



HAL
open science

CAPRE: A New Methodology for Product Recommendation Based on Customer Actionability and Profitability

Thomas Piton, Julien Blanchard, Fabrice Guillet

► **To cite this version:**

Thomas Piton, Julien Blanchard, Fabrice Guillet. CAPRE: A New Methodology for Product Recommendation Based on Customer Actionability and Profitability. ICDM Workshops, 2011, Vancouver, Canada. pp.466-473, 10.1109/ICDMW.2011.42 . hal-00698502

HAL Id: hal-00698502

<https://hal.science/hal-00698502v1>

Submitted on 26 Sep 2013

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.

CAPRE: A New Methodology for Product Recommendation Based on Customer Actionability and Profitability

Thomas Piton, Julien Blanchard and Fabrice Guillet

KnOwledge & Decision Team

LINA UMR CNRS 6241, Computer Science Laboratory of Nantes-Atlantique

Polytechnic School of Nantes University, France

{thomas.piton, julien.blanchard, fabrice.guillet}@univ-nantes.fr

Abstract—Recommender systems can apply knowledge discovery techniques to the problem of making product recommendations. This aims to establish a customer loyalty strategy and thus to optimize the customer life time value. In this paper we propose CAPRE, a data-mining based methodology for recommender systems based on the analysis of turnover for customers of specific products. Contrary to classical recommender systems, CAPRE does not aspire to predict a customer’s behavior but to influence that behavior. By measuring the actionability and profitability of customers, we have the ability to focus on customers that can afford to spend larger sums of money in the target business. CAPRE aggregates rules to extract characteristic purchasing behaviors, and then analyzes the counter-examples to detect the most actionable and profitable customers. We measure the effectiveness of CAPRE by performing a cross-validation on the *MovieLens* benchmark. The methodology is applied to over 10,000 individual customers and 100,000 products for the customer relationship management of *VM Matériaux* company, thus assisting the salespersons’ objective to increase the customer value.

Keywords-customer data mining; data mining for marketing; recommender system; actionable knowledge discovery; industrial case-based application;

I. INTRODUCTION

In a world where the number of product and service offers can be overwhelming, Recommender Systems (RS) help customers find and evaluate products or services. RS are software tools and techniques providing suggestions for products to be use by a customer [1] [2]. They connect customers with potential products to purchase by associating either the content of products or the opinions of customers, or both. In recent years, RS have found numerous applications in industry, especially in e-commerce. They help to develop and maintain profitable relationships with customers [3], for example through cross or up-selling.

With most RS, the recommendations depend on the products that are currently bought by the customer [4]. This principle is perfectly appropriate to e-commerce websites, where the customer is alone for driving the product exploration. However, it is hardly adaptable to salespersons’ visits to customers. We propose to adopt a different strategy: when salespersons visit customers, which product(s) can they recommend to trigger purchasing? Contrary to classical RS,

our goal is not to predict customer behaviors but to influence customer behaviors. Making such recommendations requires (i) to identify characteristic purchasing behaviors, and (ii) to apply them on customers while respecting their own behaviors. This last point is crucial because the cost of an inappropriate recommendation is higher in the case of salespersons’ visits than on e-commerce websites [5]. The salespeople may even refuse to use the RS if they consider the recommendations not sufficiently relevant [6].

In this paper, we propose CAPRE (*Customer Actionability and Profitability Recommendation*), a new methodology for recommender systems based on the analysis of turnover for customers of specific products. More precisely, CAPRE aggregates rules to extract characteristic purchasing behaviors, and then analyzes the counter-examples to detect the most actionable and profitable customers. Recommendations are made by targeting the *actionable* counter-examples with the most *profitable* rules. Actually, we do not use rules as naive implications of purchasing behaviors. We use the rules as a model to detect counter-examples, i.e. customers with deviant purchasing behaviors that have the potential to spend larger sums of money in the target business. Our methodology is applied over 10,000 customers and 100,000 products for the customer relationship management of *VM Matériaux*.

A. Contributions of the Paper

The main contributions of this paper can be resumed as follows:

- We develop a new *methodology* for recommending products from a turnover database.
- We propose a *measure of actionability* of the recommendations, based on the similarity between the examples and counter-examples.
- We offer an original *measure of the profitability* of the recommendations, based on criteria defined by experts.
- We apply CAPRE on an actual *VM Matériaux* trading group database composed of more than 10,000 customers and 100,000 products.
- We measure the effectiveness of our system on the *MovieLens* benchmark with a cross-validation.

B. Organization

The rest of the paper is organized as follows: section II presents a state of the art in the recommender systems and rule mining fields. Then, section III discusses our methodology CAPRE with the two measures of actionability and profitability. We apply CAPRE in order to make recommendations on an actual dataset of *VM Matériaux* trading group¹ in section IV. We analyze the effectiveness of CAPRE in section V by performing a cross-validation on the *MovieLens* benchmark. The final section concludes and describes our future work.

