

# A calendar based Internet content pre-caching agent for small computing devices

Andreas Komninos · Mark D. Dunlop

Received: 14 March 2006 / Accepted: 17 January 2007  
© Springer-Verlag London Limited 2007

**Abstract** We described in earlier publications the principles of a system where Internet content would be pre-cached, based on contextual information obtained from a user's electronic calendar. The model for such a system envisioned a set of cooperating agents, distributed on a user's desktop and mobile device, which would be responsible for making decisions on the context and preferences of the user, and downloading the relevant internet content through a land-based broadband connection and storing it on the mobile device. This paper presents and discusses established pre-caching techniques and their suitability for use on mobile information access scenarios. It proceeds in describing the implementation details of an alternative approach, a calendar-based pre-caching system and presents the findings of tests that were made with human subjects on such a system.

**Keywords** Mobile information access · Electronic calendars · Internet content pre-caching

## 1 Introduction

Motivated by the disparity of desktop and mobile Internet access, both in terms of available bandwidth and in terms

of cost, this paper presents research into an alternative method of making Internet content available for mobile users. This method is based on the extraction of contextual information regarding the user's activities and interests using their electronic calendar as a main source, and pre-loading their mobile device with Internet content, using a land-based connection.

The main aim of the research presented in this paper is to investigate whether calendars can indeed provide information that can be used to pre-fetch useful Internet content for mobile users. While it is expected that such an approach cannot fulfill the entirety of Internet content needs for a user, the work presented here provides evidence to the extent to which a mobile cache can be populated with relevant documents that the user could find of interest.

Further to this, this research is concerned with the potential of calendar entries to be used as sources for web query generation, independently of the entry brevity and without the direct involvement of the user. This is an essential step for the investigation of the aforementioned aims, given that an appropriately formulated web query would have a better chance of retrieving relevant documents and thus populate the mobile cache with more appropriate results.

Finally, this paper shows that it is indeed possible for a predictive pre-caching system to efficiently adjust itself to the preferences and circumstances of the user as an individual, in order to obtain optimal retrieval performance.

While not directly related to the main aims of this research, we report further results and findings which concern the usability and interaction patterns within electronic calendars, the document reading behaviour on mobile devices and the suitability of implicit interest indicators for information retrieval on mobile devices.

---

A. Komninos (✉)  
Glasgow Caledonian University,  
70 Cowcaddens Rd., Glasgow G4 0BJ, UK  
e-mail: andreas.komninos@gcal.ac.uk

M. D. Dunlop  
University of Strathclyde, 26 Richmond St.,  
Glasgow G1 1XH, UK  
e-mail: mark.dunlop@cis.strath.ac.uk

## 2 Motivation and current pre-caching techniques

One of the basic issues in the problem of effectively pre-caching internet content needs, whether on a large scale, such as in servers or proxies, or on a personal level, is the determination of exactly which documents should be pre-cached. Taking this decision at a level that allows for maximum personalization, involves the observation of a user's behaviors, in order to create suitable models that would encompass these behaviors and allow accurate predictions to be made.

One of the earliest attempts at the automatic prediction and retrieval of Internet content was described by Balabanovic et al. [1]. The system described there forms a model of each user's preferences and continuously adapts itself to reflect the user's opinions of the content that is prefetched. The user is presented with a collection of hyperlinks to documents that the system has identified as potentially interesting. There is an option for the user to explicitly rate each link (from +5 to -5), therefore providing the system with simple relevance feedback. Even though the scheme employed by the researchers is a relatively straightforward approach, they succeeded in proving that there are significant gains that can be made through the personal profiling of users.

Wang and Crowcroft [2] discuss the some tradeoffs between pre-fetching and the improvement of latency in the www. Also, they present an implementation of a deterministic pre-fetching approach, called Coolist. Their system is layered between the client and the proxy server and organises websites in folders. These folders can then be assigned three methods of pre-fetching. Batch pre-fetching is the first method, where a site is scheduled for downloading at a given date or time. Another method of pre-fetching is described by the term "start-up" and means that a site will be pre-fetched when Coolist is invoked. Finally, their third proposed method is pipeline pre-fetching, where sites are grouped for pre-fetching. When the first page in a group is requested, the next one will be automatically pre-fetched, regardless of the fact that a user may have not requested it.

Another discussion of the advantages of pre-fetching was carried out by Cunha and Jaccud [3], who proposed two algorithms for the prediction of the user's next action while browsing the web. Their first algorithm, using Random Walk approximation, projects the long-term interaction trend, while a second algorithm focuses on the short-term trends. Using a model described by Thiebaut [4], which relates the accumulated number of cache misses to a program's random walk range, the researchers show that it can be successfully applied to characterise users' strategies, under the hypothesis that these relate to an infinite browser's cache. This model is mathematically described as follows:

$$N(r) = Ar^{1/\theta}, r > 1, \theta \geq 1.$$

In this equation,  $r$  is the number of references,  $N(r)$  is the accumulated number of misses,  $\theta$  sets the curve growth pace, and  $A$  is a constant. A second method is described within the same report, which uses an algorithm of two phases: firstly, a preparation phase computes the first order difference of the envelope of the user's profile curve, displaced by a factor of 0.5 (for ease of detecting behaviour changes). Secondly, the prediction phase determines how conservative the user was in the last  $t$  accesses. Also, a determination of how much history is made, based on that count, in order to compute the desired set of coefficients that minimise the short-time prediction error, around a vicinity of size  $n$ , for a sample at virtual time  $r$ . A routine, based on Durbin's method to calculate the linear prediction coefficients is then called, and lastly, the predicted value is computed as a linear combination of the past NCOEF terms. The authors show that both user models manage to achieve a degree of accuracy around 85%, which can be applied in conjunction with pre-fetching techniques.

In his technical report, Palpanas [5], investigated the feasibility of using a model based on the partial-match prediction algorithm, for pre-fetching documents from the web. In his model, a pre-caching agent acts as an intermediary between the client and the server(s) that a user is connected to in a session. Having taken into consideration the special characteristics of the web and after tailoring the algorithm to accommodate those, the author concludes that his proposed scheme's implementation is feasible and that it would be assistive to users who "consistently follow regular access patterns, when searching for information". This conclusion is reached through simulations, run on the access log files of the web server of the department of computer science, at the University of Toronto.

Jiang and Kleinrock [6] presented a system in which pre-fetching is decided by the client, based on usage statistics about embedded HREF tag attributes. In their work, the client monitors its available bandwidth continuously and pre-fetches web content, choosing however, not to pre-fetch images, in order to save bandwidth. An algorithm to decide which pages should be pre-fetched is used, based on the client's access history combined with the server's access histories for each file they hold. Further filtering on the decision process is placed by placing an upper bound on the pre-fetch threshold, which is a function of the system load, capacity and cost of a time unit and a system resource unit. This two-tier decision process allows the system to maximise the performance gaining that can be achieved through pre-fetching.

Further application of Jiang and Kleinrock's work is found in another paper that investigates pre-fetching for mobile users [7]. Interestingly, in this paper, the authors

extend their prediction algorithm to achieve higher hits, by assigning users to a category (such as those interested in database research), amongst other things. The second component of their scheme is a server threshold model, which judges whether a page should be pre-fetched based on:

- The amount of time that may be saved by pre-fetching a file that may be needed
- The amount of bandwidth that will be wasted if the file is not used;
- The impact of the pre-fetch request on other users whose normal requests may be delayed by the pre-fetch request.

The latter is a necessary consideration in order to ensure that the overall system performance can be improved by pre-fetching the file. The authors proceed to conclude that their approach is well suited to mobile users, who may need to switch between different network connection methods (modem, broadband, satellite, wireless). This is because the separate server threshold module allows the adaptability of the prediction algorithm, ensuring the best possible performance under each network condition.

A simpler implementation than the one by Ziang and Kleinrock was proposed with the WCOL system, by Chinen and Yamaguchi [8]. Their system is a research prototype that pre-fetches embedded hyperlinks top-to-bottom without regard to likelihood of use. Embedded images of pre-fetched pages are also pre-fetched. Bandwidth waste can be capped by configuring WCOL to pre-fetch no more than a certain number of hyperlinks, and no more than a certain number of images embedded within pre-fetched hyperlinks.

In 1999, Dan Duchamp [9] presented his own work on pre-fetching hyperlinks, based on a predictive algorithm at a client level, which is however, able to communicate usage statistics from the server. Because the client is unable to form an objective view of the usage for a given web page, unless that page is visited often by the user, it passes on to the server the current usage statistics it has obtained, but also demands the aggregated statistics for that page, as held by the server. The performance results obtained by an implementation of the aforementioned idea were strongly encouraging. For example, of all the pre-fetched pages, a figure of 62.5% was eventually used. An improvement in latency of the order of 52.3% was observed, while notes for the consideration of network overhead due to the usage reports are being addressed. Further concerns regarding the size of the modifications necessary to the browser (Mozilla) and the execution time overhead due to these, are eased, since these do not appear to be significant.

