Recommending Insurance Riders

Lior Rokach and Guy Shani and Bracha Shapira and Eyal Chapnik and Gali Siboni
Information Systems Engineering
Ben Gurion University
Israel
{liorrk/shanigu/bshapira/eyalch/siboniga}@bgu.ac.il

ABSTRACT

Insurance riders are optional addendum to base insurance policies. In this paper we discuss the application of recommender systems to the task of matching riders to clients. This task is difficult because of the variety of possible riders, as well as the poor knowledge of the client over these riders. We focus on call centers where the agent also has limited knowledge and expertise. For such agents, discovering appropriate riders for the current client is very difficult, and automated tools that suggest such riders can play an important role in the agent-client dialogue, and may influence considerably the outcome of the interaction.

This paper presents and discusses in detail the problem of recommending insurance riders to clients in call centers, comparing it to other, classic, recommendation system applications. In addition, we present an analysis of customer purchase behavior, showing that simple item-item recommendation algorithms provide good recommendations for riders given a base policy.

1. INTRODUCTION

In the insurance industry, insurance policies are typically constructed of a base policy and a collection of possible addendum, or *riders*. For example, in health insurance, a base policy may be an extended health insurance, containing coverage for some procedures (e.g. surgeries) and medicine not covered under the basic mandatory policy[1]. Examples of possible riders for that base policy are coverage for surgeries and transplants out of the country, coverage for special medical consultation, or for constant monitoring for heart problems. These riders can be added to the base package to expand its coverage.

Traditionally, insurance sales are mainly done by professional independent agents that meet and discuss the optional policies in person with potential clients. There are several problems with such a process; First, in order for a third party to profit from the transaction, the insurance price must be raised. As such, the prices of packages sold by professional independent agents are typically higher than packages sold directly by the insurance company. Another problem is that such agents are in many cases motivated not by the welfare of the clients but rather by the commission that they earn. Thus, an agent may sell a suboptimal package because the commission he receives for such a package is higher. If insurance companies could sell insurance products directly to clients, all these problems could be avoided.

Many insurance companies have built large call centers, originally designed to help existing clients with various questions that they may have. Such call centers provide an opportunity for selling insurance products directly to clients. That is, a client that contacts the call center can be offered to expand her current insurance coverage by adding appropriate riders. The staff of the call center, however, typically has a low level of knowledge over the various packages and riders, due to the high cost of employing professional insurance agents. A typical employee of a call center may be a student in a part-time position. The problem, hence, is for the non-professional call center team to make good suggestions to clients who call the call center, thus increasing the likelihood of a transaction.

Recommender systems, that actively suggest items to users based on their profiles, can provide an adequate solution to this problem. Such systems may match appropriate riders to clients based on the clients demographics (e.g. location, income, gender), based on their interaction history, stated preferences (e.g. low risk vs. high risk), or on similar clients, and utilizing relevant domain knowledge, such as the similarity between the content of policies or riders. Insurance companies typically have a database containing relatively accurate demographic information over their clients, allowing us to choose any combination of the above sources for generating a recommendation.

Insurance policy recommendations have several key differences from traditional recommendation tasks, such as recommending movies, news articles, or electronic gadgets. For example, the relatively wide and accurate information over the clients, the complexity of the packages, the lack of information of the clients over the items, and many more. We discuss below differences and similarities between the insurance domain and other traditional tasks in domain size, item complexity, customer expertise, recommendation constraints, and interaction with the system.

The main contributions of this paper is the through description of this new and vastly different domain for rec-

[1] In Israel, for example, there is a mandatory minimal health insurance for everybody.

SAC'13 March 18-22, 2013, Coimbra, Portugal

ommendation systems. We define the tasks, highlight the differences from existing applications, and suggest methods for providing recommendations in insurance call centers. In addition, we have obtained a large dataset of insurance policies, and we report a set of offline experiments on this data. We are currently developing a recommender system for an Israeli finance company, and we intend to conduct a large scale study of the effects of our system on the sales process of the call center.

2. RECOMMENDER SYSTEMS

Recommender systems [19] have become an important tool for guiding users towards interesting items in large and complex item spaces. Perhaps the most widely used recommendation technique is the collaborative filtering approach [2] where the recommendations rely on the behavior of similar users to the current active user. Collaborative-filtering algorithms find behavior patterns within a dataset of user-item interactions, such as purchases or ratings. Such methods usually find users that have purchased the same items and predict future purchases based on those users. In many cases, collaborative-filtering methods detect correlations between items that are not obvious given the items' properties. A collaborative-filtering algorithm typically requires a user-item usage matrix R, where $R_{u,i} = 1$ iff user u has used item i. We denote the set of users with U and the set of items with I.