II. RELATED WORK

The growth of Internet and the emergence of e-commerce has led the development of RS [7]. These allow companies to filter information, and recommend products to customers according to their preferences. Recommend products and services can enhance the relationship between customer and seller, and thus increase profits [8]. Actually recommendation is not a new concept, it is already used on the Web through movies, books and music. RS have become an important research area since the appearance of the first papers on collaborative filtering [7]. RS are doubly useful: on one hand, they help customers filter through enormous numbers of products, and focus on the few ones that match their preferences; on the other hand, recommender systems help companies increase their sales on the long tail.

Many techniques have been developed, and several classifications have been proposed [2]. Montaner et al. [9] provide a taxonomy and classify existing RS applications to specific application domains: *entertainment* (movies, music and IPTV), *content* (personalized newspapers, Web pages, e-learning applications and e-mail filters), *e-commerce* (products to purchase such as books, cameras, PCs, etc.) and *services* (travel services, experts for consultation, houses to rent, matchmaking services). Examples of such applications include recommending books at Amazon.com, movies by *MovieLens* or music by Yahoo! Music. Schafer et al. [10] present a detailed classification of recommender systems in e-commerce, and elucidate how they can be used to provide personalized service customer loyalty. RS are usually classified into the following categories:

- *Content-based recommendations*: the customer will be recommended products similar to the ones the customers preferred in the past;
- *Collaborative recommendations*: the customer will be recommended products that customers with similar preferences and behaviors liked in the past;
- *Hybrid approaches*: this approach combine collaborative and content-based methods.

In the following subsections, we will focus on collaborative filtering and rule mining for RS.

¹<http://www.vm-materiaux.com>

A. Recommender Systems Based on Collaborative Filtering

Collaborative recommender systems [7] produce customer specific recommendations of products based on patterns of ratings or purchases without need for exogenous information about either products or customers. In other words, the main idea is to exploit information about the past behavior of an existing customer for predicting which products the current customer will most probably like or be interested in. These types of systems are in widespread industrial use today, in particular as a tool in online retail sites.

There have been many collaborative systems developed in the academia and the industry. It can be argued that the Grundy system [11] was the first recommender system, which proposed using stereotypes as a mechanism for building models. Later on, the Tapestry system relied on each customer to identify like-minded customers manually [12]. GroupLens, Video Recommender and Ringo use collaborative filtering algorithms to automate prediction. In the recommendation field, the difficulty to collect descriptions customers and content is a recurrent problem.

B. Rule Mining for Recommender Systems

In the context of RS, a transaction could be viewed as the set of all the previous ratings or purchases of a customer. Association rules can be used to develop *Top-N* recommendations for each customer depending on past purchases. Frequently, products are sorted by the confidence of the rules [13], using multi-level rules [14] or depending on the antecedent length of rules [15].

Recently, rule mining has become a current application of RS. Leung et al. [16] present a model based on fuzzy rules by exploiting similarities in product taxonomies. Liu et al. [17] propose to extract customer purchasing behaviors using sequential rules. Finally, Smyth et al. [18] present two different case studies using association rules for RS.

III. CAPRE METHODOLOGY

In this section, we propose CAPRE, a new methodology to apply rules to recommendation systems. A real-world application is presented in section IV. CAPRE is divided into five steps (see figure 1).




Figure 1. The Five Key Steps of CAPRE

A. Purchasing Behavior Rule Mining

1) *Data*: Let C be the set of all customers $\{c_1, c_2, \dots, c_n\}$ and P be the set of all products $\{p_1, p_2, \dots, p_m\}$ that can be recommended. Let u be a *utility function* that measures the usefulness of product p for customer c :

$$\begin{cases} u : P \times C \rightarrow \mathbf{R} \\ p, c \rightarrow u(p, c) \end{cases} \quad (1)$$

where \mathbf{R} is a part of \mathbb{R} , e.g. positive integers or numbers. In our methodology, u concerns a target business and has a financial meaning. For instance, it can be the turnover of product p for customer c . The greater $u(p, c)$, the more the purchase of p by c is profitable for the company. The function u depends on customer transaction history and can be materialized by a matrix customers \times products. This context is directly applicable in the usual framework of rating datasets (see section V).

Then, we apply two transformations (see figure 2) on the values of the utility function, to make the customers commensurable and to extract rule behaviors on discretized data.




