

Is selection optimal for scale-free small worlds?

Zs. Palotai¹, Cs. Farkas², A. Lőrincz^{1*}

August 29, 2005

¹Dept. of Information Systems, Eötvös Loránd University, Budapest, Hungary

²Dept. of Computer Science and Engr., University of South Carolina, SC, USA

Abstract

The ‘No Free Lunch Theorem’ claims that for the set of all problems no algorithm performs better than random search and thus, selection can be advantageous only on a limited set of problems. In this paper we investigate how the topological structure of the environment influences algorithmic efficiency. We study the performances of algorithms, using selective learning, reinforcement learning, and their combinations, in random, scale-free, and scale-free small world (SFSW) environments. The learning problem is to search for novel, not-yet-found information. We ran our experiments on a large news site and on its downloaded portion. Controlled experiments were performed on this downloaded portion: we modified the topology, but preserved the publication time of the news. Our empirical results show that the selective learning is the most efficient in SFSW topology. In non-small world topologies, however, the combination of the selective and reinforcement learning algorithms performs the best.

Keywords: scale-free small world, selection, NFL theorem, evolution, Internet

Running head: Selection in SFSW

*Corresponding author: Eötvös Loránd University, Department of Information Systems, Pázmány Péter sétány 1/c, Budapest, H1117, Hungary. Email: andras.lorincz@elte.hu, tel.: +36 1 209 0555 / 8473, fax: + 36 1 381 2140

1 Introduction

According to the ‘No Free Lunch Theorem’ (NFL Theorem) [1] there is no performance difference between optimization or search algorithms if we test the algorithms on every possible problem. Therefore, it seems that improved performances of specific algorithms, such as selection, are consequences of specific problem properties. Finding these properties will aid the development of optimized solutions. Recent research shows that evolving structures, both natural and artificial (like the Web), exhibit scale-free or scale-free small world properties [2, 3]. In [4] the authors show that a particular evolutionary algorithm has better performance in artificial scale-free environment than in lattice, small world, or random environments. That is, the structure of the environment has an impact on the efficiency of the algorithms.

We have designed a controllable experiment and compared the performances of different algorithms on different structures. We aimed to keep the complexity of experiments and the number of designer specifiable parameters minimal. We considered the structure of the environment, the algorithms, and the fitness values. Our choice for the structure was the World Wide Web (WWW). WWW is considered the largest source of rapidly changing data. The WWW has a scale-free small world (SFSW) structure [2, 5].

The everyday usage of the Web is the search for novel information. Therefore, we have a natural reward function; the number of novel documents that the algorithms can find. We consider this property as one of the most important components of our work: *We did not specify the temporal and structural details of the reward system.*

We ran controlled experiments on a time stamped and downloaded portion of a large WWW news site. This allowed us to preserve the temporal structure of the rewards, but also supported the modification of the underlying connectivity structure, i.e., how the novel information can be found.

Our agents were Web crawlers, or foragers. Crawlers travel from link to link foraging new, not-yet-seen information. Our agents used the simplest possible algorithms. Selective learning concerned the selection and memorization of good links. We have also used a simple version of reinforcement learning (RL). For a review on RL see, e.g., [6]. RL was used alone, and was also combined with selective learning. Our choice of RL is motivated by its structural and algorithmic simplicity and that RL is concerned with the optimization of the expected value of long-term cumulated profit.

For a recent review on evolutionary computing, see [7]. For reviews on related evolutionary theories and the dynamics of self-modifying systems see [8, 9] and [10, 11], respectively. Similar concepts have been studied in other evolutionary systems where organisms compete for space and resources and cooperate through direct interaction (see, e.g., [12] and references therein.)

Hybrid algorithms have a long history. A well-known example is the TD-Gammon program of Tesauro [13]. The author applied MLP function approximators for value estimation in RL. Reinforcement learning has also been used in concurrent learning problems like ours: robots had to learn to forage together

via direct interaction [14]. Another combinations of the algorithms concerns evolutionary learning embedded into the framework of RL to improve decision making [15, 16, 17, 18].

It is important to note, that communication and competition among our foragers are indirect. Only the first submitter of a document may receive positive reinforcement and this is the only interaction among the crawlers. Our work is different from other studies using combinations of genetic, evolutionary, function approximation, and reinforcement learning algorithms, in that i) it does not require explicit fitness function, ii) we do not have control over the original environment, iii) we can change the environment in a reproducible fashion, iv) collaborating individuals use value estimation under ‘evolutionary pressure’, and v) individuals work without direct interaction with each other. The crawler system is a self-assembling system, which is made of adaptive components, and the communication between components is kept as little as possible.

The paper is organized as follows. We review the related web crawler tools, including those [19, 20, 21] that our work is based upon, in Section 2. We describe our algorithms and the forager architecture in Section 3. This section contains the necessary algorithmic details imposed by the task, the search of the Web. We present our experiments on the Web and the controlled simulations in Section 4. Discussions can be found in Section 5. Conclusions are drawn in Section 6.

