Inf Syst E-Bus Manage (2012) 10:367–393
DOI 10.1007/s10257-011-0170-8

ORIGINAL ARTICLE

Moving recommender systems from on-line commerce to retail stores

Frank E. Walter · Stefano Battiston · Mahir Yildirim · Frank Schweitzer

Received: 11 May 2010 / Revised: 30 September 2010 / Accepted: 24 January 2011 / Published online: 30 March 2011
© Springer-Verlag 2011

Abstract The increasing diversity of consumers' demand, as documented by the debate on the long tail of the distribution of sales volume across products, represents a challenge for retail stores. Recommender systems offer a tool to cope with this challenge. The recent developments in information technology and ubiquitous computing makes it feasible to move recommender systems from the on-line commerce, where they are widely used, to retail stores. In this paper, we aim to bridge the management literature and the computer science literature by analysing a number of issues that arise when applying recommender systems to retail stores: these range from the format of the stores that would benefit most from recommender systems to the impact of coverage and control of recommender systems on customer loyalty and competition among retail stores.

Keywords Recommender systems · On-line commerce · Retail stores · Long tail · Information overload

F. E. Walter (✉) · S. Battiston · M. Yildirim · F. Schweitzer
Chair of Systems Design, ETH Zurich, Kreuzplatz 5, 8032 Zurich, Switzerland
e-mail: fewalter@ethz.ch

S. Battiston
e-mail: sbattiston@ethz.ch

F. Schweitzer
e-mail: fschweitzer@ethz.ch

M. Yildirim
Institut für Informatik, Universität Freiburg, Hermann-Herder-Str. 10, 79106 Freiburg, Germany
e-mail: yildirim@informatik.uni-freiburg.de

1 Introduction

Information technologies have transformed the way people carry out purchases. This has lead to the rise of a number of challenges for both retail stores and consumers: on one hand, retail stores have to face the competition of on-line stores such as Amazon.com. Additionally, there is an increasing diversity in the demand (Anderson 2006; Brynjolfsson et al. 2006) meaning that there are many products with low sales volume which, collectively, make up significant part of the whole share of the turnover of a retail store. In other words, the distribution of sales volume across products has a long tail (Anderson 2006; Brynjolfsson et al. 2006). This is due to the increasing number of niches in the market and can already be seen in the on-line market, but is affecting the retail market as well. Retail stores are more vulnerable to this due to the fact that it is very costly for them to keep a large variety of seldomly-sold products on their shelves. On the other hand, consumers are exposed to an increasing variety of products the information about which has to be processed—this is a situation of *information overload* (Kim and Lee 2005; Lam and Riedl 2004). One reason for this situation is the increasing concentration of stores in huge malls (Ernst & Young 2003, 2005; KPMG 2006), another reason is the integration of worldwide markets which has lead to a plethora of goods from distant countries that are unknown to customers. However, in addition to these challenges, the same information technologies also provide a number of opportunities, for example the possibility to empower consumers to carry out more informed decisions about purchases (Walter et al. 2008). One of these technologies are recommender systems, which are already widespread in on-line commerce, but are not yet available in retail stores.

Recommender systems can provide a tool that both retail stores and customers could benefit from. In this paper, we focus on the side of retail stores and analyse the main business as well as technical aspects arising from moving recommender systems from on-line stores and applying them to retail stores. Our intention is to give a not-too-technical overview of the challenges of moving recommender systems to retail stores, providing non-experts with an overview of the subject matter.

Recommender systems allow businesses to leverage knowledge about their customers. This knowledge can be used to create a personalised experience tailored to each customer's individual needs. Such recommender systems are an inherent part of today's e-commerce web sites, helping customers to identify products of interest. Their success in the "virtual world" has had significant impact on e-commerce as a whole. In the following, we investigate how "real world" retailers can also benefit from recommender systems.

The purpose of recommender systems is to help customers to identify a set of products that might be of interest to them. Nowadays, recommender systems assist millions of consumers each day. Since large e-commerce web sites offer millions of distinct products for sale, this form of guidance is of utmost importance to consumers. However, not only customers benefit from recommender systems. They also help e-commerce websites to offer greater convenience to their customers, but also to increase their sales. One of the biggest all-purpose on-line retailers in Germany reported that, two months after the deployment of a recommender system, "the rate of [requests for] cross-selling articles in the marketplace […] has increased ten-fold" (Prudsys 2006) (meanwhile, this particular retailer went—for reasons not related to its recommender system—bankrupt and is in Chap. 11 at the moment). Other e-commerce web sites such as Avitos (also on-line retailer in Germany, specialised on electronics) experienced an increase of 20% in sales in the product category "digital cameras" after it launched an on-line sales assistant for that product group (Mentasys 2007). There are many other examples that demonstrate that recommender systems play a crucial role in deciding where and what to shop on-line (Schafer et al. 2001).

So far, this scenario has only taken place in the virtual world. Facilitated by recent developments in ubiquitous computing, it seems reasonable that the success of such recommender systems in the virtual world can be replicated in the real world.

In Sect. 2, we start by putting our work into the context of the related work. We then go on to discuss the business aspects of applying recommender systems to retail stores in Sect. 3. In particular, we investigate which types of retail store would benefit from a personalisation strategy and to which extent and discuss how personalisation affects sales. In Sect. 4, we consider personalisation from a more technical point of view. We present a recommender system tailored to the retail environment, show how it can be represented by a graph model, and then illustrate a layered architecture for recommender systems which can act as a blueprint for real-world implementations. Concluding remarks and a summary of the work are given in Sect. 5.

2 Related work

The analysis carried out in this paper relies on research from several disciplines, among them recommender systems, ubiquitous computing, and market analysis.

2.1 Recommender systems

Recommender systems research is a field within computer science closely related to computational statistics, but it is also using techniques from information retrieval and machine learning. The main purpose of recommender systems is to assist users in their decision making process. They accomplish this goal by either recommending new items of interest to users or by filtering the possible choices according to users' preferences. In other words, recommender systems are a tool to cope with information overload; properly implemented, they allow users to filter the relevant from the irrelevant (Kim and Lee 2005; Lam and Riedl 2004).

Collaborative filtering (CF) is one recommendation technique that has become an industry standard and is an inherent part of many e-commerce sites (Linden et al. 2003; Resnick et al. 1994; Sarwar et al. 2000). CF systems use data on user preferences to infer similarities between users and items. Once the similarities are determined, new items can be recommended for a particular user by identifying items liked by users similar to that user (Breese et al. 1998; Herlocker et al. 1999; Resnick et al. 1994).

Since CF works well in many cases, an increasing number of online communities for product and service reviews provide personalised recommendations based on collaborative filtering algorithms. there are also some limitations. A first problem is that the predictions are often poor when users are seeking recommendations for seldom rated items. The reason is that the system may not be able to estimate the similarity between the user who requests the recommendation and the few users who actually rated the item. A second, and more fundamental problem, is that often two persons have similar mind settings and tastes although they have completely different domains of expertise and activity (e.g. two old friends of whom one became a medical doctor and the other a computer scientist). It is then possible that two persons do not have a single co-rating and still they would rate many items in the same way, were they asked to experience them. Moreover, preferences of people are not fixed: in many cases people want to *learn* from their friends how to appreciate things (e.g. about fine food or fine wines). Finally, in other cases preferences are not the main dimension in question: people want simply to *be advised* on particular topics (e.g. in health or financial investment).

Thus, trust-based recommender systems have been suggested as a means to cope with such limitations of existing recommender systems. A trust-based recommender system generates recommendations for users not from the similarity of a user to other users computed based on the actual co-ratings, but from the *trust* that they have in other users. In this context, trust is defined as the "expectancy of an agent to be able to rely on some other agent's recommendations" (Walter et al. 2008). "The small-world property of social networks (Newman et al. 2002) allows to potentially reach a lot of information, while the trust allows to filter out the relevant pieces" (Walter et al. 2009). Users can define their web of trust in a similar fashion than people define friends in social networks such as Facebook, LinkedIn, and Xing. On the one hand, research has focused on such "trust webs" themselves (Abdul-Rahman and Hailes 2000; Grandison and Sloman 2000; Marsh 1994; Sabater and Sierra 2005), and on the other hand, on their application to recommender systems (Golbeck 2005; Massa 2006; Montaner et al. 2002; O'Donovan and Smyth 2005).

