

A Testbed System for Mobile Point-of-Care Information Delivery

Susan E. Hauser, Dina Demner-Fushman, Glenn Ford, George R. Thoma
*Communications Engineering Branch, Lister Hill National Center for Biomedical
Communications, National Library of Medicine, National Institutes of Health,
Department of Health and Human Services, Bethesda, MD 20894
hauser@nlm.nih.gov, <http://archive.nlm.nih.gov/proj/pmot/pmot.php>*

Abstract

PubMed on Tap is a testbed system that supports search of and retrieval from the National Library of Medicine's MEDLINE® database from a PDA. The goal of the PubMed on Tap project is to discover and implement design principles for point-of-care delivery of clinical support information. The project explores user interface issues, information content and organization, and system performance. Here we present our progress in these areas.

1. Introduction

The Lister Hill National Center for Biomedical Communications, a research and development division of the National Library of Medicine (NLM), has undertaken a project to discover and implement design principles for point-of-care delivery of clinical support information. This implies supporting mobile healthcare professionals, which in turn implies a mobile information delivery system. Modern technology has supplied components of such a system in the form of wireless networks, wireless handheld computers (PDAs) and tablet PCs. Although these small computing devices are capable of both sophisticated computation and high-level communication, they are inherently limited by a small screen and by relatively narrow bandwidth, as compared to desktop computers. Our research and development is motivated by the hypothesis that information systems for small wireless computers require special consideration of both user interface design (display, navigation), and information content (greater precision): it is not sufficient to merely scale the desktop computer interface to a smaller screen.

Research in this area is particularly timely. Interest in Evidence Based Medicine (EBM) has grown with the technology to support point-of-care access to information [1]. Many health care professionals regularly use PDAs for access to pharmacology databases, infectious disease databases, dosage calculators and other applications. As wireless technology has become more secure, less costly and easier to administer, health care centers and medical schools have installed or plan to install wireless networks. The synergy of technology and professional interest provides an opportunity to combine emerging technologies to better serve mobile users of NLM databases.

2. Testbed system design

Our research approach is to build a working system as a testbed, recruit users for the system, solicit feedback, and compile aggregate usage statistics to discover usage patterns and preferences. Selecting from several designs [2], we developed a testbed consisting of a client program for PDAs with the Palm operating system and a Unix-based proxy server. A block diagram of the testbed system, called PubMed on Tap, is shown in Figure 1.

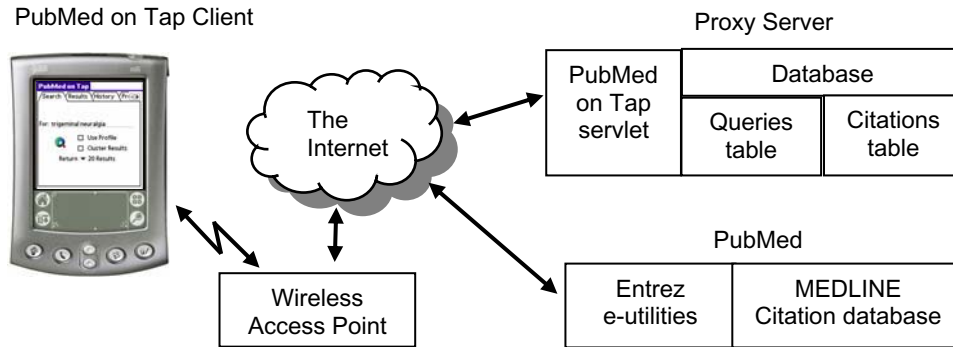


Figure 1. Testbed system overview.

We chose MEDLINE®, NLM’s premier database of bibliographic citations, as the data source because of its value to EBM and because of third party access to NLM’s PubMed search engine and XML-formatted search results that are available via the Entrez e-utilities [3]. Our client software for the PDA manages the user interface, stores user preferences on the PDA, and communicates via http and the PDA’s Internet connection to the proxy server. The proxy server submits queries to PubMed and returns search results to the client. It also stores queries and citations in a local database to improve performance for subsequent queries. A separate Update module runs nightly to refresh the citations with the current results from PubMed for the stored queries. Because the PubMed on Tap client manages the user interface, Search results and citation data are sent to the client as XML-like tagged data rather than in html format, thus saving the overhead associated with html data delivery.

