

Towards Portable Shopping Histories: Using GoodRelations to Expose Ownership Information to E-Commerce Sites

László Török and Martin Hepp

Universität der Bundeswehr
Werner-Heisenberg-Weg 39.
D-85579 Neubiberg, Germany
{laszlo.toeroek,martin.hepp}@unibw.de

Abstract. Recommender systems are an important technology component for many e-commerce applications. In short, they are technical means that suggest potentially relevant products and services to the users of a Web site, typically a shop. The recommendations are computed in advance or during the actual visit and use various types of data as input, in particular past purchases and the purchasing behavior of other users with similar preferences. One major problem with recommender systems is that the quality of recommendations depends on the amount, quality, and representativeness of the information about items already owned by the visitor, e.g. from past purchases at that particular shop. For first-time visitors and customers migrating from other merchants, the amount of available information is often too small to generate good recommendations. Today, shopping history data for a single user is fragmented and spread over multiple sites, and cannot be actively exposed by the user to additional shops.

In this paper, we propose to use Semantic Web technology, namely GoodRelations and schema.org, to empower e-commerce customers to (1) collect and manage ownership information about products, (2) detect if a shop site is interested in such information in exchange for better recommendations or other incentives, and (3) expose the information to such shop sites directly from their browser. We then sketch how a shop site could use the ownership information to recommend relevant products.

Keywords. #eswc2014Torok, Semantic Web, Recommender Systems, E-Commerce, schema.org, GoodRelations, RDF, RDFa, Microdata

1 Introduction

Recommender systems are an established part of e-commerce systems. They help prospective customers to navigate the myriad of products offered in an online shop and aid them with finding a product that best matches their needs and preferences. Many recommender systems require or benefit from data about past purchases or items ratings. This means that they perform well if the user has already a shopping history in that particular system, and perform poorly when facing a first-time user that

has not yet rated or purchased an item in a given shop. From the user's perspective, her rating and purchase history is scattered across many shop systems, which she has visited in the past. Accessing and sharing one's entire shopping history, or a subset thereof, could likely provide better personalization and better recommendations.

We also argue that purchase records are not necessarily representative when it comes to preferences. For instance, people do not always buy products for themselves. A one-time purchase of a cosmetic product for the wife does not mean that a husband has a longer-term personal interest in it. Therefore, we focus on collecting, managing and sharing **actual item ownership information**, as we believe it is a better indicator of personal preferences.

One of the key problems that prohibit a wider availability of ownership information is that parts of one's shopping history are locked up in multiple, proprietary data representations. Our goal is to develop a common data representation and exchange protocol that could improve existing recommender systems and open up new possibilities for shop systems for better understanding their customers.

1.1 Role of Item Ownership Information in Recommender Systems

Most contemporary recommender systems do not rely on actual data about items being owned by users; instead, they use past purchases, or items viewed previously, to infer likely ownership; for an overview, see e.g. [11]. In particular, Collaborative Filtering (CF) algorithms have been successfully applied to the problem of product recommendation. The numerous existing CF algorithms share two common properties: (1) they maintain a matrix of item ratings \mathbf{R} or purchases where a $R_{i,j}$ entry represents that a $User_i$ rated/purchased $Item_j$ and based on this (2) calculate a ranked list of top-N items that might appeal to a given user [1]. Our focus in this paper is the use of ownership information, hence we limit our discussion to past purchases, although CF algorithms apply to item ratings and other forms of user-item interaction as well.

Let $User_c$ be the current user for whom the recommended list of items should be computed. The user-based CF algorithms employ clustering techniques over \mathbf{R} to find other users that have similar preferences to $User_c$, that is, their purchase record has a significant overlap with that of $User_c$ [cf. 1, 11].

Another flavor of CF algorithms are item-based CF algorithms. The core of these algorithms is the item-similarity matrix \mathbf{S} , where $S_{i,j}$ denotes a similarity between $Item_i$ and $Item_j$ [1]. The item-similarity matrix is derived from the frequency of purchasing or positively rating two particular items for each item pair [1, 11]. The output is a ranked list of items likely to be purchased by $User_c$.

