Volume 378, issue 1 1 May 2007 ISSN 0378-4371

ELSEVIER

PHYSICA A
STATISTICAL MECHANICS AND ITS APPLICATIONS

Recognized by the European Physical Society

Social network analysis: Measuring tools, structures and dynamics

Proceedings of the International Conference "Social Network Analysis and Complexity", 31 July – 2 August 2006 at Collegium Budapest in Budapest, Hungary

Guest Editors:
András Lőrincz,
Nigel Gilbert
Rebecca Goolsby

Available online at
ScienceDirect
www.sciencedirect.com

http://www.elsevier.com/locate/physa

Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Physica A 378 (2007) 127–134
PHYSICA A
www.elsevier.com/locate/physa

Efficiency of goal-oriented communicating agents in different graph topologies: A study with Internet crawlers

András Lőrincz*, Katalin A. Lázár, Zsolt Palotai

Department of Information Systems, Eötvös Loránd University, Pázmány P. sétány 1/C, Budapest H-1117, Hungary

Available online 11 December 2006

Abstract

To what extent does the communication make a goal-oriented community efficient in different topologies? In order to gain insight into this problem, we study the influence of learning method as well as that of the topology of the environment on the communication efficiency of crawlers in quest of novel information in different topics on the Internet. Individual crawlers employ selective learning, function approximation-based reinforcement learning (RL), and their combination. Selective learning, in effect, modifies the starting URL lists of the crawlers, whilst RL alters the URL orderings. Real data have been collected from the web and scale-free worlds, scale-free small world (SFSW), and random world environments (RWEs) have been created by link reorganization. In our previous experiments [Zs. Palotai, Cs. Farkas, A. Lőrincz, Is selection optimal in scale-free small worlds?, ComPlexUs 3 (2006) 158–168], the crawlers searched for novel, genuine documents and direct communication was not possible. Herein, our finding is reproduced: selective learning performs the best and RL the worst in SFSW, whereas the combined, i.e., selective learning coupled with RL is the best—by a slight margin—in scale-free worlds. This effect is demonstrated to be more pronounced when the crawlers search for different topic-specific documents: the relative performance of the combined learning algorithm improves in all worlds, i.e., in SFSW, in SFW, and in RWE. If the tasks are more complex and the work sharing is enforced by the environment then the combined learning algorithm becomes at least equal, even superior to both the selective and the RL algorithms in most cases, irrespective of the efficiency of communication. Furthermore, communication improves the performance by a large margin and adaptive communication is advantageous in the majority of the cases.

Keywords: Communicating agents; Internet; Scale-free world; Small world

1. Introduction

Analysis of social networks has posed a serious challenge to a multitude of disciplines, including mathematics, physics, and economics. Methods of this highly interdisciplinary research area embrace data collection, (data) visualization and (data) explanation, development of verbal and/or mathematical models, experimental or theoretical proofs with regard to these models, simulations of analytically intractable interactions, and their combinations. For excellent (historical) reviews of this field, consult, e.g., [1,2] and the long list of the references therein. In this paper, the methods of machine learning are employed (for a recent

*Corresponding author.
E-mail addresses: andras.lorincz@elte.hu (A. Lőrincz), kati@elte.hu (K.A. Lázár), zspalotai@vnet.hu (Z. Palotai).

0378-4371/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.physa.2006.11.052

collection of similar works see [3]), and the effects of adaptivity/learning by experiments in real and simulated worlds are studied.

Our interacting agents (crawlers/Internet foragers) are able to learn by collecting food (documents from the Internet). The crawler fleet may communicate, and if they do so then they may optimize the communication. The efficiency of the fleet on a downloaded part of a vast news portal of the Internet will be investigated. Under this condition, the link structure will be manipulated and the efficiency of the fleet in different worlds, i.e., in random world environments (RWEs), in scale-free worlds (SFW), and in scale-free small worlds (SFSWs) will be investigated. The crawlers have to find novel documents, either within a time limit, in different topics, without any further restriction imposed on them or novel documents characterized by some combinations of the previous requirements. This work extends our previous study in which the crawlers were not allowed to communicate in a direct manner and the topics were not specified [4].

