

Multiagent Learning and Adaptation in an Information Filtering Market

Innes A. Ferguson

Interactive Information Group
Institute for Information Technology
National Research Council
Ottawa ON, Canada K1A 0R6
innes@ai.iit.nrc.ca

Grigoris J. Karakoulas

Global Analytics Group
Canadian Imperial Bank of Commerce
161 Bay St., BCE-8, P.O. Box 500
Toronto ON, Canada M5J 2S8
karakoul@cibc.ca

Abstract

This paper presents an adaptive model for multiagent coordination based on the metaphor of economic markets. This model has been used to develop SIGMA, a system for filtering Usenet netnews which is able to cope with the non-stationary and partially observable nature of the information filtering task at hand. SIGMA integrates a number of different learning and adaptation techniques, including reinforcement learning, bidding price adjustment, and relevance feedback. Aspects of these are discussed below.

Introduction

With the widespread availability of electronically stored information in such environments as the World Wide Web (WWW), it is becoming increasingly important to use automatic methods for filtering such information (Belkin and Croft 1992). Traditional off-line approaches to information filtering (IF) consist of building database indices upon which traditional search and retrieval algorithms are subsequently applied. Because of the need to manipulate vast numbers of heterogeneous information sources whose thematic, structural, and/or locational characteristics are subject to constant change, the database indices created by such off-line approaches become increasingly difficult, if not ultimately impossible, to update in a timely or cost-efficient manner. As such, it becomes crucial for these filtering methods to be able to adapt on-line to the problem at hand.

A number of machine learning techniques have been applied to a variety of IF tasks, including genetic algorithms and case-based reasoning methods for learning user profiles (Sheth 1994; Hammond, Burke, and Schmitt 1994) and information-theoretic classification for categorizing documents based on users' preferences and usage pattern analysis (Lang 1995). A potential shortcoming in these approaches is that, by employing learning techniques which

are cast into single agent frameworks, they will likely be unable to exploit the distributed asynchronous flow of information which is characteristic of dynamic environments such as the WWW. As some recent work suggests, learning can be sped up by extending such techniques to systems comprising multiple collaborative agents (Lashkari, Metral, and Maes 1994), but also, and perhaps more interestingly, to systems composed of multiple non-cooperative agents (Sen, Sekaran, and Hale 1994).

A number of other adaptation techniques have been proposed for dynamic multiagent coordination, including ones based on metaphors from artificial life (Menczer, Belew, and Willuhn 1995) and from computational markets (Wellman 1994), which do not use pre-fabricated coordination strategies. As such, these adaptive, distributed algorithms appear more likely to scale with problem size and so be better suited to handle the increasing complexity of today's heterogeneous information networks. In our own work, we propose an adaptive model for multiagent learning based on the metaphor of economic markets.

In this paper we briefly present the market model of SIGMA¹ and its application to a particular class of IF problems (Karakoulas and Ferguson 1995). The model defines a computational economy which we have used as the basis for constructing a multi-user IF system. We have implemented the system as a market comprising multiple specialized CALVIN² agents (Ferguson and Davlourous 1995) and have been applying the system to the task of personalized filtering of Usenet netnews. By augmenting a particular collection of these IF agents with learning and adaptation techniques – including reinforcement learning, bidding price adjustment, and relevance feedback – we have created a robust network-based application which

¹System of Information Gathering Market-based Agents.

²Communicating Agents Living Vicariously In Networks.

adapts to both changes in the characteristics of the information available on the network, as well as to changes in individual users' information delivery requirements.

The SIGMA Market Model

The activities in an IF system are characterized by competition for resources: users that arrive in the system are endowed with limited time and money to pose queries and so aim to spend their resources as efficiently as possible in order that their interests be matched as satisfactorily as possible against dynamic information streams. At the same time, the IF system itself contains agents which represent automated indexing and filtering procedures, informational resources, and human-computer interaction procedures, all of which need to be allocated computation time in order to serve as efficiently as possible the various demands for information within the system.

In the field of economics, the problem of allocating limited or scarce resources among competing agents has been studied extensively in the context of *markets*. The latter are models of decentralized decision making wherein each agent processes only local information in order to evaluate and make decisions regarding goods and services. The backbone of this decentralized decision making process is the *price* mechanism. Prices are inherently local to each market and they impose low information requirements on agents since they effectively reflect the efficiency of resource allocation within the market.

Similarly to economic markets, the computational economy of SIGMA can be defined in terms of goods and competitive agents that produce and consume these goods. The goods traded here are information items (e.g. news articles). The agents involved are of three general types: (i) *producer* agents, (ii) *consumer* agents, and (iii) *matchmaking* or *broker* agents (Kuokka and Harada 1995). Producer agents transform goods from an input form into an output according to their different technologies (Wellman 1994) and can themselves be of two types: *originators* and *intermediaries* (Brown *et al.* 1995).