The first version of the PubMed on Tap testbed system was released for testing in late April, 2003. Through our Web presence [4], demonstrations at conferences and small meetings, registration at freeware Web sites and Web search engines, and NLM’s network of libraries, we have recruited over 100 registered users plus many more who have downloaded the PubMed on Tap client software. Through personal contact, emails from our users and comments on our online forum, we have begun to learn some of the design features that are useful and some that are not. A small usability study conducted in July, 2003, was a particularly valuable tool for revealing the best and worst aspects of our initial design [2, 5]. In response to these feedback channels, a few most-requested features have been added to the initial design, and the user interface has been modified to address difficult or confusing aspects of navigation.

3. Client design

Our first client has been created for PDAs running the Palm operating system, because Palm devices are currently the most popular among physicians [6]. The client is written in C/C++ using the Code Warrior development environment. Each aspect of the user interface required consideration of the space available on a standard 2" x 2" PDA screen, the readability of text, and the number of taps required to execute a choice or navigate to a new location. Although vertical scrolling is sometimes necessary, horizontal scrolling is eliminated. Small evocative icons are used in favor of text-filled buttons whenever possible. In some cases, we relied on users’ computer experience and analytical nature to discover the required tap.

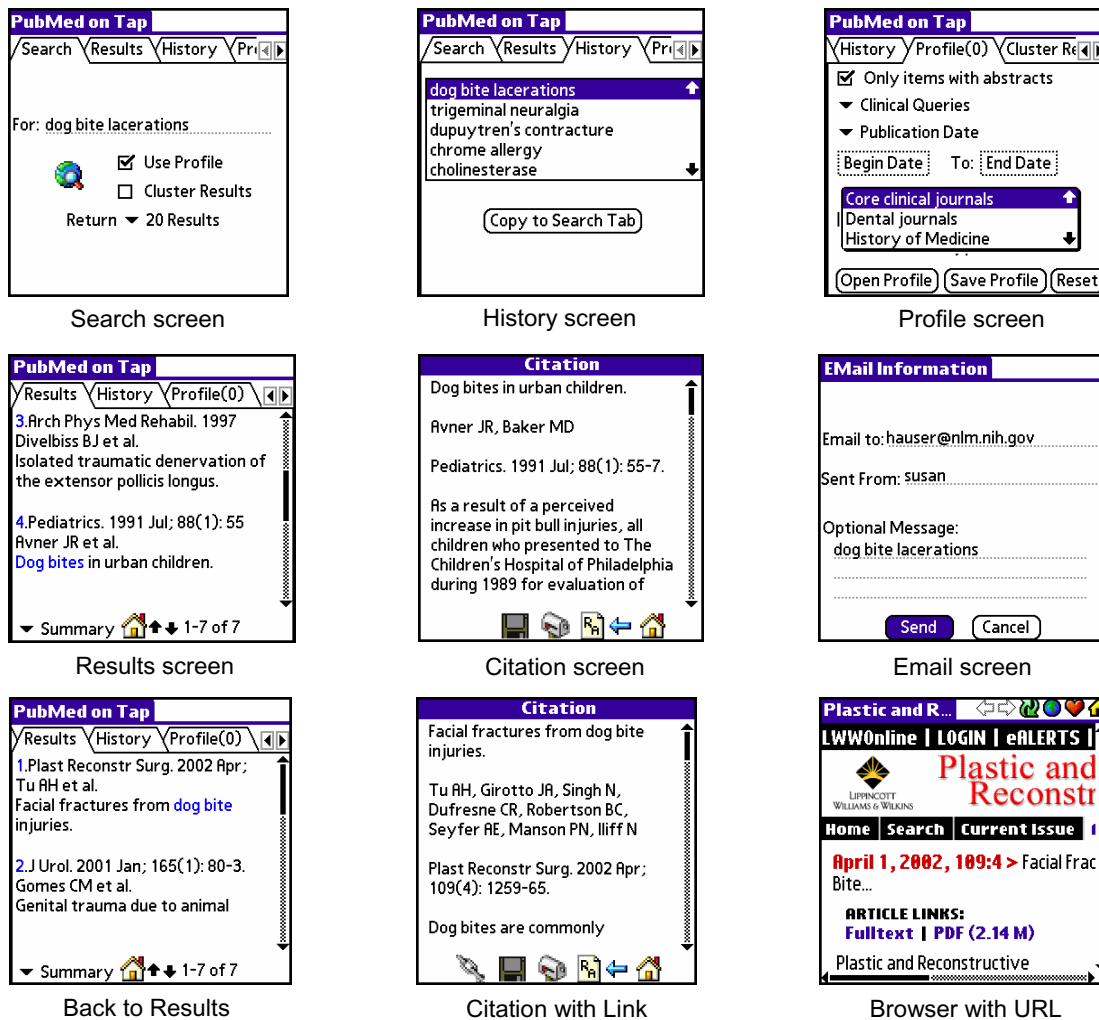


Figure 2. PubMed on Tap user interface and features.

Using “dog bite lacerations” as an example query, Figure 2 shows aspects of the current client user interface and several system features. The interface is organized into screens that the user can navigate among by tapping on the associated screen tab or by using a drop down menu available from the “PubMed on Tap” label in the upper left corner. The application opens with the Search screen, where the user enters query terms in the search field. Alternatively, the user may select from previous queries, stored locally and displayed on the History screen, and edit the query on the search screen. Search limits are selected from the Profile screen and implemented by tapping “Use Profile” on the Search screen. By default, 20 results are returned at a time, but the user can select 5 or 10 from the drop down box below the search field. The search is executed by tapping the globe, at which point the client sends the query, including search limits, to the proxy