In selective learning, alternative solutions coexist while organisms compete for space and resources and only good solutions are maintained [5]. Selection may occur, e.g., at the level of individuals, or at the level of the solutions applied by individuals. For reviews on evolutionary theories and the dynamics of self-modifying systems consult [6–8], respectively.

In a typical reinforcement learning (RL) problem what motivates the learning process [9] is the optimization of the expected value of long-term cumulated profit of the actual state (state value) or state-action pair (state-action value). A well-known RL example is Tesauro's TD-Gammon program [10], which employs MLP function approximators for value estimation. The RL has also been utilized in concurrent multi-robot learning, see, e.g., [11], the robots have to learn to forage in concert via direct interaction.

In this paper a selective learning algorithm is compared to linear function approximation-based RL algorithm. In our problem domain, fitness is influenced by the ever changing external world and by the competing individuals. Communication can either occur in direct or indirect manners. Indirect communication is invoked by the reward system: positive reward is delivered only to the first sender of a news item. What differentiates this work from the other studies addressing issues of the same nature are as follows: (i) no control is exercised over the environment; (ii) the collaborating individuals employ selective learning and value estimations under 'evolutionary pressure'; (iii) direct and indirect communications may occur simultaneously; and (iv) different topologies are considered.

First, the architecture of the crawlers and the crawler fleet (see Section 2) is briefly reviewed. Secondly, in Section 3 the results of our simulations in different worlds are presented. Finally, Section 4 discusses and summarizes our findings.

2. Forager architecture

In the sequel, our agents and their working mechanisms are delineated. Two types of agents, i.e., foragers, or crawlers and reinforcing agents (RA) can be distinguished. The foragers crawl the web and send back the addresses (URLs, uniform resource locators) of the selected documents to the RA. The RA is a simple machine that schedules the work of the foragers, launching one after the other an equal number of times. In effect, it acts as a central agent able to determine the forager finding the new information first and then it reinforces that forager based on the received URL.

The forager 'moves to' or 'visits' linked URLs in a predefined order so as to find novel documents. Indeed, regions abundant in information may be separated from each other by those lacking in information, which form the 'barriers'. This segregation motivates reinforcement-based value estimation enabling the crawler to overcome such short-term penalties in the hope of gaining long-term profits [12,13].

Algorithms, pseudo-codes, and parameters of the algorithms are detailed in [4]. Herein, the core features of the different agent–environment–communication procedures are reviewed.

2.1. Weblog algorithm

Each forager possesses a weblog consisting of the URLs and their associated weblog values. The weblog is utilized for the periodical recommencement of the activity of the forager it belongs to. At the beginning of a period, the forager selects randomly a URL from the best elements of the weblog. The sequence of visited URLs between two restarts forms a path. In effect, the weblog value of a URL estimates the expected sum of rewards along the path after a visit to the underlying URL. Estimation proceeds by simple averaging of the measured sum of rewards along the path after the visit. Averaging, in effect, takes into consideration a given number of recent paths after the visit. The average is the value of the URL in the weblog. The weblog value of a new URL is the actual sum of rewards that can be collected along the rest of the path after the visit to the new URL. The high weblog value of a URL indicates an abundance of relevant documents around it. Consequently, it is advisable to launch a search from that URL.

2.1.1. Crawling applying the weblog algorithm

Each document is characterized by an $N(=50)$ dimensional vector, with each components mapped non-linearly to the interval [0,1]. The vector is computed by the well-known probabilistic term-frequency inverse document-frequency (PrTFIDF) classifier method [14] generated on a previously downloaded portion of the Internet. Each forager has a randomly chosen N dimensional 'weight vector'. At a given URL, the crawler calculates the scalar products of its weight vector and all vectors that belong to the documents at the frontier (the list of linked but not yet visited URLs found during the crawling along the actual path). Then the crawler 'moves' to (one of) the URL(s) of the maximum scalar product value.