Continuing on the theme of document pre-fetching on a personalised level, more related work was carried out by

Fan et al. [10]. They propose pre-fetching at a proxy level, under the argument that because proxies can collect access histories for limited numbers users, this would present significant advantages over server-level pre-caching, since servers would be able to only collect access histories for the entire www population. An investigation of prediction by partial match algorithms follows in their work, followed by a simulation which shows an improvement in latency between low-bandwidth clients and proxies, of the order of ~23%.

Pitkow and Pirolli [11] explore predictive modelling techniques that attempt to reduce model complexity while retaining predictive accuracy. The techniques merge two methods: a web-mining method that extracts significant surfing patterns by the identification of longest repeating subsequences (LRS) and a pattern-matching method that embodies the principle of weighted specificity. Their work is largely motivated by previous studies by Schechter et al. [12], who utilized path and point profiles generated from the analysis of Web server logs to predict HTTP requests. Also, much reference is made to the work by Padmanabhan and Mogul [13], who describe the efficiency of Markov models for pre-fetching. The authors use the definition of longest repeating subsequence by Crow and Smith [14], which contains the following terms:

- Subsequence means a set of consecutive items
- Repeated means the item occurs more than some threshold  $T$ , where  $T$  typically equals 1
- Longest means that although a subsequence may be part of another repeated subsequence, there is at least one occurrence of this subsequence where this is the longest repeating.

Two models are proposed, firstly a hybrid LRS model, which extracts LRS patterns from a training set and uses them to estimate a first order Markov model. This model is compared against a first order Markov model estimated from all the paths in the training data set. The second hybrid model proposed is one that decomposes the extracted LRS subsequences into all possible  $n$ -grams of various lengths. This is called the All  $K$ th Order LRS model, as all orders of  $k$  are able to make predictions. This model is compared against an All- $K$ th-order Markov model derived from all the possible subsequences, decomposed into varying length  $n$ -grams. Further simplification is added to their web surfing model, by assuming that surfing paths have an average branching factor  $b$ . Surfers may start in  $b$  places and from each page, they move on to one of  $b$  pages on average. By assuming that surfing paths of length  $S$  can be divided into  $S/k$  subpath partitions ( $0 < k \leq S$ ), the complexity cost  $C(k)$ , in terms of the number of patterns as a function of  $k$  can be described as

$$C(S) = \sum_{i=1}^S (S/i)b^i$$

through analysis of the applied methodology, the authors showed that in the case of modeling paths under the first hybrid model, the reduced LRS model was able to match the performance accuracy of the first order Markov model while reducing the complexity by nearly a third. They also then showed that overall hit rates could be raised by including the principle of specificity, with the All-Kth-Order LRS model almost equaling the performance of the All-Kth-order Markov model while reducing the complexity by over an order of magnitude. Finally, within their findings, it was further shown that increasing the prediction set has a dramatic impact on predictive power, with the predictive power of each method nearly doubling by increasing the set size to four elements.

Swaminathan and Raghavan [15] presented a study of web pre-fetching, which is based on the characterisation of the web client alone, without depending on server or proxy side algorithms (Fig. 1).

The above figure shows the principle of the system as described by the authors. The input stream consists of symbols representing the URLs. The learning module learns the trend in visit counts associated with the past  $n$  URLs and the order of URLs visited and it is implemented using genetic algorithms. The comparison module compares the predicted and the actual streams and provides feed back of the error in prediction to the learning module. The generative module attempts to predict the next  $k$  URLs that might be visited for possible pre-fetching. Finally, the pre-fetching module decides which URLs to actually pre-fetch based on constrains such as available bandwidth, recommendation from the server and the state of the web client. In their research, the authors highlight the problem of dynamically generated Internet content, which renders pre-fetching approaches useless. Indeed, given the trend to deliver highly customisable content to users and the generic structure that is “populated” by dynamic articles,




Fig. 1 Operational principles diagram [15]

which now forms the basis of many major sites, the validity of the pre-fetched content becomes an important issue. An interesting point in this research was that the authors manage to prove, through simulation on actual client traces, that their proposed pre-fetching technique allows the maintenance of client cache hit ratio of around 13% on average, even when all the visited URLs are dynamic.

Interesting research on proxy cache was also presented by Foygel and Strelow [16]. The authors propose a system of hierarchical proxy caches, where their algorithm observes requests to a cache and its ancestors, before initiating pre-fetching for the predicted future requests. This would only happen if the pre-fetching action is deemed likely to reduce the overall latency experienced by the cache’s clients. Their algorithm is based on the continuous evaluation of the usefulness of each document in the cache, but also of documents that are not in the cache, but are likely to be needed in future requests. The set of documents (fetched or un-fetched), which has the greatest esteemed value is kept in the caches. In their conclusions, the authors argue that a hierarchical cache network structure is the ideal foundation on which pre-fetching can yield significant performance gains. Again, however, they highlight the concern over increased network utilisation, although they argue that it is often the case that because traffic is added to under-utilised networks, the performance gains can be obtained without significant cost.

Brian Davidson [17] presented an article, which relates to predicting web actions from HTML content. In his work, he compared the simplistic approaches so far taken for pre-fetching based on HTML content, with an information retrieval-based one that ranks the list of links using a measure of textual similarity to the set of pages recently accessed by the user. These simple approaches vary, but examples are namely pre-fetching all hyperlinks in a page, or pre-fetching all hyperlinks in a serial manner, as time allows.

The algorithm used for measuring the similarity between two text documents ( $D_1, D_2$ ) is

$$\text{TextSim}(D_1, D_2) = \sum_{\text{all}_w} TF(w, D_1) TF(w, D_2)$$

( $TF(w, D_1)$  = the number of times term  $w$  appears in  $D_1$ )

having compared text-similarity-based ranking methods to simple original link ordering and a baseline random ordering, the author found that similarity-based rankings performed 29% better than random link selection for prediction, and 40% better than no pre-fetching in a system with an infinite cache.

The ideas of Davison are implemented in Mozilla, an open-source web browser and a technical report is made by Zhang et al. [18], who incorporated the aforementioned



content-based prediction algorithm with the history-based prediction described in another work by Davison [19].

Further related research, however, not immediately a document pre-caching technique in its own right, was presented by Cohen and Kaplan [20]. Their proposal is that in order to overcome potential problems in the validity of cached documents, and other problems that relate to the increased network utilisation, which in turn might cause clients to experience even more latency, one could pre-fetch (rather, pre-execute, one might add) the means of getting a document, rather than the document itself. This is because of the observation that the actual steps required for the setup of a connection, are relatively costly in terms of time. A suggestion is made that this pre-transfer prefetching could be accomplished by

- Pre-resolving, which means that the browser or a proxy could perform a DNS lookup before a request to a server is issued, therefore eliminating the DNS query time from user-perceived latency.
- Pre-connecting, where the browser or a proxy establishes a TCP connection to a server, prior to a user's request. This should address the problem of connection establishment time, which is significant compared to HTTP request response times.
- Pre-warming, or, in other terms, sending a dummy HTTP HEAD request prior to the actual request, in order to address the problem of start-of-session latency at the server, which tends to be larger for first-time requests than follow-up requests (the server is referred to as being "cold" or "hot", once a request has already been issued).

The tests conducted in this research show a significant decrease in average latency times, proving that pre-fetching the means for getting a document is a useful technique that can be applied to the problem of reducing overall latency.

### 3 Limitations of current research

The techniques presented in these research papers describe the process of pre-caching documents in a dynamic manner, whilst however, presuming that the user has a currently active connection and is using (surfing) the Internet. Although all of the concepts are largely relevant and indeed of great interest, they suffer from the disadvantage that network connectivity is a pre-requisite and is essential to their operation.

In mobile devices such as PDAs and Smartphones, network connectivity is not a transparent service which is normally available under typical operating circumstances, such as on modern broadband-connected desktop computers. Indeed, network connectivity depends heavily on the

location of the user and the strength of wireless network signals, the interference of nearby radio-wave emitting devices with network antennae, the subscription to co-operating networks, the present load of the wireless network and the type of wireless network connectivity supported by the mobile device hardware. All of these factors are commonly present in everyday situations and make the likelihood of network unavailability rather high.

Further to this, most of the aforementioned pre-caching models make no, or very little, consideration of the bandwidth used to pre-fetch files which will never be used and assume a practically unrestricted local (or proxy) cache. With mobile devices, one has to consider that the available bandwidth is generally very limited, as well as extremely costly, as the charging mechanism is per unit of data (typically KB), rather than time, when using connections that provide acceptable surfing speeds (GPRS, 3G). Additionally, the devices themselves offer very little memory capacities, thus placing significant constraints on the size of caches that can be created for web surfing. It can therefore be generally concluded that current predictive pre-fetching models cannot be directly applied.