Figure 2. Normalization and discretization of the utility function

- The first step is to make the turnover of a customer c comparable with another, regardless of their financial potential. For each customer c , the sales are transformed as a percentage of sales, representing the usefulness ratio of a product p to a customer c .

$$\begin{cases} u^{\%} : P \times C \rightarrow [0; 100\%] \\ p, c \rightarrow u^{\%}(p, c) \end{cases} \quad (2)$$

with, $u^{\%}(p, c) = \frac{u(p, c)}{\sum_{p \in P} u(p, c)}$

- The second step is to normalize the proportions $u^{\%}(p, c)$, that we denote $u^*(p, c)$, by taking into account their distribution for each product. This normalization results in a discretization of proportions in several intervals $\mu_1 < \mu_2 < \dots < \mu_t$.

$$\begin{cases} u^* : P \times C \rightarrow U = \{\mu_1, \mu_2, \dots, \mu_t\} \\ p, c \rightarrow u^*(p, c) \end{cases} \quad (3)$$

The normalization and discretization depend on the application. Each customer $c \in C$ is described by various characteristic variables (numerical or categorical), such as age, income, etc.

For example, the utility function u is the turnover of customers for product families such as plaster, concrete, etc. $u^*(p, c) = \mu_1$ (respectively μ_2 and μ_3), means that the turnover of customer c for product p is low (respectively medium and large). Each building contractor customer c is described by characteristic variables composed of the number of employees, the professional category, etc.

2) *Mining Cohorts of Rules (step 1 on figure 1)*: We apply an association rule mining algorithm [19] to extract purchasing behavior rules of the form:

$$u^*(p_X, \cdot) = \mu_X, u^*(p_Y, \cdot) = \mu_Y, \dots \rightarrow u^*(p_Z, \cdot) = \mu_Z \quad (4)$$

where $\forall_i p_i \in P$ and $\forall_j \mu_j \in U$. The examples of the rule are the customers who satisfy the antecedent and the consequent (corresponding to the dots "·" in (4)). By contrast, counter-example customers satisfy the antecedent and not the consequent². Such a rule means that the customers who satisfy the antecedent tend to satisfy the consequent. In our methodology, as our goal is to make recommendations, we use rules with only one product in the consequent. For simplicity, we will denote $u^*(p_X, c) = \mu_X$ by $p_X = \mu_X$ (where $p_X \in P$ and $\mu_X \in U$). So, we denote the rules in (4) as follows:

$$p_X = \mu_X, p_Y = \mu_Y \dots \rightarrow p_Z = \mu_Z \quad (5)$$

In the classical association rule terminology, $p_X = \mu_X$ is called an item. The set of rules is filtered by support and confidence thresholds. The resulting ruleset is denoted R . Given an item $p_X = \mu_X$, a cohort is the set of all the rules in R concluding on this item:

$$Ct(p_X = \mu_X) = \{r \in R \mid consequent(r) = (p_X = \mu_X)\} \quad (6)$$

R is partitioned into as many cohorts as consequents of rules. For the cohorts to be robust, we retain only the cohorts with a sufficient number of rules. The sets of examples and counter-examples of the cohort, noted $Ct^+(p_X = \mu_X)$ and $Ct^-(p_X = \mu_X)$ are the union of the sets of examples and counter-examples of the rules respectively. Detecting counter-examples provides us with a precise group of customers who can theoretically allow the company to run at a healthy profit.

The interest of the cohorts is to merge all the rules which provide the same recommendation. By exploiting rule redundancy, we are able to extract more robust purchasing behavior models (with more examples). Given a cohort Ct , a Ct-Recommendation is to propose the consequent of Ct to a counter-example.

For example, let us consider the cohort $Ct(Gravel = c)$ composed of two rules:

- $Cement = c$ AND $Sand = c \rightarrow Gravel = c$
has 190 examples and 70 counter-examples.
- $Sand = c$ AND $Concrete_Mixer = c \rightarrow Gravel = c$
has 135 examples and 40 counter-examples.

The Ct-recommendation consists in proposing some gravel to customers who are counter-examples of at least one of the rules.

²In CAPRE, the eligible counter-examples are only the customers who purchase p_Z at a level lower than μ_Z .

B. Actionability of Counter-Examples

Let us now consider a cohort. With their partially unexploited purchase potential, the counter-examples can theoretically help the company to develop its turnover. However, all the counter-examples cannot be equally receptive to our recommendations. This is why our methodology focuses on the most actionable counter-examples, i.e. the counter-examples that are similar to examples according to the customer characteristic variables (age, gender, income, etc.). The more a counter-example is similar to the examples, the more it is likely to behave like an example, i.e. to purchase more of the consequent of the cohort.