There have also been attempts to design recommender systems for retail stores (Kitts et al. 2000; Liu et al. 2004; Wang et al. 2004). For example, Lawrence et al. have presented a personalisation approach that is designed for supermarkets (Lawrence et al. 2001). However, their approach needs a product taxonomy in order to be applied. In Sect. 4, we will introduce a personalisation approach that does not require any taxonomy but leverages the data that is obtained in the store and provided by the customers. Beyond computer science, recommender systems are also studied in marketing and, consequently, in customer behavioural science—there, of course, the focus is on how recommender systems affect customers.

2.2 Ubiquitous computing

The driving force which makes the transition of technologies from the virtual world to the real world possible lies in the fast pace of developments in information and communication technologies. More than fifteen years ago, Mark Weiser has framed the concept of "ubiquitous computing", a scenario in which computing is an essential and invisible part in people's everyday life (Weiser 1991).

Because of these developments, it is, nowadays, possible to equip retail stores with a highly mobile infrastructure that allows automated data recording in real-time and improved communication with customers (Decker et al. 2003; Fleisch 2001; Kourouthanassis and Roussos 2003; Krohn et al. 2005; Sackmann et al. 2006). Technologies such as radio-frequency ID (RFID) tags improve the supply chain and inter-firm cooperation because they allow the consistent integration of real and virtual worlds (Fleisch 2001). In principle, such technologies could also be used to provide additional information to the customers of a retail store, as already proposed by Cinicioglu et al. (2007).

Many other applications are already being explored in practice by big retail chains such as Walmart or Metro. "Smart shelves" that are able to track basic actions of the customers, such as removing items from and returning them to a shelf have already been implemented. Also, experiments with mobile shopping assistants in form of personal digital assistant (PDA) computers with interactive screens on shopping carts and in-store information terminals are on the brink of becoming commercial practice. It can be useful to combine ubiquitous computing with recommender systems to personalise the experience of customers in retail stores.

2.3 Market analysis

Market analysis is concerned with providing corporate decision-makers with the necessary information to plan the activities of a firm. Some of the aspects of a company that are adjusted based on the information provided by a market analysis are its inventory, workforce, and services.

Recent market analyses of the retail market have shown that, as response to dropping margins and lower sales, retailers have opened larger stores with a wider product assortment, less sales clerks, and much higher spending on advertising. These measures have further intensified the competition giving rise to the typical structural problems of the retail market: shrinking productivity per space, indifferent stores, extremely low margins, and a more selective customer base (Ernst & Young 2003, 2005; KPMG 2006; Mundt et al. 2002; Schröder et al. 2003).

Could personalisation realized by recommender systems offer a competitive edge to those retailers who take advantage of this technology? This would be achieved by reducing the customers' search cost and by better matching their needs with the retailers' assortments.

3 Business perspectives

We will start this section by depicting the main factors influencing cost and revenue for the retail stores and analyse different levels of personalisation within this context. Then, we will point out the benefits, evident in e-commerce, that would also boost personalised retail stores. In fact, there already is a trend towards personalisation in the retail market and recommender systems would be a method to make the personalisation

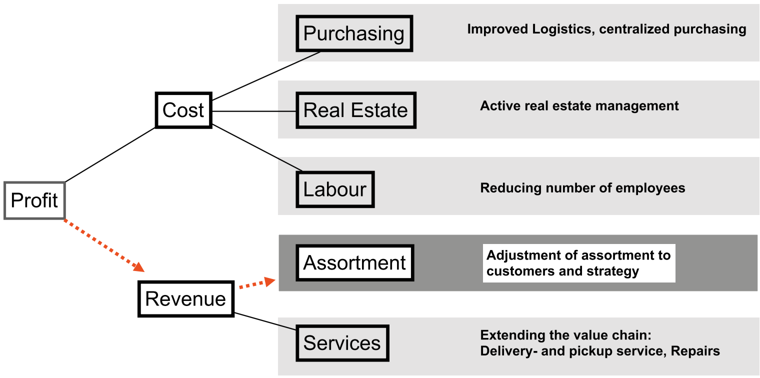

Fig. 1 Tree of the factors impacting profit of retailers

more sophisticated and effective. Interestingly, according to our analysis, not all retailers would benefit equally from personalisation; rather, this depends on the variety of their product assortment and the diversity of their customer base. We will conclude this section by arguing that the main added value that recommender systems can bring to retailers is to improve their assortment and placement.

3.1 Factors that affect the profit of retailers

In order to determine the effect of recommender systems on profit of retailers, we have to determine what factors affect the profit of retailers and how this, in turn, is affected by recommender systems. Figure 1 illustrates several factors that impact the profit of retailers. Please note that the tree is not supposed to be exhaustive; rather, its purpose is to highlight some of the prominent drivers that have an impact on profit and the typical approaches taken by retailers to deal with these challenges. An increase in profits can either be achieved by decreasing costs or increasing sales.

On the cost side, the factors are purchasing (which includes logistics and stock management), real estate, and labour costs. On the revenue side, the factors are assortment and services. At this stage, we can see a possible impact of recommender systems on the assortment factor.

The assortment and placement provides a means for retailers to distinguish themselves from other competitors beyond the price of products and services. However, the wider the assortment, the greater is the cost of stock management. Additionally, choosing the assortment and placement wisely is not trivial and requires good knowledge of the customer base. The assortment of a retail store can be improved by making certain products and services better reachable and visible to the customers as discussed in more detail in Sect. 3.5. Recommender systems and personalisation are key to these improvements. However, in the following, we will see that there are also additional benefits for personalised retail stores.

3.2 What retailers can learn from e-commerce

Over the last decade, Europe has witnessed the spread of bigger shopping centres with less sales clerks, but a wider range of products. Having the choice among so many products not only increases the amount of information that has to be processed by customers, but also creates the need to guide the customer towards the products.

E-commerce sites address this issue of their customers facing information overload by utilising recommender systems that identify products in which the customer may be interested. The recommender system can suggest products based, for instance, on the demographics of the consumer, on the best-selling or highest rated items, or by utilising the purchasing history of each individual customer. In this way, the e-commerce web sites adapt themselves to each customer's personal needs, wants, and demands.

3.2.1 Developing one-to-one marketing

One-to-one marketing is a customer relationship management strategy that attempts to overcome the impersonal nature of mass-marketing by treating each customer individually. In other words, it implies 'treating different people differently'. Such a one-to-one approach is the natural manner for a shopkeeper to interact with one of their customers: they remember details about each customer's preferences or characteristics and use that knowledge to provide better service. However, most businesses of today cannot maintain a one-to-one interaction with their customers because of rationalisation effects such as larger retail stores with a wider range of goods and fewer employees. Recommender systems provide a technology that helps businesses in implementing the one-to-one marketing strategy to treat each customer individually at a reasonable cost.

3.2.2 Improving customer loyalty

There is no consensus on the definition of customer loyalty and therefore there are several proposed indicators of loyalty: for instance, the customer retention rate, the degree of customer satisfaction, the share of money that customers spend in a particular store, or whether the customer would recommend the product or service to a friend (Reichheld 2003).

Empirical results of studies conducted in the literature show that "only about 10% of the buyers for many types of frequently purchased consumer goods are 100% loyal to a particular brand over a one-year period" (Dowling and Uncles 1997). Furthermore, these buyers usually are not heavy, but rather light buyers of a product. Thus, about 90 % of the consumers of a particular type of good are not completely loyal to one brand, and many of these are hopping between various brands.

Hopping, in this context, means that customers change frequently between products or services to find the one that they like best. The utility of hopping is to—eventually—find the product which matches the consumer's preferences, but there also is an inevitable cost of hopping in form of the time spent on trying different products and the risk that a particular product does not suit the customer's preferences. Thus, consumers try to find the best trade-off between exploring products and exploiting the ones they know to match their preferences. This issue has also been investigated in economics (Weisbuch et al. 2000).

In this context, recommender systems are a way of significantly reducing the costs associated with finding a suitable product. They suggest the customer products and services that better match their preferences. Thereby, to customers, trying different brands is not as costly as without a recommender system—the risk to choose something that the customer would not like is much lower as the recommender system acts as a filtering mechanism that only suggests the customer products and services that are similar to their preferences.