server. If the query is identical to a query in the Query table, current results are already in the Query and Citations tables, so the proxy server sends the results to the client without communicating with PubMed. If the query is new, the proxy server submits an esearch and then an efetch to PubMed to obtain search results. The results are stored in the proxy server’s database, and short summaries are extracted and sent to the client. The client automatically switches to the Results screen where these are displayed. To minimize the need to scroll, only the publication, first author name

and up to two lines of title are displayed for each article. The user can tap the down arrow at the bottom of the screen to obtain the next 20 (or 5 or 10) results. To view a complete citation, the user can tap anywhere on the summary text. The client then sends a request for the citation to the proxy server. The proxy server retrieves the citation from the local database and sends it. The client displays all available citation data: title, author, publication, abstract, affiliation, MeSH terms, substances and PMID. In addition to the Home and Back icons, the bottom of the Citation screen has three icons for cases where the user may have found this citation to be of interest: The Related Articles icons retrieves the list of PubMed-selected Related Articles; the floppy disk icon saves the citation to the PDA's Memo Pad from where it can be hotsync'd to the user's desktop computer or beamed to a colleague; the mailbox icon lets the user email the citation. The citation is emailed from the proxy server, so it is not necessary for the user's PDA to have email capability.

High on the list of features requested by our users is the ability to access the full text of articles from PubMed on Tap. Link-out to full text capability has been developed for the Palm OS client and will be released once it is fully tested, probably in April 2004. Support for this feature is possible through a recent Entrez e-utility that returns the URL to the full text of an article, if it exists. Because a separate request to PubMed is needed to obtain the link, there is additional overhead associated with retrieving a citation. The URL information is returned to the client with the citation and also becomes part of the citation record in the proxy server database. If a link exists, the client displays an additional link icon at the bottom of the citation screen, as shown in Figure 2. When the user taps the icon, the PDA's browser is launched with the given URL. At this point, the browser and the publication's server take over communications and data layout, and the user is responsible for logging into the publication's site, if needed.

