality”, such as novelty and generality. However, we observe no correlation between these properties and the groups of measures discovered using CAPA.

These 38 measures are compared in [21]. Authors consider the case of extracting and ranking temporal rules (*event A* \rightarrow *event B*) from the execution traces of programs. Each measure is evaluated in its ability to rank highly rules known from a ground truth (library specification). We observe that the measures scoring the highest are all from the groups identified in this work as G_1 and G_3 , which are also favored by our analysts. There are however some counterexamples, with measures from G_1 scoring poorly. The main difference between CAPA and [21] is the absence of a ground truth

of interesting rules for our dataset. Thus, our evaluation of measures is first comparative, with 4 correlations measures covering both the top of the ranking and the entire ranked list. The differences in the results obtained also highlight the importance of performing domain-specific studies, as properties of data and expectations of analysts vary significantly.

The closest work to ours is HERBS [6]. HERBS relies on a different and smaller set of measures to cluster rule rankings. Authors perform an analysis of the properties of measures, in addition to an experimental study. The datasets used are from the health and astronomy domains. Each of them contains at most 1,728 transactions and leads to the extraction of 49 to 6,312 rules. Rankings are then compared between all pairs of measures using Kendall's τ correlation measure averaged over all datasets. The largest group of measures identified, which includes confidence, is quite similar to G_1 . However, there are also significant differences. For instance, we find G_2 and G_6 to be very different, while [6] considers the measures of this group similar. The authors observe a weak resemblance between the theoretical and experimental analysis of the measures. CAPA is entirely focused on retail data, which has different properties and contains millions of transactions and rules. We also consider more interestingness measures, and 4 different ranking correlation measures instead of 1.

Our use of the p -value (via Pearson's χ^2 test) in the evaluation of rule interestingness is borrowed from [12]. A low p -value shows a correlation between a rule's antecedent and consequent. The use of Fisher's exact test on association rules is inspired by [11]. Both of these works aim at finding highly-correlated itemsets, which requires the analyst to set a threshold on the p -value. This is common practice in biology, but less meaningful in the retail industry. In [12], Liu *et al.* also propose an exploration framework where rules are grouped by consequent, then traversed by progressively adding items to the antecedent. The framework provides hints to help guess how each additional item would make a difference. Such a framework is suitable to some of the scenarios we consider and could be integrated in a future version of CAPA.

VII. SUMMARY AND EVOLUTIONS

In this paper, we present CAPA, a framework for mining association rules from large-scale retail data. We defined 3 mining scenarios allowing analysts to extract associations between user segments and product categories, or products themselves. Given a scenario, CAPA builds a dataset of transactions and mines in parallel association rules containing target products chosen by analysts. Our main contribution is the study of 34 interestingness measures for ranking rules. We first performed an analytical and an empirical comparison between different rule rankings and grouped measures into 6 families. Resulting groups were then evaluated in a user study involving retail experts. We concluded that *lift* and *Piatetsky-Shapiro* best fit the needs of analysts, as they ensure high confidence. Our user study also led us to thinking about the extraction of negative results (*anti*-rules). We are currently studying how to add and rank negative rules.

ACKNOWLEDGMENT

This work was partially funded by the Datalyse PIA project.

REFERENCES

- [1] R. Agrawal, T. Imieliński, and A. Swami, "Mining Association Rules between Sets of Items in Large Databases," in *Proc. SIGMOD*, 1993, pp. 207–216.
- [2] P. S. G.C., C. Sun, K. G. K., H. Zhang, F. Yang, N. Rampalli, S. Prasad, E. Arcaute, G. Krishnan, R. Deep, V. Raghavendra, and A. Doan, "Why Big Data Industrial Systems Need Rules and What We Can Do About It," in *Proc. SIGMOD*, 2015, pp. 265–276.
- [3] L. Geng and H. J. Hamilton, "Interestingness Measures for Data Mining: A Survey," *ACM Comput. Surv.*, vol. 38, no. 3, 2006.
- [4] M. Kirchgessner, V. Leroy, A. Termier, S. Amer-Yahia, and M.-C. Rousset, "jLCM," <https://github.com/slide-lig/jlcm>, [Online; accessed 27-May-2016].
- [5] G. Piatetsky-Shapiro, *Knowledge Discovery in Databases*. Menlo Park, CA: AAI/MIT, 1991.
- [6] P. Lenca, B. Vaillant, P. Meyer, and S. Lallich, "Association Rule Interestingness Measures: Experimental and Theoretical Studies," in *Quality Measures in Data Mining*. Springer, 2007, pp. 51–76.
- [7] N. Pasquier, Y. Bastide, R. Taouil, and L. Lakhal, "Discovering Frequent Closed Itemsets for Association Rules," in *Proc. ICDT*, 1999, pp. 398–416.
- [8] J. Pei, J. Han, and R. Mao, "Closet: An Efficient Algorithm for Mining Frequent Closed Itemsets," in *Proc. SIGMOD*, 2000, pp. 21–30.
- [9] T. Uno, M. Kiyomi, and H. Arimura, "LCM ver. 2: Efficient Mining Algorithms for Frequent/Closed/Maximal Itemsets," in *Proc. ICDM Workshop FIMI*, 2004.
- [10] P.-N. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining, (First Edition)*. W. W. Norton & Company, 2007.
- [11] S. ichi Minato, T. Uno, K. Tsuda, A. Terada, and J. Sese, "A Fast Method of Statistical Assessment for Combinatorial Hypotheses based on Frequent Itemset Enumeration," *Lect. Notes Artif. Int.*, vol. 8725, pp. 422–436, 2014.
- [12] G. Liu, M. Feng, Y. Wang, L. Wong, S.-K. Ng, T. L. Mah, and E. J. D. Lee, "Towards Exploratory Hypothesis Testing and Analysis," in *Proc. ICDE*, 2011, pp. 745–756.
- [13] V. K. Vavilapalli, A. C. Murthy, C. Douglas, S. Agarwal, M. Konar, R. Evans, T. Graves, J. Lowe, H. Shah, S. Seth, B. Saha, C. Curino, O. O'Malley, S. Radia, B. Reed, and E. Baldeschwieler, "Apache Hadoop YARN: Yet Another Resource Negotiator," in *Proc. SOCC*, 2013, pp. 5:1–5:16.
- [14] The Apache Software Foundation, "HBase," <http://hbase.apache.org>, [Online; accessed 27-May-2016].
- [15] J. Dean and S. Ghemawat, "Mapreduce: Simplified Data Processing on Large Clusters," *Commun. ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [16] S. Mishra, V. Leroy, and S. Amer-Yahia, "Discovering characterizing regions for consumer products," in *Proc. DSAA*, 2015, pp. 1–10.
- [17] W. Daniel, *Applied Nonparametric Statistics*. Houghton Mifflin, 1978.
- [18] M. G. Kendall, "A New Measure of Rank Correlation," *Biometrika*, vol. 30, no. 1/2, pp. 81–93, 1938.
- [19] K. Järvelin and J. Kekäläinen, "Cumulated Gain-based Evaluation of IR Techniques," *ACM Trans. Inf. Syst.*, vol. 20, no. 4, pp. 422–446, 2002.
- [20] R. R. Sokal and C. D. Michener, "A Statistical Method for Evaluating Systematic Relationships," *Univ. Kans. Sci. Bull.*, vol. 38, pp. 1409–1438, 1958.
- [21] T.-D. Le and D. Lo, "Beyond Support and Confidence: Exploring Interestingness Measures for Rule-Based Specification Mining," in *Proc. SANER*, 2015, pp. 331–340.