There are many well-known collaborative-filtering algorithms, including decision trees [2], latent semantics indexing [8], and the recently popular base models, also known as matrix factorization models [11]. Given any such models, the ratings for other items that "similar" users liked (as computed by the model) are aggregated to create a list of recommendations [20, 4]. Algorithms that directly work with the user-item rating matrix are known as *memory based*, while algorithms that construct some model (e.g., a decision tree) are known as *model based*.

One of the simplest ways to provide collaborative-filtering recommendations is the pairwise item-item setting, where given a single item, we recommend other related items [21, 12]. A well-known method to evaluate the relevance of two items i_1 and i_2 is the conditional probability of selecting item i given that the user already selected item j that can be estimated using:

$$\Pr(i|j) = \frac{count(i,j)}{count(j)} \tag{1}$$

where $count(i)$ is the number of users who used item i, and $count(i,j)$ is the number of users who used both item i and item j.

Given an item j that was selected by the active user, we recommend items i with the largest $\Pr(i|j)$ that have not yet been selected by the active user. Even though this item-to-item recommendation scheme is simple, it has been employed successfully in large-scale commercial recommender systems (e.g., [12]).

3. INSURANCE PRODUCTS AND RIDERS

Insurance companies (or simply *Insurers*) protect people or organizations from the risk of a loss, in exchange for a payment. The insurer sells protection packages, often referred to as policies. A *policy* is a written contract effecting insurance, and including all clause, riders and endorsements. Insurance companies usually design several base-policies templates which target different audiences. Each *base-policy* template contains a basic set of mandatory *coverage* which determines the scope of protection provided under the insurance policy. The policy can be extended by attaching an *insurance rider*. A rider or *addendum* is an endorsement to an insurance policy that modifies clauses and provisions of the policy, including or excluding coverage. When a rider is added to the policy it becomes an integral part of it and it is subject to the same general conditions. In most cases the rider extends an insurance policy with more coverage and thus may increase the premium. A *policy-owner* (or *policy-holder*) is the person who owns the insurance policy. This is usually the insured person, but it may also be a relative of the insured, a partnership, or a corporation.

Insurance companies that serve the end-customers (i.e. people) can be involved in two main lines of business: 1. Life and Annuity; 2. Property and Casualty. In this paper we focus on Life and Annuity insurance, where policies are based on mortality or morbidity risk. This area includes various types of policies including life insurance, annuity, disability, health and long-term care.

Insurers market various insurance covers either directly (using call centers or online) or through various distribution channels such as agents. Dumm and Hoyt [6] identify the following main policy distribution channels:

1. Agent-led channels: An *agent* is a licensed insurance company representative who solicits and negotiates contracts of insurance, and provides follow-up service including helping the policyholders making changes in the policy in response to new needs. An agent can be independent, i.e. representing at least two insurance companies, or a direct agent who represents and sells policies for one company only.

2. Bank-led channels - Some banks market policies from various insurers to their customers. In Europe this type is playing a major role in the distribution of insurance while in the United States it does not.

3. Company-led channels - Certain companies market insurance policies through mediums such as direct mail or telephone call centers. The practice in this case does not require a meeting with the insured an thus cannot be used for all policy types. In many cases the company is mostly owned by the insurer, but due to regulations it still operates as a separate business. This type of channel has seen increasing growth in the last two decades.

4. Internet led channels: this type of channel enables customers to compare, customize and buy insurance policies online.

While the sales of insurance policies via the internet or via company-led channels are constantly rising, most of the life and annuity policies are still sold by an agent. Specifically agents yet account for 90% of first-year premiums [17]. Competing channels capture a relatively small share in the first year. Nevertheless company-led channels begin to play a major role in the follow up sales and cross selling which is the focus of this paper.

4. RECOMMENDER SYSTEMS FOR CALL CENTERS

A call center is a communication channel between customers and companies from which customers buy products and services, and is also used by companies for promotion of new products. A call center employs a relatively large staff, enabling a direct human interaction between the customer and the company [1]. These employees have two main tasks; answering incoming phone calls from clients, and issuing outgoing calls to clients. Being an important and popular tool for CRM (Customer Relation Management), many aspects of call centers were studied. For example, a large study conducted at Cornell University (Holman, 2007) examined the differences between management and employment of call centers across countries and looked at questions like "What strategies contribute to better operations, job quality, turnover, and absenteeism?" or the adaptation of new technologies, workforce characteristics and more.