2.2. Reinforcement learning

A forager can estimate the long-term cumulated value/profit of a URL according to the reinforcements obtained after the URL has been visited. The (immediate) profit is the difference between the rewards and the penalties received at any given step. In fact, the immediate profit characterizes a step to a URL in a myopic manner. Foragers use adaptive linear value estimator (ALE) [9] to overcome this short sightedness. They follow the policy maximizing the expected long-term cumulated profit (LTP) in lieu of the immediate. The policy and the profit estimation are related concepts: the profit estimation determines the policy, whereas the policy influences the choices and, in turn, the expected LTP. (For an excellent introduction to this subject, see [9].)

2.2.1. Crawling employing RL

The crawler performs a step according to ALE. ALE trains the weight vector to improve the crawling. The LTP of the actual URL is estimated as it is described below. At a given URL, the crawler computes the scalar products of its weight vector and all vectors that belong to the documents at the frontier (the list of linked but not yet visited URLs found during the crawling). Then the crawler 'moves' greedily to (one of) the URL(s) having the highest scalar product value. After the step has been performed and the sum of the components of the immediate reward (cost of a step, cost of sending a document, rewards received for novel and topic relevant documents) has been calculated the following terms are obtained: a denotes the LTP of the previous URL, b the LTP of the actual URL, and c the immediate reward. The error of our state value estimation δ is as follows:

$$\delta = \mathrm{a} - \gamma(\mathrm{b} + \mathrm{c}),$$

where $0 < \gamma < 1$ is the discount factor. Should there be no error, then the estimation may be regarded as good. Otherwise, the weight vector has to be modified, in proportion to the sign and the magnitude of the error. This method is called temporal differencing (TD). For further details about the TD method and an in-depth description of our algorithms consult, [4,9], respectively. The procedure gives rise to adaptive crawling. At the beginning of each period, the forager continues the previous path.

2.3. Weblog URL selection and reinforcement learning-based crawling

The two methods can be combined, the weblog algorithm can be used for the evaluation and the selection of the start points, whereas and weight vector of the crawler can be updated by RL.

2.4. Sending documents

The crawlers apply a threshold for document sending. Not all downloaded documents are novel and not in all of the cases their value can pass the sending threshold. The document will be sent provided that it is novel and its value passes the threshold. Neither do all sent documents invoke positive rewards. If positive reward is delivered for a sent document, then it is alleged to be a relevant sent document, or a relevant document, for short. If documents can be sent to the RA only, then there is no direct communication between the crawlers. Nonetheless, crawlers are coupled: they are adaptive and work in different domains and/or they tend to follow different paths [15].

The value of a document, in fact, is the scalar product of the sending weight vector and its PrTFIDF vector. Sending can be adaptive, for instance, if the crawlers use the PrTFIDF vector of a document and adjust the components of the sending weight vector, i.e., the sending weights, by averaging the weights of the relevant documents. In case of changing worlds, the moving window averaging might improve performance. Concerning the estimation of the sending weights two cases can be distinguished: the crawler finds a document and sends it (i) to the RA or (ii) to another crawler. Thus, if a crawler receives a document from another crawler and sends it to the RA then the value of this document must have passed the sending thresholds of the two crawlers. This condition can be weakened provided that the sending weights of a crawler are communicated to its partners.

2.5. Topic-specific crawlers

A crawler is topic specific if the sent document for which it is rewarded belongs only to a given topic. In the experiments, crawlers seek either novel information, irrespective of topic, or novel information in different topics. The topics are defined by means of keyphrases. The keyphrases of a document are determined by the application of the keyphrase-extracting algorithm (KEA, for short) [16]. The reward condition is fulfilled in the simulations, in case the extracted keyphrase set of the document sent by the crawler contains all keyphrases of a given topic. Keyphrases were selected meticulously: at least 100 and at most 2000 documents belong to each topic in our downloaded database.

3. Experiments

An 18-day experiment amassing real data has been conducted on the web in order to compare the weblog update and the RL algorithms. Crawlers have collected the data from the Internet. Some of them utilized the weblog algorithm, the others the RL algorithm or the combined RL-based crawling and weblog update-based restart algorithm, so as to eliminate the biases from the gathered data.