2 Related work

There are important problems that have been studied in the context of crawlers. Angkawattanawit, Rungsawang [28], and Menczer [29] study topic specific crawlers. Risvik et al. [30] address research issues related to the exponential growth of the Web. Cho and Gracia-Molina [31], Menczer [29] and Edwards et. al [32] study the problem of different refresh rates of URLs (possibly as high as hourly or as low as yearly).

An introduction to and a broad overview of topic specific crawlers are provided in [28]. They propose to learn starting URLs, topic keywords and URL ordering through consecutive crawling attempts. They show that the learning of starting URLs and the use of consecutive crawling attempts can increase the efficiency of the crawlers. The used heuristic is similar to the weblog algorithm [21], which also finds good starting URLs and periodically restarts the crawling from the newly learned ones. The main limitation of this work is that it is incapable of addressing the freshness (i.e., modification) of already visited Web pages.

Menczer [29] describes some disadvantages of current Web search engines on the dynamic Web, e.g., the low ratio of fresh or relevant documents. He proposes to complement the search engines with intelligent crawlers, or web mining agents to overcome those disadvantages. He introduces the InfoSpider architecture that uses genetic algorithm and reinforcement learning, also describes the MySpider implementation of it, which starts from the 100 top pages of AltaVista. Our

weblog algorithm uses local selection for finding good starting URLs for searches, thus not depending on any search engines. Dependence on a search engine can be a suffer limitation of most existing search agents, like MySpiders. Note however, that it is an easy matter to combine the present algorithm with URLs offered by search engines.

Risvik and Michelsen [30] overview different dimensions of web dynamics and show the arising problems in a search engine model. The main part of the paper focuses on the problems that crawlers need to overcome on the dynamic Web. As a possible solution the authors propose a heterogenous crawling architecture. The main limitation of their crawling architecture is that they must divide the web to be crawled into distinct portions manually before the crawling starts. A weblog like distributed algorithm – as suggested here – may be used in that architecture to overcome this limitation.

Cho and Garcia-Molina [31] define mathematically the freshness and age of documents of search engines. They propose the Poisson process as a model for page refreshment. The authors also propose various refresh policies and study their effectiveness both theoretically and on real data. They present the optimal refresh policies for their freshness and age metrics under the Poisson page refresh model. The authors show that these policies are superior to others on real data, too. Although they show that in their database more than 20 percent of the documents are changed each day, they disclosed these documents from their studies. Their crawler visited the documents once each day for 5 months, thus can not measure the exact change rate of those documents. While in our work we definitely concentrate on these frequently changing documents.

3 Forager architecture

There are two different kinds of agents: the foragers and the reinforcing agent (RA). The fleet of foragers crawls the web and sends the URLs of the selected documents to the reinforcing agent. The RA determines which forager should work for the RA and how long a forager should work. The RA sends reinforcements to the foragers based on the received URLs.

We employ a fleet of foragers to study the competition among individual foragers. A forager has simple and limited capabilities, like a stack for a limited number of starting URLs and a simple, content based URL ordering. The foragers compete with each other for finding the most relevant documents. In this way they efficiently and quickly collect new relevant documents without direct interaction.

At first we present the basic algorithms, followed by the algorithms for the reinforcing agents and the foragers.

3.1 Algorithms

Our constraints on finding the minimal set of algorithms were as follows: The algorithms should (i) allow the identification of unimportant parameters, (ii)

support the specialization of the individuals (the foragers), (iii) allow the joining of evolutionary learning and individual learning, (iv) minimize communication as much as possible. We shall return to these points in Section 5.

Weblog algorithm and starting URL selection

A forager periodically restarts from a URL randomly selected from the list of starting URLs. The sequence of visited URLs between two restarts forms a path. The starting URL list is formed from the 10 first URLs of the weblog. In the weblog there are 100 URLs with their associated weblog values in descending order. The weblog value of a URL estimates the expected sum of rewards during a path after visiting that URL. The weblog update algorithm modifies the weblog before a new path is started. The weblog value of a URL already in the weblog is modified toward the sum of rewards ($sumR$) in the remaining part of the path after that URL:

$$newValue = (1 - \beta)oldValue + \beta sumR,$$

where β was set to 0.3. A new URL has the value of actual sum of rewards in the remaining part of the path. If a URL has a high weblog value it means that around that URL there are many relevant documents. Therefore it may worth it to start a search from that URL.

Without the weblog update algorithm the weblog and thus the starting URL list remains the same throughout the searches. The weblog algorithm is a very simple version of evolutionary algorithms. Here, evolution may occur at two different levels: the list of URLs of the forager is evolving by the reordering of the weblog. Also, a forager may multiply, and its weblog, or part of it may spread through inheritance. This way, the weblog algorithm incorporates the basic features of evolutionary algorithms. This simple form shall be satisfactory for our purposes.