Other studies tried to make use of the information gathered during the interaction between the call center employee and the customer, for improving the satisfaction of customers or the call centers functionality [16]. The analysis of the collected data can be improved by utilizing data-mining techniques. Some studies (e.g., [16]) use performance data such as records filled by customer representatives during and after the call, while other studies analyze the actual conversation text between the customer and the representative [23, 7, 15]). Many studies focused on the text mining aspects of the analysis. Researchers examined, for example, the feasibility of identifying and extracting procedural data or important segments from the conversation [15, 23], and the practicability of classifying the data to pre-defined topics [7]. Most of the research in this area is focused currently on the extraction of important data from these interactions. An obvious next step is the utilization of this information for meeting the call center's goals.

Still, to date, very little research has been conducted on improving the call centers functionality or users satisfaction using the collected data. For example, an obvious idea is to identify cross-sell or up-sell opportunities during the call based on data collected during previous calls [7], or build a predictive model for identifying future marketing targets [13]. However, studies of this type are scarce, and most only suggest ideas without evaluating them on real users or through an on-line study to examine the effects of the analysis results on the call centers perfromance.

The idea of integrating recommender engines into call centers dates back to Driskill and Riedel [5]. The authors suggest using recommendations for inbound calls where agents would receive a list of additional products to suggest to a customer based on her current basket, and for outbound calls where the recommender would suggest a list of customers to call based on previous customer activities. Since then, to the best of our knowledge, there is little public research studying the application of recommender systems to call center for selling items. We are unaware of any previous research examining insurance call centers recommendations.

Debnath [3] describes an algorithm for a help desk at an IT company designed to help people with technical problems. The recommender system suggests to the agent a procedure for resolving the problem described by the current call that would satisfy the caller, based on an analysis of former calls. The system was evaluated offline.

To conclude, call centers are becoming a premier tool for contacting customers directly. Still, most of the employed customer representatives for the insurance companies hold temporary positions [9] and are not knowledgeable with the complexity of the insurance domain. It thus seems necessary to explore options of improving call centers functionality and customers satisfaction by applying recommender systems techniques to the insurance domain. To truly show the value of recommender systems to that domain it is important to conduct a real-world study to examine their effect.

5. RECOMMENDING INSURANCE RIDERS THROUGH CALL CENTERS

Call centers are now common in many businesses, including cell-phone companies, banks, and insurance companies. These call centers provide multiple services, such as answering questions about products, updating customer details, handling the first stage of claim processing, and many more. In some cases, call centers provide the main access point for insurance customers in contacting the insurance companies. As such, an obvious addition to the call center tasks is to suggest more products for customers that already called the company with a different goal in mind.

While this sales channel can certainly provide more revenue to the company, there are several difficulties in selling insurance items through the call center. First, the staff of these call centers are typically non-professionals, as employing professional insurance agents in call-center seems too costly. As certain insurance types, such as pensions, require professional training to sell[2], the call-center staff is legally prohibited to offer them to clients. Thus, the number of possible products that could be sold is rather limited. Furthermore, due to the high transition in call-center staff, most of the staff does not know all the insurance products well enough to be able to properly recommend an appropriate product to a customer.

It seems that the products that fit the most this sales channel is riders, or addendum, to basic insurance packages that were previously purchased by the customer. Such products do not require a license for sales, and are also relatively simple to explain. However, it is still needed to match the appropriate rider to the current customer.

The task of finding appropriate riders for customers with already purchased base insurance plans is a perfect match for recommender systems. One of the most common task of recommender systems is the personalized recommendation of items to users based on the user profile. In our case the items are riders, and the users are customers who call the call-center for some reason. An alternative view identifies the call center staff as the users of the system. This can be true in other call center applications, where the call center employee is presented with a recommendation concerning, e.g., the discussion protocol for identifying a technical problem. In our case, though, we consider the call center employee as the communication channel between the company and the end customer. That is, we expect the call center employee to attempt to sell the riders presented to him by the recommender system to the client, without making any judgment calls.

[2]At least in Israel.