Given the high probability of network unavailability in the daily environment, Internet content pre-caching seems a logical solution to partially solving the problem of mobile information access. However, pre-caching, as described in aforementioned research, is a technique that is largely inapplicable (or even unacceptable), due to the significant cost and the charging mechanism for accessing the Internet over a wireless connection. In fact, given the charging per kilobyte that most service providers adopt for GPRS and 3G, it would seem obvious why a user might want to initiate data transfer from the Internet, as only they deem necessary. Wi-Fi networks on the other hand are a lot less expensive but are very limited in range, thus preventing true mobility.

Another type of solution to pre-caching Internet content is common to several types of mobile devices are not equipped with wireless connection (actually even to devices that have integrated wireless capabilities), but have software for browsing the web already installed. Users of devices such as the ones we described, tend to resort to commercial programs such as AvantGo, that provide a means of pre-caching selected or interesting pages based on declared user preferences and storing them for off-line browsing. However, such services offer static, rather than dynamic personalization, and depend heavily on explicit user instruction on the (limited) type of content that should be pre-cached.

### 4 A predictive pre-caching system for mobile devices

Information access relates strongly to the individual user's interests, which rise from general, always valid (perma-

ment) preferences (e.g. aviation news will always interest an aviation enthusiast, regardless of their profession). Information access also is strongly dependent on preferences which can vary, according to the task the users are, or anticipate to be engaged in. While the knowledge of both contexts under which information access might be desirable are necessary for pre-caching useful information, in our research, we focus more on the automatic prediction and comprehension of the activities a user might be engaged in, in the near future.

Thus we hypothesized that it would be possible for a personal pre-caching algorithm to examine a user's electronic calendar as a source of information that it could use, in order to make dynamic informed predictions about the user's future tasks and provide useful content to the user. While it may be unrealistic to expect a calendar-based pre-fetching algorithm to describe all of the user's preferences and content needs or desires in a perfect manner, we expect that content that is relevant to the user's daily interests can be picked up. Further more, we were interested in investigating whether permanent user preferences can be captured in order to tailor the quantity and quality of the content better to suit the user.

## 5 The pre-caching system

### 5.1 Overview

To explain the principles of operation behind the system we propose, it might be useful to present an example scenario. Let us consider the following example: each entry in the user's calendar should contain at least one keyword, which will help us understand the nature of the activity denoted by the entry. If a calendar entry contains a keyword such as the name of a city (e.g. Edinburgh), of which the user is not a resident, it can be assumed that the user will be present at that location for some purpose which is possibly described elsewhere in the entry. Typical information about cities could then be downloaded through web searches such as "Edinburgh map", "Edinburgh accommodation", "Edinburgh Museums", etc. Also, further and more specialised searches, formed through combination with other information retrieved from the calendar entry, can be conducted, e.g. "Edinburgh Holiday Inn" or "Edinburgh HCI2004 conference".

The predictive system can then be enhanced over time through use of implicit or explicit feedback to learn and remember a users' preference for different categories of information, e.g. more interested in transport links than hotels. Such a predictive system should be able to obtain information from the user directly and indirectly. The majority of information should be obtained through

indirect means, in order to minimise interference with the user's other activities. However, the system should maintain its ability to directly interact with the user, in order to resolve any possible uncertainties.

In previous work [21], we described the overview principles for our pre-caching system, whose software comprises of two core modules, each of which encompasses a number of cooperating agents. The first module resides on a desktop computer and works in order to extract entries from the user's electronic calendar. It then uses the desktop's connection to pre-cache documents. The second module is responsible for presenting the pre-cached content on the user's mobile device and therefore resides therein. The second module is also responsible for evaluating the user's interactions with the pre-cached content and passes back this information to the desktop module, which then uses it to make more informed guesses at what it needs to pre-cache (Fig. 2).

### 5.2 Identifying keywords

One of the most important problems that the system needs to solve is the identification of keywords within the calendar entries, which can be used to initiate the process of web query formulation. To achieve this, a database of known keywords (identifiers) could be kept, against which the system would compare the content of the calendar entries.

In order to train the system with the ability to recognise potential keywords, it was decided that an analysis of the contents of a sample of real world calendars would be necessary. The scope of this analysis would be to determine firstly which suitable words and of what types, would be frequently encountered in a calendar. Subsequently, in order to confirm previous early research on calendars, an attempt would be made to identify calendar entry categories, so the list of keywords for these could be supplemented with other words that belonged to similar contexts. For example, if the analysis was to show that placenames

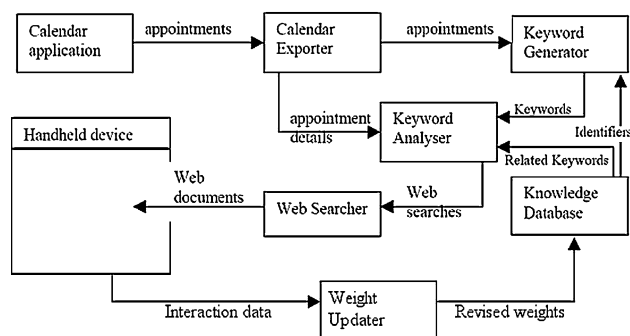


Fig. 2 Overview model of the system components

(“Glasgow”, “Edinburgh”) were common occurrences, then an appropriately extensive list of placenames would need to be compiled. These keywords in essence are identifiers for the categories they represent.

Although early research by Kincaid and Dupont [22] had highlighted some of the categories of entries commonly encountered in calendars, it was felt necessary to re-investigate this matter, firstly due to the age of the preceding research (which was not based entirely on electronic calendars), and secondly to tune the system performance for use by individuals in an academic context. Because the subjects of our final experiments were expected to be from an academic environment, it appeared reasonable to attempt to gear the system towards the particularities of the academic community. The system could have similarly been geared towards other professional areas or could be tailored towards a general population. This would require the selection of appropriate test subjects and since the comparison of performance on different target groups is not part of the scope of this research, the academic sample and subjects were judged to be appropriate and adequate.

From our research in the categorization of entries, it immediately became clear that the calendar entries tend to fall within specific categories. The categorisation of the entries was done manually, with guidance from the original entry authors, where ambiguity made it necessary. This categorisation comes as a confirmation of the findings of previous research, and especially those of Kincaid and Dupont, who discovered that users tend to use their calendars mainly to keep a record of meetings, appointments, events, travel, reminders, notes and as “to do” lists. This, in turn, suggests that there is a consensus to the items that form appropriate calendar entries and that there is a common mental representation model for the organization of these entries amongst humans. There is a slight variation in the categorisation of entries that was made with these new findings, although this can be considered natural, due to the particularities of the academic environment. However, the categories with the highest frequency are the same as those in previous research. Table 1 shows the categories and entry frequencies as derived from this research’s sample.

It is important to stress here that these categories reflect the opinions of the calendar users, who are highly familiar with the context of their entries. One can observe that there is some overlap between the calendar entries; for example, birthday could be considered a subset of Social. However, where such overlap is maintained, it is because there was a strong indication from the users that such a low-level category is significant and should exist separately from its high-level parent.

The database of category identifiers which would be compiled would not exclusively contain words, but also short keyphrases, such as “travel to”. The reason for this

**Table 1** Calendar entry categories and their frequencies

Category	Frequency
Meeting (group)	53
Meeting (another person)	25
Reminder	18
Travel	13
Social event	13
Work task	10
Class (to attend)	10
General task (to-do)	7
Miscellaneous	7
Birthday	5

choice was that it would be almost impossible to compile a list of every possible associated identifier for every category. Therefore by including keyphrases, we could infer that an unknown word which would follow these might actually be a good identifier candidate. For example, “trip to Garnethill<sup>1</sup>” is a sentence where a human can immediately identify Garnethill as the name of a place, even though they might never have heard of it. It is also very unlikely that the name of Garnethill would come up in any general placename list. It was desired that same functionality for inferred knowledge should be implemented for the system as well.

It might be argued that using pre-compiled lists and rules as a basis might not be an optimal solution. However, the process of recognising keywords and inferring their meaning based on pre-compiled rules and lists is well-established practice in the field of Information Extraction. This practice is one of the two approaches generally available for solving the problem of machine text analysis and understanding, the other one being automatically trained systems that do not rely on handcrafted rules and databases. Given the recommendations of Appelt and Israel [23] on the choice of approach and based on the good availability of resource lists (e.g. name and place lists) and the lack of extensive training data (calendar entries), the manual approach was elected. Further detail on the automatic categorization of calendar entries the performance of our algorithm are described in [24].

### 5.3 Formulating queries based on keywords

With the appropriate keywords identified by the keyword generator module, the next step the system should take is to generate appropriate queries for these keywords. Sometimes a keyword on its own might be a good candidate for a query. However, in order to obtain information that is

<sup>1</sup> Garnethill, a region in the Glasgow city centre area (UK).

closer to the needs or preferences of the user, the keyword will have to be combined with other keywords to form longer queries. Those additional keywords can be obtained from the calendar entry itself, or could be retrieved from a separate database of additional keywords which are known to be typically included in queries that are based around the original keyword.

