Learning User Preferences in a Multiagent System

Adorjan Kiss, Joël Quinqueton

► To cite this version:

Adorjan Kiss, Joël Quinqueton. Learning User Preferences in a Multiagent System. CEEMAS: Central and Eastern Europe on Multi-Agent Systems, Sep 2001, Cracow, Poland. pp.169-178, 10.1007/3-540-45941-3_18 . lirmm-00268632

HAL Id: lirmm-00268632

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-00268632>

Submitted on 4 Oct 2019

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

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

Learning user preferences in a multiagent system

Adorjan Kiss and Joël Quinqueton

LIRMM, Multiagent Team, 161 rue Ada, F-34392 Montpellier Cedex, France
{kiss,jq}@lirmm.fr,
WWW home page: <http://www.lirmm.fr/~kiss>

Abstract. We present in this paper some attempts to design a Machine Learning method to predict preference knowledge in a multi-agents context. This approach is applied to a corporate knowledge management system.

1 Introduction

In this paper we will present some attempts to design a Machine Learning method to predict preference knowledge in a multi-agents context. Here we define preference knowledge as knowledge about a preference between elements of a set.

For instance, the documents found by a search engine on the web are ordered according to a preference function computed from the user request. Thus, they can be considered as ordered according to a preference relation.

The framework that gave birth to this work is a joint research project, CoMMA¹, dedicated to corporate memory management in an intranet. The main objective of the project is to implement and test a Corporate Memory management framework integrating several emerging technologies in order to optimize its maintenance and ease the search inside it and the use of its content by the members of the organization.

The main challenge is to create a coherent system that relies upon several promising new technologies which are in the middle of their struggle to become standards:

- Multi-agent architecture: it is well suited to the heterogeneity of the Corporate Memory; its flexibility eases the system maintenance and keeps the rhythm with the dynamics and evolution of the Corporate Memory; cooperating and adaptive agents assure a better working together with the user in his pursuit to more effectively achieve his goals. The FIPA standard, supported by the CoMMA project, offers the specifications for interoperable intelligent multi-agent systems.

¹ This work was supported by the CoMMA (Corporate Memory Management through Agents) project [Con00] funded by the European Commission under Grant IST-1999-12217, which started beginning of February 2000.

- XML: is a standard recommended by the World Wide Web Consortium intended to offer a human and machine understandable description language: a good choice if it is important to ensure an easy maintenance, and seamless flow through various information processing systems that evolve in time.
- RDF/RDFS another W3C recommendation, that creates a semantic level on the top of XML formal description. RDF annotations allow having an integrated, global view of the Corporate Memory keeping untouched (in terms of storage, maintenance) the heterogeneous and distributed nature of the actual info sources. RDF also allows us to create a common ontology to represent the enterprise model. The ontological commitment, a fundamental choice in our approach to design the Corporate Memory, is motivated by our belief that the community of corporate stakeholders is sharing some common global views of the world that needs to be unified and formalized (RDFS) to form the basis of the entire information system.
- Machine Learning Techniques make the system adaptive to the user, and comes even more naturally due to the previous choices, as presented in the following section.

2 The role of preference in Knowledge Management

The purpose of this section is to discuss how our work (preference learning) fits into the Knowledge Management system.

2.1 Getting the user profile

One of the advantages of an enterprise that should be exploited by such a corporate information management system is that the users (i.e. the employees) can be known (their domains of interest/competence, their current activities/tasks). This can be especially useful in some cases where users are likely to be overwhelmed by the quantity of information to process and navigate themselves through (new employees during accommodation, technology monitoring scientists) who would appreciate personalized, automated help in their process of information retrieval.

2.2 Using semantic annotations

On the other hand, we have “human and machine understandable” semantic information upon the corporate knowledge offered by the RDF formalization, based upon an “enterprise ontology” (RDF schema).