In SIGMA, raw information is obtained from an originator agent, the Usenet Wrapper (UW) agent, whose output goods are then transformed by an intermediary, a Feature Extractor (FE) agent, into a feature-based knowledge representation formalism that can be used to reason about the news article (see Figure 1). The vector space model (VSM) is used for implementing this formalism (Belkin and Croft 1992). FE agents, in addition, use dynamically compiled stop lists (of content-free words) together with the Wordnet on-line thesaurus to discard noise words and construct a subset of discriminating keywords. Weights are then calculated for each keyword by using dynamically compiled term-frequency/ inverse-document frequency (tf-idf) factors. The VSM and any other intermediate representations of processed articles are cached in repositories which are subsequently accessed by FE agents through appropriate Repository Manager (intermediary) agents RM_1 and RM_2 .

A key group of agents in the SIGMA market are the Profile Selector or PS agents. Each of these represents a different query by the user for a filtering activity over a period of time. At each stage of this activity, a PS agent is endowed with a budget M which is a constraint on how much it can spend on buying goods (news articles) from producers, Profile Generator (PG) agents, at prices that depend on a bidding process (which is brokered by a Bid Manager or BM agent) and on the relevance of each good. The behavior of a PS agent is based on *reinforcement learning*. The agent initially allocates its budget uniformly to the PG agents that have won the bidding and selects from each of them the respective top-ranked articles. For each article that the user accepts or rejects the PS agent receives (from this user) a positive or negative reinforcement, the size of the reinforcement being a function of the price that the agent has paid for the respective article. Such reinforcement signals enable the agent to learn over time a policy of buying the best articles (i.e. maximizing the user's long term cumulative benefit) from the most efficient producers. We have developed a Q-learning algorithm for implementing the learning behavior of the PS agent, based on earlier work by Karakoulas (1995).

PG agents are initialized with a profile that contains the keywords specified by the user.³ When selected at bidding, each PG agent explores the relevant newsgroups and selects articles for which its marginal profit is greater than a given threshold. The filtering technology that is used for producing the selection of articles is based on similarity metrics of the VSM representation; marginal profit is defined as a function of the similarity of the article and the profile of the agent. To ensure proactivity in the agents' behavior we allow the agents to follow this greedy policy with probability p and to select any other article with probability $1-p$. Given the articles that have been accepted by the user, the agent applies relevance feedback to update its profile. This is done by adjusting the weights of existing keywords and by including new terms, the latter selected from term/tf-idf weightings tables which are dynamically compiled by the FE agents. Finally, articles selected by the PS agents are presented to the user via the Interface Module (IM) agent, which is also responsible for propagating the user's feedback into the system (see Figure 1).

Discussion

Agents in the SIGMA market model exhibit two levels of learning and adaptation. At the intra-agent level, PS and PG agents learn and adapt their behavior via Q-learning and relevance feedback using users' local evaluation of presented articles. At the inter-agent level, the price mechanism and the purchasing of articles from any subset of a population of stochastically diverse and evolving agents combine to make the system, as a whole, more proactive and anticipatory with respect to the personalized interests of its users; in addition, these strategies combine to