Reinforcement Learning based URL ordering

A forager can modify its URL ordering based on the received reinforcements of the sent URLs. The (immediate) profit is the difference of received rewards and penalties at any given step. Immediate profit is a myopic characterization of a step to a URL. Foragers have an adaptive continuous value estimator and follow the *policy* that maximizes the expected long term cumulated profit (LTP) instead of the immediate profit. Such estimators can be easily realized in neural systems [6, 22, 23]. Policy and profit estimation are interlinked concepts: profit estimation determines the policy, whereas policy influences choices and, in turn, the expected LTP. (For a review, see [6].) Here, choices are based on the greedy LTP policy: The forager visits the URL, which belongs to the *frontier* (the list of linked but not yet visited URLs, see later) and has the highest estimated LTP.

In the particular simulation each forager has a $k(= 50)$ dimensional probabilistic term-frequency inverse document-frequency (PrTFIDF) text classifier

[24], generated on a previously downloaded portion of the Geocities database. Fifty clusters were created by Boley’s clustering algorithm [25] from the downloaded documents. The PrTFIDF classifiers were trained on these clusters plus an additional one, the $(k + 1)^{th}$, representing general texts from the internet. The PrTFIDF outputs were non-linearly mapped to the interval $[-1, +1]$ by a hyperbolic-tangent function. The classifier was applied to reduce the texts to a small dimensional representation. The output vector of the classifier for the page of URL A is **state**(\mathbf{A}) = $(state(A)_1, \dots, state(A)_k)$. (The $(k + 1)^{th}$ output was dismissed.) This output vector is stored for each URL.

A linear function approximator is used for LTP estimation. It encompasses k parameters, the *weight vector* **weight** = $(weight_1, \dots, weight_k)$. The LTP of document of URL A is estimated as the scalar product of **state**(\mathbf{A}) and **weight**: $value(A) = \sum_{i=1}^k weight_i state(A)_i$. During URL ordering the URL with highest LTP estimation is selected. (For more details, see, [21].)

The weight vector of each forager is tuned by temporal difference learning (TD-learning) [26, 22, 23]. Let us denote the current URL by URL_n , the next URL to be visited by URL_{n+1} , the output of the classifier for URL_j by **state**(\mathbf{URL}_j) and the estimated LTP of a URL URL_j by $value(URL_j) = \sum_{i=1}^k weight_i state(URL_j)_i$. Assume that leaving URL_n to URL_{n+1} the immediate profit is r_{n+1} . Our estimation is perfect if $value(URL_n) = value(URL_{n+1}) + r_{n+1}$. Future profits are typically discounted in such estimations as $value(URL_n) = \gamma value(URL_{n+1}) + r_{n+1}$, where $0 < \gamma < 1$. The error of value estimation is

$$\delta(n, n + 1) = r_{n+1} + \gamma value(URL_{n+1}) - value(URL_n).$$

We used throughout the simulations $\gamma = 0.9$. For each step $URL_n \rightarrow URL_{n+1}$ the weights of the value function were tuned to decrease the error of value estimation based on the received immediate profit r_{n+1} . The $\delta(n, n + 1)$ estimation error was used to correct the parameters. The i^{th} component of the weight vector, $weight_i$, was corrected by

$$\Delta weight_i = \alpha \delta(n, n + 1) state(URL_n)_i$$

with $\alpha = 0.1$ and $i = 1, \dots, k$. These modified weights would improve value estimation in stationary and observable environments (see, e.g, [6] and references therein), but were also found efficient in large Web environments [21].

Without the reinforcement learning based URL ordering update algorithm the weight vector remains the same throughout the search.

Document relevancy

A document or page is possibly relevant for a forager if it is not older than 24 hours and the forager has not marked it previously. The selected documents are sent to the RA for further evaluation.

Multiplication of a forager

During multiplication the weblog is randomly divided into two equal sized parts (one for the original and one for the new forager). The parameters of the URL ordering algorithm (the weight vector of the value estimation) are either copied or new random parameters are generated. If the forager has a URL ordering update algorithm then the parameters are copied. If the forager does not have any URL ordering update algorithm then new random parameters are generated.

3.2 Reinforcing agent

A reinforcing agent controls the ‘life’ of foragers. It can start, stop, multiply or delete foragers. RA receives the URLs of documents selected by the foragers, and responds with reinforcements for the received URLs. The response is 100 (a.u.) for a relevant document and -1 (a.u.) for a not relevant document. A document is relevant if it is not yet seen by the reinforcing agent and it is not older than 24 hours. The reinforcing agent maintains the score of each forager working for it. Initially each forager has 100 (a.u.) score. When a forager sends a URL to the RA, the forager’s score is decreased by 0.05. After each relevant page sent by the forager, the forager’s score is increased by 1.