The combination of these two sources of information can provide a rich ground to infer knowledge about the users probable/possible preferences. This combination is made possible due to the fact that we use the same RDF standard for formalizing the user profile; the same base ontology for the enterprise and user models.

It can be imagined that the information combined from these sources will form sets of attributes that will be used as input for a Machine Learning (ML) mechanism.

In order to set up such a ML mechanism, there are two main tasks to carry out:

1. Getting and formalizing the information to be decomposed as attributes to feed the ML mechanism.
2. Defining the ML methodology to process this info

2.3 Collecting the information to create a set of most meaningful attributes

We will need to answer the following question: *Why does a user prefer a document?*

In our attempt to give an example of some possible answers, we are gradually going deeper and deeper into details in case of complex answers: *The document is interesting.*

- Because it has been stated so:
 - By the user himself (the user has already seen the document, and “told” the system, that he is interested in)
 - By someone else (someone, maybe “close” to the user, wanted to share a favorable opinion about a document)
- Because it concerns a topic close to the users *interest fields*:
 - by the relation with the user:
 - * Personal interest fields
 - * Professional interest fields (known by his role in the enterprise)
 - by the way they are obtained:
 - * Declared interest fields (the user has stated his interest in documents concerning a topic)
 - * Implied interest fields (the user is included in a community of interest which is close to a topic, like in [PMB01])

The second question, that introduces the notion of temporality into the preference: *Why does a user prefer a document at a given moment?*

In other words, to make the difference from the first question: *Why does a user prefer a document at a given moment, and does not prefer it at another moment?*

The document is interesting only if seen the first time (or the first few times)

- It is interesting during a certain period:
- When the user performs a certain activity
- Etc.

These answers are just some samples, one can think of many other possible reasons. Though, it is a very important to find the right questions and answers, that include the majority of possible situations. Indeed, getting the right questions and answers and translating them into quantifiable attributes, and making sure that the highest number of possible situations are observed is a key to the success of such a learning mechanism, that may even outclass in importance the chosen learning technique.

Nevertheless, we will present our approach in the Comma project to choose some typical answers and attributes, but we will focus more on the second issue: the preference learning methodology.

3 The design of the CoMMA system

The design of CoMMA, presented in this section, can be viewed as our attempt to implement a corporate knowledge management framework.

3.1 The agent architecture

The chosen MAS consists of a society of coarse-grained agents, that fulfill in general multiple roles, and are organized in a small number of functional sub-societies. The MAS architecture was designed in order to optimize task-division, flexibility and robustness of the system, and network layout (extensibility, scalability, traffic optimization).

For the implementation of the prototype system, the Jade agent platform was chosen, which is an Open Source Project developed by project partners, University of Parma and CSELT. Jade is an agent platform implemented in Java, which is Fipa compliant, thus having the advantage of a wide opening towards Internet and the Web, interoperability with other MAS-s, and future systems.

In the current status of implementation, the CoMMA system will help the user in three main tasks:

- insertion and RDF annotation of documents,
- search of existing documents, and
- autonomous document delivery in a push fashion to provide the user with information about new documents that the system predicts interesting for him.

The agent categories present in the system can be classified into four main areas:

1. Document and annotation management.
2. Ontology (Enterprise and User Models) management.
3. User management.
4. Agent interconnection and matchmaking.

The agents from the document dedicated sub-society are concerned with the exploitation of the documents and annotations composing the corporate memory, they will search and retrieve the references matching the query of the user with the help of the ontological agents.

The agents from the ontology dedicated sub-society are concerned with the management of the ontological aspects of the information retrieval activity especially the queries about the hierarchy of concepts and the different views.

The agents from user dedicated sub-society are concerned with the interface, the monitoring, the assistance and the adaptation to the user. Finally the agents from the interconnection dedicated sub-society are in charge of the matchmaking of the other agents based upon their respective needs.

We have already experimented such an architecture in Network Supervision [EDQ96], with Machine Learning abilities [QEN97]. More recently, an interaction-based strategy has been experimented [SS01].