3.1. Web

The experiment has been performed on a single personal computer. The forager architecture is implemented in Java. In course of the experiment, a fixed number of foragers compete against each other to collect news items from the CNN web site. They crawl along their path in quest of novel information for equal time intervals according to a predefined order. After the given period has elapsed, the forager has to finish the step that it has started.

A fleet of foragers employing the weblog-based restarts, the RL-based crawling, and their combination, are deployed in the simulation.

The link structure of the gathered web pages follows the power-law distribution ($P(k) = k^{\gamma}$): γ equals to -1.45 for the outgoing links and -3.04 for the incoming ones, as it is demonstrated by Fig. 1. Thus, it can be concluded that the link structure has the scale-free property. The clustering coefficient [17] of the link structure is 0.02 and the diameter of the graph is 7.2893. The link structure has been reordered as follows: the random permutations have been applied to the starting and to the end points of the links. The reordered links have the same distribution. The magnitude of the clustering coefficient of the new generated graph is lower by an order

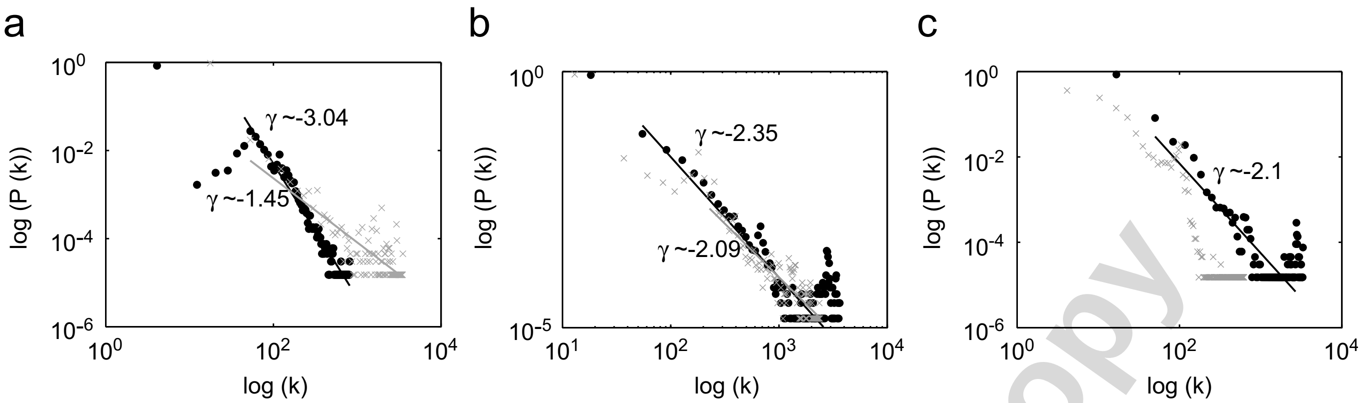

Fig. 1. Scale-free properties of the environments. Log–log scale distribution of the number of (incoming and outgoing) links of all URLs found in course of the investigation. Horizontal axis: number of edges (log k). Vertical axis: relative frequency of the number of edges at different URLs (log P(k)). Dots and dark line correspond to outgoing links, crosses and gray line to incoming links: (a) downloaded portion of the Internet (SFSW environment); (b) rewiring of outgoing links of the new nodes with preferential attachment algorithm (SFW environment); (c) random rewiring of outgoing links of new nodes (RWE). For further details, see the text.

of magnitude than that of the original graph, whilst the diameter of the new graph remains the same. It implies that the links of the amassed pages form a small world structure.

The pages for the experiment are stored in a centralized component with two indices (and time stamps): the URL and the page index. Multiple instances of a page may have the same URL index provided that it has been downloaded from the same URL and the content of the page has been changed. The page index, in effect, uniquely identifies the page content and the URL belonging to it. For more details, see [4,15].