If a retail store offers a recommender system to its customers for these purposes, it can be assumed that this effort is rewarded by customers by shopping more frequently or by purchasing more items because they feel well-treated and satisfied. Thus, recommender systems may eventually lead to improve all of the metrics for customer loyalty. We will come back to this issue in Sect. 3.6.

3.2.3 Increasing cross-selling

By suggesting additional products for the customer to purchase, recommender systems can significantly improve cross-selling—i.e., selling related or complementary products in addition to a product that the customer originally wanted to purchase. An example for cross-selling are the often used phrases in fast-food restaurants: "Would you like French fries?" or "Would you like a drink with your meal?". E-commerce sites such as Amazon.com put this marketing strategy of cross-selling into practice by displaying related or complementary products under the label "Customers who bought this … also bought …" or "You might also find these … interesting". Similarly, recommender systems provide more opportunities for retail stores to increase cross-selling. Often, cross-selling is realized by putting related items close to each other. Recommender systems relax this requirement because they allow to advertise related or complementary products without the need for these products to be in close proximity. For example, based on what a customer has in their shopping cart and their location in the store, the recommender system could suggest products that are related and in proximity.

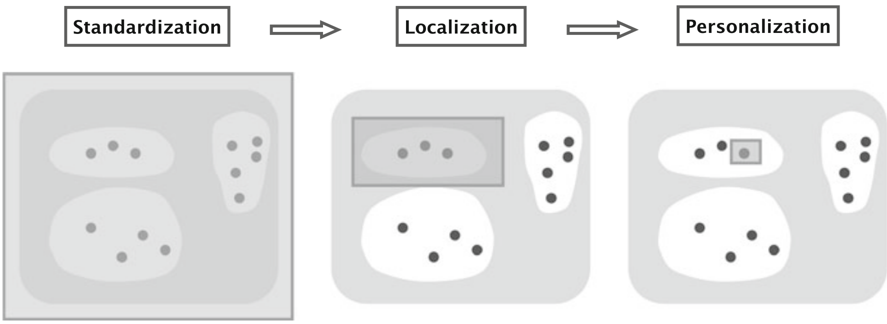

Fig. 2 Strategies for retailers for interacting with customers

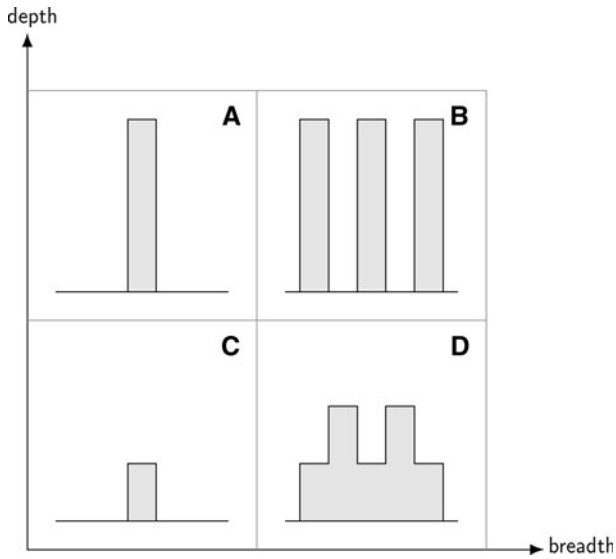

Fig. 3 Diagram illustrating how the product mix of a retail store can vary along the dimensions of breadth and depth

3.3 Strategies for retailers

In this section, we describe how increasing levels of personalisation are reflected in different strategies for retailers to interact with their customers. These three possible strategies as depicted in Fig. 2 differ in the granularity of interaction between retailers and customers: standardisation is used to address mass markets, localisation creates market segments, and personalisation addresses customers individually. Each of these strategies can benefit from the lessons learnt in e-commerce, even though each in a different way and to a different extent (Fig. 3).

- *Standardisation* Standardisation initiated the first revolution in the retail market. Big chains such as Walmart, McDonald's, and Tesco have all pursued strategies of standardisation (Rigby and Vishwanath 2006). They have developed uniform businesses in terms of store format, merchandise mix, and marketing efforts in order to take advantage of *economies of scale*. Although standardisation has had and still has an important role, with the increasing diversity of demand (Anderson 2006; Brynjolfsson et al. 2006) and the decreasing number of places to introduce new stores, standardisation has reached certain limits.
- *Localisation* A current trend in the retail market is to customise the stores and their offerings to the local customer base (Rigby and Vishwanath 2006). This localisation strategy is inferred from analysing differences in the regional customer base. The goal of localisation is to split the current store format into several modules that are targeted towards the local customer base, but that are still taking advantage of economies of scale. For example, Best Buy has identified five customer-centric store formats, such as one that appeals to busy

mothers who want a quick and personalised help through the world of technology, or another that is targeted at experienced users of technology interested in the latest gaming and entertainment gear (Rigby and Vishwanath 2006).

- *Personalisation* The next possible strategy for retailers is to build up a one-to-one relationship with their customers. It has been shown (Anderson 2006; Brynjolfsson et al. 2006) that in on-line retail stores, the distribution of sales per product has a long tail, meaning that there are very many products with very few sales which, however, sum up to significant total sales volume. In other words, there is a growing number of market niches, which can be targeted only by personalised customer relationship management. Since consumers are increasingly used to the possibilities offered by the on-line market, it is crucial for retail stores to address this issue and offer similar services.

As in the e-commerce world, recommender systems provide a scalable technology to establish a personalised communication channel between customers and retail stores. With this new form of customer communication, retail stores can capture the preferences of their customers at a greater level of granularity and guide customers through the overwhelming set of products by identifying items that best match the customers' current needs. This greatly reduces the search cost for customers and thus improves their ability to cope with the information overload that they face when being confronted with thousands of potential products. Furthermore, by making the appropriate products better reachable and visible for their customers, retailers make the best use of their assortment (see Sect. 3.5). However, not all retail stores will benefit from introducing recommender systems in order to realize personalisation equally, as we will illustrate in the next section.

3.4 Mapping strategies to retail formats

In management science, it is customary to classify the format of retail stores based on *breadth* (i.e. the number of product lines they cover) and *depth* (i.e. the number and variety of items within a given product line) of their product mix (Kotler and Armstrong 2004). For instance, a shop specialised ties with a large variety of brands, quality and tastes, has a deep and narrow product mix.

In our context, it is important to discuss for which of these formats recommender systems could be most useful. Because a recommender system processes information on the whole customer base to make predictions on a given customer preferences, we need to combine the breadth-depth classification with the level of heterogenity of the customer base to which the store is targeted. This leads to the following considerations.

1. *Narrow-Shallow mix (C)* For instance, little convenience stores or grocery shops (typically open until late in metropolitan areas of the US and some EU countries) offer just basic products in few categories. The customer base is heterogeneous. Essentially all types of customers now and then shop there, regardless of their income and taste, simply out of convenience and or

necessity. Customers do not expect to find in such a store a good match for their taste. In this case, a recommender system is not very useful.

2. *Broad-Shallow mix (D)* For instance, chain stores or discount stores with limited choice over many categories (e.g. food, houseware, cloths). If the shop is visited by a heterogenous customer base, a recommender system is useful to customers to decide whether the quality of the products offered matches their requirements and whether they do or do not need to go to a more specialized shop.
   However, an alternative strategy for the retailer would be the localisation (as see above), consisting of targeting a specific segment of the potential customer base. A recommender system is useful in this case to help the retailer to adapt the assortment to the targeted customer base.
3. *Narrow-deep mix (A)* These are speciality shops. An example is the tie shop mentioned above. In this case, the sales staff play a crucial role. They can establish a relationship to their customers by remembering details about their preferences or, based on their specialised knowledge of products, guide the customers to products that best suit them. In particular, personalisation is a valid strategy for retailers that trade products and services in the premium sector. Thus, recommender systems would be applicable in such stores, too, for example as an assisting tool for the sales staff.
4. *Broad-deep mix (B)* Retail stores of this type offer extensive assortment in several product lines. This is the case in which recommender systems operating within a single shop are most useful (but see also the discussion on coverage in Sect. 3.6). Because these are necessarily large shops with many categories and the customer base is likely to be very heterogeneous, recommender systems are very helpful for customers to orient themselves in the store and to compare products.

It is important to mention at this point that the personalisation strategies are crucially influenced by whether the recommender system is in the control of the customer or the retail store and whether it can cover one or several retail stores. We discuss these issues of control and coverage of recommender systems in more detail in Sect. 3.6.