3.2 The learning agent

The first context to assess preference learning was chosen to be the document retrieval scenario, via semantic annotations. The search engine used for document retrieval in the CoMMA system is an inference engine called CORESE [CDH00] developed by INRIA, one of the partners of the project. CORESE uses Conceptual Graphs and combines the advantages of using the RDF language for expressing and exchanging metadata, and the query and inference mechanisms available in CG formalism. In order to produce inferences, CORESE exploits the common aspects between CG and RDF: it defined a mapping from annotation statements (RDF triples) to Conceptual Graphs and vice-versa.

One of the shortcomings of such a query retrieval engine is that there is no standard method to sort the information returned, such as keyword frequency in keyword-based search engines. The returned data set must be post-processed, filtered and sorted to present the user with the relevant information. Here comes the aid offered by our ML mechanism.

In the CoMMA system, information that feeds the ML comes from several sources: The document sub-society (the annotations accompanying a query response), the user sub-society (user monitoring and explicit user feedback), and ontology sub-society (to help getting the meaning of the results). Therefore, the learning behavior was awarded to the User Profile Manager agent, which belongs to the connection dedicated sub-society, and performs notably a role of middleman between agents. This decision was justified also by network traffic optimization reasons, especially because in reaction to a user action (query), several interactions can be triggered between agents of different roles.

For example, during a query retrieval, the main interactions are as described in the following figure.

Fig. 1. The main interactions in CoMMA during a query

3.3 The learning cycle

The goal of the ML component is to produce a set of rules that will be used to produce predictions about user preferences. It can be supposed that the system starts with a set of predefined rules, that will be gradually improved during the process of adaptation to users. Otherwise, the system would start with an empty knowledge base, and will undergo a training period, to accumulate sufficient knowledge to allow its deployment.

In our approach, user adaptability comes from both implicit observation of the user (user monitoring subsystem), and explicit user feedback.

Within the ML subsystem, we use a "predict-or-learn" type protocol, that is, when the input is coming from query answers, the system tries to use its knowledge to predict, otherwise, when the input comes from the user in the form of negative feedback, the system tries to update its rules set.

3.4 A sample set of attributes

A sample set of attributes is used to create instances from the answers we gave as examples to the question of preference. We tried to make it diverse and for most

of them restricted the scope as much as possible for the sake of simplicity of our prototype. We also tried to stick close to the concepts defined in the prototype ontology, that we used as a basis for our first enterprise model.

The first set of attributes relates the document to knowledge recorded in the user profile. Concepts defined in the ontology which are used for this purpose are *Topic*, *Concerns*, *Interestedby*, *CommunitiesofInterest*, etc .

We have also used a number of attributes automatically generated by *user monitoring*. In our example we have imagined a simple user monitoring scenario, that supposes tracking each consultation of a document by the user, and building a navigation history (or consultation trace). These attributes refer to statistics about document and document category visits.

Finally we have a set of attributes containing specific information about the particular document being analysed (ratings for the document of the user in question, public ratings of other users eventually close to the user).

The issue regarding the correctness of the choice of these concepts, how to make them more generic or how they can be adapted to a specific enterprise, we are addressing in the next stage of our project: there we will intend to create a generic process for attribute collection, instead of analyzing how these attributes are fit for their purpose.

4 The use of Preference knowledge

As mentioned in the introduction we define preference knowledge as knowledge about a preference between elements of a set. Such knowledge can be stated in various forms: a numerical value assigned for each item, a total ordering relation, a partial ordering relation or even a preordering of the set.

Logical models have been proposed to deal with such knowledge, some dealing directly with the comparison abilities [Sch96].

It is generally admitted that a preference stands for an order that maybe partial, even a preorder, but that it is often convenient to represent it by a linear extension (which is a total order) or a numeric value compatible with the known orderings.

Then, in terms of Machine Learning, different strategies may be used, depending on the form of the preference knowledge.