3.2. Tasks, crawlers, and simulated web environment

Five types of experiments have been conducted:

(1) Experiments without communication:
  (a) there is no communication between the crawlers. The documents are claimed to be relevant if they are found within 24 h of their respective time stamps. In this experiment, our previous results [4] have been reproduced ('+' signs in Fig. 2, case 'reproduced').
  (b) topic-specific experiments without communication ('○' signs in Fig. 2, case 'no comm').
(2) Topic-specific experiments with communication:
  (a) both types of document sending, i.e., sending to the RA and sending to other crawlers, are adaptive ('∇' signs in Fig. 2, case 'learn all')
  (b) sending of a document to the RA is adaptive, but the situation is more straightforward: each crawler sends its learned weight to the other crawlers, and crawlers receiving the weights employ them in the direct communication of the documents ('△' signs in Fig. 2, case 'send learned')
  (c) averaged weights of the previous experiment are used in both types of communications ('×' signs in Fig. 2, case 'good all').

Three different environments have been generated for the simulations:

(1) SFSW: each simulated page has exactly the same links as the original page on the web does.
(2) SFW: the number of links of the new simulated pages equals to that of the original pages on the web. New documents are added according to their time stamps. A new document keeps the number of its links. The targets of these links are selected from the already existing simulated web via the preferential attachment algorithm.
(3) RWE: similar to the SFW except that the targets of the links of a new document are selected randomly from the already existing simulated web according to the uniform distribution.