1) *Pre-actionability on Purchasing Variables (step 2 on figure 1)*: To be close to the examples Ct^+ of a cohort, a counter-example e^- should not display an extreme purchasing behavior on the antecedents of the rules of the cohort Ct . Given the purchases of example distribution on the antecedents of the cohort, we define the median M and $(Q_3 - Q_1)$ the interquartile range of the distribution. A counter-example $e^- \in Ct^-$ is pre-actionable if and only if:

$$M - (Q_3 - Q_1) \leq \sum_{\substack{p \in \text{antecedents} \\ \text{of the cohort } Ct}} u\%(p, e^-) \leq M + (Q_3 - Q_1) \quad (7)$$

For example, given the cohort $Ct(\text{Gravel} = c)$, the purchases of examples on the antecedents of the cohort are characterized by $M = 40\%$ and $Q_3 - Q_1 = 20\%$. To be pre-actionable, a counter-example $e^- \in Ct^-$ must have spent between 20% and 60% of its turnover on cement, sand and concrete mixer.

2) *Actionability on Characteristic Variables (step 3 on figure 1)*: In the multidimensional space of the customer characteristics, we need a distance measure to assess similarities between customers. As customer characteristics are numerical and categorical variables, we consider the framework of Factor Analysis of Mixed Data (FAMD) introduced by Escofier in 1979 (see [20] and [21] for further details). This framework provides the distance measure for mixed data that we use in CAPRE. The customer characteristics consist of numerical and categorical variables:

- Let N be the set of centered and reduced numerical variables.
- By applying a complete disjunctive coding on categorical variables, we generate a set I of indicators³ (a binary variable by modality).

The FAMD distance between customers c_1 and c_2 is defined by:

$$d^2(c_1, c_2) = \sum_{i=1}^v (N_{i,c_1} - N_{i,c_2})^2 + \sum_{j=1}^w (I_{j,c_1} - I_{j,c_2})^2 \quad (8)$$

³The values of these indicators are divided by the square root of frequency of each indicator.

For example, roofers and tilers purchase many common products. They can be close by their purchases (see section III-B1) but remote by their characteristic variables. For the rule $\text{Cement_Bag} \rightarrow \text{Tile}$, most examples are tilers. However, roofers tend to regularly purchase 25kg bags of cement to sink a fork or install a ridge roof. Even if the counter-examples are tilers and roofers, you have to recommend purchasing tiles only to tilers. Computing actionability allows to avoid recommend purchasing tiles to roofers.

To sum-up, a counter-example e^- is actionable if it is close enough to the examples. Given two thresholds minDist and δ , e^- is considered as δ -actionable for a cohort Ct if and only if:

- 1) e^- is pre-actionable (see section III-B1);
- 2) e^- respects the following inequality:

$$\frac{|\{e^+ \in Ct^+ \mid d(e^-, e^+) \leq \text{minDist}\}|}{|Ct^+|} \geq \delta \quad (9)$$

with minDist a threshold defining a neighborhood around the counter-example e^- , and δ a threshold defining a minimum ratio of neighbor examples to be actionable.

Given the distribution of the examples in the space of customer characteristics, the distance from a counter-example to the barycenter of the examples is not necessarily representative. That is why we do not use the barycenter in our measure.

C. Profitability of the Cohorts (step 4 on figure 1)

The profitability of a counter-example corresponds to the value in Dollars of the Ct-Recommendation. For example, for each actionable counter-example (see section III-B), we compute the total turnover in Dollars that should have been spent if the counter-example had behaved as an example of the cohort. We denote by $u(\cdot, c)$ the sum $\sum_{p \in P} u(p, c)$. Given a threshold $\theta \in \mathbb{R}^+$, a counter-example e^- is defined as θ -profitable for the product p_Z , i.e. for the recommendation $Ct(p_Z)$ if and only if:

$$\text{Profit}(Ct(p_Z), e^-) = (u(\cdot, e^-) * \alpha) - u(p_Z, e^-) \geq \theta$$

with the mean purchase ratio $\alpha = \frac{\sum_{c \in Ct^+} u(p_Z, c)}{\sum_{c \in Ct^+} u(\cdot, c)}$ (10)

Therefore, the profitability of a cohort of rules corresponds to the sum of the profitabilities of the θ -profitable counter-examples of the cohort. From this value, we deduce fixed and variable costs incurred by a business approach. Moreover, we suggest using a *scoring* to refine the profitability expected for each customer. The learning phase is guided by the examples of the cohort and the application phase by the profitable counter-examples.

D. Triggering the Most Actionable and Profitable Cohorts (step 5 on figure 1)

1) *Marketing Recommendation Campaigns*: Several scenarios exist for decision makers to profitably action the recommender system:

- Sorting cohorts in order to identify the most profitable Ct-Recommendations. Therefore, commercial actions can be triggered to promote products.
- For each customer, sorting the cohorts in which the customer is a θ -profitable counter-example. Salespersons can trigger to one or more recommendations, creating a personalized relationship with the target customer.