When the forager’s score reaches 200 and the number of foragers is smaller than 16 then the forager is multiplied. That is a new forager is created with the same algorithms as the original one has, but with slightly different parameters. When the forager’s score goes below 0 and the number of foragers is larger than 2 then the forager is deleted. Note that a forager can be multiplied or deleted immediately after it has been stopped by the RA and before the next forager is activated.

Foragers on the same computer are working in time slices one after each other. Each forager works for some amount of time determined by the RA. Then the RA stops that forager and starts the next one selected by the RA.

3.3 Foragers

A forager is initialized with parameters defining the URL ordering, and either with a weblog or with a seed of URLs. After its initialization a forager crawls in search paths, that is after a given number of steps the search restarts and the steps between two restarts form a path. During each path the forager takes 100 steps, i.e., selects the next URL to be visited with a URL ordering algorithm. At the beginning of a path a URL is selected randomly from the starting URL list. This list is formed from the 10 first URLs of the weblog. The weblog contains the possibly good starting URLs with their associated weblog values in descending order. The weblog algorithm modifies the weblog and so thus the starting URL list before a new path is started. When a forager is restarted by the RA, after the RA has stopped it, the forager continues from the internal state in which it was stopped.

The URL ordering algorithm selects a URL to be the next step from the frontier URL set. The selected URL is removed from the frontier and added to the visited URL set to avoid loops. After downloading the pages, only those URLs (linked from the visited URL) are added to the frontier which are not in the visited set.

In each step the forager downloads the page of the selected URL and all of the pages linked from the page of selected URL. It sends the URLs of the possibly relevant pages to the reinforcing agent. The forager receives reinforcements on any previously sent but not yet reinforced URLs and calls the URL ordering update algorithm with the received reinforcements.

4 Experiments

We conducted an 18 day long experiment on the Web to gather realistic data. We used the gathered data in simulations to compare the weblog update (Section 3.1) and reinforcement learning algorithms (Section 3.1). In the Web experiment we used a fleet of foragers using combination of reinforcement learning and weblog update algorithms to eliminate possible biases on the gathered data. First we describe the experiment on the Web then the simulations. We analyze our results in the next section.

4.1 Data collection on the Web

We ran the experiment on the Web on a single personal computer with Celeron 1000 MHz processor and 512 MB RAM. We implemented the forager architecture (described in Section 3) in Java programming language.

In this experiment a fixed number of foragers were competing with each other to collect news at the CNN web site. The foragers were running in equal time intervals in a predefined order. Each forager had a 3 minute time interval and after that interval the forager was allowed to finish the step started before the end of the time interval. We deployed 8 foragers using the weblog update and the reinforcement learning based URL ordering update algorithms (8 WR foragers). We also deployed 8 other foragers using the weblog update algorithm but without reinforcement learning (8 WL foragers). The predefined order of foragers was the following: 8 WR foragers were followed by the 8 WL foragers.

We investigated the link structure of the gathered Web pages. As it is shown in Fig. 1 the links have a power-law distribution ($P(k) = k^\gamma$) with $\gamma = -1.3$ for outgoing links and $\gamma = -2.57$ for incoming links. That is the link structure has the scale-free property. The clustering coefficient [27] of the link structure is 0.02 and the diameter of the graph is 7.2893. We applied two different random permutations to the origin and to the endpoint of the links, keeping the edge distribution unchanged but randomly rewiring the links. The new graph had 0.003 clustering coefficient and 8.2163 diameter. That is the clustering coefficient was smaller than the original value by an order of magnitude, but the diameter is

almost the same. Therefore we can conclude that the links of gathered pages form *scale-free small world* structure.




Figure 1: **Scale-free property of the Internet domain.** Log-log scale distribution of the number of (incoming and outgoing) links of all URLs found during the time course of investigation. Horizontal axis: number of edges ($\log k$). Vertical axis: relative frequency of number of edges at different URLs ($\log P(k)$). Dots and dark line correspond to outgoing links, crosses and gray line correspond to incoming links.

The data storage for simulation is a central issue in our experiments. Pages are stored with 2 indices (and time stamps). One index is the URL index, the other is the page index. Multiple pages can have the same URL index if they were downloaded from the same URL. The page index uniquely identifies a page content and the URL from where the page was downloaded. For any foragers, at each page download we stored the followings (with a time stamp containing the time of page download):

1. if the page is relevant according to the RA then store 'relevant'
2. if the page is from a new URL then store the new URL with a new URL index and the page's state vector with a new page index
3. if the content of the page is changed since the last download then store the page's state vector with a new page index but keep the URL index
4. in both previous cases store the links of the page as links to page indices of the linked pages
 - (a) if a linked page is from a new URL then store the new URL with a new URL index and the linked page's state vector with a new page index
 - (b) if the content of the linked page is changed since the last check then store the page's state vector with a new page index but same URL index

4.2 Simulations

For the simulations we implemented the forager architecture in Matlab. The foragers were simulated as if they were running on one computer as described in the previous section.