Markov-chain based (a.k.a next-basket prediction) methods [6] model purchasing as a probabilistic sequential process and attempt to predict the next set of items likely to be purchased by the user. In other words, they consider the temporal relations between purchases of items. The explicit cause of the next-item relevance remain typically undefined even in semantically augmented CF approaches, such as [4], where a product taxonomy is employed in order to better capture consecutive purchases of related items, for example a camera, followed by a tripod and an additional camera lens. In other words, these advanced models may be able learn that customers who

purchase a camera are likely to purchase a tripod, from training data. However, they do not reveal the fact that the reason behind the two being frequently purchased together is that the tripod is an accessory of a video camera.

More complete or more detailed ownership information could pave the way to novel product recommendation techniques that will rely on semantic relations between products. Recommender systems could e.g. utilize functional dependencies between items in certain product recommendation scenarios, for example, directing the user to a compatible toner cartridge for the laser printer that she purchased previously. Alternatively, a recommender system informed by popular preference structures could infer that someone who owns a yoga mattress is likely to be interested in organic food, because both is related to an interest in a healthy lifestyle.

1.2 Impact of Additional Item Ownership Information

As we have shown in 1.1, most personalized recommender systems are driven by past purchases. However, a Web shop only holds a fraction of a user's online purchase record, as it is only aware of items purchase there (see Figure 1).

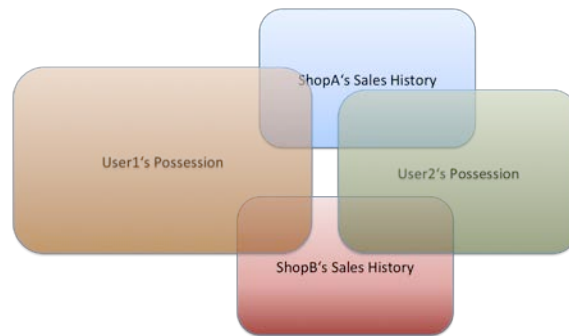


Fig. 1. Owned items versus purchase records

The ability alone to actively share purchases made at other shops could improve the performance of recommender systems. For example, they could help mitigate the cold-start problem [cf. 5] for new customers or alleviate the data sparsity problem in collaborative filtering approaches [cf. 1]. In this paper, we argue that ownership information offers additional insight into the user's preferences than pure data about past purchases. Specifically, we will develop a data model to express which items the user currently owns and which items she did own in the past. This level of detail allows one to filter out items purchased for someone else and offers hints on how long an item or a group of items were in possession of the user. Given this level of detail, someone making a one-time purchase for his spouse can be spared from constantly receiving recommendations about further products of a similar kind.

1.3 Portability Barriers on Item Ownership Information

Currently, the exchange of information about owned items between users and Web sites is hampered by the following limitations:

- (1) Purchasing records as structured data are mostly available only within shop site applications but not on the machines of users. While users receive purchasing confirmation and invoices, those are typically only unstructured text¹.
- (2) There is no common data model for representing and sharing ownership information.
- (3) There is no common protocol for initiating and governing the exchange of item ownership information between users and Web sites.

Currently, only product model master data is exchanged between partners in value chains, mostly via XML-based product catalogs, and product models of commodities are identified via product identifiers like the standardized Global Trade Item Number² (GTIN). There is, however, no standard way of representing, managing, and sharing item ownership information from a user's perspective.

1.4 Our contribution

In this paper, we describe a **conceptual model, an RDF-based syntax, and a protocol** that allow users to actively share information about items owned with Web applications, thus empowering the latter to provide better recommendations even for first-time users. Our paper is structured as follows: In Section 2, we present a motivating scenario, and develop the conceptual model and its machine-processable representation based on schema.org in combination with an exchange protocol. In Section 3, we describe two relevant product recommender scenarios based on ownership information. In section 4, we provide preliminary evidence for the viability and relevance of our proposed method.