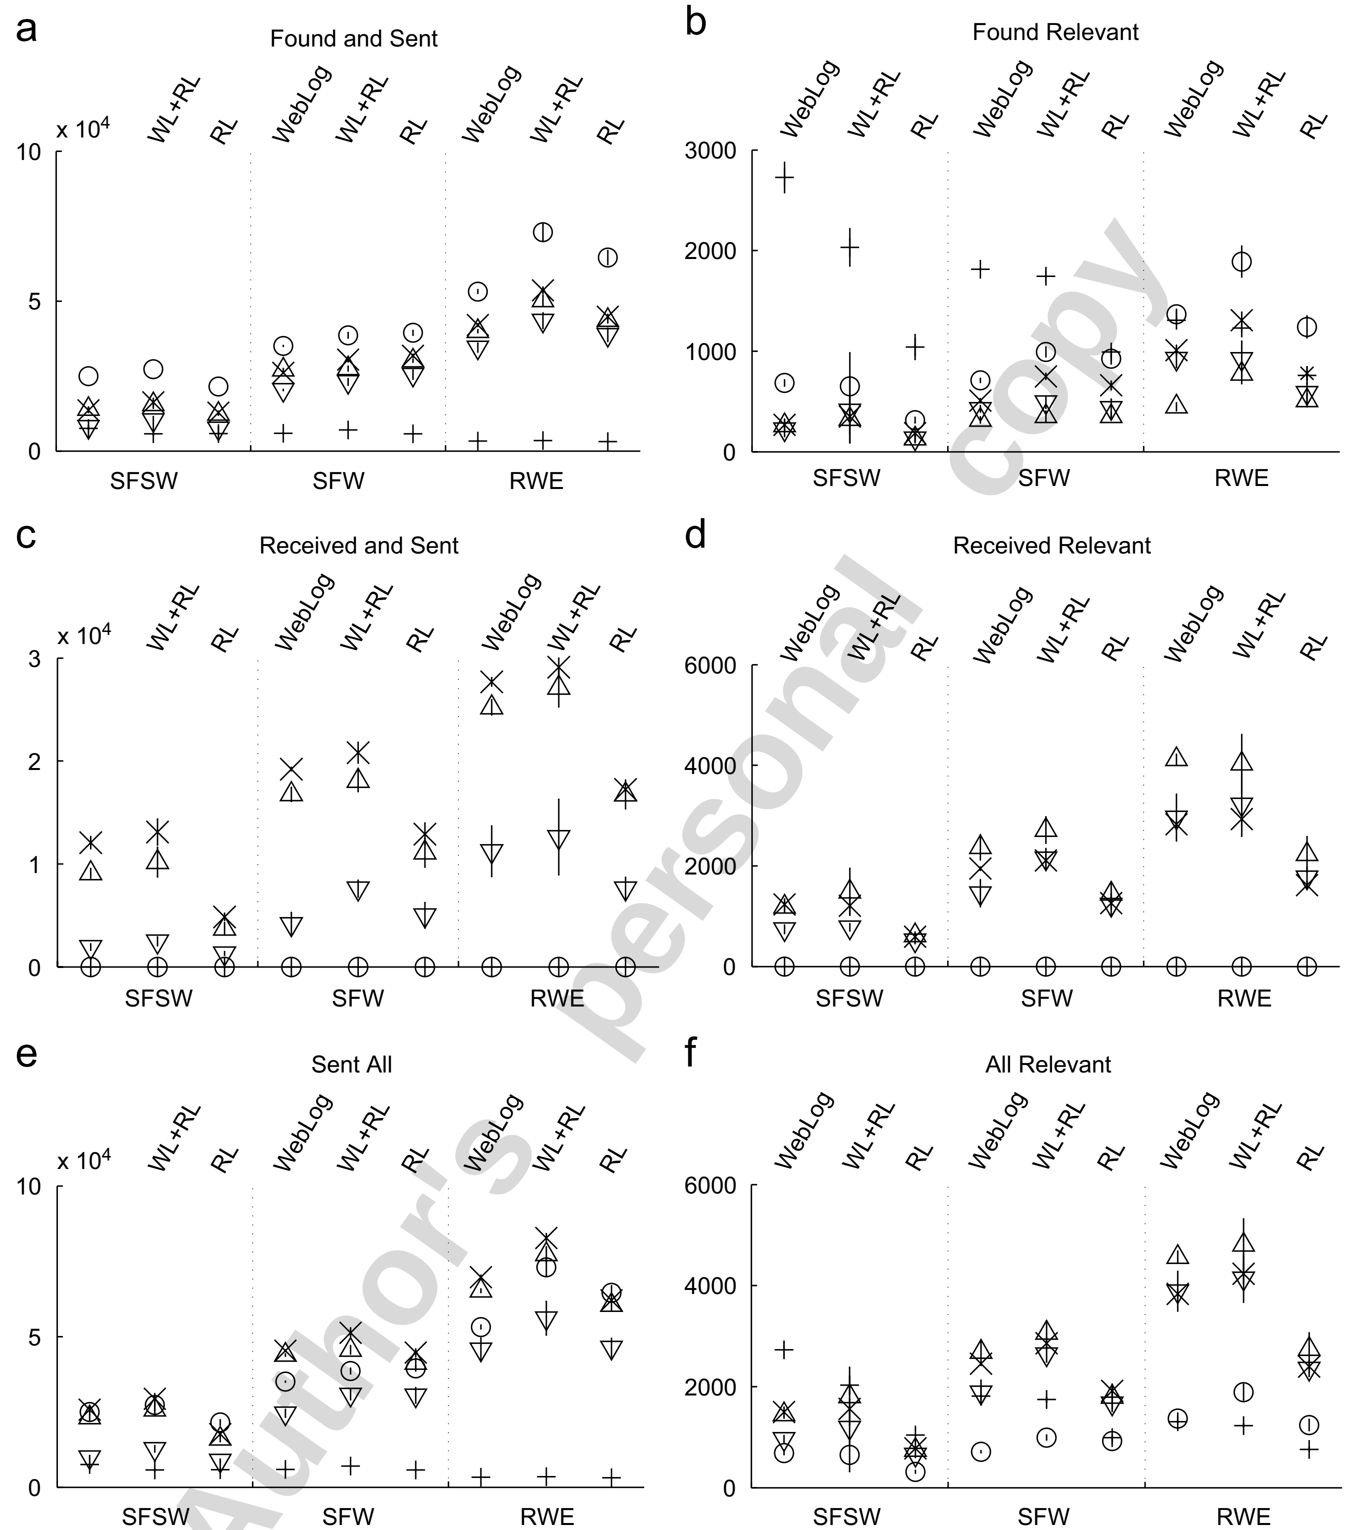

Fig. 2. Results in different worlds, with and without communication, and with different learning mechanisms. Notations: '+': 'reproduced', '○': 'no comm', '∇': 'learn all', '△': 'send learned', and '×': 'good all', SFSW: scale-free small world, SFW: scale-free world, RWE: random world environment, weblog: crawling applying the weblog algorithm, WL: abbreviation for weblog, RL: crawling utilizing reinforcement learning, WL + RL: weblog employing good restarts, crawling using RL. Vertical lines: standard deviations of 12 simulations for each case. For more details, see the text.