3.5 Improving the assortment

Space on the shelves is a fundamental limitation to increasing the product assortment of a store. However, as discussed in Sect. 3.1, recommender systems can be used to adjust the assortment and placement through personalisation: on one hand, customers are guided to the products and services that correspond to their preferences; on the other hand, retailers are guided to adjust their assortment and placement such that it reflects the preferences of their customers. For example, items which are perceived as similar or associated—a property that can be discovered through the recommender system—can be positioned in proximity to each other. Furthermore, recommender systems allow retail stores to learn about the preferences of customers and therefore it allows them to better select the target segment of customers by adjusting their assortment to these customers. However, an efficient information system is required in order to evaluate shopping behaviour and preferences of customers.

We identify three different ways to better adjust the assortment and placement when retailers make use of recommender systems:

- *Reachability* A large number of products stocked by retailers on their shelves makes it hard for customers to find their favourite products. An application that guides the consumer to the desired product reduces search costs and thus make shopping more convenient.
- *Visibility* About two thirds of customer purchases are unplanned (Kaufman 2000). There is an optimisation potential for retailers that can best adapt and recommend products according to the customer's current preferences.
- *Trust* Consumers are often facing the situation of having found a product that appears interesting, but they are not sure whether to purchase it or not. In this scenario, recommender systems can provide reviews of trusted peers and thereby increase customer satisfaction by reducing the risk of mispurchases.

3.6 Control and coverage of recommender systems

Following (Sackmann and Strüker 2004), there are two important dimensions affecting recommender systems: the control over the recommender system and the coverage of the recommender system. Note that (Sackmann and Strüker 2004) refer to control as "power of disposal over the end device" and to coverage as "range of provision".

- *Control* At one extreme, it could be that the recommender system is owned by the retail store and supplied to the customer; at the other extreme, it could be that users have their own recommender system which allows them to gather and process information (such as quality, price, etc.) of products.
- *Coverage* At one extreme, it could be that the recommender system covers only one retail store (or several ones of the same chain); at the other extreme, it could be that the recommender system covers several, competing retail stores.

From the point of view of the consumer, it is obviously more desirable to have a recommender system that is under their own control (and thus not susceptible to any form of bias) as well as that covers all sorts of retail stores. For the retail store it is just the opposite. In this respect, consumers and stores have opposed interests which can not be reconciled in a simply way.

Recommender systems work well if customer preferences can be estimated with enough accuracy. It is easier to achieve this goal if customers are willing to reveal their preferences or at least if they allow retailers to collect enough information about their purchasing habits. This requires (1) a certain level of trust of the customer towards the retail store, and (2) the perception of a potential benefit—the expectation to be better served—. Obtaining customer trust and providing the incentive to reveal information are thus preconditions to be able to use succesfully a recommender system.

At this point, one might ask, "Why would a supermarket shopper want to reveal his preferences or provide recommendations to other shoppers?".

The answer is twofold: first, many shoppers already do happily reveal their preferences, as can be witnessed by the countless number of product reviews written online (e.g., on websites dedicated to reviews such as Epinions or Dooyoo, but also on retail websites such as Amazon, etc.—there are countless other examples). Often, people derive pleasure from communicating to others how they experienced a product—be it a positive or a negative experience. Second, it is possible for retailers to provide incentives to shoppers to reveal their preferences. For example, this is common practice in form of loyalty cards where all purchases are recorded. Moreover, supermarkets may join privacy programmes (e.g., by regularly letting an independent party do an audit and be awarded a privacy seal by the auditor if successful) to alleviate the concerns of privacy-aware shoppers.

To sum up, obtaining the preferences of customers is quite a challenge for the retailer, but it represents also a potential competitive advantage over other stores. The issue is controversial. It has been argued that, in the situation of consumer control and large coverage, "market transparency is expected to increase", but also "customer's loyalty is expected to decrease and their varying choice of store can possibly lead to higher intensity of competition and lower margins" (Sackmann and Strüker 2004). However, one could also argue instead that recommender systems—even in the case of customer control and large coverage—could increase customer loyalty. For example, for products which have associated services, such as mobile phones, price may not be the sole decision criterion for consumers. In this case, retailers could better shape the bundles of services (in the example, the combination of mobile phones and types of contracts) that they offer according to the niches of customers identified by the recommender system. In our opinion, this issue remains open and could be addressed, for example, by simulations of multi-agent models and—possibly in the future—by empirical validations.

4 Technical perspectives

Recommender systems have reshaped the world of e-commerce by enabling on-line retailers to provide services tailored to each customer. Having investigated the business aspects of personalising retail stores, we will now turn our focus to the technical challenges that come along with using recommender systems in retail stores and present an approach to tackle these technical challenges (Walter et al. 2008).

Throughout the past, people always have had visions of how the future could look. In particular, progress in technology has played a core role in some of such predictions. A realistic vision of how shopping could look like in a few years from now is given by the "Extra Future Store" of the Metro Group in Germany (Metro 2007).

> In this grocery store visited by real customers, the Metro Group tests, among others, mini-computers with a touch screen attached to a shopping cart as personal assistants. These devices are equipped with wireless connectivity and a bar-code reader. Future forms of customer interaction will include customers

using mobile electronic devices to communicate with products, enhanced with electronic product code tags and the store's computer system (Sackmann et al. 2006).

The technical infrastructure described in this vision allows retail stores to take the individual situation of each customer into account, i.e. their current position within the store, the products in their shopping carts, and their preferences based on purchasing histories. The retail store can use all this information to offer personalised services through electronic interaction channels such as mobile computers on shopping carts.

Before going into technical details, we will first present our own vision of a personalised retail store which is centred around the use of a recommender system. Subsequently, we will design, in a step-by-step approach, a recommender system tailored to the retail environment and following the *Personalisation Pyramid* (Sackmann et al. 2006). We will match each layer of the personalisation pyramid to an appropriate recommendation technique. Finally, we present a coherent graph model which incorporates all of the different recommendation techniques used at different layers in the personalisation pyramid in one single framework.

4.1 Vision: personalised retail stores

In the following, we will illustrate a scenario in the future which describes, from a customer's perspective, the functionalities of a recommender system:

- Customers use "smart" shopping carts connected to a hand-held computer with a touch-screen interface—a personal shopping assistant (PSA) which assists customers in their shopping tasks. In the store, all products are tagged with radio-frequency identification (RFID) chips that work similar to bar-codes, but can be read wirelessly. The PSA can read RFID chips and thus identify products in the shopping cart and also authenticate customers through their loyalty cards of the retail store. These features can also be used to allow an automatic checkout when the customers are done shopping.
- Customers store their shopping lists on the PSA, which also tries to intelligently complete lists and detect products that customers might have forgotten. The PSA alerts customers when a product that they might like is on sale. Based on the shopping list, the fastest route through the store is computed by the PSA.
- Based on the items that are currently in the cart, the location of the customer, and their purchasing history, the PSA suggests related products.
- Customers are able to rate products and set their preferences with respect to particular types of products, brands, or labels in the PSA. For instance, they could state their trusted fair-trade labels or leave comments if they particularly like/dislike a product.
- Customers may send their friends and family recommendations and share their shopping lists among each other.
- If requested, the PSA can assist customers in finding products with some particular features. It can also provide descriptions of products, such as their

nutrition facts, origin and manufacturing details. This information can be used by customers to compare products.

This illustrates a vision of how a personalised recommender system in a retail store can change the way people shop. In the following, we will discuss the technical building blocks behind an implementation of this vision.

4.2 Design: recommender systems for retail stores

One way of classifying personalisation is the *Personalisation Pyramid*, introduced by Sackmann et al. (2006). It identifies three different levels of personalisation of products and services based on the available information, such as context and personal data about customers: universal services, individualised services, and personalised services. In this section, we show how to map each layer of the personalisation pyramid to an appropriate recommendation technique.

4.2.1 Universal services

The *universal services* require neither personal nor context information about customers. These basic functionalities can be, for example, a lookup of other customers' reviews or further product information. This could be accomplished by an application that stores reviews about products written by fellow customers and a product information database. Such services can be realized without any recommender system and without taking information about the customer's current context into account.