2 User-managed Ownership Information

2.1 Our Approach

Our goal is to augment the typical interaction pattern between a user and a Web shop by the ability to share information about items owned by the user in exchange for a more personalized shopping experience or other incentives, as sketched in Figure 2. To enable this interaction pattern, both the user's Web browser and the shop system must support (1) the common data model and (2) the exchange protocol described in 2.4.

¹ This may improve by the availability of the support for transactions in schema.org, <http://schema.org/Order>, and schema.org markup in JSON-LD in email messages, see <https://developers.google.com/gmail/actions/>.

² Standardized by the GS1 standards body

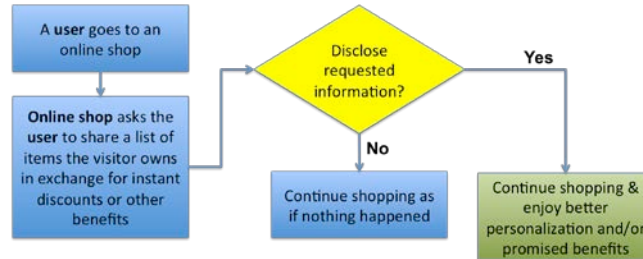


Fig. 2. High-level interaction pattern for exchanging item ownership information

2.2 Modeling Ownership Information

In the following, we first develop a conceptual model for capturing item ownership information, and then show how it can be represented in various RDF syntaxes.

Requirements. We express the requirements on the conceptual model in the form of competency questions, which is a common approach in ontology development [7]. The competency questions arose from our study of recommender system literature and novel rule-based recommendation techniques that use on granular semantic product relationships, as discussed in Section 3.

- Q1: Which items does the user currently own?
- Q2: Which items did the user ever own in the past?
- Q3: During which period of time has a certain user owned a certain item?
- Q4: Why has a user ceased to own a certain item? (e.g. sold, broken, superseded)
- Q5: From whom did a user acquire a certain item?
- Q6: How/for how much was the item acquired?
- Q7: What class/category does the item belong to?

Conceptual Model for Item Ownership Information. We propose the following set of entity and relationship types: An *Agent* represents a real person or an organization that can possess arbitrary items. An *Item* stands for arbitrary physical or electronic goods (e.g. music) that can be owned by an *Agent*. A natural approach is to employ a binary relation “owns” to express basic ownership information, however, such simplistic approach has a very limited expressivity.

Following our competency questions, we extend the binary owns relation to capture more granular information of an ownership relation (see Figure 3). The attributes *from* and *to* cover the time interval during which an *Item* was owned by an *Agent*. These are particularly valuable for non-perishable goods that are intended for long-term use. In certain scenarios, it can be valuable to know from whom the *Agent* did acquire a given *Item*. This can be directly expressed via the *acquired from* relation. As the most common scenario for an acquisition is accepting a commercial offer of a merchant, we will reuse some of the conceptual elements of the GoodRelations and schema.org vocabulary to provide further details (e.g. purchase price) for the acquisition. In order to express the circumstances of ceasing ownership of an item, we define

the relation *disposal*, which can point to one or more elements from a set of predefined values: *Sell*, *Broken*, *Donation*, and *Superseded*.