The user profile can be constructed of the following available information:

- The base plans and riders already purchased by the customer.
- The client demographic data; Insurance companies typically maintain a wide and accurate database of customer demographics, including age, gender, marital status, income level, address of residence, and many more. This is in contrast to typical e-commerce recommendation applications where the recommender knows little about the users.
- The reason that the customer called the call center to begin with. For example, customers who called because they feel that their current coverage is insufficient may be more willing to accept certain types of products than clients who called to update their personal data.

The recommender system suggested above is to be used by the call-center staff for deciding which packages to suggest to customers. This is again a slight deviation from the traditional setting in which users directly communicate with the recommender system. One obvious result is that the recommendation system interface becomes less important, as displaying recommendations in an attractive manner has no effect on the success of the recommender system. In fact, one could consider the call-center personal as the user interface of the recommender system.

6. THE UNIQUE PROPERTIES OF RECOMMENDATIONS IN THE INSURANCE DOMAIN

There are several classic applications for recommender systems [22, 14], including movies (e.g., MovieLens[3] and Netflix[4]), books [12, 24], news stories [10], and electronic gadgets. Another domain that is less explored, yet interesting to compare against, is tourism recommendations, such as vacation packages [18]. In this section we focus on the differences between these applications and recommendations for the insurance domain.

6.1 Domain size

As opposed to retail recommendations, vacation packages, and also news stories, the number of possible items to recommend is typically very small. There are perhaps several dozens of possible items for each customer. An important goal of recommender systems for classic domains is to ease the information overload of browsing through numerous possible items. In the insurance domain it is feasible, although unlikely, that a customer will observe all possible items.

6.2 Item complexity

While there are a limited number of possible items (insurance riders), their complexity is typically non-trivial. As such, understanding the items may require a considerable cognitive overload. In the news stories domain reading the title and potentially a few lines of summary can reveal much of the content. In movies, or books, consumers also typically read a short summary before deciding whether to purchase the item.

Electronic gadgets are also somewhat similar to insurance packages in this regard. While some gadgets, such as memory sticks, have only a few properties, other items in this family, such as laptops or desktops have a complex configurations, in many cases with a possibility to replace certain parts of the configuration (e.g. upgrade the memory). Such items, like insurance policies have a relatively high complexity. However, it may be that the expertise required to understand the configuration of electronic gadgets is more commonly found in the general public than the ability to understand insurance policies.

Vacation packages are also complex items, typically constructed of a multitude of optional sub-items, such as flights, hotels, rentals, and so forth. The number of possible sub-items is vast in comparison with the possible insurance riders. These sub-items, however, are typically easier to understand than insurance riders, and can be exchanged in many cases without any real change to the final product (e.g. replacing one 3 star hotel with another, nearby 3 star hotel).

6.3 Customer expertise

In some classic domains people often browse the item space for the pure enjoyment of learning about items. For example, many people want to know about the new movies or books, even if they do not intend to read or watch them. Users in online news websites often read many of the headlines, even if they will never read the full story. To a lesser degree, there are many people who like to know of the recent development in electronic gadgets, even if they do not intend to purchase them.

Insurance products are vastly different in that regard. It is unlikely that people visit the websites of insurance companies just to read about the insurance packages that they offer, unless they are explicitly shopping for a particular insurance. As such, it is unlikely that a typical customer is familiar with the items that the insurance company is offering.

For example, in books or movies or even laptops usually most people are familiar with the most popular items, due to the massive media coverage, and it is in many cases wasteful to provide recommendations for them. In the insurance domain, even insurance products that are bought by most of the customers are unlikely to be known by new customers, and recommendations for such items can have an important value.

In vacation packages judging the quality of a flight, an hotel, or a rented car, is typically easy. Users can easily decide whether they prefer morning flights to evening flights, or whether an hotel closer to attractions is preferred to a higher quality hotel further from the attractions. Making similar judgment in the insurance domain requires an expertise that most people do not possess.

6.4 Constraints

In the insurance domain many constraints may apply when matching an insurance package to a customer. For example, certain health plans are not available for people under or over a certain age. Other insurance items may be offered only to people who already have other insurance items. It is also possible that certain items have a significant overlap and should not be purchased together. These constraints

[3] www.movielens.org
[4] www.Netflix.com

	Movies	Books	News	Gadgets	Vacation	Insurance
Domain size	Medium	Medium	Large	Medium	Medium	Small
Item complexity	Low	Low	Low	Medium	High	High
Customer expertise	High	High	High	Medium	Medium	Low
Constraints	None	None	None	Low	High	High
Interactions	Direct	Direct	Direct	Direct/Indirect	Direct/Indirect	Indirect
Attention span	High	High	Very high	Medium	Medium	Low

Table 1: Summary of domain properties comparison.

reduce the amount of available items, thus making the recommendation task easier. On the other hand, customers are sometimes required to plan ahead, and purchase a certain item in order to be eligible for other items in the future.