To be able to adapt the system to the user's preferences and needs, it becomes apparent that all additional keywords will need to be weighted in order of importance to the user. Because some of the additional keywords that will initially be provided might be irrelevant to the user's context, the system should be able to either negatively weigh these, or increase the weight of additional keywords for which the user was presented queries that retrieved "good" documents, while leaving other weights intact. A combination of positive and negative weighting can also be used.

One approach to the problem of constructing an additional keyword database would be to create a list of potential associations for each keyword contained in the identifier database. In this manner, the level of relevance between identifiers and additional keywords would be high, although this would come at a considerable cost. The costs of this approach would firstly be the large degree of duplication of data, as an additional keyword might be relevant to a multitude of identifiers. Secondly the construction of appropriate additional keyword lists for every single keyword would have to be done manually and there is no guarantee that every single identifier would have an additional keyword to reflect all context scenarios. Finally, because of the specificity of the additional keyword associations, a sudden or temporary change of context for the user might be completely ignored or take a long time to adapt to, which is of course undesirable.

An alternative approach would be to cluster the additional keywords into smaller databases, which will reflect the association of each keyword category, rather than each keyword individually, with additional keywords. The term of clustering is borrowed from the Information Retrieval field, where it is defined as follows:

"...We define the organisation as the grouping together of items (e.g. documents, representations of documents) which are then handled as a unit and lose, to that extent, their individual identities. In other words, classification of a document into a classification slot, to all intents and purposes identifies the document with that slot. Thereafter, it and other documents in the slot are treated as identical until they are examined individually. It would appear, therefore, that documents are grouped because they are in some sense related to each other; but more basically, they are grouped because they are likely to

be wanted together, and logical relationship is the means of measuring this likelihood..."<sup>2</sup>

This is a more natural step to take, considering that calendar entries are already clustered into categories by nature. Therefore web queries will be made by combining the original keyword (category identifier) with either all or the top  $n$  additional keywords that are associated with the category. In fact, as the system becomes increasingly accustomed to the preferences the user, the user-defined  $n$ -sized window of additional keywords ranks should reflect the true preferences of the user by returning a smaller amount of keywords within the same window.

The additional keywords can initially be set to have the same score, which would indicate an equal opportunity for each additional keyword to rise up in the ranks and therefore reflect a user's preference, independently of that preference's commonality. Alternatively, additional keywords can all be given different initial scores which would reflect the general preferences and assumptions that could be made for an average user. This would impede the training process for users that have very particular requirements, although for the majority of the users this method should provide adequately promising performance even from the early phases of use. We employed the second option in our implementation.

In order to obtain additional keywords for each category, three different methods were used in conjunction:

- Interviews where people were asked to give details of typical searches they might have performed for calendar entries of each category (during the calendar entry sample collection)
- Calendar entry samples were given to independent subjects who were asked to produce as many web queries as they could for each entry sample
- Google's keyword suggestion tool<sup>3</sup>, where category identifiers were entered and a list of potential associated keywords was returned.

In our implementation, the keyword analyser module forms the following two types of web queries:

- Original keyword (or keywords, e.g. name + surname) only
- Original keyword + additional keyword

The system can currently be set-up to fetch the top  $N$  rated queries for each user, according to their rank value, although for the main experiment, the amount of generated

<sup>2</sup> HAYES, R.M., 'Mathematical models in information retrieval', in *Natural Language and the Computer* (Edited by P.L. Garvin), McGraw-Hill, New York, 287 (1963).

<sup>3</sup> Google Adwords keyword suggestion tool: <https://www.adwords.google.com/select/main?cmd=KeywordSandbox>



queries was limited to the top five rated items. When two or more queries happen to be ranked equally, they are both included in the generated query list.

#### 5.4 Pre-caching documents

Having described the methodology for formulating queries, the next logical step in the sequence of operation events for the system is to submit these queries into one or more search engines on the Internet, which in turn should be able to retrieve links to relevant documents.

Those links should be followed in order to retrieve the documents. From within the documents, further links can be followed to documents within the same web site, or to documents that are external to the site. If those external documents were fetched and analysed, the system would encounter yet more links, which could also be retrieved. The retrieval process might be likened to a multi-tier tree, where a document from the original query can be considered a root, further linked from it are second-level nodes and links to further documents from these form edges that lead to even lower-level layers of nodes.

It becomes apparent that the system could be locked into a fetching mode, which will exponentially increase the number of documents retrieved, without guarantees that all of the lower-level nodes will be relevant to the original query. In fact, it is logical to expect that the degree of relevance will reduce with the increase of the depth of the retrieval tree. Considering Jansen's [25] observation that only 50% of the original web documents retrieved from web queries are classed as "interesting" by users, the lack of any guarantee of relativity of all the supplementary retrieved documents is apparent. For this reason a control mechanism should exist, which will either limit the amount of levels the retrieval tree can have, or be able to intelligently predict whether a linked document might be worth pre-fetching. Based on the pre-fetching methods researched earlier for desktop computers and proxy caches, it becomes obvious that such decisions have to be based either on data from server request statistics or an analysis of the target document, in which case it needs to be fetched so it can be analysed. Such an implementation technique as the latter might be considered an interesting exploration for the adjustment of performance levels, however, it is beyond the scope of this investigation. Therefore a logical solution for this problem under these circumstances would be to implement a mechanism that would limit the depth of the retrieval tree levels for each document.

Further considerations would have to include the type of Internet content that the system should be able to store. Further from HTML documents which contain formatted text, other elements such as images, .pdf or .doc files, audio

and video, might be desirable to have. All of these however, place a considerable stretch of the memory limitations of current handheld devices. For this purpose, we carried out interviews with users in order to discover exactly what kind of content should be prioritised for pre-fetching. The results of the interviews showed an overwhelming preference to text (HTML, Word, PDF), while other media such as images, audio and video were considered largely irrelevant, unless a user was specifically looking for such types.

Of the search engines available today on the Internet, the most successful is the Google engine<sup>4</sup>. This engine was chosen for the submission of the formulated web queries. An observation by Jansen and Spink [26] indicated that the majority of users (80%) only view between 10 and 20 results for each query, i.e. no more than one or two pages of results. Because the Google search engine can retrieve thousands of relevant results, a choice was made to limit the returned results to the top ten, as ranked by Google. It was felt that the choice to restrict the results to the top ten was justified, bearing in mind the memory constraints that are present on handheld devices. Furthermore, because Google is widely considered to be the best Internet search engine currently available, the choice was made not to submit the queries to other search engines. The comparison between the performance of Google and other search engines was not of interest for the purposes of this research and beyond this, such an attempt would greatly increase the amount of retrieved documents without any expected significant increase in the amount of retrieved relevant documents. This would place a significant and unnecessary burden on the subjects of the main experiment. Although collective filtering could be applied to perhaps include only the top K results from a number of search engines, it should be kept in mind that the optimisation of the cache, while desirable, was beyond the purposes of this research. Thus no services other than Google were employed in this instance.

The Google result page is trimmed from all unnecessary HTML elements in order to remove graphics, advertisement and unnecessary links. Subdirectories are created to store all the documents of the desired tree depth levels. The retrieval procedure then works for each document of each tree level, including the Google result page, in the following manner: firstly, the document is parsed and searched for links. Once a link is encountered, the linked document is retrieved and stored in the appropriate subdirectory. The link in the document under analysis is changed to reflect the local relative URL of the newly fetched document. The process continues until no further links can be found in the document, in which case, the system moves

<sup>4</sup> <http://www.google.com>

on to the next document of the same depth level. In this manner, one can envisage a horizontal breadth-first traversal of the document traversal tree, in order to generate the lower levels of pre-fetched document. The entire process ends when the user-specified retrieval depth has been reached.

Once the system is finished retrieving the top ten documents for each Google page, further transformation of the Google page is made. A second-level parsing is performed in order to separate the document titles, summaries and URLs contained therein, and store them in an XML structure. This XML structure will later be passed to the handheld device, along with the relevant documents, and will be used to display the retrieval results to the user.

## 5.5 Mobile component

### 5.5.1 Presenting results

Once the documents have been retrieved and sent to the handheld device, a separate software component therein should be responsible firstly for displaying the documents to the user, and secondly for observing the interactions of the user with these documents.

It is clear that for the first task, the design needs to consider the physical characteristics of the handheld device and especially the constraints placed by the dimensions of the device screen. Taking in mind Nielsen's observations that users tend to dislike long pages which require lots of scrolling [27], an implementation of the results browser should be considerate of this natural tendency and contain facilities that will allow the users to minimise the scrolling needed. A collapsible tree-structured list of calendar entries, identified keywords, searches and retrieved document titles/summaries is potentially a good way of addressing the scrolling problem. Unfortunately, due to the nature of web pages, scrolling to view their content is inevitable. However, the quality of the browsing is beyond the system's control and will be fully dependent on the device's integrated web browser. As this issue is beyond the scope of this research, no attempt to write a dedicated web browser was made.