Simulation specification

During simulations we used the Web pages that we gathered previously to generate different environments (note that the links of pages point to local pages (not to pages on the Web) since a link was stored as a link to a local page index):

- Simulated documents had the same state vector representation for URL ordering as the real pages had
- Simulated relevant documents were the same as the relevant documents on the Web
- Pages and links appeared at the same (relative) time when they were found in the Web experiment - using the new URL indices and their time stamps
- Pages and links are refreshed or changed at the same relative time as the changes were detected in the Web experiment – using the new page indices for existing URL indices and their time stamps
- Simulated time of a page download was the average download time of a real page during the Web experiment.

We generated 4 different environments for the simulations:

1. **SFSW**: each simulated page had exactly the same links as the original page had on the Web (a simulated page linked those simulated pages, page indices of those pages, which were linked by the original Web page).
2. **SF1**: in each second the new simulated pages had the same number of links as the original pages on the Web. A new simulated page linked to simulated pages selected by the preferential attachment algorithm from the existing simulated pages.
3. **SF2**: the previous algorithm applied for the **SF1** environment.
4. **SFRandom**: similar to the **SF1**, but the linked simulated pages are selected from a uniform distribution of the pages.

The SFSW environment has exactly the same scale-free and small world properties as the web environment downloaded by the web foragers. The SF1 and SF2 environments have 10 times smaller clustering coefficient than the SFSW environment has. These environments have scale-free degree distributions, although those are slightly different from the web environment (see Fig. 2(a)).

The SFRandom environment, also, has 10 times smaller clustering coefficient then SFSW. The SFRandom environment has scale-free outgoing link degree distribution. But because of the uniform selection of linked documents the incoming link degree distribution is exponential (see Fig. 2(b)). With the above given constraints this environment is the most random in the sense that all of the free parameters (linked documents) were selected from the uniform random distribution.




Figure 2: **Degree distribution of the environments.**

Dots and dark line correspond to outgoing link distribution. X-s and gray line correspond to incoming link distribution. (a) Upper: degree distributions of SF1 environment. Lower: degree distributions of SF2 environment. (b) : degree distributions of SFRandom environment.

We conducted simulations with four different kinds of foragers in each environment:

1. **WR foragers** used both the weblog update and the reinforcement learning based URL ordering update algorithms.
2. **WL foragers** used only the weblog update algorithm without URL ordering update. Each WL forager had a different weight vector for URL value estimation – during multiplication the new forager got a new random weight vector.
3. **RL foragers** used only the reinforcement learning based URL ordering update algorithm without the weblog update algorithm. RL foragers had the same weblog with the first 10 URLs of the gathered pages – that is the starting URL of the Web experiment and the first 9 visited URLs during that experiment.
4. **Fix foragers** did not use the weblog update and the reinforcement learning based URL ordering update algorithms. These foragers had fixed starting URLs and fixed weight vectors, but the latter was different for each Fixed forager.

In each case, initially there were 2 foragers and they were allowed to multiply until reaching the population of 16 foragers. The simulation for each type of foragers were repeated 3 times with different initial weight vectors for each forager. The variance of the results show that there is only a small difference between simulations using the same kind of foragers, even if the foragers were started with different random weight vectors in each simulation.

Simulation measurements

The first thing that we should note concerns the efficiency as a function of the number of crawlers. On a single computer, and under the time sharing method we applied, and without direct competition between the different crawlers, we found that bipartition gives rise to a transient decrease of the efficiency of the new crawlers, but it quickly recovers. Within the limits of the number of crawlers that we studied (between 2 and 24), performance of the fleet is a slowly increasing function of the number of crawlers. We fixed the number of the crawlers and this slow dependence did not enter our considerations. Table 1 shows the investigated parameters during simulations.

Table 1: Investigated parameters

downloaded	number of downloaded documents
sent	number of documents sent to the RA
relevant	number of found relevant documents
found URLs	number of found URLs
download efficiency	ratio of relevant to downloaded documents in 3 hour time window throughout the simulation.
sent efficiency	ratio of relevant to sent documents in 3 hour time window throughout the simulation.
exploration	ratio of found URLs to downloaded at the end of the simulation
freshness	ratio of the number of current found relevant documents and the number of all found relevant documents [31]. A stored document is current, up-to-date, if its content is exactly the same as the content of the corresponding URL in the environment.
age	A stored current document has 0 age, the age of an obsolete page is the time since the last refresh of the page on the Web [31].

Parameter ‘download efficiency’ is relevant for the site where the foragers should be deployed to gather the new information. Parameter ‘sent efficiency’ is relevant for the RA. Note that during simulations we are able to immediately and precisely calculate freshness and age values. In a real Web experiment this is impossible, because of the time needed to download and compare the contents of all of the real Web pages to the stored ones.