In movies, books, or news items there are no such constraints. It is quite common, for example, for people to watch two action movies with Bruce Willis, even though they know that there will be a significant overlap between the movies. the same applies to books of the same author, or news stories that cover the same event. Electronic gadgets have a similar constraint to insurance packages in that people typically buy only one laptop, memory stick, or MP3 player (at least for themselves).

Vacation packages possess similar constraints. For example, one cannot assign hotels for a trip disregarding the flights. Suggesting two hotels for the same night is also useless. In this regard, the two domains are quite similar.

Demographic constraints, however, hold to a certain degree in all domains through user tastes. For example, certain books or movies are intended for teen-agers and are rarely read or watched outside that crowd. Certain vacations are designed for families, while others are for the newlywed. While these are not hard constraints, the recommendation system is expected to take such trends into account.

6.5 Interacting with the recommender

In the classic applications we have mentioned, in a typical interaction scenario, users browse online the space of items, assisted by the recommender system. One obvious example of such an interaction process is users browsing a news website, such as cnn.com, where the webpage contains, in addition to the displayed story, links to the "most popular stories", or "people who read this also read" links. In such cases there is a direct interaction between the user and the recommendation system. The same applies for people selecting movies for their watching queues in Netflix, or customers buying books at Amazon.

In the insurance domain, it is rarely the case that people buy insurance products directly online, the obvious exception being car insurance. In most cases customers buy insurance from an insurance agent, either in a face-to-face meeting, or over the phone. One reason for this is the high complexity of insurance products, and the low expertise of customers, resulting in the need to receive many explanations over the packages.

As such, the recommender system in the insurance domain can be designed to be used by the insurance agents themselves, either in the call center or in a face-to-face meeting as a helping tool in suggesting which products are more suitable for the current customer. That is, we believe that the interaction between the customer and the recommender system is done indirectly through the sales agent.

6.6 Attention span

Most people interact only rarely with insurance companies, and buy a relatively small number of insurance products. Furthermore, insurance products typically have a relatively high price, and are bought to be used for a long time period (from one year in car insurance to a lifetime in many health insurance products). As such, people expect that purchasing an insurance will be a tedious chore, requiring a long dialogue of questions and answers from both sides.

People who shop for an explicit insurance product, are rarely willing to hear about other insurance products. That is, a person buying a car insurance may be willing to hear about items related to car insurance, such as additional riders for the package that she has bought, but may not be interested in hearing about health insurance plans.

In other domains, such as books or movies people interact many times with the recommender system (in the case of Netflix, the minimal package is one movie per month, which suggests a monthly interaction). People may read hundreds of news stories, watch hundreds of movies, read many dozens of books, and buy several dozens of electronic gadgets. It is also quite likely for people who enter Amazon to buy a certain book to be persuaded to add a completely different book to the transaction, and the same holds for electronic gadgets. In fact, the main purpose of recommender systems in retail websites is to promote cross-sell — cause customers to buy more than they originally intended.

People also buy vacation packages only once or twice an year. As such, interaction with the system is not done very often, although probably more often than in the insurance domain. It is likely, though, that most people may find exploring their vacation options more enjoyable than reading about insurance plans. Therefore, we may expect longer interactions and a longer attention span in the case of purchasing vacation packages.

In domains other than insurance, people can be expected to be willing to explore a larger set of items, both because the shopping experience itself in these domains can be enjoying, and because items are simpler to understand without much explanations.

7. EXPERIMENTAL EVALUATION

We now report the results of an offline study over a dataset of customer purchasing insurance riders. We then briefly discuss an online study that was conducted with real customers.

7.1 Properties of the dataset

We have obtained a dataset containing customers with base insurance policies in the health domain, and a number of riders, that purchased additional riders for their policies. The data was collected during the last quarter of 2010, and contains 30,000 customers, 13 base policies, and 64 riders. There were 73,565 transactions sold through multiple channels — through insurance agents, through the call-center, and through other channels. We removed from the dataset customers who purchased a collective package, typically sold to large customer groups, such as a large organization employees, as a fixed, non-configurable package.