### 5.5.2 Processing relevance feedback

Observing the interactions of the users poses several questions that need to be answered. Firstly, which actions of a user do reflect interest and therefore should be monitored? Secondly, once such actions are identified, do they all indicate the same amounts of interest or should their importance be weighted with different measures? Based on previous related research, an instinctive negative answer to the last question is probably the right one. Indeed, as

mentioned previously in chapter 2, it has been shown that not all kinds of interaction can provide dependable implicit information on relevance.

On the handheld device, a log is kept of the user interactions with the content. Based on this information, an attempt should be made to judge the relevance of a given keyword from the knowledge base to the user's context and determine which of these are likely to be wanted as part of a query in the future. Given the collapsible presentation structure, a log can be kept for:

- The viewing of the document index for a given search
- The viewing of the summary of a given document.
- The viewing of the document.
- The amount of time spent on a document that has been opened.
- Any explicit feedback rating that a user might provide for the document.

The incorporation of further heuristics, such as pointer movement, highlighting of text, bookmarking and scrolling in the document, would also be desirable. However, given the implementation was on a PocketPC platform, programming pocket internet explorer to trap such behaviours was not feasible. In addition, some of these heuristics (e.g. scrolling) would not be reliably applicable, given the lack of previous research on small screen devices for such measures.

Upon loading, the handheld component examines the XML structure passed to it and the XML-formatted Google pages, in order to load the necessary details for presentation to the user. The pre-fetching activity details are presented using a collapsible tree structure list, which gives details of the appointments, keywords identified, web queries formed and document titles of each search (see Fig. 3).

When a document title is highlighted, the user is given the option via two buttons to either launch a descriptive summary of the document, or open the document with pocket internet explorer. The system monitors the collapse

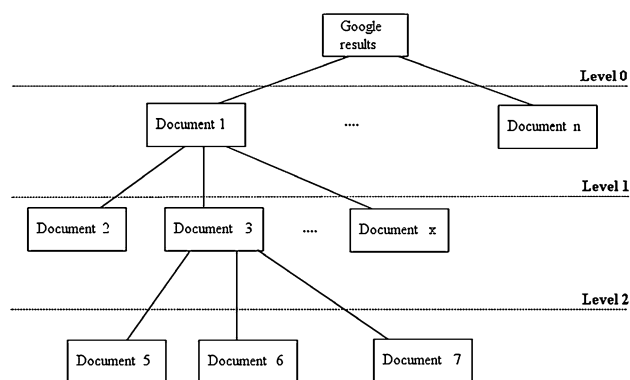


Fig. 3 The multi-tier retrieval tree

of the document list for a web query, the viewing of a summary, the launch of pocket internet explorer and the duration for which it is active, i.e. in the foreground. It is assumed that the user will be reading the document for that time. Also, when pocket internet explorer exits and the user returns to the handheld software that is running in the background, an option is given to explicitly rate the quality of the document just read on a scale of 1–5 (Fig. 4).

To implement the document summary function, the short document narrative that Google provides directly under the document title in its results page is used. This is displayed as a pop-up dialog box to the user, upon request, therefore imposing an interaction cost on its viewing (Fig. 5).

### 5.6 Processing relevance feedback

For the heuristics mentioned above, we had to devise appropriate weights and a method to combine their values to form a score for each keyword. We define the importance  $I(k)$  of a keyword  $k$  to be

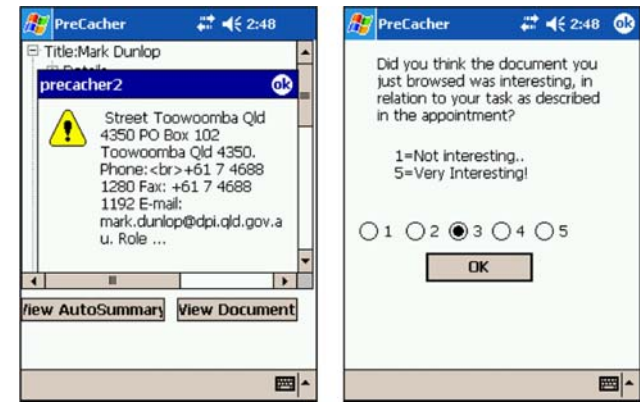
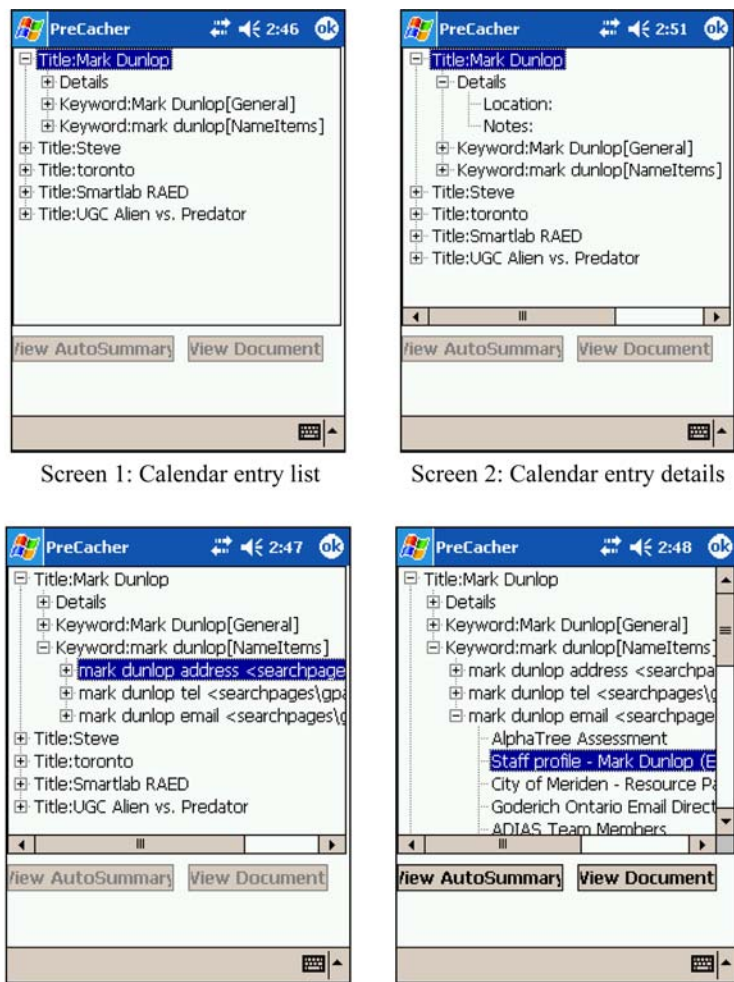


Fig. 5 The auto-summary function and the explicit relevance feedback screens

$$I(k) = w_{E(k)} + \sum D(i)$$

where  $w$  is the weight associated with a viewed document index for the keyword  $E(k) \in \{0,1\}$ , and  $D(i)$  is the

Fig. 4 The mobile user interface, showing the collapsible tree list. Each keyword under a given calendar entry (title), can be expanded to show the queries formed for it. For each query, a list of retrieved documents is provided



Screen 1: Calendar entry list      Screen 2: Calendar entry details      Screen 3: Queries formed for a keyword      Screen 4: Document list for a query

importance of each document  $i$  that has been retrieved for keyword  $k$ . Further more, we defined  $D_{(i)}$  as follows:

$$D_{(i)} = \alpha O_{(i)} + \beta S_{(i)} + \gamma F_{(i)} + \varepsilon T_{(i)}$$

where  $O_{(i)}$  indicates the viewing of document  $i$ ,  $S_{(i)}$  indicates the viewing of its summary,  $F_{(i)}$  is the explicit feedback given by a user to the document,  $T_{(i)}$  is the time spent reading the document and  $\alpha, \beta, \gamma, \varepsilon$ , are the respective weights for each of these measures.

The weights  $w, \alpha, \beta, \gamma$ , and  $\varepsilon$  take positive or negative values. We wanted a mechanism that would promote the appearance of preferred keywords in the queries. However, the negative marking for undesirable keywords would not only allow the improved promotion of desirable keywords, but also, should one of the latter become undesirable, due to perhaps the change of context of the user, it would not take too long for it to start disappearing from the queries (Table 1).

We experimented with several weights for our system and we arrived to the conclusion that it is not only important to consider the relationship between the weights, but also the bias towards positive or negative marking. Current research [8] shows that users will view only two or three (on average) documents per web query and the vast majority will visit at most two pages or results (approximately 20 results in all). The same research shows that an estimate of 50% of documents viewed from these results are expected to be relevant to the query. Therefore we decided that a bias of approximately 1:7 in favor of positive marking was reasonable. Because our queries will not generate more than ten documents each, this means that two documents with a positive overall rating will indicate a successful and relevant query was made. The weights that were used were as shown in Table 2.