5 Discussion

We observed, that the efficiency of the algorithms depends strongly on the weight vectors. As we have mentioned above, the number of foragers had slight effects on the efficiency. This observation is supported by the fact that upon bipartition the weight vectors of the descendant foragers are similar, causing the descendant foragers to follow similar paths and to spoil the performance of the other. This is the reason that efficiency shows a transient decrease, but as a result of adaptation it quickly disappears.

The dependence of the efficiency on the number of starting points is not too much different. Two foragers could use 20 different starting points, whereas 16 foragers could use 160 different starting points, but the efficiency was only slightly influenced by this order of magnitude difference. That is, the number of starting points and the number of foragers are less important than the weight vectors in our simulations indicating that 20 starting points or so, which can adapt, can efficiently shatter the structure. Note that other parameters, e.g., larger number of crawlers, more than one computer can change this picture. The weak dependence, however, is most advantageous for our purpose namely to be confident that we can study structural dependencies.

The measured parameter values are presented in Fig. 3. The figure contains the values for each measured parameter for each type of forager in each type of environment. From the subfigures we can conclude the followings:

It can be seen in the top left and top middle subfigures that freshness and age values of different foragers are changing in the SFSW environment (marker x) while in the other 3 environments the values are almost the same for the different type of foragers. RL and Fix foragers can get trapped in clustered environments, can not easily escape (find outgoing links leaving the cluster) from clusters containing less relevant documents. The weblog algorithm provides a way to go to the found best clusters and in this way to escape from the worse clusters. This can be the reason why the RL and Fix foragers in the SFSW environment performs worse than WL and WR foragers.

In the top right subfigure it can be seen that finding new documents is the hardest in the SFRandom environment for all foragers. In this environment the pages has exponential incoming link degree distribution which means that there are relatively more small degree pages than in the other 3 scale-free environments. It is harder to get to pages which links to many pages. Therefore it is harder for the foragers to find not yet seen documents.

It can be seen in the middle left and middle subfigures that finding relevant or possibly relevant documents is the easiest in the SFSW environment for all foragers. This environment is more clustered than the other three environments. When a forager finds a relevant document in a cluster then it finds other relevant documents in that cluster while it can not escape from the cluster. The other 3 environments are less clustered. Foragers can get out from clusters easier by following the links and therefore have to forage the entire environment for new relevant documents.

In the bottom left subfigure it can be seen that the download efficiency is

the best in the SFSW environment for all foragers. This is because of the high number of found relevant documents compared to the other three environments.

In the bottom middle subfigure it can be seen that the WR forager's sent efficiency is the same in all environments. Although these foragers found less relevant documents in not SFSW environments but also sent back less documents to the Reinforcing Agent. The other three foragers sent more less documents to Reinforcing Agent in the not SFSW environments, therefore their sent efficiencies are the worst in the SFSW environment. Although these foragers also found the most relevant documents in the SFSW environment, compared to the other 3 environments.

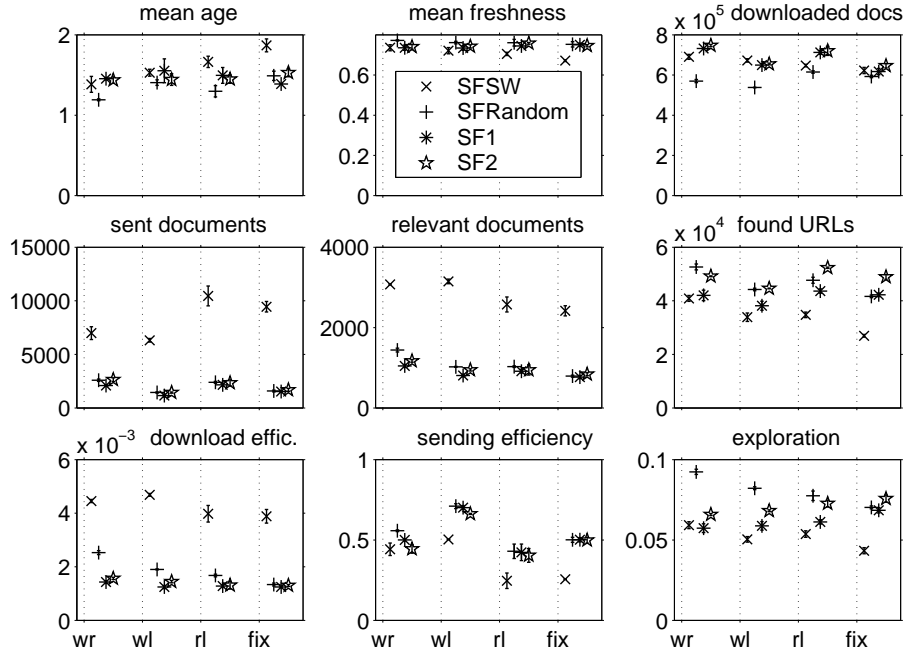


Figure 3: **Measured parameter values.**

The subfigures show the investigated parameters. Each subfigure contains the parameter values for the four type of foragers in four columns as shown below the bottom subfigures. The 4 different markers correspond to the measured parameter values in the 4 environments as shown in the top middle figure legend. On each marker an error bar shows the standard deviation of the corresponding parameter values for the 3 simulations. Mean age is in hours in the upper left subfigure.