4.2.2 Individualised services

The next layer in the personalisation hierarchy consists of *individualised services*. These services take the current context of the customer into account, but do not require them to reveal their identity. Consider a retail store where the products are tagged with RFID chips and the current position of the customer is available through a device on the shopping cart. In this kind of setting, the supermarket could offer individualised services such as the running total of the items in their shopping cart or recommendations for products on sale that are in proximity to the current location of the customer.

At this level of personalisation, the recommendation technique to be applied is an *item-to-item algorithm*, for example the one presented by Sarwar et al. (2001). Based on the items that the customer is looking at, that are in the shopping cart, or that are on the shopping list, this algorithm recommends items to a customer which all customers tend to purchase together with those items. Figure 4 shows a graph in which each node represents a product and the weight of a link represents the number of co-purchases between the two products it links. For instance, each time a customer has purchased a beer and chips together, the weight of the link between these two products is increased by one. Over time, strong links will develop between products that are purchased together frequently and weak links will develop between products that are not. The recommender system then can recommend items which are strongly linked to one or more items that the customer already knows or considers. This approach makes recommendations based on the assumption that the purchase of one item makes it likely that the customer is willing to purchase another item that is frequently co-purchased with that item. Nowadays, algorithms of this type are applied in many recommender systems, for example in the Amazon.com on-line store: there, recommendations of the form "customers who bought … also bought …" are presented to the customer (Linden et al. 2003).

4.2.2.1 Item-to-item collaborative filtering Item-to-item collaborative filtering (Sarwar et al. 2001) is based on the number of co-purchases between two items. A co-purchasing matrix with the dimensions [|*items*|, |*items*|] contains in each entry (i, j) the number of co-purchases between item i and item j. An example of a subset of a co-purchasing-matrix is presented in Fig. 5; this can be used to generate a list of items that a user might be interested in. For instance, a recommendation algorithm based on the item-to-item approach would recommend a customer that has beer and shampoo in their shopping cart to also purchase chips. This result is obtained by selecting the rows corresponding to items already in the shopping cart (in the example, beer and shampoo) and summing up the values along the columns which represent products not yet purchased (in the example, chips and milk). The column that yields the maximum sum corresponds to the item to be recommended (in this case, chips with 14 co-purchases with beer and shampoo).

Note that the weighted graph illustrated in Fig. 4 is equivalent to the co-purchasing matrix in Fig. 5. An entry in the matrix indicating the co-purchasing between product i and j with value v can be mapped to the graph model by drawing a weighted, undirected edge with value v between the two vertices i and j, and vice versa.

Furthermore, it is straightforward to extend this approach to greater granularity, i.e. in which products are specific brands instead of generic categories as assumed in the example.

4.2.3 Personalised services

A personalised service requires more personal data about the customer to be available. Once one knows a customer's identity, it is possible to match their

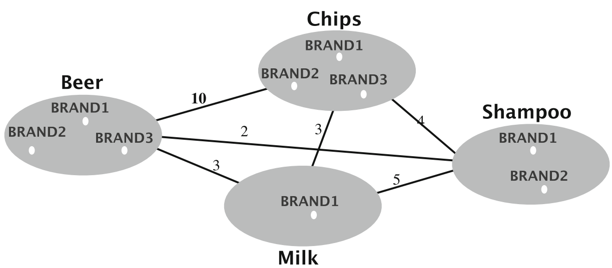

Fig. 4 Identifying Customer Behaviour: graph in which *each node* represents a product and the weight of a link represents the number of co-purchases between the two products it links

	Beer	Chips	Shampoo	Milk
Beer	0	10	2	3
Chips	10	0	4	3
Shampoo	2	4	0	5
Milk	3	3	5	0
	-	14	-	8

Fig. 5 Item-to-item algorithm: co-purchasing matrix and computation of the recommendation (see text for more details)

purchasing profile with those of other customers, form a neighbourhood of similar customers, and finally recommend products which are popular in a target customer's neighbourhood. This is a typical collaborative filtering approach that is also utilised by many e-commerce sites (Linden et al. 2003; Sarwar et al. 2000, 2002).

4.2.3.1 User-to-user collaborative filtering Collaborative filtering algorithms apply the "*word of mouth*" principle to help users filter information. Figure 6 illustrates this analogy in the context of movies:

1. Before seeing a movie, people may ask their friends that have seen the movie how they liked it. Usually, they would ask friends with whom they share a similar taste, under the assumption that their opinion on the unseen movie is likely to be similar, too. In terms of the collaborative filtering paradigm, this implies to find users that highly correlate in taste so that they benefit from each other's opinions.
2. Once, for each member, such a group of users that highly correlates in taste is identified, it is possible to infer an opinion for any user based on the degree of similarity to their group and the group's opinion about the movie.

Note that in collaborative filtering, information is propagated among similar users only. The similarity measure thus represents a filtering step, since the opinions of users that are not similar will not be taken into account. This, of course, requires some criterion on which the similarity between two users can be estimated. Typically, such criteria are established based on items that two users have both rated. Once a neighbourhood of similar users is built, predictions for items that are not known to the active user can be computed by aggregating the ratings of the users in their neighbourhood. Besides generating predictions, new items can be recommended to the active user by scanning the neighbourhood and identifying items that were rated well by neighbours, but not yet used by the active user.

Collecting rating tuples

A rating tuple is a triple containing a user, an item, and the value of the rating of the item by the user. Collecting enough rating tuples is crucial for the neighbourhood formation step. The more rating tuples are available, the better we can match the active user with other users in the dataset. Building neighbourhoods based on only few ratings leads to inaccurate recommendations. The data from the rating tuples can be combined in a $[n, m]$ rating matrix of the ratings of n customers

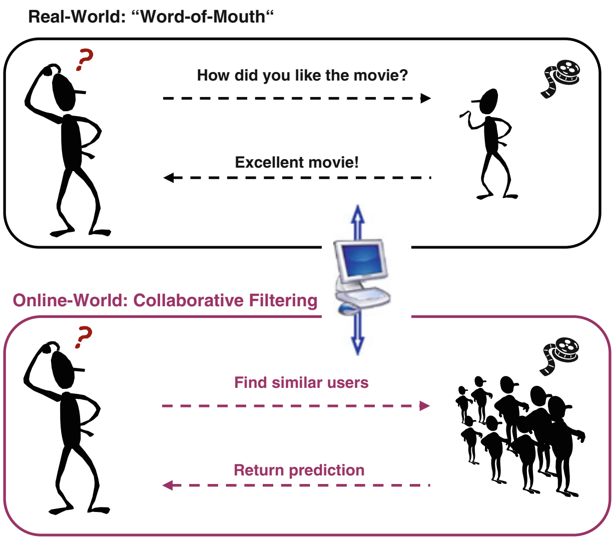

Fig. 6 Mapping of the Real World to the On-Line World: using CF in the on-line world to implement word-of-mouth principle from the real world

on m products, such that the entry $r_{i,j}$ represents the rating that customer i gave to product j.

Building the recommendation model

The next step in collaborative filtering is the formation of a neighbourhood that represents a group of similar users. First, we build a *neighbourhood* N_i of a user i, where the neighbourhood is the set of those other users which rated one or more objects that user i also rated. In that way, one does not consider the whole set of all users in the system, but rather a subset of users which are *comparable* in terms of similarity. Typically, statistical techniques are used to determine similarities between users. Once we have, for each user, identified a group of similar users, different strategies can be employed to generate recommendations.

Next, we can compute the *similarity* ω_{ik} between a user i and its neighbours $k \in N_i$. There are several possible measures of similarity, e.g. based on Pearson correlation, Manhattan distance, or Tanimoto coefficients. For many purposes, though, Pearson correlation, which is defined as follows, is considered one of the most appropriate choices:

$$\omega_{ik} = \frac{\sum_{j=1}^{n}(r_{i,j} - \bar{r}_i)(r_{k,j} - \bar{r}_k)}{\sqrt{\sum_{j=1}^{n}(r_{i,j} - \bar{r}_i)^2 \sum_{j=1}^{n}(r_{k,j} - \bar{r}_k)^2}}, \tag{1}$$

where $\bar{r}_i$ indicates the average rating of user i to objects over all its ratings. Note that $\omega_{ik} \in [-1,1]$, but only ω_{ik} such that $\omega_{ik} \in [0,1]$ are used in the next step. In this way, the values of similarity range from 0 to 1 where a value of 1 means that there is a perfect match of the given ratings by two users a and b (Breese et al. 1998; Herlocker et al. 1999).