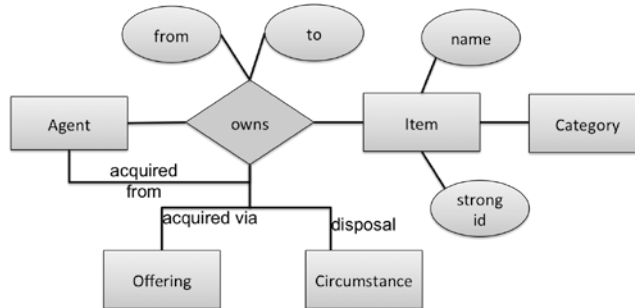


Fig. 3. Conceptual model of item ownership information

Representation Using Semantic Web Formalisms. As discussed in 2.2, we need a machine-processable representation of our conceptual model to make data adhering to our model suitable for computer-based processing. We choose RDF³ and established Semantic Web vocabularies to implement our model. The utility of common data schemas and protocols depends on the number parties that support it. For this reason, we adopt established Semantic Web vocabularies such as **Schema.org**⁴ and **GoodRelations**⁵ [3]. Schema.org is a Web vocabulary covering common topics in Web publishing, such as events, news and product offers that can be used to express simple semantic structures of Web page content. The GoodRelations Web vocabulary has been designed to express fine-grained facts in the e-commerce domain and has been recently integrated into the schema.org effort, backed by major search engines. RDF's unit of information is an entity-attribute-value tuple, commonly referred to as triple. Multiple standardized concrete syntaxes have been developed for RDF that are conceptually equivalent. In this paper, we settled upon using the human-friendly, legible Turtle⁶ notation.

Basic Ownership Information. Initially, GoodRelations (now part of schema.org) only allowed expressing ownership information as a binary relation between an agent and an item. For instance,

```

<http://alice.me/#i> <http://schema.org/owns>
<http://dbpedia.org/resource/Mona_Lisa>.
  
```

encodes the statement that Alice, identified as `<http://alice.me/#i>` owns the Mona Lisa painting. The advantage of this approach is its simplicity, as it requires only a single binary attribute, yet it has very limited expressive power.

³ Resource Description Framework (<http://www.w3.org/standards/techs/rdf>)

⁴ <http://schema.org/>

⁵ <http://purl.org/goodrelations/>

⁶ <http://www.w3.org/TeamSubmission/turtle/>

Granular Ownership Record. As RDF only allows binary relations, for N-ary relations we need to include an additional element `s:OwnershipInfo`. This element has been recently added to schema.org following a proposal by the authors of this paper, along with the attributes `s:typeOfGood` referring to the owned Product; `s:ownedFrom`, `s:ownedTo` denoting temporal bounds of ownership and the attribute `s:acquiredFrom` pointing to the source of the item.

```
@prefix s: <http://schema.org>.
<http://alice.me/#i> s:owns [
  a s:OwnershipInfo;
  s:typeOfGood <http://alice.com/mylaptop>;
  s:acquiredFrom <http://amazon.com/#company>;
  s:ownedFrom "2011-11-09T00:00:00";
  s:ownedThrough "2013-10-01T00:00:00" ].
```

2.3 Data Management

One open issue with handling ownership information is that we need to make sure that the information about items owned stays in sync with the purchasing and disposal of items, and we need to organize the initial exchange, and update of previously shared, information with sites. In other words, we need to address how data is acquired, stored and edited, and shared. These are responsibilities of the client implementing the data model and the exchange protocol described in this paper. Our **prototype client**⁷ currently supports manual form-based entry and the import of Amazon.de purchase history data. Other viable sources are extracting data from purchase receipts. These are usually poorly structured and manual intervention is likely necessary. As for storing the data, multiple options are available, either storing it locally or storing it remotely. Our reference implementation uses the browser's local storage, since this can be accessed from a browser extension and does not involve tackling access control, as it would be in a remote-storage scenario. Other capabilities, such as sharing ownership information selectively are also part of the client implementation, hence independent from our data model and protocol.

2.4 Protocol for Exchanging Ownership Information

We define a minimal protocol for exchanging over HTTP. Our only concern here is to support data exchange. How the recipient will actually use ownership information (OI) is not in the scope of our protocol. The protocol consists of three abstract phases:

1. **Discovery:** A Web shop advertises interest in receiving ownership info and a capable client detects this intent.
2. **Authorization:** The user decides whether item information should be shared.

⁷ <http://demo.portable-shopping-history.info/>

3. **Sharing:** If step 2 succeeds, the browser sends ownership information to the shop.