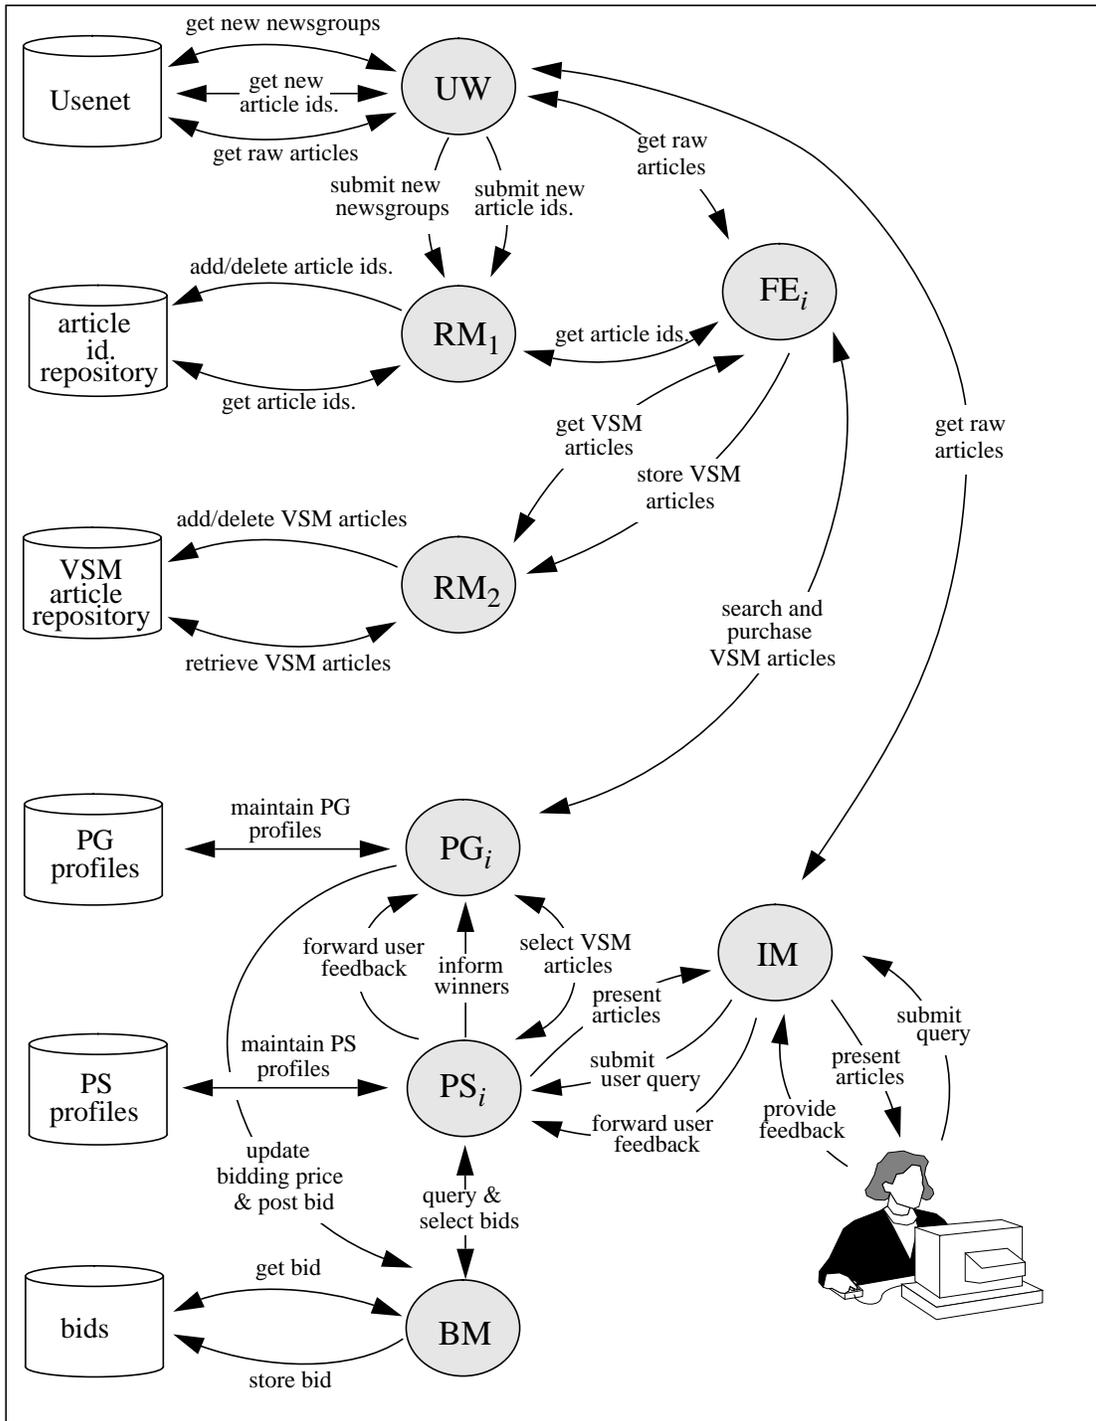


Figure 1: SIGMA agent-level system description.

enhance the robustness of the system with respect to the uncertainties and dynamics of the environment.

The SIGMA market model described above has been implemented as a collection of specialized CALVIN agents. An extension of the TouringMachine architecture (Ferguson 1992), the CALVIN agent framework is an open architecture for facilitating the development of highly concurrent, embedded agent-oriented applications. As such, the framework provides application developers with a powerful set of agent programming tools including libraries of intra- and inter-agent protocols (in the latter category, we are currently using KAPI (Kuokka and Harada 1995)), sensory and effector apparatus, internal agent behavior APIs, persistent storage management, and support for pre-emptive (anytime) scheduling of behaviors, among others.