4.1 Numeric labelling

A numeric labeling, i.e. a mapping of our set of examples into a set of real numbers, is a convenient way to summarize a preference relation. Some Machine Learning methods are available to learn numerical variables [Gas89,Bre96b,Bre96a].

Generally, the methods for learning to predict a numerical variable v measure the quality of a predictive rule R by the standard deviation $Q(R, v) = \sigma^2(v|R)$ of the value of the variable among the set of objects verifying the concept R to be tested.

The lower $Q(R, v)$ is, the better R is to predict v , because the mean value of v can be used with less error. With such criteria, any learning method will lead to grouping several neighbour values around their mean. Then, the learnt rules will not be very different from rules learnt from examples roughly rated with a finite set of values.

4.2 The order relation

By definition, a binary relation, which we note $\triangleleft$, is an order if it has the following properties:

- reflexive $x \triangleleft x$,
- transitive: if $x \triangleleft y$ and $y \triangleleft z$, then $x \triangleleft z$,
- antisymmetric: if $x \triangleleft y$ and $y \triangleleft x$, then $x = y$.

Then, we can imagine to learn the binary relation by learning each of its elements, that is, learn on each couple of objects (a, b) such that $a \triangleleft b$. Then, let us summarize the suitable properties of such a learning set for this approach to work correctly.

First, if (a, b) with $a \triangleleft b$ is an example, then (b, a) is a counter-example. Then, what happens to (a, a) ? We can see that they would be both examples and counter-examples, then it is better to consider the strict order relation, and eliminate diagonal elements.

With these hypotheses, the description of an example is made of 2 parts: the attributes which are modified between a and b , and those which keep the same value. We can notice here that these attributes are the same as those involved in the sorting tree of our examples.

Finally, our method appears to be “half lazy”, in comparison with lazy learning methods, like kNN or LWR [Aha92]. Our learned knowledge is partly explicit, but in the classification step, we need to compare a new instance with several elements of the learning set (maybe in a dichotomic way) to put it in the right place.

4.3 Statistical evaluation criteria

Usually in Machine Learning, particularly for building of decision trees, the learned classification rules are evaluated by their similarity to the desired classification.

We can use the same principle here, and we have two possible families of criteria. If we can compute a rank for each element, the similarity is computed by measuring the rank correlation to the expected ranking. Otherwise, each pair must be given: then we use a pairwise comparison between the expected order and the learnt order.

Several measures of similarity between 2 different orderings of the same data have been proposed. In each case, one has to deal with tied elements.

The Spearman rank order correlation The Spearman rank order correlation r_s is a correlational measure that is used when both variables are ordinal. The traditional formula for calculating the Spearman rank-order correlation is

$$Corr(r, r') = 1 - \frac{6 \sum_{i=1}^n (r_i - r'_i)^2}{n(n^2 - 1)}$$

where r and r' are the ranks to compare of paired ranks. When there are tied cases they should be assigned the mean of their ranks. The mean of the ranks from $p+1$ to $p+n$ is $\frac{1}{n}(\frac{(n+p)(n+p+1)}{2} - \frac{p(p+1)}{2})$, which become after simplification $\frac{n+2p+1}{2}$.

The Kendall pairwise τ criterion When we have to compare each pair of data, they can be classified as either tied (T), concordant (P), or discordant (Q).

The best measure for this case is Kendall's τ_b which takes into account a correction for tied pairs. Its formula is

$$\tau_b = \frac{P - Q}{\sqrt{(P + Q + T_x)(P + Q + T_y)}}$$

where T_x is the number of pairs tied on X but not Y, and T_y is the number of pairs tied on Y but not X.

4.4 Verifying the consistency of the method

In the case we perform a pairwise learning of the order relation, we can notice that a fundamental property, the transitivity, can be guaranteed by the learning process itself, as we show below for a version space method [Mit97].