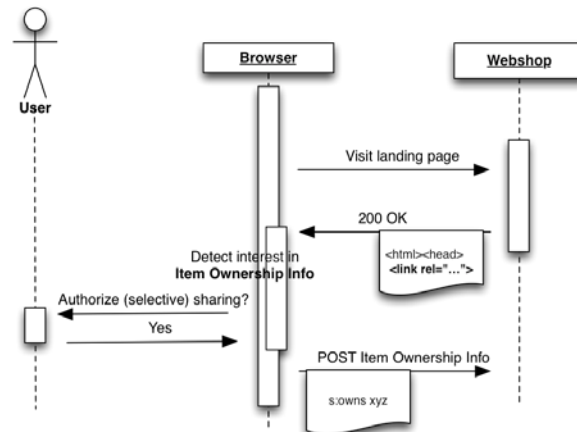


Fig. 4. Exchange of item ownership information

A Web shop may advertise its interest in OI by adding the following link element to the `<head>..</head>` of the landing page

```

<link rel="http://purl.org/xventory/sink"
media="application/json"
href="http://laptopcase.biz/personalize">
  
```

where `<http://purl.org/xventory/sink>` is a unique identifier which indicates that the shop is able to receive OI sent via HTTP to the URL denoted by the `href` attribute. The interaction can only take place if both parties adhere to the simple rules above. It will not break or impede the usual browsing sequence if one of the parties does not support the protocol.

3 Product Recommendation Using Item Ownership Information

In Section 1.2 we have argued that additional item information received from external sources can complement current recommender systems. The "new-user" problem and the data sparsity present in many recommender systems from a low number of past purchases can be mitigated in a straightforward way, should the user choose to share his list of belongings with the shop. Let `ItemsUser` denote the set of items in the user's possession and `Itemsshop` the items known to the shop. Taking `ItemsUser` `Itemsshop` serves as a valuable additional input for the next time the item-similarity matrix S is recomputed.

In order to provide some evidence for the further utility of item ownership information for a more personalized shopping experience, we will present two rule-based scenarios that rely on the availability of rich product data in a shop system. In the following, we will implement the recommender rules as SPARQL queries to demonstrate interesting inferences, as they are high level, declarative, yet directly executable. Depending on the problem scale, i.e. the number of items involved, real-time requirements, a real-world implementation may have to use optimizations.

Our process for obtaining useful recommendations is as follows:

1. Request product ownership information.
2. Add it to the shop's database.
3. Run recommender rules formulated as SPARQL queries over the new, integrated knowledge base.

3.1 Interpreting Transmitted Item Information

Although our data model places both the client and the shop system in a common frame of reference with respect to data schema, there is no single canonical reference to items that would serve as identity check. Product names or labels (given by `s:name`) prove rather unreliable for this purpose due to their variability and language dependence. There could be two slightly different product labels or two semantically equivalent labels in two different languages referring to the same product. These are the same issues that arise when merging data originated from two different sources.

Therefore, in our following scenarios, we only consider items received from the client for which a strong identifier is provided, such as GTIN13, which is globally unique for commodities. If such a strong identifier is available, establishing a link to a known product in the database is efficient and trivial.

Alternatively, product classification information referring to a widely deployed taxonomy, such as the Google Product Taxonomy⁸ can provide valuable insights into the user's interest.

3.2 Related Product Recommendation

Due to the increasing number of specific products available in a single shop, it is not always easy for the non-expert user to choose the correct spare part or compatible product for an already owned product. High quality product catalogs typically define these relationships, so they can be used to aid the user in her search for the right product. `schema.org` defines `s:isConsumableFor` and `s:isAccessoryOrSparePart` attributes to denote such relationships between two products, which we use in our rule.