Using this, the matrix of similarity values can be constructed as shown in Fig. 7. The matrix has dimensions [|*users*|, |*users*|] where each entry [*i*, *j*] indicates the similarity between the users *i* and *j*. A neighbourhood of size two is obtained by considering the two most similar users in the model. For instance, for Rick, the two most similar users are Jason and Karen, with similarity values of 0.9 and 0.74, respectively. This results in the following neighbourhoods of size two: Jason (Rick, Karen), Karen (Mary, Rick), Mary (Karen, Jason), and Rick (Jason, Karen).

Generating recommendations

Once the similarity neighbourhoods are computed, different techniques can be used to derive a recommendation. A scalable method that delivers good results to come up with a list of recommendable items is to scan the neighbourhood of the current user by performing a frequency count for each item (Sarwar et al. 2000). Once all neighbours are accounted for, the system sorts the items according to their frequency in the neighbourhood, eliminates the ones already known and returns a recommendation list of the remaining most frequent items. A *prediction* $p_{i,j}$ of how a user *i* might rate item *j* based on the ratings of the neighbours *k* and their similarity to the user *i* under consideration can be computed as follows:

$$p_{i,j} = \sum_{k \in N_i} \omega_{ik} r_{k,j}, \quad \omega_{ik} > 0. \tag{2}$$

Since $\omega_{ik} \in (0, 1]$, $p_{k,j}$ will be in the same range as $r_{k,j}$. If there is no neighbour *k* for which $\omega_{ik} > 0$ that has rated product *j*, no prediction can be made for that product *j*.

4.2.4 Trust-based recommender approaches

The discussed methods all were based on collaborative filtering and established similarities between users. Recent research, however, has focussed on also exploiting the fact that users have social network in order to generate recommendations (O'Donovan and Smyth 2005; Walter et al. 2009). This has fuelled research on trust-based recommender systems. Assuming that we can assign a level of trust between a user and some of its peers, either because users themsleves declare it (e.g., as implemented in Epinions.com or Dooyoo.co.uk) or because it can be computed (e.g., as done in Walter et al. 2009), a trust-based recommender system computes predictions as

$$p_{i,j} = \sum_{k \in N_i} T_{ik} r_{k,j}, \quad T_{ik} > 0, \tag{3}$$

(1) User-Item Matrix

	Flipper	Rambo	Star Wars	Titanic
Jason	2	5	5	1
Karen	5	2	4	4
Mary	4	1	4	4
Rick	1	5	4	2

(2) Similarity Matrix

	Jason	Karen	Mary	Rick
Jason	0	0.76	0.71	0.9
Karen	0.76	0	0.99	0.74
Mary	0.71	0.99	0	0.70
Rick	0.9	0.74	0.70	0

Fig. 7 CF model building process

where $T_{ik} \in [0,1]$ reflects the trust between users i and k and N_i the neighbourhood of i defined by trust. Of course, if there is no neighbour k for which $T_{ik} > 0$ who rated product j, a prediction cannot be made for that product j. Notice that formally, this approach leads exactly to the same equation as in collaborative filtering with the essential difference that instead of using similarity between users, this method uses trust between users to generate predictions for the recommender system.

4.2.5 Further issues: importance of providing explanations for recommendations

Herlocker et al. (2000) have presented experimental evidence that users are more likely to accept recommendations when the mechanism behind these is explained to them beforehand. They propose several explanation interfaces that provide transparency about the reasoning behind a recommendation. For example, a histogram of neighbours' ratings for a recommended item can give users more insight into how many of their neighbours rated an item in a particular way than just a number with the average rating.

In particular, such explanations help users to cope with recommendations that appear to be wrong: since collaborative filtering uses statistical mechanisms to compute recommendations based on heuristic approximations of human behaviour, there may be inaccuracies. Even though, usually, most of the recommendations are correct, there may be cases where they are not when the recommender systems is operating on a sparse rating matrix. Explanations provide a means for the customer to put the recommendation in context and to decide when to trust a recommendation and when to doubt one. Thus, explanations help build trust towards the recommender system. Nonetheless, based on the assumption that the rating matrix is not too sparse, collaborative filtering usually is accurate enough to be deployed on a wide scale. Additionally, at the same step in which explanations are generated, recommendations can also be filtered by a set of rules to prevent, for example, that alcohol is recommended to under-age customers.

In the next section, we will introduce a graph-based model to implement all of the presented recommendation techniques in one coherent framework.

4.3 Graph models of recommender systems

In recent years, there have been several attempts to provide a general framework for recommender systems based on graph theory. Graph models provide an abstraction to implement different recommendation techniques within one framework (Aggarwal et al. 1999; Huang et al. 2004; Mirza et al. 2003) by representing relationships between users and items as networks. In this section, we briefly summarise the main concepts of such graph models of recommender systems.

In a recommender system, the rating matrix corresponds to a *rating network*, linking users and items, a link representing the rating of an item by a user. This graph is bipartite, meaning that there are two classes of nodes—users and items—and links can only connect two nodes of different class. In deed, an item can not rate another item or a user another user.

From this graph, it is also possible to build two so-called one-mode projections, one containing only users as nodes and the other containing only items as nodes. In the user projection—in the following, called *user-user network*—, there is a link between two users if and only if they have co-rated at least one item. The weight of the link represents the similarity of ratings of these users according to, for instance, the formula given in Eq. 1. In the item projection—in the following, called *item-item network*—, there is a link between two items if and only if there is one user which has rated both of them. The weight of the link represents the similarity of the ratings across users on the two items; the formula would be analogous to the one given in Eq. 1, but with users and items interchanged.

Additionally, there is a *social network* and a *content network*. These are, again, graphs between only users and items, respectively, but these graphs are not established based on similarity measures, but based on a priori knowledge. In the social network, one could establish a link between two users if they, for example, know each other, trust each other, or have similar demographic features. In the content network, one could establish a link between two items if they, for example, belong to the same category of products (they are substitutes) or if one product goes well with the other (they are complements).

Each of these graphs can be used by a recommender system to generate recommendations. For example, collaborative filtering is based on the rating network and the user-user or item-item network (see Sect. 4.2.3); trust-based recommender systems (Walter et al. 2008) use social networks; content-based recommender systems use the content network (Huang et al. 2004). Which graph to use can also be an issue of performance: the item-item network usually is significantly smaller than the rating network and still can provide quite good performance as long as no personalisation is required (see Sect. 4.2.2). Hybrid approaches that combine two or more of these graphs are possible (Huang et al. 2004) and, in some cases, lead to better performance. Figure 8 shows the social network, the rating network, and the content network of the running example used in the paper so far; the item-item network and the user-user network are not represented.

In the next section, we present an architecture which could be used as a blueprint to implement such graph models of recommender systems.

4.4 Layered architecture for recommender systems

Given this abstraction of recommender systems as models operating on different graphs, let us give an idea of a potential architecture for recommender systems that enables the application of all prevalent recommendation techniques within one single framework: LARS, a layered architecture for recommender systems (Yildirim et al. 2011), and a step towards a unified framework for recommender systems.

As we have seen in the previous sections, there are several approaches for generating recommendations, each with its own benefits and shortcomings. To overcome the drawbacks of using only a single recommendation algorithm we have developed LARS, a hybrid architecture for recommender systems that is organized in layers. LARS enables us to unify various recommendation methods in one single

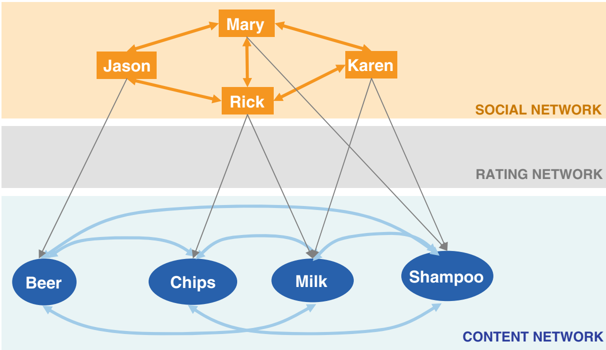

Fig. 8 Graph model of recommender systems: example of a social network, a rating network, and a content network. The item-item network and the user-user network are not represented

framework. The key feature of LARS is that it is completely *modularized*, i.e. the system is decomposed into small components, each delivering a particular, atomic piece of functionality. These modules are hierarchically ordered into *layers*, where each layer provides services to the layer above it by using services from the layers below it. Each layer is independent of the other layers; as long as a commonly defined interface to other layers remains the same, the implementation of a layer can be changed as necessary.