As for the downloaded portion of the environment it is a small world and it demonstrates the scale-free property (Fig. 1). Consequently, it can be considered as an SFSW environment. The clustering coefficient of the SF environment (Fig. 1), on the other hand, is 10 times smaller than that of the SFSW environment. Nevertheless, it is also scale free concerning both the incoming and the outgoing degree distributions. The clustering coefficient of the random RWE environment is also 10 times smaller than that of the SFSW environment. The outgoing link degree distribution of this environment is scale free, whilst the incoming link degree distribution is exponential (Fig. 1) owing to the uniform selection of the linked documents. This environment can be regarded as the most random in the sense that all of the free parameters (linked documents) have been selected according to the uniform random distribution.

3.3. Simulations

The results of the simulations are summarized in Fig. 2.

In all cases, the number of the downloaded documents are approximately the same ($\sim 2 \times 10^6$ downloads). Fig. 2(a) shows the number of found documents that have been sent to the RA. Fig. 2(b) illustrates the number of documents that have been come across by the crawlers, sent to, and reinforced by the RA. Fig. 2(c) depicts the number of found documents that have been received by the crawlers. Fig. 2(d) demonstrates the number of documents received by the crawlers ,forwarded to and reinforced by the RA. Fig. 2(e) and (f) represent the sums of values of Fig. 2(a) and (c), and Fig. 2(b) and (e), respectively.

4. Discussion and conclusions

It can be seen easily that the performance of different methods may differ considerably in different environments.

Our previous results have been reproduced (sign '+'): in SFSW the selective weblog (WL) algorithm is superior to RL, and its performance will be deteriorated if RL is added to it (case WL + RL) . On the other hand, in SFW RL does not ruin WL. Should the crawlers not communicate, the SFSW environment will perform the best, the RWE the least.

In the topic-specific cases, studies are restricted by the sparsity of data and time limits have not been imposed: all novel documents have been rewarded, irrespective of their time stamps. Consider first the case without communication (sign $\circ$). The trend concerning RL remains the same: in SFSW, RL is much worse than WL, whereas supplying WL with RL (case WL + RL) slightly influences the performance of WL. On the contrary, in SFW RL + WL is the best, followed by RL and WL. In RWE, it is WL + RL that performs the best and this combined method collects twice as much reward in this environment as in SFW.

As for the methods employing communication, it can be observed that for all algorithms, RWE is the easiest, followed by SFW and SFSW. The original database is the hardest, although the SFSW environment is seemingly not randomly organized (Fig. 2(f)). Furthermore, in all worlds, the combined algorithm accomplishes the search the best. The learning and the communication of the weights are trivially advantageous (sign '$\triangle$'). Nevertheless, the averaged weights perform similarly well in SFW worlds (sign '$\times$'). If the crawlers do not communicate their weights, nor are the averaged weights available, then the transfer of (a fraction of the) reward received from the RA makes it possible for the crawlers to learn the sending weights. This extra learning worsens the performance somehow (sign '∇', see (Fig. 2(c) and (d))), but the learning of these weights is (i) highly advantageous in most worlds and (ii) its performance is slightly better than that of the average weights (sign '$\times$'). The weblog algorithm in SFSW may be regarded as an extreme, the extra learning of the sending weights does not contribute considerably to the performance of the crawlers regarding the number of collected documents.

It is difficult to compare the different worlds to each other as for the topic-specific case, since the topics are well organized in the original news site and this arrangement will be destroyed when the topology is changed. Furthermore, the comparison of the topic-specific and the non-topic-specific cases requires strenuous efforts owing to the application of the time limit in the latter case: the novel documents become obsolete beyond that limit. It cannot be inferred from our experiments to what extent the lack of time constraint influences the efficiency of topic-specific crawlers in RWE.