Figure 1: Base policy popularity distribution

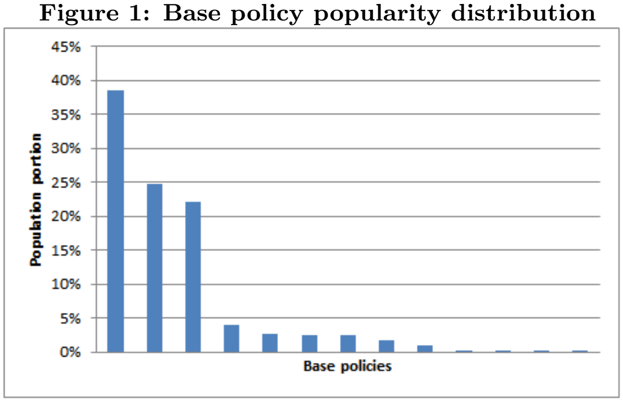

Figure 1 shows the distribution of base package popularity. There is one highly popular policy, purchased by almost 40% of the customers, and 2 more base policies with 22% and 25% popularity. Figure 2 shows the distribution of rider popularity. There is one rider that is purchased by more than 70% of the customers, and 5 riders purchased by about 10% of the customers or more, but about 30% of the purchases are for unpopular riders (riders purchased by 8% or less of the customers). These two distribution present the long tail behavior often found in item selection datasets in other recommendation system applications, where the number of popular items is relatively low, and many of the purchases are for non-popular items.

Figure 2: Riders popularity distribution

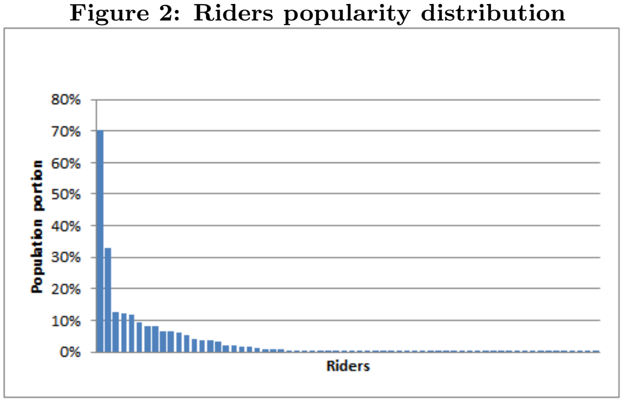

Figure 3 shows the distribution of the number of purchased riders per customer. Customers purchase 5.42 riders on average. About 20% of the customers purchased only a base policy with no riders.

In addition, we have data over customers demographics, including their income level, marital status, gender, zip code, and so forth. Due to privacy issues we do not provide a detailed description of customers statistics over this data.

Figure 3: Riders per customer distribution

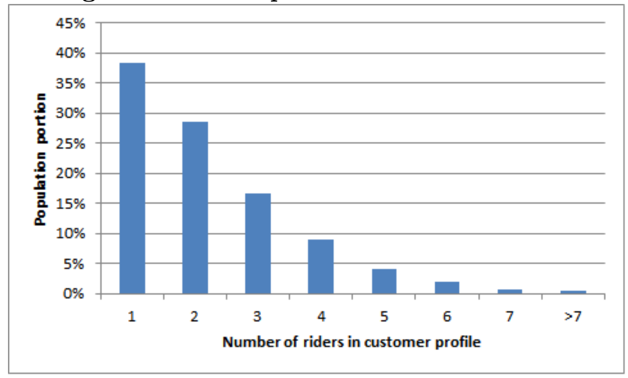

7.2 Offline Experimental results

We conduct an experiment to analyze the performance of simple collaborative filtering techniques on the insurance dataset that we have. The goal of the experiment is to study a simple and fast, yet effective, recommendation technique that could be easily integrated into the call-center system.

We implemented a simple item-item collaborative filtering approach[12], based on the conditional probability estimation $\Pr(i_1|i_2) = \frac{count(i1,i2)}{count(i_2)}$. The system's task is the recommendation of additional riders. We therefore compute conditional probabilities estimations for riders given base packages, and for riders given other riders, but not for base packages given riders. Given a customer u with a base package and several riders, we define $\Pr(r|u) = max_{i \in u} \Pr(r|i)$ for every rider r that u has not yet purchased, and item i (base policy or rider) that u has purchased, and order these recommendations by decreasing conditional probability.