Figure 9 gives an overview of the architecture of LARS. There are four layers:

- A *rating layer* that is responsible for capturing the ratings that users give to items. The rating layer acts as a data storage and retrieval mechanism; its main purpose is to allow to save and access rating information in a straightforward and uniform manner. For example, the rating layer could transparently store and retrieve ratings through a database without the other layers being aware that a database is being used.
- The *statistics layer* provides the statistical quantities that are needed by the layers upstream. This include for each user the average rating he gave across all rated items, a quantity that is then used in the computation of the weights in Eq. (1). Similarly, it is also useful for each item its average rating across the users. Thus, the output of the statistics layer can be represented as two sets of tuples: [$UserID_i$, $UserStat_i$], where $UserStat_i$ is a vector of statistical quantities related to user i. [$ItemID_k$, $ItemStat_k$] is the analog tuple for item k.
- A *weight layer* that establishes weights between entities in the system: for example, weights between users (items) could reflect the similarity between these users (items) and can be used to establish neighbourhoods of similar users (items). Thus, for each relation between two objects, users or items, the output is the tuple $ObjectID_i$, $ObjectID_j$, $Weight_{ij}$, where $Weight_{ij}$ is a vector. Each component $Weight_{ij}^m$ of this vector is a value of the weight of the relation between objects i and j computed according to method m. For instance, method 1 could be the user similarity described by Eq. (1), while method 2 could be the level between two users as described in Sect. 4.2.4. Depending on the type of

relation and the type of weight that needs to be computed, a different input maybe necessary. For instance, method 1 needs only to access the rating tuples from the rating layer. In contrast, method 2 needs to access trust data in a social network.
- A *prediction layer* that uses the information provided by the rating, statistics, and weight layers to generate recommendations. For example, to make a recommendation on how a particular user *u* might like an item *i*, the prediction layer can aggregate the ratings of other similar users in the neighbourhood of *u* weighted by the degree of similarity between *u* and the neighbours.

Given the structure above, the prediction layer is able to provide predictions according to various methods. It is obviously possible to modify the formulas and the methods as long as there no changes in the interface, i.e. in the structure of the variables provided as output by each layer. Of course, in case a new method requires a new variable to be computed, this change propagates backwards in the lower level layers, which have to be adapted. However, the overall structure remains the same and the system is able to provide recommendations no matter where the rating data are coming from, nor what items are rated, as long as the inputs to the rating layer and the weight layer are in the correct format.

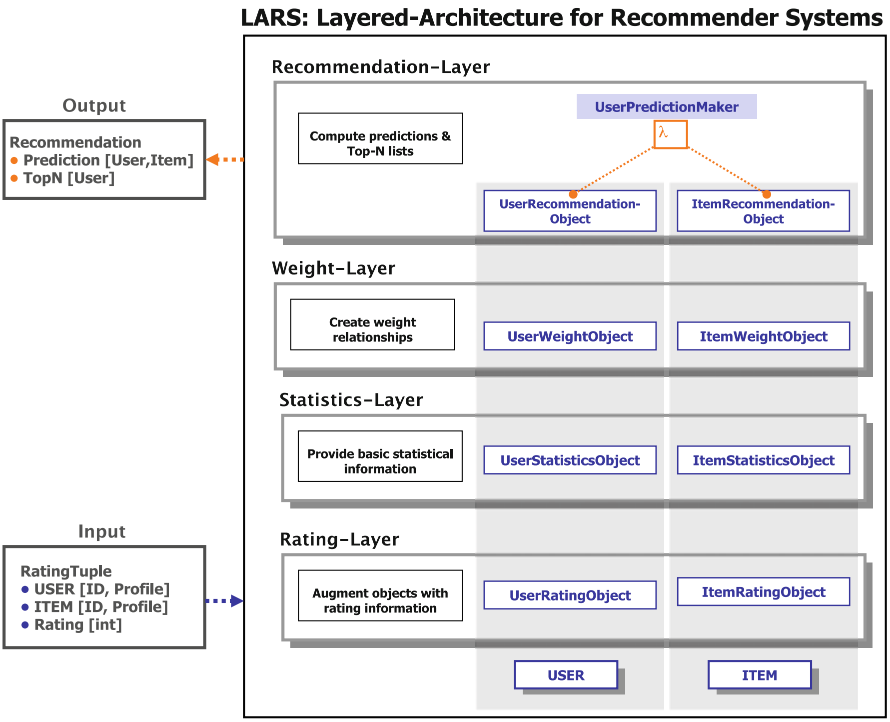

Fig. 9 Architecture of LARS

Thus, the core benefit of the layered architecture is that it is very flexible with respect to changes to the recommender system: for example, employing different prediction strategies would imply to only replace the prediction layer (while all other layers remain the same). This allows us to easily change the prediction strategy and at the same time to re-use existing code from other layers. Furthermore, the weight layer hides the information how the weights are computed among users or items. This is also appropriate, since there are many ways of establishing similarity between users or items. For instance, in CF, the two most popular similarity measures are the Cosine-similarity metric and the Pearson-similarity metric, but many more are suitable depending on the context of the recommender system (Adomavicius and Tuzhilin 2005; Herlocker et al. 1999, 2004; Sarwar et al. 2000). We could provide a customised module for each similarity metric and use them, no matter what is done in other layers. Equally, the statistics layer, can provide all sorts of statistical information that might be useful for other modules. Finally, the rating layer encapsulates access to the rating information which has to be aggregated in some form in a central location—otherwise, all other layers would need to have their own ways of storing and retrieving rating information which would lead to duplicate functionalities inside the system. Thus, in summary, a change of a particular aspect of the implementation of the recommender system usually corresponds to a change in one particular module. The layered architecture allows to make such changes to particular components of the system independently of the rest of the system, providing a great deal of flexibility. Overall, this architecture encompasses virtually any recommender system method which is based on a rating matrix and some similarity measure among users.

Certainly, when implementing a recommender system for a specific purpose, one has to address the issue of designing the appropriate interface the database model and the rating layer, as well as between the prediction layer and the final user. Here we do not provide a detailed implementation scheme but a proof-of-concept.

5 Conclusion

The continuous progress in information technologies and ubiquitous computing has made it feasible to move applications from the virtual world to the real world. Nowadays, it is, for instance, possible to equip traditional retail stores with technology that can serve as infrastructure for recommender systems similar to the ones used in on-line stores. In particular, personalisation strategies can be put into practice by using the building blocks that we have outlined in this article.

Recommender systems offer a tool to cope with the challenges that today's retail stores have to face. Both retail stores and customers can benefit from recommender systems, but, in this paper, we have focused on the side of retail stores and analysed the main business as well as technical aspects arising from moving recommender systems from on-line stores and applying them to retail stores. There are a number of issues related to consumers, such as those arising from the coverage of and the control over recommender systems, which we will investigate in future work. For example, an open issue is whether recommender systems increase or decrease loyalty among customers. We have investigated which types of retail store would benefit from a personalisation strategy and to which extent. From a more technical perspective, we have shown how to build a recommender system tailored to the retail environment. We have also illustrated recommender systems in the framework of graph models, a useful abstraction that encompasses consumer ratings, product features, and social networks. Embarking from this perspective, we have presented LARS, a layered architecture for recommender systems which can serve as a blueprint for real-world implementations.

In the past, recommender systems have successfully been deployed in e-commerce settings. In the future, they have the potential to also be successfully deployed in the retail world, potentially leading to the next generation of retail stores.

Acknowledgments We would like to thank Elgar Fleisch, Florian Michahelles, and Dirk Martignoni for their fruitful suggestions on drafts of this paper.