Find accessories of products owned by the user (R1)

```
PREFIX s: <http://schema.org/>
```

⁸ <http://www.google.com/basepages/producttype/taxonomy.en-US.txt>

```

PREFIX : <http://mysemanticshop.com>
SELECT ?itemName ?relatedItemName WHERE {
# try all properties that refer to common
# globally unique ids
VALUES ?strongIdProperty { s:gtin8 s:gtin13 s:gtin14 }
# using the basic binary ownership property
?customer s:owns ?itemOwned.
?itemOwned ?strongIdProperty ?productId;
s:name ?itemName.
# find item in shop
?itemShop ?strongIdProperty ?productId.
?relatedItem
s:isAccessoryOrSparePartFor ?itemShop;
s:name ?relatedItemName;
:rating ?rating.
# rank by popularity, :rating is a simple
# numeric attribute
} ORDER BY DESC(?rating)

```

R1 presents a SPARQL query that retrieves all products that are suitable accessories or spare parts for any recognized item in the user's possession. For example, it will retrieve paper bags that are compatible with the user's vacuum cleaner or recommend a laptop bag for the user's 13" laptop. Ranking the result set can be performed using arbitrary criteria. In our example, we resort to ranking the results by item popularity.

3.3 Successor Product Recommendation

Recommending successor products is a common marketing tactic. In schema.org, the `s:successorOf` property denotes a successor relationship between two products. Assuming the user is in possession of a Phone123 smartphone and the shop system has the fact

```
:phone124 s:successorOf :phone123.
```

available, recommending an upgrade to Phone124 may prove valuable to the user. Collaborative filtering system will typically capture the correlation between a Phone123 and Phone124, as users interested in Phone123 may likely be interested in Phone124, too. However, a user, who has just purchased a Phone123 device is less likely to find a Phone124 recommendation useful [cf. 8].

Our semantic recommender rule is able to make a better decision based on knowing how long a given item has been in the possession of the user. Assuming that the shop system has data on the average product lifetime, that is, the average period after a product is replaced, represented as

```
:phone123 :avgProductLTMonths 24.
```

Then, the R2 rule presented below finds all successor products to the ones owned by the user and which are already past their average life-time.

Find successors of products owned by the user (R2)

```

PREFIX s: <http://schema.org/>
PREFIX : <http://myshop.com>

```

```

SELECT ?relatedItemName WHERE {
  VALUES ?strongIdProperty { s:gtin8 s:gtin13 s:gtin14 }
# use granular ownership information
?customer s:owns ?ownershipInfo.
?ownershipInfo s:typeOfGood ?itemOwned.
              s:ownedFrom ?dateOfAcq.
?itemOwned s:name ?itemName;
           ?strongIdProperty ?productId.

?itemShop ?strongIdProperty ?productId.

?relatedItem s:successorOf ?itemShop;
             s:name ?relatedItemName;
             :rating ?rating.
             :avgProductLTMonths ?avgPLTMonths.
# recommend a successor product if the
# currently owned is "old"
BIND (YEAR(NOW())*12+MONTH(NOW())-
      YEAR(?dateOfAcq)*12 AS ?age)
FILTER (?avgPLTMonths < ?age) }
ORDER BY DESC(?rating)

```

4 Evaluation and Future Work

We indicated earlier that the current recommender systems can benefit from access to ownership information in order to mitigate the cold-start problem and the data sparsity issue. For the future, we expect that precise and granular ownership information can be best leveraged by novel rule-based expert systems. These systems will capture expert-level domain-specific product recommendation knowledge and in effect act as automated sales assistants.

4.1 Expert Survey

During our initial investigation, we interviewed six e-commerce experts who possess domain knowledge in multiple product domains. The goal was to assess the relevance of information about items owned by the customer in order to improve recommendations in different product categories. For example, we wanted to know whether these human experts think they can provide more relevant recommendations for a customer in the area of consumer electronics if they knew about all the furniture owned. They were asked to rate all combinations of five item category pairs on a 5-level Likert scale. The categories were picked from the set of categories most frequently bought online [2]. Apart from three cases, our experts reached consensus, which means that the majority settled on two adjacent scores in all cases but three. The diagonal of the matrix naturally received the highest utility assessment. However, some categories can be useful to slightly improve product recommendations in other categories as well (see Table 1).