There was some confusion as to the weights  $\varepsilon$ , which should be used for association with the reading time

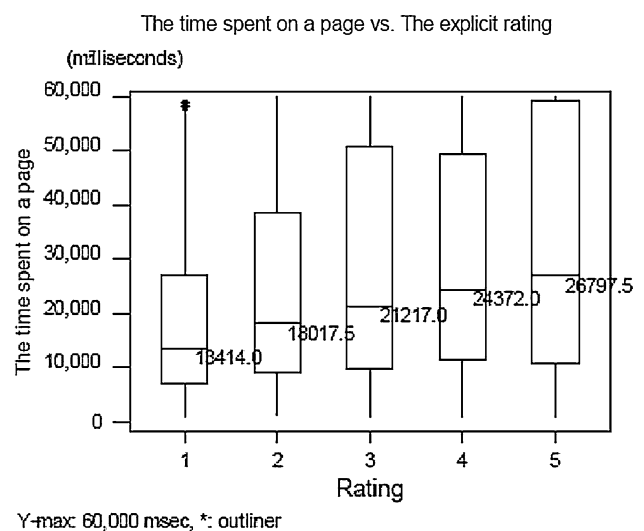
**Table 2** Keyword score adjustment weights as used in the system implementation

Weight	Value
$\alpha$	-0.03 (document not opened) +0.03 × 7 (document opened)
$\beta$	-0.021 (summary not viewed) +0.021 (summary viewed)
$\gamma$	-0.15 (feedback = 1) -0.075 (feedback = 2) 0.00 (feedback = 3) +0.15 × 6 (feedback = 4) +0.15 × 7 (feedback = 5)
$\varepsilon$	Inconclusive

interest indicator. As mentioned in Sect. 2, previous research highlights a possible correlation between average reading time and relevance of a document. However, all previous research had been conducted on desktop computers, where screen readability issues were not as much of problem. The initial inclination was to apply a tiered weighting system according to the reading time averages reported by Claypool [28] (Fig. 6). The weights for each metric were chosen on an ad-hoc basis, mainly through consultation with the main experiment’s control group subjects and their observed interactions with the system. From the figure below, it becomes immediately obvious that there is quite some overlap between the distributions of reading times for each explicit rating, something which makes it difficult to accurately infer document relevance from reading time alone, in most cases. After experimentation (described in Sect. 7), it was decided also that these average reading times were completely irrelevant to the handheld device environment. In fact, it was discovered that no distinct correlation could be made between reading time and document relevance on a handheld device, so the decision was made to abandon this metric completely.

### 6 Limitations of the system

We would have liked to be able to incorporate further heuristics, such as pointer movement, highlighting of text, book marking and scrolling in the document. However, given the implementation was on a pocketPC platform, programming pocket Internet explorer to trap such behaviors was not feasible. In addition, some of these heuristics (e.g. scrolling) would not be reliably applicable, given the lack of previous research on small screen devices for such



**Fig. 6** Claypool’s findings on relevance and reading time correlation



measures. Indeed, further on, we describe how we were forced to also exclude reading time from our heuristics. Further more, we decided not to pre-cache images or other media (pdf, word), purely for reasons of storage space constraints.

## 7 Experimentation

### 7.1 Experiment design

Having implemented a full system, as detailed in the previous section, it was time to test its performance and determine whether the hypotheses proposed by this thesis could actually be met. Ideally the system should be given to several users and they should be allowed to run it for a period of time which should be as long as possible. However, given the lack of volunteers that would be willing to run the experiment as part of their everyday routines and also given disproportionately large timescale the experiment would take, a decision was made to test the final system under supervised conditions which would simulate real world scenarios as closely as possible.

Two groups of users were given the same scenario with some details of their imaginary living location, job and a list of some names of people and how they would be related to them. Furthermore, over the duration of 3 weeks, the users would be given five tasks per week that form their hypothetical schedule of activities for that week. These activities were given in the form of a calendar entry that contained a title, location and notes for each one. Some activities did not contain items in the location or notes fields, as they were based on real-world entries that we had collected in previous studies. The users were also given clear instructions on the exact meaning of each entry, through the provision of accurate descriptions of the entries.

The users were then allowed to freely navigate through the pre-cached content that was fetched for these hypothetical schedules, and try to locate content that they thought might be helpful to them. We would also ask the users to give an indication of whether they found the provided content for each activity useful.

It was decided that the users should not be told that their behaviours would be logged. Also, one of the groups would have their logs analyzed and we would attempt to provide them with content that was personalized on the basis of these logs. Again, the groups would not be made aware of this discrepancy until after the experiment had ended. The analysis of the logs was done automatically by the system, for each of the monitored subjects individually. Their respective profiles were maintained and updated at the end of each session, therefore influencing the retrieval process to personally match each of the monitored group's subjects.

Finally, another factor which was considered in the experiment, was that the physical storage limitations for the devices used led to a choice to pre-cache only HTML documents, and furthermore, these were restricted to the documents proposed by Google for each search. Effectively this meant that the retrieval tree was limited to just one level. For a first-level document that contains three hyperlinks, a two-order retrieval means a total of three documents retrieved. The amount of generated documents from just the one-level tree was very large and this would only grow further with the implementation of additional levels. Therefore, in order to avoid overloading the user with documents and to overcome storage limitation problems, the choice was made to restrict the retrieval tree to just one level. Further to this, the focus was placed on the four most popular categories, according to the findings of the query test. Therefore the calendar sets given contained only entries of type Meeting, Travel and Social (including birthday)

### 7.2 Initial experiment setup

An initial group of ten subjects volunteered to test the system before we proceeded with the actual experiment. All of the subjects were from a similar background and considered themselves computer literate, although most did not have previous experience with a PDA. The initial group was given different, but similar in context, data than those that would participate in the actual experiment; however, the rules of the experiment were the same, apart from the duration of the experiment, which would only encompass the virtual timeframe of 1 week. The goal of this initial experiment was to ensure the system ran smoothly with users that were unfamiliar with it. A further, and more important goal, was to observe the average reading times for the web documents and their relation to explicit feedback, as we planned to use this metric for implicit relevance feedback.

Having analysed the results of these initial groups, it was observed that the average reading times were not what we expected, and were certainly in contrast with previous research such as that mentioned in Sect. 5.6 (Fig. 7).

It is clear from this graph that the users take, on average, the same amount of time to distinguish between either relevant or largely irrelevant documents. Therefore it is apparent that the use of time as a metric is not a reliable source of information, since there is not much significant discrepancy between the average reading times for each feedback score. This observation brought about the decision to eliminate this metric from the weight recalculation formula, as it is in contrast with other findings, such as those by Morita [29] and Claypool [28], but seem to confirm Kelly's [30] conclusion that reading time is an unre-

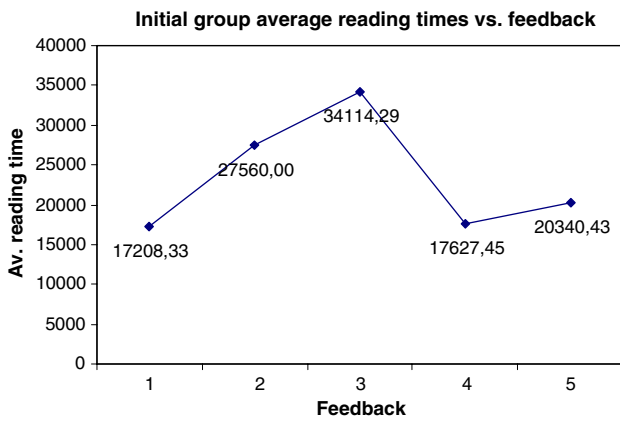


Fig. 7 Initial group average reading times versus feedback ratings

liable source for implicit relevance feedback. While our research and that of Claypool’s use scales of 1–5, Kelly uses a scale of 1–7. However, what is more interesting than the direct comparison of observations is the fluctuation between these (Fig. 8).

### 7.3 Actual experiment

For our actual experiment, two groups of ten people each were used. The rules were applied in full this time and we were able to obtain some interesting results at the end of the experiment. Unfortunately, due to data corruption on the logs of two members of one group, we were forced to exclude them from the analysis, removing also two random members from the other group, to make the figures directly comparable.

In the following graphs, a representation of the average reading times for both groups, over the three experiment weeks is depicted (Fig. 9).