2) *Return On Investment*: The methodology presented above allows us to propose an equation to calculate the expected ROI before triggering the most actionable and profitable cohorts, assuming that for each θ -profitable counter-example e^- , the profitability $Profit(Ct(p), e^-)$ is computed for the most profitable cohorts.

$$ROI = \sum_{\substack{c \in C \\ p \in P}} Profit(Ct(p), c) - \omega \quad \text{with } \omega = \alpha + \beta - \gamma \quad (11)$$

- α : operation fixed costs (salespersons' visits, communication, gifts, meals, advertising, etc.).
- β : data mining costs: time for preprocessing, modeling and reporting.
- γ : cost reduction due to time spent gain in the marketing department to establish recommendations.

IV. APPLICATION ON REAL DATA

A. Context

The *VM Matériaux* trading group offers more than 100,000 product references for professional building contractors whose have different characteristic variables and often atypical purchasing behaviors. Therefore, the use of our methodology is appropriate to profitably recommend products for target customers, creating a personalized relationship and assisting the salespersons' objective to increase the customer value.

B. Purchasing Behavior Rule Mining

1) *Data*: To address the problem of *over-specialization*, we use the product taxonomy of *VM Matériaux*. The advantage of this approach is to obtain relevant rules, in which the most common products do not hide the less common products by their frequency. For example, the rarest products, such as photovoltaic panels, will be codified in a more general level, while the most common products, such as plaster boards, will be codified at a finer level. We decided to collect data from our datawarehouse composed of approximately 91 tera-bytes and 180 tables. 9,575 customers having achieved a turnover threshold (defined by decision makers) are selected and described through variables divided into three groups:

- 20 internal variables (in the datawarehouse): address, loyalty, main salesperson, authorized outstanding, etc.
- 9 external variables (obtained from a Coface⁴ file): number of employees, professional category, etc.
- 414 aggregates (computed from levels of the taxonomy): sum of turnover for each product family (cement, insulation, etc.) within the year before a datetime.

This data preprocessing produces a model with 9,575 observations and 443 variables. In agreement with business requirements, the purchases are made commensurable and are discretized (see section III-A1) into three intervals a , b and c of the same frequency $\frac{1}{3}$, to segment the small, medium and large customer purchases, respectively.

2) *Mining Cohorts of Rules*: The extraction of association rules is performed using the algorithm CHARM based on in-depth research of closed itemsets [22] with support threshold of 0.005 (i.e. about 50 customers) and a confidence threshold of 75 %. A set of 7,788 rules is extracted (see table I) and partitioned into 57 cohorts of rules.

Table I
SOME RULES EXTRACTED BY THE CHARM ALGORITHM

Sup.(%)	Conf.(%)	Rules
0.006	78.13	Plaster = b & Metal = c \rightarrow Plaster_Board = c
0.008	78.22	Tile = a & Parpen = c \rightarrow Cement = c
0.011	76.81	Adhesive_Coating = c \rightarrow Plaster = c

To illustrate CAPRE, we focus on the cohort to recommend the purchase of plaster $Ct(Plaster = c)$. The cohort is composed of 12 rules, 150 distinct examples and 50 distinct counter-examples.

C. Actionability of Counter-Examples

To select the actionable counter-examples of $Ct(Plaster = c)$, we compute the parameter $minDist = 28.43$ which is the average distance between all customers (see section III-B2) and we define with the commercial direction (trading director, purchasing director and sales force team) the threshold $\delta = 10$ %. Therefore, a counter-example will be actionable if and only if:

- It is pre-actionable, i.e. its purchasing behavior for the antecedents of the rules of the cohort are in the interquartile range defined in section III-B1.
- It has 10 % of the examples Ct^+ of the cohort in its neighborhood at a maximal distance of 28.43.

To apply these two steps, let us visualize on figure 3 the distributions of examples Ct^+ and counter-examples Ct^- on purchases from the antecedents of the cohort.

Antecedent purchases of Ct^+ are characterized by a median $M = 37.32$ % and an interquartile range $Q3 - Q1 = 28.05$ %. Therefore, the counter-examples outside the interval [9.27 %; 65.37 %] are considered *non pre-actionable*. This filtering allows to prune 4 counter-examples.

⁴<http://www.coface.com>




Figure 3. Distributions of Examples and Counter-Examples on Purchases from the Antecedents of a rule of the Cohort $Ct(Plaster = c)$

Secondly, to be *actionable*, the 46 counter-examples must have 10 % of Ct^+ in their neighborhood at a distance less than or equal to 28.43 (see section III-B2). This second filter allows to prune 4 counter-examples. In the end, we select 42 actionable counter-example customers.

D. Profitability of the Cohorts

With the help of decision makers of *VM Matériaux*, we considered that the solicitation of a customer by a salesman were approximately \$2 000. That is why, to run at a healthy profit, decision makers decide to set the threshold $\theta = \$2,000$ from which a counter-example is profitable for a cohort. When applying (10) of our methodology, 38 actionable counter-examples are \$2000-profitable for the purchase of plaster.