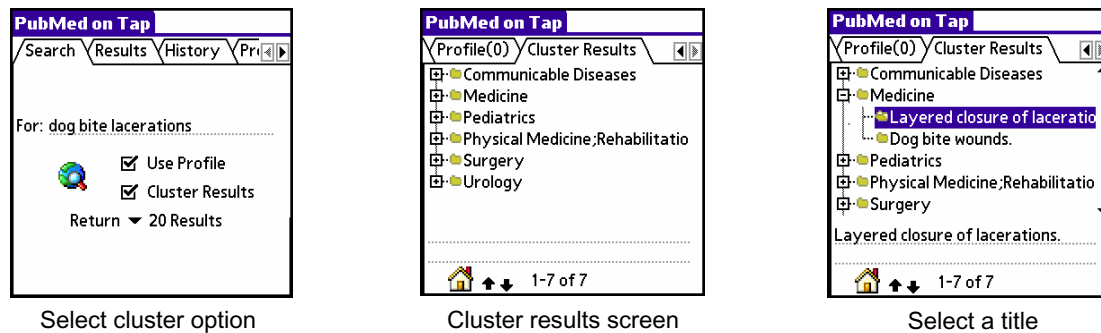


Figure 3. Results clustered by journal subject area.

As an option to viewing results in PubMed's traditional method of listing citations by date, PubMed on Tap users may choose to have search results clustered by subject area by tapping the Cluster Results box prior to executing the search. For the approximately 4500 journals indexed by NLM, subject area is determined by the subject field of the journal in which the article appears. Subject areas for more than 4000 additional journals have been determined following rules established by the Lister Hill Center's Indexing Initiative research project. Citations from PubMed include a journal title identifier that the proxy server uses in conjunction with a local table of journal title identifiers and subject areas to rapidly assign a subject area to retrieved citations on the fly. The server returns only titles and subject areas to the client in response to a search with clustered results. As shown in Figure 3, the PubMed on Tap client displays the subject areas for the set of returned citations in a hierarchical tree format on the Cluster Results screen, with each cluster of citations identified by its subject area. The user can expand the tree to see abbreviated titles within each subject. When an

abbreviated title is tapped, the complete title is shown at the bottom of the screen. Tapping anywhere on the complete title causes the client to retrieve the full citation from the proxy server and display it on the citation screen. Although this technique is fast, journal subject areas themselves are uneven and do not always reflect the main topic of individual articles in the cluster. Part of our ongoing research is to explore natural language and statistical techniques to cluster citations into medical categories (Findings, Diagnosis, Intervention, Outcomes, Prevention) and extract other citation characteristics (Scope, Weight) based upon the text of the citation. In parallel with these efforts, we are exploring visualization techniques to provide cues regarding important aspects of the publication profile. For example, EBM recommends seeking information from clinical trials; therefore a recognizable symbol for the clinical trial publication type might significantly reduce the time taken to browse search results.

4. Proxy Server Performance

Because the proxy server is an extra step between the user and the source data (MEDLINE), it is important that the response time of the proxy server itself be negligible. The servlet, written in Java 1.4.2, runs in a Tomcat 5.0.19 container on an Apache http server in a UNIX environment. Queries and citations are stored in a mySQL database on the same platform as the servlet. Database transactions performed by the servlet are optimized through the use of connection pooling, a technique whereby a pool of database connections are created and are ready for use by any thread that needs one. Connection pooling allows for more controlled use of resources and more predictable behavior under loads from concurrent users. Database performance is further optimized by executing batch inserts and retrievals whenever possible. We test performance of the proxy server from desktop computers with a tool called Jmeter. Following user-defined scripts for content and timing, Jmeter issues http requests to the proxy server and records response time. When all the requested data is in the proxy server database, average response time to a request for a query (i.e. a search) is 33 ms, and average response time to a request for a citation is 23 ms. In both cases, the standard deviation = 19 ms.

Of course, most user sessions will require search and retrieval from PubMed, whose resources are heavily used. Over one day, the response time measured by JMeter for 800 requests routed to PubMed varied between 125 ms and 56 seconds, with an average of 804 ms. The proxy server also records the amount of time to service each request, time that does not include data transmission between the client and the proxy server. 4707 requests from PubMed on Tap clients outside of the NLM network were received between mid-November, 2003 and mid-March, 2004. Presumably, most of these required a request to PubMed from the proxy server. The average internal service time for these requests was 1.7 seconds; the median internal service time was 50 ms.