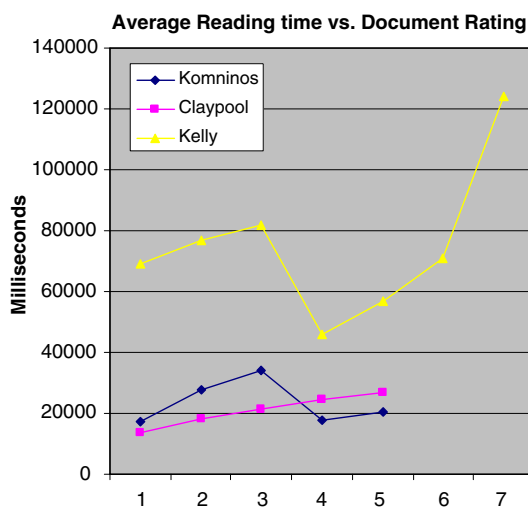


Fig. 8 Studies on average document reading times (ms) versus perceived document usefulness

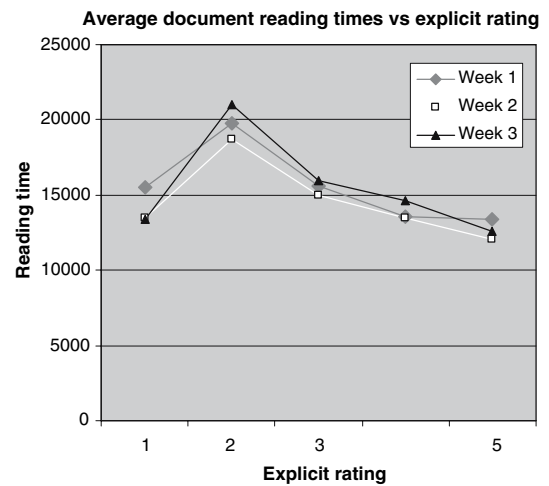


Fig. 9 Experiment group average reading times versus feedback ratings

The trend shown here is slightly different from the results of the initial group. This is expected as the data for the two groups were not the same. However, again from this graph one can clearly see that determining a relationship between feedback and reading time is not feasible. The average reading times for the top two and the worst mark are very close, making any secure distinction between the two almost impossible. Further to this, it is surprising that the group behaviour for all 3 weeks shows similar trends. This finding seems to confirm our suspicion that average reading times are strongly dependent on each individual test group characteristics and cannot thus be easily generalized.

Further analysis of the results should show whether there is a trend in the improvement of cache hits for the group whose data was tuned according to their previous logs. The following two tables show the numerical and percentile quantities of the opened documents (cache hits) versus the total documents offered.

From Table 3, 4 one can clearly see an almost linear trend developing for group 1, who were the group that had their data adjusted according to their logs. This is a strong indication that for these users, the system provides an improvement in potential efficiency, if one considers that opening a document indicates the user’s potential interest in it. For group 2, a solid conclusion cannot be made, as the percentage rates seem to fluctuate almost randomly,

Table 3 Total documents versus opened documents (%)

	Group 1 (%)	Group 2 (%)	Joint (%)
Week 1	8.6	24.4	16.5
Week 2	16.5	18	17.7
Week 3	24.8	30.5	27.8

**Table 4** Total documents versus opened documents (absolute values)

	Group 1		Group 2		Joint	
	Total	Opened	Total	Opened	Total	Opened
Week 1	1464	127	1,464	357	2,928	484
Week 2	1,879	31	1,880	357	3,759	667
Week 3	1,065	264	1,248	381	2,313	645

affecting of course the joint outcome as well. It would appear that for this group, the nature of the entries in their given schedules is the only determinant in the percentage of documents opened. For group 1 however, it appears that the application of interest indicators found in previous logs, has a restraining effect to the fluctuation of the variation in percentages and an overall effect that shows increased performance of the system.

As mentioned in Sect. 5, a further indicator of interest in a document is its summary. A look at how many documents were judged by the summary reveals the following results (Table 5):

The percentages compare each figure with the total number of opened documents. From these trends we see that approximately only one in four times did the users consult the summary before making a decision. This suggests that a user will be inclined to navigate to a website based on the information contained in its title solely. A higher percentage was expected in this situation, especially since visiting a document is costly (in terms of loading and reading time) and also because in the implementation, explicit feedback was requested after each user had finished reading. According to Nielsen [27] also, users tend to like summaries and will read them before resuming with the rest of the text. However, it must be noted here that the summaries were only displayed on demand, while the document titles were immediately available. This was a

**Table 5** Summary viewing as a deciding factor for opening a document

	Group1		Group2	
	Immediate Open	After summary	Immediate Open	After summary
Week 1	90 (65.6%)	37 (34.4%)	255 (71.4%)	102 (28.6%)
Week 2	282 (90.9%)	28 (9.1%)	252 (70.6%)	105 (29.4%)
Week 3	209 (79.1%)	55 (20.9%)	205 (53.8%)	86 (46.2%)
Joint	Immediate open		After summary	
Week 1	345 (71.3%)		139 (28.7%)	
Week 2	534 (80%)		133 (20%)	
Week 3	504 (78.1%)		141 (21.9%)	

necessary tradeoff in order to reduce the scrolling required for the retrieval results overview, although the action of opening a summary incurs an additional cost to the subjects.

Finally, the following Table 6 takes a look at the average document scores for each session:

It would appear from this table that the scores remain at a constant level and in fact, at around the middle of the scoring table. This in turn is consistent with the results by Jansen [26], where it is mentioned that users should expect one of two web documents they view to be relevant.

The choice was made not to measure the individual scores attained by each group in order to establish a trend. These would be actually just measuring the ability of Google to return relevant results, where as this research is only concerned with measuring the relevance of the web query in the context of the calendar entry and the user’s needs.

### 7.4 Further discussion

#### 7.4.1 Experimental environment

The experiment was performed in a quiet room. This setting might not appear to be realistic in the sense that there were no external distractions for the subjects, although they were given food and drink and were allowed to communicate and interact with each other. Mobile devices are used in both mobile and stationary environments and since the experiment shows that reading times are not long (around 25 s), it is expected that a mobile user could easily dedicate such small times to interact undisturbed. I perceive the notion of ‘mobile’ to mean ‘out of office’ rather than ‘walking’ or ‘driving’, therefore the setting seems adequate for the purposes of the experiment. In any case, a maximum time limit of 2 min, based on the observations from an initial test group, is imposed on the measuring to eliminate gross inaccuracies due to user distraction. Therefore it can be concluded that the environment settings for the experiment were appropriate for its purposes and did not deduct from its credibility.

#### 7.4.2 Statistical confidence

Further to the results of the test, a two-paired *t* test was conducted in order to investigate the statistical significance

**Table 6** Average document scores (0 lowest, 5 maximum)

	Group 1	Group 2	Joint
Week 1	2.26	2.37	2.44
Week 2	2.60	2.38	2.48
Week 3	2.49	2.36	2.42

of the findings that were observed, in relation to the improvement of cache hit-rate improvement for the two groups. The  $t$  test was the recommended approach as the experiment dealt with two groups of different subjects, who came however, from a homogenous background, for one of which an external factor was applied and its effect was observed. This external factor was the monitoring and consideration of interaction and feedback behaviour, and its implication in the retrieval process.

The cache hit rates between week 1 and 3 were measured for each individual subjects and their difference was analysed. The findings of the  $t$  test are summarised in Table 7. With a statistical probability of error of approximately 1.2% when considering whether the external factor was indeed responsible for the cache hit-rate improvement, the credibility of the results is further enhanced.

### 7.5 Summary of findings

Several important conclusions were reached by this experiment. Firstly, the system shows that useful Internet content can indeed be pre-cached based on calendar information alone. This is shown by the cache hit rates, which rose close to 30%. Another important finding was that the reading behaviour of the subjects when faced with documents on a small screen, showed that the time spent on a document does not accurately reflect the quality of the document. A correlation between these two cannot be established, therefore the use of reading time for implicit relevance feedback on small screen devices is not recommended. Finally, and perhaps most importantly, the results of this experiment show a gradual, almost linear improvement of the retrieval performance for the group whose behaviours were taken into account. Although the duration of the experiment could have been longer, the statistical confidence is such that it can be argued that the results are solid enough to provide adequate confirmation of a promising learning curve performance.

## 8 Discussion and future work

We described a pre-caching system which is based on the information found in electronic calendars, in order to provide useful content for a user with a small mobile computing device. While such a system in its own right

would not be able to completely satisfy all of a user's internet content needs or desires, we show that this system it can indeed provide useful content for the appropriate entry categories. Even in the case of entries where the information contained therein comprises of a single word, the automatic generation of web queries based on common knowledge and the users's preferences proves to be able to provide meaningful and useful content.

While in this instance we chose to focus on the contextual information that could be obtained from electronic calendars, it is possible to foresee an expansion of the system which may make use of additional sources of information. Some examples could be the scanning of a user's email inbox, SMS inbox or instant messaging backlogs for information on upcoming activities. It is common these days that meetings are arranged through email and tasks are also delegated through this medium. Thorough analysis of email messages would require a careful natural language processing, we believe that such a process would be a very valuable source of information.