The current focus of our work is on system evaluation. We propose to evaluate the IF application through two sets of experiments. The first of these concerns performance evaluation of the computational economy with respect to parameters such as initial number of PG agents assigned to each user query, number and initial weights of terms in the VSM representation, producers' marginal profit thresholds and consumers' budgets, declaration of bankruptcy status for PG agents, and the PG and PS agents' exploration factors. The second set of experiments involves measuring the (hopefully) improved access to Usenet that is provided to users. We have identified a number of quantitative measures to be applied to this task including, among others, the number and ratio of articles accepted/rejected by a user, the average length of time spent searching for articles to read in a newsgroup, the average length of time spent reading a given article, and the number of different newsgroups articles are sought from.⁴ On the assumption that it should become easier for users to find relevant articles and that they should have more time available for reading articles from other newsgroups, we hope to see an increase in the number of newsgroups that articles are selected from, as well as an increase in the proportion of relevant articles the user actually reads. Alternatively, for a user with fixed interests, we would expect this user to have to spend less time reading news.

While our current interest in computational markets has largely been focused on the design and application of adaptive information filtering techniques to WWW-based news, it is our belief that as the Internet and WWW become increasingly commercialized, the need for effective profit-based agents which can act on behalf of their owners and

³Optionally, should a user's latest query be deemed similar enough to any queries previously made by the same or any other users of the system, the new query can be assigned a clone of the appropriate matching PG profile as a starting point for bidding and selecting articles. This method of bootstrapping new queries (and new users) in effect permits a degree of collaborative filtering through sharing of learned information about users' existing information interests. And since agents are modular entities, sharing is merely a matter of copying whole agents whose profiles satisfy a pre-determined similarity metric.

seek payment for services rendered will increase dramatically (Genesereth and Ketchpel 1994). At the same time, we also believe that in order for such sophisticated agents to materialize, much work remains to be done in defining and deploying automated techniques for brokering of information and for linking requests of demand (consumer) agents to the ability of supply (producer) agents to provide the required information items (Brown *et al.* 1995; Kuokka and Harada 1995). Aspects of a number of these issues are already addressed by our current work. We are confident that our experiences gained in furthering broadening the scope and applicability of the SIGMA market model to other information access tasks will provide other useful insights in these directions.

References

- Belkin, N.J. and Croft, W.B. 1992. Information filtering and information retrieval: Two sides of the same coin? *Communications of the ACM*, 35(12):29–38.
- Brown, C.; Gasser, L.; O'Leary, D.; and Sangster, A. 1995. AI on the WWW. *IEEE Expert*, 10(4):50–55.
- Ferguson, I.A. and Davlourous, J.D. 1995. On establishing multi-sensory multi-channel communications among networked users. In *Proceedings IJCAI Workshop on AI in Distributed Information Networks*, pp. 103–109.
- Ferguson, I.A. 1992. *TouringMachines: An architecture for dynamic, rational, mobile agents*. Ph.D. Thesis, Computer Laboratory, University of Cambridge, Cambridge, UK.
- Genesereth, M.R. and Ketchpel, S.P. 1994. Software agents, *Communications of the ACM*, 37(7):48–53.
- Hammond, K.; Burke, R.; and Schmitt, K. 1994. A case-based approach to knowledge navigation. In *Proceedings AAAI Workshop on Indexing and Reuse in Multimedia Systems*, pp. 46–57.
- Karakoulas, G.J. and Ferguson, I.A. 1995. A computational market for information filtering in multi-dimensional spaces. In *Proceedings AAAI Fall Symposium on AI Applications in Knowledge Navigation and Retrieval*. pp. 78–83.
- Karakoulas, G.J. 1995. A Q-Learning approach to cost-effective classification. Technical Report, Knowledge Systems Laboratory, National Research Council, Ottawa, ON, Canada.
- Kuokka, D. and Harada, L. 1995. Matchmaking for information agents. In *Proceedings International Joint Conference on Artificial Intelligence*, pp. 672–678.
- Lang, K. 1995. NewsWeeder: Learning to filter netnews. In *Proceedings International Conference on Machine Learning*, pp. 331–339.

⁴Several of these or related measures were employed by Maltz in his evaluation of a collaborative Usenet filtering application (Maltz 1994).

Maltz, D.A. 1994. *Distributing information for collaborative filtering on Usenet net news*. SM Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Menczer, F.; Belew, R.K.; and Willuhn, W. 1995. Artificial life applied to adaptive information agents. In *Proceedings AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments*.

Sen, S.; Sekaran, M.; and Hale, J. 1994. Learning to coordinate without sharing information. In *Proceedings Conference of the American Association for Artificial Intelligence*, pp. 426–431.

Sheth, B.D. 1994. *A learning approach to personalized information filtering*, SM Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Wellman, M.P. 1994. A computational market model for distributed configuration design. In *Proceedings Conference of the American Association for Artificial Intelligence*, pp. 401–407.