5. User preferences

As the PubMed on Tap system becomes more widely used, we anticipate that analysis of user queries retained in the proxy server database will yield useful information regarding user preferences. For example, of the 4707 requests mentioned in section 4, 2032 were original queries, 498 were continuations of a query (i.e. beyond the first 20 articles), 2009 were requests for citations, 234 were requests for related articles and 188 were requests for citations to be emailed. Of the 2032 original queries, 9.3% resulted in an emailed citation and 11.5% resulted in a request for related articles, suggesting that these are useful features. It is also encouraging to find that since

clustered results were introduced and monitored, 10.4% of original and continuation queries requested clustered results. Although the use of search limits (i.e. Profiles) was not the norm, 17.8% of original queries included one or more search limits. The most frequently used limits were Abstract Only (5.8%), Publication Date (4.5%), and Clinical Queries (3.2%).

Assuming that emailed citations were of interest to the user, and possibly directly related to the user's query, we chose the first 33 citations sent in an email in March 2004 for closer analysis. Of these, 18 were published in 2002 through 2004, and 15 published between 1979 and 2001. 10 articles were not yet indexed. Of the 23 indexed articles, 14 (61%) were labeled with a publication type other than journal article: 4 Clinical Trials, 3 Reviews, 3 Case Reports and 3 Letters. The wide variety of dates and publication types indicates that common search limits alone do not adequately screen for articles of interest. Our research in clustering techniques [8] addresses this need.

6. Summary

The PubMed on Tap system is a fast and stable system for exploring information content and user interface issues associated with information delivery to mobile healthcare providers. User feedback through our user forum and from personal contact at meetings and demonstrations heavily influences features that are modified or added. In addition, we learn which features are actually used by analysis of queries and emailed citations. Our experience to date shows that all features are used to some extent. Although greater system utilization will provide more data toward understanding user preferences, we may discover that there is no "typical" information seeker among mobile healthcare providers, and the best strategy will be a feature-rich but flexible system.

7. References

- [1] Lewis M. Evidence-based-medicine tools for your Palm-top computer. *Family Practice Management*, American Academy of Family Physicians, May, 2003, pp. 73-6.
Available at: <http://www.aafp.org/fpm/20030500/73evid.html>
- [2] Hauser SE, Demner-Fushman D, Ford G, Thoma GR. PubMed on Tap: Discovering design principles for online information delivery to handheld computers. Accepted for publication in *Proceedings of the 11th World Congress on Medical Informatics*, San Francisco, CA, Sep 7-11, 2004.
Available at: <http://archive.nlm.nih.gov/pubs/biblio/biblio.php>
- [3] Entrez Programming Utilities. February 3, 2004.
Available at: http://www.ncbi.nlm.nih.gov/entrez/query/static/eutils_help.html
- [4] PubMed on Tap. March 2004. Available at: <http://archive.nlm.nih.gov/proj/pmot/pmot.php>
- [5] Alexander G, Hauser SE, Steely K, Ford G, Demner-Fushman D. A usability study of the PubMed on Tap user interface for PDAs. Accepted for publication in *Proceedings of the 11th World Congress on Medical Informatics*, San Francisco, CA, Sep 7-11, 2004.
- [6] Criswell DF, Parchman ML. Handheld computer use in U.S. family practice residency programs. *JAMIA*, 2002, 9, pp. 80-6.
- [7] Humphrey SM, Rindflesch TC, Aronson AR. Automatic indexing by discipline and high-level categories: Methodology and potential applications. *Proceedings of the 11th ASIST SIG/CR Classification Research Workshop*, Chicago, IL, Nov 12, 2000, pp. 103-16.
- [8] Demner-Fushman D, Hauser SE, Ford G, Thoma GR. Organizing literature information for clinical decision support. Accepted for publication in *Proceedings of the 11th World Congress on Medical Informatics*, San Francisco, CA, Sep 7-11, 2004. Available at: <http://archive.nlm.nih.gov/pubs/biblio/biblio.php>