Further studies are required to establish the cause of the high performances of the topic-specific cases. It is possible that the overlaps between the paths are smaller, i.e., compartmentalization is more pronounced, for the topic-specific case.

In conclusion, we found that selective learning can be improved by RL both with and without communication. The SFSW has proven to be the only exception, due to the fact that in the absence of communication, RL decreases the efficiency of the selective algorithm. In our previous studies [4,15], it was shown that this diminution may be explained by the fact that compartmentalization (niche formation) occurs easily in SFSW. The extension to topic-specific searches (to environments that require work sharing) and to the communication diminishes the advantages of selective learning in all worlds, including SFSW.

References

[1] L.C. Freeman, The Development of Social Network Analysis: A Study in the Sociology of Science, BookSurge, LLC, North Charlston, SC, 2004.
[2] M. Newman, A.-L. Barabasi, D.J. Watts, The Structure and Dynamics of Networks, Princeton University Press, Princeton, NJ, 2006.
[3] R. Sun (Ed.), From Cognitive Modeling to Social Simulation, Cambridge University Press, Cambridge, UK, 2006.
[4] Zs. Palotai, Cs. Farkas, A. Lőrincz, Is selection optimal in scale-free small worlds?, ComPlexUs 3 (2006) 158–168.
[5] E. Pachepsky, T. Taylor, S. Jones, Mutualism promotes diversity and stability in a simple artificial ecosystem, Artif. Life 8 (1) (2002) 5–24.
[6] J.M. Fryxell, P. Lundberg, Individual Behavior and Community Dynamics, Chapman & Hall, London, 1998.
[7] C.W. Clark, M. Mangel, Dynamic State Variable Models in Ecology: Methods and Applications, Oxford University Press, Oxford UK, 2000.
[8] G. Kampis, Self-Modifying Systems in Biology and Cognitive Science: A New Framework for Dynamics, Information and Complexity, Pergamon, Oxford, UK, 1991.
[9] R. Sutton, A.G. Barto, Reinforcement Learning: An Introduction, MIT Press, Cambridge, MA, 1998.
[10] G.J. Tesauro, Temporal difference learning and TD-Gammon, Commun. ACM 38 (1995) 58–68.
[11] M.J. Mataric, Reinforcement learning in the multi-robot domain, Autonomous Robots 4 (1) (1997) 73–83.
[12] I. Kókai, A. Lőrincz, Fast adapting value estimation based hybrid architecture for searching the world-wide web, Appl. Soft Comput. 2 (2002) 11–23.
[13] A. Lőrincz, I. Kókai, A. Meretei, Intelligent high-performance crawlers used to reveal topic-specific structure of the WWW, Int. J. Found. Comp. Sci. 13 (2002) 477–495.
[14] T. Joachims, A probabilistic analysis of the Rocchio algorithm with TFIDF for text categorization, in: D.H. Fisher (Ed.), Proceedings of ICML-97, the 14th International Conference on Machine Learning, Nashville, USA, Morgan Kaufmann, San Francisco, USA, 1997, pp. 143–151.
[15] Zs. Palotai, S. Mandusitz, A. Lőrincz, Computer study of the evolution of 'news foragers' on the Internet, in: C. Grosan, A. Abraham, V. Ramos (Eds.), Swarm Intelligence and Data Mining, Springer SCI, vol. 34, Springer, Berlin, Germany, 2006, pp. 203–220.
[16] E. Frank, G.W. Paynter, I H. Witten, C. Gutwin, C.G. Nevill-Manning, Domain-specific keyphrase extraction, 16th International Joint Conference on Artificial Intelligence, San Francisco, CA, Morgan Kaufmann, Los Altos, CA 1999, pp. 668–673.
[17] D.J. Watts, S.H. Strogatz, Collective dynamics of 'small-world' networks, Nature 393 (6684) (1998) 440–442.