Table 1. Relevance of item ownership information for improving product recommendations (expert opinion). Likert scale scores used in the study: Helpful (4), Often helpful (3), Sometimes helpful (2), Rarely helpful (1), Not helpful (0), No consensus (X)

		Owned items				
		<i>Furniture</i>	<i>Consumer electronics</i>	<i>Apparel</i>	<i>Sportswear & gear</i>	<i>Cosmetics</i>
Item to recommend	<i>Furniture</i>	4	3	2	1	0
	<i>Consumer electronics</i>	1	4	X	1	0
	<i>Apparel</i>	0	1	4	2	X
	<i>Sportswear & gear</i>	0	1	2	3	0
	<i>Cosmetics</i>	0	0	X	0	4

4.2 User study

Another important factor of the viability of our approach is how willing the users are to share information about their belongings. There are situations where the perceived breach to privacy may outweigh the promised additional personalization effects and the users will be likely to refuse to reveal sensitive information. A user study described in [9] focused on personal data, e.g. name, email address, whereas we focus on revealing information about one's belongings. Also, users may be more likely to reveal certain categories of items than others. The online shop's reputation is expected to be a very important trust factor, when deciding whether to share any information with it. Online shops can encourage sharing by offering various incentives, such as discounts or free shipping. They can also explicitly declare how they intend to use the acquired information to alleviate users concerns.

In general, measuring privacy concerns of users is very difficult [see e.g. 10]. Therefore, in our work, we resorted to simulated decision making in practical scenarios. In order to provide some preliminary evidence that users are indeed willing to reveal information about their belongings in certain situations, we conducted a study, in which users were asked to complete five simulated decision-making situations in an online shopping context⁹.

⁹ <http://help.portable-shopping-history.info/>

Table 2. Model situations

#	Buy	Asked to reveal	Additional incentive (p=0.5)
1	Paper bag for vacuum cleaner	List of household appliances	Free shipping
2	Case for one's smartphone	Computer and smartphone	Free shipping
3	Pair of running shoes	Sportswear and gear	Free shipping
4	Belt	All clothes	10% discount
5	French cookbook	List of all book titles	10% discount and free shipping

In each situation, the user was told that she or he is looking for an item on the Web and just discovered a promising online shop that she or he decided to visit. After being taken to the site, a dialog box appeared asking the user to reveal some of her or his items of a certain category, which were related to the shopping task at hand. The standard incentive always offered in exchange for revealing the information was “reduced search time and better personalized service”. Additionally, monetary incentives such as free shipping or a 10% discount on the next purchase were offered with a 0.5 probability. Table 2 contains the situations in the order of appearance, the information requested, and the additional incentive.




Fig. 5. Tendency to share categories of items with and without incentive

We collected responses from 31 individuals (men and women, aged 20-50) over the Web who declared themselves as knowledgeable in matters of online shopping. All respondents revealed the requested item category at least in one of the model situations (see Figure 5). In Figure 6 one can see that for almost all item categories, 50% or more of the respondents were willing to share information about owned items. It seems that in our model situations, the incentives actually had a slightly negative impact on encouraging sharing. We also see that people are happy to reveal their sportswear or sports gear, incentivized or not, however they are wary of exposing information about their cellphone and computer.




Fig. 6. User distribution per number of shared items

4.3 Conclusion and Future Work

We have presented a conceptual model (that has already been added to schema.org based on our input), a protocol, and a reference implementation for sharing item ownership information between users and Web sites. Our first evaluation provides very preliminary evidence that human experts consider such information helpful for product recommendations, and that typical users are generally willing to expose a part of their personal ownership information in turn for better recommendations or additional incentives. We plan to submit the proposed protocol to a standardization body so that it can be widely used in real-world implementations.