Figure 4. Summary of Counter-Example Filtering for $Ct(Plaster = c)$

For the 38 profitable counter-examples, the expected profit for the cohort is estimated at about \$420,000 of additional turnover, i.e. over \$11,000 per customer. To refine the profitability expected for each customer, we decide to use a scoring with the data mining software program KXEN. To predict the continuous target of turnover per customer for the consequent $plaster = c$, we choose to calculate a tendency score with a ridge regression. The learning phase is guided by the 150 examples of the cohort and the application phase by the 38 \$2000-profitable counter-examples for the purchase of plaster. The learning dataset consists of a few customers. However, they represent a very homogeneous population, i.e. with a strong purchasing potential on the consequent of the cohort. Therefore, the accuracy of the model rises. For the 38 profitable counter-examples, the expected profit for the cohort is estimated at about \$264,373 of additional turnover, i.e. over \$7,000 per customer.

E. Triggering the Most Actionable and Profitable Cohorts

Our methodology has enabled the sales team to analyze the behavior of their customers in order to understand why

some of them do not purchase certain products (price?, competition?, etc.). Therefore, other marketing measures will be implemented:

- 1) To promote products on targeted customers (free samples) by selecting the most profitable cohorts.
- 2) To target customers in a region or a store and recommend some products based on cohorts in which these customers are the most profitable.

Finally, based on profitable recommendations, supplier partnerships may be agreed on to highlight products.

V. EXPERIMENTAL VALIDATION WITH *MovieLens*

In this section, we present the experimental results on the movie-rating benchmark. *MovieLens* benchmark was collected by GroupLens Research Project at the University of Minnesota and contains 100,000 movie-rating data from 943 users on 1,682 movies. Each user has rated at least 20 movies from 1 (strongly unsatisfactory) to 5 (strongly satisfactory). The dataset was converted into a binary user-movie matrix $C \times P$ that has 943 customers and 1,682 movies that were rated by at least one of the users. We apply our methodology on the five usual learning datasets named $U[1-5].base$ and the five usual application datasets named $U[1-5].test$. Training and validation datasets contain 80 % and 20 % of all ratings respectively. Customer characteristic variables, i.e. $N \cup I$, are the age, the gender and the occupation. We gathered the ratings into three classes $a = [1, 2]$, $b = [3]$ and $c = [4, 5]$.

A. Cross-Validation and Comparison

With a support threshold of 0.05 and a confidence threshold of 50 %, we extract rules and generate the cohorts summarized in table II. For robustness considerations, only the cohorts with at least 20 rules are retained. To evaluate the *Top-N* recommendations, we use the two classical metrics widely used in the recommender system community namely, *precision* and *recall*. We privilege the precision measure to minimize the number of false positive recommendations. We present in table III the *Top-10* and *Bottom-10* recommendations sorted by average precision \overline{Pr} .

Table II
CAPRE EXPERIMENT ON *MovieLens* BENCHMARK

	U1.base	U2.base	U3.base	U4.base	U5.base
# of rules	26,658	12,487	12,456	9,516	13,028
# of cohorts	99	95	92	88	92

We notice the good results reached with CAPRE on average precision \overline{Pr} on the *Top-10* recommendations (even the *Bottom-10* results are acceptable). The errors measured by precision and recall are similar to those generally measured on *MovieLens* with RMSE (Root Mean Square Error). Aggregating rules into cohorts allows to build a more robust model than single rules, this reduces the errors of recommendation.

Table III
PRECISION (PR) AND RECALL (RC) OF RECOMMENDATIONS

<i>Top-10</i>	Average	
	Pr	Rc
Rear Window	91,50	78,77
The Shawshank Redemption	91,07	87,67
Casablanca	90,02	88,16
Schindlers List	89,97	90,97
To Kill a Mockingbird	89,03	78,13
One Flew Over Cuckoos Nest	88,28	90,39
The Silence of the Lambs	87,98	89,90
The Usual Suspects	86,76	87,32
Star Wars	86,58	89,70
The Godfather	85,01	86,32
<i>Bottom-10</i>		
The Terminator	72,87	90,57
2001: A Space Odyssey	72,32	85,15
Toy Story	71,48	87,61
Die Hard	71,36	78,05
Contact	70,85	66,68
Indiana Jones	69,81	89,40
Forrest Gump	69,73	89,14
Twelve Monkeys	68,04	84,35
Back to the Future	67,83	92,48
Jaws	67,29	84,40

B. Example and Discussion

Let us consider the rule $Men\ in\ Black = c \rightarrow Independence_Day = c$, 102 customers who watched the action movie *Men_in_Black* and voted positively (4 or 5) have also watched the action movie *Independence_Day* by voting positively. However, 43 counter-examples have watched and voted positively *Men_in_Black* but did not appreciate *Independence_Day* (vote 1 or 2 or 3). CAPRE allows to obtain 99 cohorts of rules. Let us focus on the cohort recommending the movie *Independence_Day* (see table IV). The cohort is composed of four rules.