Given the opened document trends as described previously, we have reason to believe that our system is able to adapt accordingly to the individual preferences of a user. A further trial over an extended period of time, preferably over 6 months, would be able to show the fluctuation between improvement rates and whether a peak is reached, which would indicate the system's optimum performance level. A foreseeable problem with our current system is that the current adaptation algorithm adjusts the system gradually, and not abruptly, to the needs of a user. Thus, if a dramatic change of circumstances was to occur, or if a user was to require information from a very specific and known source, it is likely the system would fail to provide the necessary information. To that extent we programmed the system so that if a calendar entry contained a website address in the Notes section, that URL would automatically be pre-cached. However, our interest focussed on the adaptability algorithm and the usefulness of calendar entries as a source of contextual information, thus our experiment did not include the usage of this convenience.

Further to this, the system currently lacks the ability to automatically add keywords to its knowledge base. Such an inclusion, we believe, would help dramatically in the improvement of the system's performance.

Apart from the findings that were part of our main target, we encountered several other interesting facts. The similarity between decision times for judging positively or extremely negatively against a document prohibits the use of such a metric from any further studies. Further to this, we were also impressed by the low reading times, which are in stark contrast with other studies that are concerned with the average reading time of a web document, such as [28, 29]. Other studies report average reading times closer to the ones

**Table 7**  $t$  Test results

Mean $a$ -Mean $b$	$t$	$df$
0.1003	2.4985	14
$P$		
One-tailed	0.01277	
Two-tailed	0.02554	



we experiences, but again higher [30, 31], although these were not based on web documents. However, all of these previous studies relate to documents viewed on a desktop, where a large monitor facilitates the viewing of documents. It is our assumption that the smaller reading times on the handheld indicate a tendency for users to “skim” through the document in order to decide on its usefulness. This should be considered normal, given that the need for immediate and full comprehension of the information in the text was not there (due to the virtual environment). Therefore the users would try to acquire a general “feel” for the quality of the searches, and refer to these later on when they have more time or immediately need the information. Nielsen argues that scanning the text in a web document is common practice. Further more, in his work, long pages that cause lots of scrolling are considered to be largely disliked by users. Since the limited size display on a handheld causes websites to appear unproportionately large and causes lots of scrolling, our findings of reduced reading times seem to be further supported by these statements.

The reduction of used bandwidth was a starting point in our thinking; however, starting from thinking about how we could reduce the need for mobile bandwidth, we became interested in examining calendars as a source of contextual information. The significance of the findings of this study do not relate solely to how bandwidth can be reduced. Rather, with the discovery of patterns in electronic calendar use and the automatic de-ciphering of calendar entries, one could proceed in solving other problems in the usability of mobile devices, or invent new services.

(Yet) unpublished work that we have carried out draws upon the categorization of entries to improve the usability of calendars, firstly by implementing entry categories that reflect those truly needed by users (most mobile devices offer a categorization which is far too simplistic to address real user needs). Also, based on the categorization, entries could be colorized to carry layers of information such as type of entry (meeting, to-do, etc.) and distance of event from current location. This would enable a mobile calendar to automatically adjust reminders, pre-fetch traffic information or building service information (e.g. lift status).

Our work shows that de-ciphering calendar entries and automatically discovering which category they fall under can be used successfully in pre-caching Internet content. However, the same technology can be used to solve other real needs in mobile device interaction, such as augmenting the usability of mobile calendars.

## References

- Balabanovic M, Shohav Y, Yun Y (1995) An adaptive agent for automated web browsing. *J Vis Commun Image Represent* 6:4
- Wang Z, Crowcroft J (1996) Prefetching in the world wide web. In: *Proceedings IEEE global internet conference*, London
- Cunha C, Jaccoud C (1997) *International symposium on computers and communication* 97, Alexandria
- Thiebaut D (1989) On the fractal dimension of computer programs and its applications to the prediction of the cache miss ratio. *IEEE Trans Comput* 38(7):1012–1026
- Palpanas T (1998) Web prefetching using partial match prediction. In: *Proceedings of the 4th international web caching workshop*, San Diego
- Jiang Z, Kleinrock L (1998) An adaptive network prefetch scheme. *IEEE J Sel Areas Commun* 16(3):358–368, 1–11
- Jiang Z, Kleinrock L (1998) Web prefetching in a mobile environment. *IEEE Pers Commun* 5:25–34
- Chinen K, Yamaguchi S (1997) An interactive prefetching proxy server for Improvement of www latency. In: *Proceedings of INET97*, June 1997
- Duchamp D (1999) Prefetching hyperlinks. In: *Proceedings of the 2nd USENIX symposium on Internet technologies and systems*, Boulder, Colorado
- Fan L, Jacobson Q, Cao P, Lin W (1999) Web prefetching between low-bandwidth clients and proxies: potential and performance. In: *Proceedings of the joint international conference on measurement and modelling of computer systems (SIGMETRICS 99)*. Atlanta, Georgia
- Pitkow J, Pirolli P (1999) Mining longest repeated subsequences to predict www surfing. In: *Proceedings of the second USENIX symposium on Internet technologies and systems*, October 1999
- Schechter S, Krishnan M, Smith MD (1998) Using path profiles to predict HTTP requests. In: *Proceedings of the seventh international www conference*. Brisbane, pp 457–467
- Padmanabhan V, Mogul JC (1996) Using predictive prefetching to improve www latency. *ACM SIGCOMM Comput Commun Rev* 26(3):22–36
- Crow D, Smith B, Habits DB (1992) Comparing minimal knowledge and knowledge-based approaches to pattern recognition in the domain of user-computer interactions, neural networks and pattern recognition in human computer interaction. Ellis Horwood, New York, pp 39–63
- Swaminathan N, Raghavan SV (2000) Intelligent pre-fetching in www using client behaviour characterization. In: *Proceedings of the eighth international symposium on modeling, analysis and simulation of computer and telecommunication systems (MAS-COTS)*
- Foygel D, Strelow D (2000) Reducing web latency with hierarchical cache-based prefetching. In: *Proceedings of the international workshop on scalable web services (in conjunction with ICPP0)*, Toronto
- Davison B (2002) Predicting web actions from HTML content. In: *Proceedings of the 13th ACM conference on hypertext and hypermedia*, College Park, pp 159–168
- Zhang W, Lewanda DB, Janneck CD, Davison BD (2003) Personalized web prefetching in Mozilla. Technical report LU-CSE-03-006, Department of Computer Science and Engineering, Lehigh University
- Davison B (2004) Learning Web Request Patterns. In: Poulavassilis A, Levene M (eds) *Web dynamics: adapting to change in content, size, topology and use*. Springer, Heidelberg
- Cohen E, Kaplan H (2000) Pre-fetching the means for document transfer: a new approach for reducing web latency. In: *Proceedings of the 2000 IEEE INFOCOM conference*. Tel-Aviv, pp 854–863
- Komninos A, Dunlop MD (2003) Towards a model for an Internet content pre-caching agent for small computing devices. In: *Proceedings of the 10th international conference on human computer interaction (HCI2003)*, Crete

22. Kincaid CM, Dupont PD, Kaye AR (1985) Electronic calendars in the office: an assessment of user needs and current technology. *ACM Trans Off Inf Syst* 3(1):89–102
23. Appelt D, Israel DJ (1999) Introduction to information extraction technology. *Int J Commun Artif Intell* 12:161–172
24. Komminos A, Dunlop MD (2004) Keyword based categorisation of calendar entries to support personal internet content pre-caching on mobile devices, 2nd international workshop on mobile and ubiquitous information access (MUIA04), in conjunction with Mobile HCI04. Glasgow
25. Jansen BJ, Spink A, Bateman J, Sarasevic T (1998) Real life information retrieval: a study of user queries on the web. *ACM SIGIR Forum* 32(1):5–17
26. Jansen B, Spink A (2003) An analysis of web documents retrieved and viewed. In: *Proceedings of the 4th international conference on Internet computing*. Las Vegas, pp 65–69
27. Nielsen J, Morkes J (1997) Concise scannable and objective: how to write for the web, <http://www.useit.com/papers/webwriting/writing.html>. Link valid January 2005
28. Claypool M, Le P, Waseda M, Brown D (2001) Implicit interest indicators. In: *Proceedings of the 6th international conference on intelligent user interfaces (IUI '01)*, USA, pp 33–40
29. Morita M, Shinoda Y (1994) Information filtering based on user behavior analysis and best match text retrieval. In: *Proceedings of the 17th annual international ACM SIGIR conference on research and development in information retrieval*. Ireland, pp 272–281
30. Kelly D, Belkin N (2001) Reading time, scrolling and interaction: exploring implicit sources of user preferences for relevance feedback. In: *Proceedings of the 24th annual international ACM conference on research and development in information retrieval*. New Orleans, pp 408–409
31. Kelly JD (2004) Understanding implicit feedback and document preference: a naturalistic user study. PhD Thesis, State University of New Jersey