It can be seen in the middle right subfigure that finding new URLs is the hardest in the SFSW environment because of its clustered nature. The foragers checks the same URLs for changed documents in the clusters therefore they can collect many new relevant documents. In the other three environments

foragers do not get trapped as much in the clusters and they search the whole environment continuously.

Now, consider Table 2, which contains some of the data of Fig. 3. The number of sent relevant documents is somewhat larger for the WL foragers than for WR foragers in SFSW environments (3149 and 3075, respectively, i.e., the difference is about 2.3%). This slight difference changes sign and becomes much more pronounced in all other environments. For example, in the SFRandom environment the numbers are 1029 and 1441, i.e., the difference is about 28%, but in the opposite direction. It is important to note that, if these foragers compete with each other, then the slight 2.3% difference or the larger 28% difference both enter the argument of an exponential because of bipartitions. Thus, even slight differences can become large, and may give rise to overwhelming population differences. Such competitive runs are under way. The issue becomes more pronounced for real situations, where time is not shared on a single computer and all foragers may search for food at all times.

Our results and the results of Annunziato et al. [4] can not be compared directly. The most obvious reason is that their investigations were restricted to SF, SW and random structures, but they did not study the SFSW structure, which plays a central role in our work. Artificial studies are, however, desirable, because in such examples one can finely gauge the different components and may find the necessary and sufficient conditions of our findings.

Table 2: Quantitative results for algorithm–structure pairs

	Weblog		Weblog and RL		RL	
	SFSW	SFR	SFSW	SFR	SFSW	SFR
No of sent docs	6313	1443	6985	2585	10455	2396
No of relevant docs	3149	1023	3075	1441	2575	1029
Sending efficiency	0.5035	0.7106	0.4425	0.5574	0.2463	0.4299
No of found URLs	33882	44217	40888	52636	34759	47668

We consider the following findings important: Bipartition gives rise to certain transient disadvantages if the descendants are similar. They will have to share the food until they start to learn. Still, selective learning can be more effective than selective learning combined with other methods, or than other methods alone if the environment is SFSW. It seems that selection fits the SFSW structure and vice versa. This could be a good reason for the abundance of emergent SFSW structures in nature. This can be understood through the NFL Theorem if it is read backwards: If we find that a particular algorithm is more efficient than random search or than any other algorithm, then this winning algorithm ‘knows’ (fits) the most the underlying problem among the investigated algorithms. Based on our computer simulations it seems that highly clustered SFSW structures and simple selective learning algorithms match each other.

6 Conclusions

We investigated algorithms using evolutionary, reinforcement learning, and combined evolutionary and reinforcement learning strategies. We experimented with environments having different degree distributions and clustering coefficients. We generated different topological environments, using data collected during real Web search. Our study focused on the task of searching for new relevant documents. We found that in the scale-free small world environment the evolutionary weblog update algorithm performs the best. It outperformed a reinforcement learning based algorithm and the combinations of these two algorithms. We conjecture that the highly clustered nature and the small diameter of the environment match simple selection over other more sophisticated learning schemes. However, when the scale-free nature of the environment was kept but the small diameter of the environment was increased simply by restructuring the environment, then other algorithms performed better than the simple selection. We found in the 3 not small world environments that the combination of the weblog and reinforcement learning algorithms are the best. That is, when the diameter of the world becomes larger, then estimation of the long-term cumulated reward becomes important. Moreover, the combined algorithm showed the smallest performance variation both on the scale-free small world and on scale-free environments.

7 Acknowledgements

This material is based upon work supported by the European Office of Aerospace Research and Development, Air Force Office of Scientific Research, Air Force Research Laboratory, under Contract No. FA8655-03-1-3036. This work is also supported by the National Science Foundation under grants No. INT-0304904 and No. IIS-0237782. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the European Office of Aerospace Research and Development, Air Force Office of Scientific Research, Air Force Research Laboratory.

References

- [1] WOLPERT DH, MACREADY WG: No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation* 1, 1 (April 1997), 67–82.
- [2] BARABÁSI A, ALBERT R, JEONG H: Scale-free characteristics of random networks: The topology of the world wide web. *Physica A* 281 (2000), 69–77.
- [3] ALBERT R, BARABÁSI A: Statistical mechanics of complex networks. *Reviews of Modern Physics* 74 (2002), 47–91.