Table IV
A COHORT OF FOUR RULES RECOMMENDING *Independence_Day*

Sup.(%)	Conf.(%)	Rules
0.13	54.12	The_Rock = c \rightarrow Ind._Day = c
0.12	52.28	Star_Trek = c \rightarrow Ind._Day = c
0.11	54.12	Men_in_Black = c \rightarrow Ind._Day = c
0.10	56.54	Star_Wars = c & The_Rock \rightarrow Ind._Day = c

The pre-actionability, actionability and profitability steps allow to prune the recommendations and target users:

- *Pre-actionability*: from 188 to 140 pre-actionable counter-examples. For instance, the user #236 is non pre-actionable. Indeed he enjoyed the movie *Star_Trek = c* but voted negatively for movies *The_Rock*, *Men_in_Black* and *Star_Wars*.
- *Actionability*: from 140 to 93 actionable counter-examples. We note that examples are mostly men between 20 and 45 years old of a particular occupation (student or engineer). Therefore, we prune many

counter-examples, such as 25-year old artist women.

- *Profitability*: we could easily assign a monetary value on each movie, considering that the recommendation of a movie like *Independence_Day* yields \$5 per streaming viewing. Therefore, the expected profitability of the cohort could be \$465.

We can note that CAPRE methodology allows (i) to prune sharply the number of recommendations and (ii) to filter efficiently the number of potential customers. Moreover, we observe customer rating behaviors throughout other cohorts:

- Some cohorts have scattered examples because their consequents are a very popular movie, such as *Star_Wars*. The behavior of the examples is difficult to characterize through ratings or characteristic variables. That is why, we little prune the counter-examples.
- However, for horror movies, examples show much more typical behaviors. Counter-example customers who view and appreciate at least one horror movie of the cohort are generally actionable and profitable because of a more restricted customer rating behavior. That is why, we widely prune the counter-examples.

Finally, in [23], Onuma et al. showed that for *Horror Movie Fan*, their system recommends *A_Nightmare_on_Elm_Street*, *The_Shining* and *Jaws*. However, in our methodology, we can note that for the cohort recommending *Jaws*, composed of 7 rules, we prune the number of counter-examples from 144 to 75 actionable customers. By this way, we target more finely the recommendation to the most receptive customers.

VI. CONCLUSIONS AND FUTURE WORKS

In this paper we propose CAPRE, a new methodology for recommender systems based on the analysis of turnover of customers for specific products. By measuring the actionability of rules based on the similarity between the examples and counter-examples, we present an original measure of the profitability of rules. Thus, the methodology allows not only to make recommendations, but also to identify and sort the most receptive customers for these recommendations, and estimate the related profit. CAPRE is applied to over 10,000 individual customers and 100,000 products in the framework of the customer relationship management of *VM Matériaux* company, thus assisting the salespersons' objective to increase the customer value. The most actionable and profitable recommendations are triggered and an expected ROI is computed, with 7,788 rules and 57 cohorts, representing about \$264,373 of additional turnover for plaster recommendation. Thus, recommender systems applied to the analysis of the turnover of customers would improve customer relationships and generate qualitative and quantitative profits for companies. The cross-validation performed on the *MovieLens* benchmark shows that CAPRE makes good recommendations with little errors (even on the *Bottom-10* results).

We are currently using CAPRE in a real business application. The goal is to use extensively CAPRE for any databases of the type customers \times products. This article offers the prospect for further work. First, the measure of profitability can be improved: the actionability of cohorts could be modulated by the number of actionable counter-examples and the net margin of recommended products. Secondly, an application of recommendations can be integrated in the corporate information system for decision makers. Finally, an OLAP-based approach can be adapted to address decision makers by trying to provide aggregated recommendations (e.g. for brands or product families) to certain *segments* of customers.

ACKNOWLEDGMENT

The authors would like to thank Pierrick Richard and Gaëtan Blain (*VM Matériaux*) for supporting this work.