We computed the conditional probabilities estimations for 75% of the users (the training set), and measured the recommendation accuracy for the rest of the users (the test set). Given a test set customer with a base policies and a list of n purchased riders, we uniformly pick a number m between 1 and n, and then pick m random riders and hide them. We then try to predict a list of varying size k from 1 through 7, and compute the precision at k.

As we can see from Figure 2, several riders are very popular. It seems reasonable, hence, to compare our recommendation technique against a fixed popularity recommender. Furthermore, one can gain additional power by recommending the popular items given the customer's base policy. We compare our algorithm also against this technique.

Figure 4 shows the precision-recall curve for the different methods. As we can see, the precision of the conditional probability method drops down after the second item recommendation. As we can see in Figure 3 only 16% of the population purchased more than 3 riders. For the rest of the customers, who buy 3 or less riders, the system cannot offer good recommendations beyond one or two riders. We have therefore allowed the recommender to truncate the list when the probability of the next item drops below a certain threshold (0.5 in the reported experiments). This is equivalent to the system prediction that the customer will not purchase more items. In a real system, this can be a valuable feature, because bothering customers with rider recommendations that they are not interested in has a cost in terms of customer satisfaction and trust in the insurance company. As

Figure 4: Precision-Recall results for recommending riders to customers.

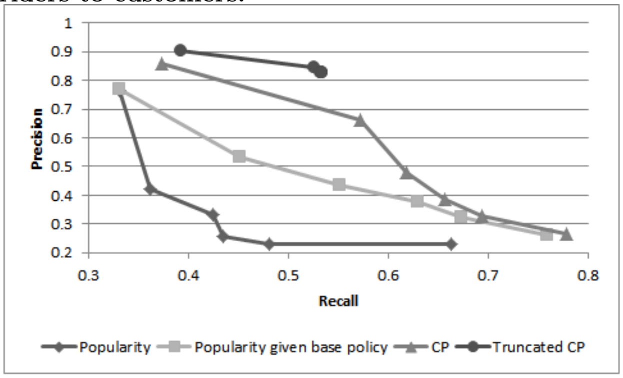

we can see, this method (denoted "Truncated CP") outperforms the other methods in terms of precision, although it does not achieve as high recall rates, as expected. In the experiment, the method never saw any value in recommending more than 3 items.

In our experiments, we tried to add available demographic information, such as age, gender, zip code, level of income, and so forth. That is, we tried a conditional probability model where the purchased item is conditioned upon the demographic properties of the customer, as well as on the purchased base policy. We did not find any positive effect of adding these factors to the recommendation quality. It seems that the already purchased base policy and riders are much more informative concerning the customer preferences than, e.g. their zip code.

These preliminary results are very encouraging because they show that very good predictions (with precision above 80%) can be made to customers based on their currently purchased packages. They also suggest that very little could be gained by employing more sophisticated collaborative filtering techniques.

We acknowledge that our dataset may contain several biases that are hard to compensate for. These are mainly because many of the insurance packages in the dataset were sold by insurance agents. As we explained above, these insurance agents often sell riders not only because they fit the client the most, but also because of other factors, such as the commissions that they earn, or the ease of convincing clients to buy them. Furthermore, we have noted distinct patterns in the way that agents from a particular agency sell products. That is, some agencies tend to sell certain products much more often than other agencies. This can be attributed to specialization in various customer populations, but can also be attributed to products that are favored by the agency due to other reasons.

7.3 Online Experiment

Following the successful offline results, the insurance company has decided to test our approach online in their call center. In the experiment we have computed offline a set of recommendations for customers using the truncated CP method. A few call center employees were then given the list of recommendations and were asked to call the clients and offer them the recommended riders. This was done while the rest of the call center staff called customers with the regular suggestions.

While we are not at liberty to provide exact numbers, the sales following the recommended riders were about 3 times more than the sales following the regular suggestions! The insurance company was very pleased by these results and is exploring an implementation of an automated system for their call center.

8. CONCLUSION

In this paper we have presented the problem of recommending insurance riders to policyholders. We have described the insurance domain and its various special properties, compared with more traditional recommender system applications.

In collaboration with a large Israeli finance company, we have obtained records of customers purchases of base policies and riders, and customer demographic properties. We use this data in an offline evaluation to measure the ability to predict, given a base policy and some riders what other riders will the person be interested in. Our experiments show that simple item-item collaborative filtering approaches provide an impressive ability to predict additional riders, far better than suggesting the most popular products.

Our system was also experimented with at the finance company call center, showing an impressive increase in sales following recommendations versus the regular call center suggestions.