In our future research, we will tackle the following issues: First, we will try to understand the necessary amount and type of item information for better recommendations. Ideally, we can find a small number categories of products (e.g. car make and model, cellphone, computer, leisure activity items, the favorite pair of shoes, ...) that would be already effective for better recommendations. Second, we will extend the data management part of the protocol and the implementation, e.g. how updates of information can be effectively shared between user and sites, and how the users can maintain their item information locally. Third, we will have to design and carry out more formal experiments on the effects of such information on state-of-the-art recommender systems. Another important concern is explicating incentives and designated use for the user to encourage sharing ownership information. Extensive sharing of personal data presents a privacy challenge; hence it's also worthwhile exploring. In the current state of our research we have not yet addressed these questions.

Our final goal is to have an experimental, end-to-end implementation of the presented vision deployed in real Web shops, browser extensions for all major browsers, and the broad adoption of the protocol and data model.

Acknowledgments: We would like to thank Bene Rodriguez for providing valuable feedback on our early prototypes. This research has been partially funded by the Eurostars program (within the EU 7th Framework Program) of the European Commis-

sion in the context of the Ontology-based Product Data Management (OPDM) project (FKZ 01QE1113D).

References

1. S. Badrul, K. George, K. Joseph, and John, R.: Item-based Collaborative Filtering Recommendation Algorithms. Proceedings of the 10th International World Wide Web Conference, Hong Kong, pp. 285-295 (2001).
2. European Commission. Bringing E-commerce Benefits to Consumers. Brussels, SEC(2011) 1640 final (2012),
http://ec.europa.eu/internal_market/e-commerce/docs/communication2012/SEC2011_1640_en.pdf.
3. Hepp, M.: GoodRelations: An Ontology for Describing Products and Services Offers on the Web, Proceedings of the 16th International Conference on Knowledge Engineering and Knowledge Management (EKAW2008), September 29 - October 3, 2008, Acitrezza, Italy, Springer LNCS, Vol. 5268, pp. 332-347.
4. Kanagal, B., Ahmed, A., Pandey, S., Josifovski, V., Yuan, J., and Garcia-Pueyo, L.: Supercharging Recommender Systems Using Taxonomies for Learning User Purchase Behavior. Proceedings of VLDB Endow. 5, 10, pp. 956-967 (2012).
5. Schein, A.I., Popescul, A., Ungar, L.H., and Pennock, D.M.: Methods and Metrics for Cold-start Recommendations. Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Tampere, Finland, 2002, ACM, pp. 253-260 (2002).
6. Steffen, R., Christoph, F., and Lars, S.-T.: Factorizing Personalized Markov Chains for Next-basket Recommendation. Proceedings of the 19th International World Wide Web Conference, ACM, Raleigh, North Carolina, USA, pp. 811-820 (2010).
7. Uschold, M. and Grünniger, M.: Ontologies: Principles, Methods, and Applications. Knowledge Engineering Review 11, 2 (June, 1996) pp. 93-155 (1996).
8. Wang, J., Sarwar, B., and Sundaresan, N.: Utilizing Related Products for Post-purchase Recommendation in E-commerce. Proceedings of the 5th ACM Conference on Recommender systems, Chicago, Illinois, USA, 2011, ACM, pp. 329-332 (2011).
9. Ackerman, M. S., et al.: Privacy in E-commerce: Examining User Scenarios and Privacy Preferences. Proceedings of the 1st ACM Conference on Electronic Commerce. Denver, Colorado, USA, ACM, pp. 1-8 (1999).
10. Preibusch, S.: Guide to Measuring Privacy Concern: Review of Survey and Observational Instruments. International Journal of Human-Computer Studies 71(12): pp. 1133-1143 (2013).
11. Adomavicius, G.; Tuzhilin, A.: Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. IEEE Trans. on Knowl. and Data Eng. 17, 6 (June 2005), pp. 734-749 (2005).