References

Abdul-Rahman A, Hailes S (2000) Supporting trust in virtual communities. In: Proceedings of the 33th annual Hawaii international conference on system sciences. IEEE Press

Adomavicius G, Tuzhilin A (2005) Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. IEEE Trans Knowl Data Eng 17(6):734–749

Aggarwal CC, Wolf JL, Wu K-L, Yu PS (1999) Horting hatches an egg: a new graph-theoretic approach to collaborative filtering. In: Proceedings of the fifth ACM SIGKDD international conference on knowledge discovery and data mining (KDD '99). ACM, New York, NY, USA, pp 201–212. doi: 10.1145/312129.312230

Anderson C (2006) The long tail: how endless choice is creating unlimited demand. Random House, New York

Breese JS, Heckerman D, Kadie C (1998) Empirical analysis of predictive algorithms for collaborative filtering. In: Proceedings of the fourteenth conference on uncertainty in artificial intelligence (UAI-98), pp 43–52

Brynjolfsson E, Hu YJ, Smith MD (2006) From niches to riches: anatomy of the long tail. Sloan Manag Rev 47(4):67–71

Cinicioglu EN, Shenoy PP, Kocabasoglu C (2007) Use of radio frequency identification for targeted advertising: a collaborative filtering approach using Bayesian networks. Lect Notes Comput Sci 4724:889

Decker C, Kubach U, Beigl M (2003) Revealing the retail black box by interaction sensing. In: ICDCSW '03: proceedings of the 23rd international conference on distributed computing systems, IEEE Computer Society, p 328

Dowling GR, Uncles M (1997) Do customer loyalty programs really work? MIT Sloan Manag Rev 38(4):71–82

Ernst & Young (2003) Händler am Scheideweg

Ernst & Young (2005) Consumer trends report

Fleisch E (2001) Business perspectives on ubiquitous computing. Technical report, M-Lab Working Paper

Golbeck J (2005) Computing and applying trust in web-based social networks. PhD thesis, University of Maryland at College Park

Grandison T, Sloman M (2000) A survey of trust in internet applications. IEEE Commun Surv Tutorials 3(4):2–16

Herlocker JL, Konstan JA, Borchers Al, Riedl J (1999) An algorithmic framework for performing collaborative filtering. In: SIGIR '99: proceedings of the 22nd annual international ACM SIGIR conference on research and development in information retrieval. ACM Press, pp 230–237

Herlocker JL, Konstan JA, Riedl J (2000) Explaining collaborative filtering recommendations. In: Proceedings of the 2000 ACM conference on computer supported cooperative work (CSCW '00). ACM Press, pp 241–250
Herlocker JL, Konstan JA, Terveen LG, Riedl JT (2004) Evaluating collaborative filtering recommender systems. ACM Trans Inf Syst 22(1):5–53
Huang Z, Chung W, Chen H (2004) A graph model for E-commerce recommender systems. J Am Soc Inf Sci Technol 55(3):259–274
Kaufman PR (2000) Consolidation in food retailing. Economic Research Service/USDA
Kim JH, Lee ES (2005) User XQuery pattern method based personalization recommender service. In: First international conference on semantics, knowledge and grid, Beijing, 99 p. doi: 10.1109/SKG.2005.137
Kitts B, Freed D, Vrieze M (2000) Cross-sell: a fast promotion-tunable customer-item recommendation method based on conditionally independent probabilities. In: KDD '00: proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining. ACM Press, pp 437–446
Kotler P, Armstrong G (2004) Principles of marketing. Prentice Hall, Englewood Cliffs
Kourouthanassis P, Roussos G (2003) Developing consumer-friendly pervasive retail systems. IEEE Pervas Comput 2(2):32–39
KPMG (2006) Trends im Handel 2010
Krohn A, Zimmer T, Beigl M, Decker C (2005) Collaborative sensing in a retail store using synchronous distributed jam signalling. In: Pervasive computing. Springer, Berlin, pp 237–254
Lam SK, Riedl J (2004) Shilling recommender systems for fun and profit. In: WWW '04: proceedings of the 13th international conference on world wide web. New York, NY, USA, ACM, pp 393–402
Lawrence RD, Almasi GS, Kotlyar V, Viveros MS, Duri SS (2001) Personalization of supermarket product recommendations. Data Minining Kowl Discov 5(1–2):11–32
Linden G, Smith B, York J (2003) Amazon.com recommendations: item-to-item collaborative filtering. IEEE Internet Comput 7(1):76–80
Liu Y-H, Yih J-S, Chieu TC (2004) A personalized offer presentation scheme for retail in-store applications. In: E-commerce and web technologies. Springer, Berlin, pp 296–304
Marsh S (1994) Formalising trust as a computational concept. PhD thesis, University of Stirling
Massa P (2006) Trust-aware decentralized recommender systems. PhD thesis, Università degli Studi di Trento
Mentasys (2007) Guided selling solutions company, case studies
Metro Group (2007) Future store initiative
Mirza BJ, Keller BJ, Ramakrishnan N (2003) Studying recommendation algorithms by graph analysis. J Intell Inf Syst 20(2):131–160
Montaner M, López B, de la Rosa JL (2002) Developing trust in recommender agents. In: Gini M, Ishida T, Castelfranchi C, Johnson WL (eds) Proceedings of the 1st international joint conference on autonomous agents and multiagent systems (AAMAS'02). ACM Press, pp 304–305
Montaner M, López B, de la Rosa JL (2002) Opinion-based filtering through trust. In: Proceedings of the 6th international workshop on cooperative information agents (CIA 2002). Springer, pp 164–178
Mundt K, Almquist E, César J (2002) Profitable retailing in a zero-sum game. Mercer Manag J (14):60–69
Newman MEJ, Watts DJ, Strogatz SH (2002) Random graph models of social networks. Proc Natl Acad Sci USA 99(90001):2566–2572
O'Donovan J, Smyth B (2005) Trust in recommender systems. In: Proceedings of the 10th international conference on intelligent user interfaces (IUI '05). ACM Press, pp 167–174
Prudsys AG (2006) Prudsys recommendation engine goes online at Quelle.de
Reichheld FF (2003) The one number you need to grow. Harvard Bus Rev 81:46–54
Resnick P, Iacovou N, Suchak M, Bergstrom P, Riedl J (1994) Grouplens: an open architecture for collaborative filtering of netnews. In: CSCW '94: proceedings of the 1994 ACM conference on computer supported cooperative work. ACM Press, pp 175–186
Rigby DK, Vishwanath V (2006) Localization—the revolution in consumer markets. Harvard Bus Rev 84:82–92
Sabater J, Sierra C (2005) Review on computational trust and reputation models. Artif Intell Rev 24(1):33–60
Sackmann S, Strüker J, Accorsi R (2006) Personalization in privacy-aware highly dynamic systems. Commun ACM 49(9):32–38

Sarwar B, Karypis G, Konstan J, Riedl J (2002) Recommender systems for large-scale e-commerce: scalable neighborhood formation using clustering. In: Proceedings of the fifth international conference on computer and information technology

Sarwar B, Karypis G, Konstan J, Riedl J (2000) Analysis of recommendation algorithms for E-commerce. In: Proceedings of the 2nd ACM conference on electronic commerce (EC-2000). ACM Press, pp 158–167

Sarwar BM, Karypis G, Konstan JA, Riedl J (2001) Item-based collaborative filtering recommendation algorithms. In: World wide web, pp 285–295

Schafer JB, Konstan JA, Riedl J (2001) E-commerce recommendation applications. Data Min Knowl Discov 5(1/2):115–153

Schröder H, Feller M, Zimmermann G (2003) Retail-Studie. Mercer Management Consulting

Sackmann S, Strüker J (2004) Success factors for electronic customer communication in brick-and-mortar retailing. In: Proceedings of the third international mobile business conference

Walter FE, Battiston S, Schweitzer F (2008) Coping with information overload through trust-based networks. In: Helbing D (ed) Managing complexity: insights, concepts, applications. Springer, Berlin, pp 273–300

Walter FE, Battiston S, Schweitzer F (2008) A model of a trust-based recommendation system on a social network. J Auton Agents Multi Agent Syst 16(1):57–74

Walter FE, Battiston S, Schweitzer F (2009) Personalised and dynamic trust in social networks. In: RecSys ’09: proceedings of the third ACM conference on recommender systems. ACM Press, pp 197–204

Wang Y-F, Chuang Y-L, Hsu M-H, Huan-Chao K (2004) A personalized recommender system for the cosmetic business. Exp Syst Appl 26(3):427–434

Weisbuch G, Kirman A, Herreiner D (2000) Market organisation and trading relationships. Econ J 110:411–436

Weiser M (1991) The computer for the 21st century. Sci Am 265(3):94–104

Yildirim M, Walter FE, Battiston S, Schweitzer F (2011) Towards a unified framework for recommender systems (under preparation)