We can check that, if, for 3 examples the transitivity holds, then it is not necessary to add the 3rd pair as example to learn the relation :

- let $(a, b) = (a_1 \dots a_n, b_1 \dots b_n)$ and $(b, c) = (b_1 \dots b_n, c_1 \dots c_n)$. Then, S is of the form $L \wedge R$, with the left part L as a generalisation of both a and b , and the right part R of both b and c . Then, as L is a generalisation of a and R of c , S is a generalisation of (a, c) .
- with the same conventions, G has a disjunctive form whose elements reject all the examples, then, if we represent any of its elements as $L \wedge R$. If (b, a) and (c, b) are rejected, it means that L rejects b or R rejects a , and L rejects c or R rejects b . But G must also be a generalisation of S .

Of course, this is only a scheme of the proof, and is, strictly speaking, only available for version-space-like learning. In a more general case, like decision tree learning, we can only make the hypothesis that it is true. We concluded that we could learn directly a sorting rule (in a greedy way, like decision trees) and evaluate the obtained rule with the τ criteria defined in section 4.3.

5 Conclusion

The Machine Learning and user adaptability component is a very important aspect in the CoMMA project. In this paper we presented the advances that we have made in this domain, especially focussing on the learning of preference data for document retrieval.

The main choice we focus on does not only present the usefulness of Machine Learning, but also tries to overcome some of the gaps and limitations of semantic based information retrieval systems.

The first implementation of our algorithm has given the expected quick results during the intermediate trials during the project, and allowed experimental validation through feedback from the users. The next stage of our research is focusing on the "generic learning problem", that instead of being built upon a predefined set of attributes, will try to offer a formal procedure to collect the attributes from "observable concepts" present in the ontology.

References

- [Aha92] D. Aha. Tolerating noisy, irrelevant, and novel attributes in instance-based learning algorithms. *International Journal of Man Machine Studies*, 36(2):267–216, 1992.
- [Bre96a] L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [Bre96b] L. Breiman. Stacked regression. *Machine Learning*, 24(1):49–64, 1996.
- [CDH00] Olivier Corby, Rose Dieng, and C. Hébert. A Conceptual Graph Model for W3C Resource Description Framework. In *the 8th International Conference on Conceptual Structures (ICCS'00)*, number LNCS 1867 in Lecture Notes in Artificial Intelligence, Darmstadt, Germany, 2000. Springer Verlag, Springer Verlag.
- [Con00] CoMMA Consortium. Corporate Memory Management through Agents. In *E-Work and E-Business conference, Madrid*, October 2000.
- [EDQ96] Babak Esfandiari, Gilles Deflandres, and Joël Quinqueton. An interface agent for network supervision. In *Intelligent Agents for Telecommunication Applications*, Budapest, Hungary, 1996. ECAI'96 Workshop IATA, IOS Press.
- [Gas89] O. Gascuel. A conceptual regression method. In Katharina Morik, editor, *EWSL-89, 4th European Working Session on Learning*, pages 81–90, Montpellier, France, Décembre 1989. Jean Sallantin and Joel Quinqueton, CRIM, Pitman, Morgan Kaufman.
- [Mit97] Tom M. Mitchell. *Machine Learning*. Mac Graw Hill, 1997.
- [PMB01] Michal Pechoucek, Vladimir Marik, and Jaroslav Barta. Cplant: An acquaintance model based coalition formation multi-agent system. In *this book*, 2001.
- [QEN97] Joël Quinqueton, Babak Esfandiari, and Richard Nock. Chronicle learning and agent oriented techniques for network management and supervision. In Dominique Gaïti, editor, *Intelligent Networks and Intelligence in Networks*. Chapman & Hall, September 1997.
- [Sch96] P-Y. Schobbens. A comparative logic for preferences. In Pierre-Yves Schobbens, editor, *Working Notes of 3rd ModelAge Workshop: Formal Models of Agents*, Sesimbra, Portugal, January 1996.
- [SS01] Nicolas Sabouret and Jean Paul Sansonnet. Learning collective behavior from local interactions. In *this book*, 2001.