- [4] ANNUNZIATO M, HUERTA R, LUCCHETTI M, TSIMRING LS: Artificial life optimization over complex networks. In *Fourth International ICSC Symposium on ENGINEERING OF INTELLIGENT SYSTEMS* (March 2004).
- [5] KLEINBERG J, LAWRENCE S: The structure of the web. *Science* 294 (2001), 1849–1850.
- [6] SUTTON R, BARTO A: *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, 1998.
- [7] EIBEN AE, SMITH JE: *Introduction to Evolutionary Computing*. Springer, 2003.
- [8] FRYXELL J, LUNDBERG P: *Individual Behavior and Community Dynamics*. Chapman and Hall, London, 1998.
- [9] CLARK C, MANGEL M: *Dynamic State Variable Models in Ecology: Methods and Applications*. Oxford University Press, Oxford UK, 2000.
- [10] KAMPIS G: *Self-modifying Systems in Biology and Cognitive Science: A New Framework for Dynamics, Information and Complexity*. Pergamon, Oxford UK, 1991.
- [11] CSÁNYI V: *Evolutionary Systems and Society: A General Theory of Life, Mind, and Culture*. Duke University Press, Durham, NC, 1989.
- [12] PACHEPSKY E, TAYLOR T, JONES S: Mutualism promotes diversity and stability in a simple artificial ecosystem. *Artificial Life* 8, 1 (2002), 5–24.
- [13] TESAUR GJ: Temporal difference learning and td-gammon. *Communication of the ACM* 38 (1995), 58–68.
- [14] MATARIC MJ: Reinforcement learning in the multi-robot domain. *Autonomous Robots* 4, 1 (1997), 73–83.
- [15] STAFYLOPATIS A, BLEKAS K: Autonomous vehicle navigation using evolutionary reinforcement learning. *European Journal of Operational Research* 108 (1998), 306–318.
- [16] MORIARTY D, SCHULTZ A, GREFENSTETTE J: Evolutionary algorithms for reinforcement learning. *Journal of Artificial Intelligence Research* 11 (1999), 199–229.
- [17] TUYLS K, HEYTENS D, NOWE A, MANDERICK B: Extended replicator dynamics as a key to reinforcement learning in multi-agent systems. In *ECML 2003, LNAI 2837* (2003), N. L. et al., Ed., Springer-Verlag, Berlin, pp. 421–431.

- [18] KONDO T, ITO K: A reinforcement learning with evolutionary state recruitment strategy for autonomous mobile robots control. *Robotics and Autonomous Systems* 46 (2004), 11–124.
- [19] KÓKAI I, LŐRINCZ A: Fast adapting value estimation based hybrid architecture for searching the world-wide web. *Applied Soft Computing* 2 (2002), 11–23.
- [20] LŐRINCZ A, KÓKAI I, MERETEI A: Intelligent high-performance crawlers used to reveal topic-specific structure of the WWW. *Int. J. Founds. Comp. Sci.* 13 (2002), 477–495.
- [21] PALOTAI Z, GÁBOR B, LŐRINCZ A: Adaptive highlighting of links to assist surfing on the internet. *Int. J. of Information Technology and Decision Making* 4 (2005), 117–139.
- [22] SZITA I, LŐRINCZ A: Kalman filter control embedded into the reinforcement learning framework. *Neural Computation* 16 (2004), 491–499.
- [23] SCHULTZ W: Multiple reward systems in the brain. *Nature Review of Neuroscience* 1 (2000), 199–207.
- [24] JOACHIMS T: A probabilistic analysis of the Rocchio algorithm with TFIDF for text categorization. In *Proceedings of ICML-97, 14th International Conference on Machine Learning* (Nashville, US, 1997), D. H. Fisher, Ed., Morgan Kaufmann Publishers, San Francisco, US, pp. 143–151.
- [25] BOLEY D: Principal direction division partitioning. *Data Mining and Knowledge Discovery* 2 (1998), 325–244.
- [26] SUTTON R: Learning to predict by the method of temporal differences. *Machine Learning* 3 (1988), 9–44.
- [27] WATTS DJ, STROGATZ SH: Collective dynamics of ‘small-world’ networks. *Nature* 393, 6684 (1998), 440–442.
- [28] ANGKAWATTANAWIT N, RUNGSAWANG A: Learnable topic-specific web crawler. In *Hybrid Intelligent Systems* (2002), A. Abraham, J. Ruiz-del-Solar, and M. Köppen, Eds., IOS Press, pp. 573–582.
- [29] MENCZER F: Complementing search engines with online web mining agents. *Decision Support Systems* 35 (2003), 195–212.
- [30] RISVIK KM, MICHELSEN R: Search engines and web dynamics. *Computer Networks* 32 (2002), 289–302.
- [31] CHO J, GARCIA-MOLINA H: Effective page refresh policies for web crawlers. *ACM Transactions on Database Systems* 28, 4 (2003), 390–426.
- [32] EDWARDS J, MCCURLEY K, TOMLIN J: An adaptive model for optimizing performance of an incremental web crawler. In *Proceedings of the tenth international conference on World Wide Web* (2001), pp. 106–113.