In the future, we would like to conduct a wider and longer study, measuring user satisfaction with the riders that were sold. Such a study could, e.g., measure whether users maintain the riders that were bought following a recommendation, or cancel them after a while. It would also be valuable to receive data over riders that were not sold through a promotion by a sales agent, in order to reduce the biases int he data, although we acknowledge that such data may be difficult to achieve.

9. REFERENCES

[1] J. Anton. The past, present and future of customer access centers. *International Journal of Service Industry Management*, 11(2):120–130, 2000.

[2] J. Breese, D. Heckerman, and C. Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *UAI '98*, pages 43–52, 1998.

[3] S. Debnath. Machine learning based recommendation systems. Master's thesis, Department of Computer Science and Engineering ,Indian Institute of Technology, 2008.

[4] C. Desrosiers and G. Karypis. A comprehensive survey of neighborhood-based recommendation methods. In *Recommender Systems Handbook*, pages 107–144. 2011.

[5] R. Driskill and J. Riedl. Recommender systems for e-commerce: Challenges and opportunities. Technical Report WS-99-01, AAAI, 1999.

[6] R. Dumm and R. Hoyt. Insurance distribution channels: markets in transition. *Journal of Insurance Regulation*, 22(1):27–48, 2003.

[7] A. Gilman, B. Narayanan, and S. Paul. Mining call center dialog data. *Information and Communication Technologies, Data Mining V*, 33:317–325, 2004.

[8] T. Hofmann. Latent semantic models for collaborative filtering. *ACM Trans. Inf. Syst.*, 22(1):89–115, Jan. 2004.

[9] R. Holman, D.; Batt and U. Holtgrewe. The global call center report: International perspectives on management and employment. Technical Report 13, Research Studies and Reports. Cornell University, 2007.

[10] J. Konstan, B. Miller, D. Maltz, J. Herlocker, L. Gordon, and J. Riedl. Grouplens: Applying collaborative filtering to usenet news. *Communications of the ACM*, 40:77–87, 1997 1997.

[11] Y. Koren and R. M. Bell. Advances in collaborative filtering. In *Recommender Systems Handbook*, pages 145–186. 2011.

[12] G. Linden, B. Smith, and J. York. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7:76–80, 2003.

[13] J. Luan, C. Summa, and M. Weiland. Use of data mining to examine an outreach call center's effectiveness and build a predictive model for classifying future marketing targets. *Data mining in E-learning , State-of-the-art in Science and Engineering*, 2005.

[14] M. Montaner, B. López, and J. L. de la Rosa. A taxonomy of recommender agents on the internet. *Artif. Intell. Rev.*, 19(4):285–330, 2003.

[15] D. Padmanabhan and K. Kummamuru. Mining conversational text for procedures with applications in contact centers. *Int. J. Doc. Anal. Recognit.*, 10:227–238, December 2007.

[16] M. Paprzycki, A. Abraham, R. Guo, and S. Mukkamala. Data mining approach for analyzing call center performance. In *Proceedings of the 17th international conference on Innovations in applied artificial intelligence*, IEA/AIE'2004, pages 1092–1101, 2004.

[17] C. Pfeil, T. Posselt, and N. Maschke. Incentives for sales agents after the advent of the internet. *Marketing Letters*, 19(1):51–63, 2008.

[18] U. Rabanser and F. Ricci. Recommender systems: Do they have a viable business model in e-tourism? In *Information and Communication Technologies in Tourism*, pages 160–171, 2005.

[19] P. Resnick and H. R. Varian. Recommender systems. *Commun. ACM*, 40(3):56–58, 1997.

[20] B. Sarwar, G. Karypis, J. Konstan, and J. Reidl. Item-based collaborative filtering recommendation algorithms. In *WWW '01: Proceedings of the 10th international conference on World Wide Web*, pages 285–295, New York, NY, USA, 2001. ACM.

[21] 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*, pages 158–167, New York, NY, USA, 2000. ACM.

[22] J. Schafer, J. Konstan, and J. Riedl. E-commerce recommendation applications. *Data Mining Knowledge Discovery*, 5:115–153, 01/2001 2001.

[23] H. Takeuchi, L. V. Subramaniam, T. Nasukawa, and S. Roy. Getting insights from the voices of customers: Conversation mining at a contact center. *Information Sciences*, 179:1584–1591, 2009.

[24] C.-N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *WWW*, pages 22–32, 2005.