REFERENCES

- [1] T. Mahmood and F. Ricci, "Improving Recommender Systems with Adaptive Conversational Strategies," in *Proceedings of the 20th ACM conference on Hypertext and hypermedia*, ser. HT '09. New York, NY, USA: ACM, 2009, pp. 73–82.
- [2] P. Resnick and H. R. Varian, "Recommender Systems," *Communications of the ACM*, vol. 40, no. 3, pp. 56–58, 1997.
- [3] M. Ren, Z. Chen, C. Liu, and G. Chen, "An Evolving Information System Based on Data Mining Knowledge to Support Customer Relationship Management," in *Advanced Management of Information for Globalized Enterprises*, September 2008, pp. 1–5.
- [4] G. Adomavicius and A. Tuzhilin, "Toward the Next Generation of Recommender Systems: A survey of the State-of-the-Art and Possible Extensions," *IEEE Transactions On Knowledge and Data Engineering*, vol. 17, no. 6, pp. 734–749, 2005.
- [5] G. Linden, B. Smith, and J. York, "Amazon.com Recommendations: Item-to-Item Collaborative Filtering," *IEEE Internet Computing*, vol. 7, pp. 76–80, January 2003.
- [6] D. Cosley, S. K. Lam, I. Albert, J. A. Konstan, and J. Riedl, "Is Seeing Believing?: How Recommender System Interfaces Affect Users' Opinions," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, ser. CHI '03. New York, NY, USA: ACM, 2003, pp. 585–592.
- [7] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl, "GroupLens: an Open Architecture for Collaborative Filtering of Netnews," in *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, ser. CSCW '94. New York, NY, USA: ACM, 1994, pp. 175–186.
- [8] Y. Zhang and J. R. Jiao, "An Associative Classification-Based Recommendation System for Personalization in B2C e-commerce Applications," *Expert Syst. Appl.*, vol. 33, no. 2, pp. 357–367, 2007.
- [9] M. Montaner, B. López, and J. L. De La Rosa, "A Taxonomy of Recommender Agents on the Internet," *Artif. Intell. Rev.*, vol. 19, pp. 285–330, June 2003.
- [10] J. B. Schafer, A. J. Konstan, and J. Riedl, "E-Commerce Recommendation Applications," *Data Mining and Knowledge Discovery*, vol. 5, no. 1 - 2, pp. 115–153, January 2001.
- [11] E. Rich, "Readings in Intelligent User Interfaces," M. T. Maybury and W. Wahlster, Eds. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1998, ch. User Modeling via Stereotypes, pp. 329–342.
- [12] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry, "Using Collaborative Filtering to Weave an Information Tapestry," *Commun. ACM*, vol. 35, no. 12, pp. 61–70, 1992.
- [13] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Analysis of Recommendation Algorithms for E-Commerce," in *EC '00: Proceedings of the 2nd ACM conference on Electronic commerce*. New York, NY, USA: ACM, 2000, pp. 158–167.
- [14] C. Kim and J. Kim, "A Recommendation Algorithm Using Multi-Level Association Rules," in *Proceedings of the 2003 IEEE/WIC Int. Conf. on Web Intelligence*, ser. WI '03. Washington, DC, USA: IEEE Computer Society, 2003, p. 524.
- [15] M. Nakagawa and B. Mobasher, "Impact of Site Characteristics on Recommendation Models Based On Association Rules and Sequential Patterns," in *Intelligent Techniques for Web Personalization*, 2003.
- [16] C. W.-k. Leung, S. C.-f. Chan, and F.-l. Chung, "A Collaborative Filtering Framework Based on Fuzzy Association Rules and Multiple-Level Similarity," *Knowledge and Inf. Syst.*, vol. 10, pp. 357–381, 2006.
- [17] D.-R. Liu, C.-H. Lai, and W.-J. Lee, "A Hybrid of Sequential Rules and Collaborative Filtering for Product Recommendation," *Inf. Sciences*, vol. 179, no. 20, pp. 3505–3519, September 2009.
- [18] B. Smyth, K. McCarthy, J. Reilly, D. O'Sullivan, L. McGinty, and D. C. Wilson, "Case Studies in Association Rule Mining for Recommender Systems," in *IC-AI*, 2005, pp. 809–815.
- [19] J. Han, H. Cheng, D. Xin, and X. Yan, "Frequent Pattern Mining: Current Status and Future Directions," *Data Mining and Knowledge Discovery*, vol. 14, no. 1, 2007.
- [20] M. Bécue-Bertaut and J. Pagès, "Multiple Factor Analysis and Clustering of a Mixture of Quantitative, Categorical and Frequency Data," *Computational Statistics & Data Analysis*, vol. 52, no. 6, pp. 3255–3268, 2008.
- [21] G. Saporta, "Simultaneous Analysis of Qualitative and Quantitative Data," in *Proceedings of the 35th Scientific Meeting of the Italian Statistical Society*, 1990, pp. 63–72.
- [22] M. J. Zaki and C.-J. Hsiao, "CHARM: An Efficient Algorithm for Closed Itemset Mining," in *Proceedings of the Second SIAM DM*. Arlington, VA: SIAM, 2002.
- [23] K. Onuma, H. Tong, and C. Faloutsos, "TANGENT: a Novel, "Surprise me", Recommendation Algorithm," in *KDD '09: Proc. of the 15th ACM SIGKDD Int. Conf. on KDD*. New York, NY, USA: ACM, 2009, pp. 657–666.