

Contextual Contact Retrieval

Jonathan Trevor, David M. Hilbert, Daniel Billsus, Jim Vaughan and Quan T. Tran¹

FX Palo Alto Laboratory
3400 Hillview Avenue, Bldg. 4
Palo Alto, CA 94304 USA
+1 650 813 7233

{trevor, hilbert, billsus, vaughan}@fxpal.com

¹College of Computing, GVU Center
Georgia Institute of Technology
Atlanta, GA 30332 USA
+1 404 385 1102
quantt@cc.gatech.edu

ABSTRACT

People routinely rely on physical and electronic systems to remind themselves of details regarding personal and organizational contacts. These systems include rolodexes, directories and contact databases. In order to access details regarding contacts, users must typically shift their attention from tasks they are performing to the contact system itself in order to manually look-up contacts. This paper presents an approach for automatically retrieving contacts based on users' current context. Results are presented to users in a manner that does not disrupt their tasks, but which allows them to access contact details with a single interaction. The approach promotes the discovery of new contacts that users may not have found otherwise and supports serendipity.

Categories and Subject Descriptors

H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces – *Computer-supported cooperative work*.
H.4.1 [Information Systems Applications]: Office Automation – *Groupware*.

General Terms

Human Factors

Keywords

Contact management, context, proactive retrieval, social networks.

1. INTRODUCTION

Personal and organizational contacts are critical to getting work done, particularly in knowledge-intensive professions [1]. However, we often forget our contacts, or must refresh our memories regarding their details. In some cases, we may be unaware of potentially relevant contacts who are well known to others in our social networks. This can reduce efficiency, collaboration, and productivity.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright is held by the author/owner(s).

IUI'04, Jan. 13–16, 2004, Madeira, Funchal, Portugal.

ACM 1-58113-815-6/04/0001.

Today there are many ways to capture information regarding contacts: we collect business cards; we record contact information; we record information regarding past interactions, such as phone calls and meetings; we may even shoot photos or video, or use sophisticated contact and relationship management software to keep track of important contacts. However, regardless of the method used to capture such information, we must manually look it up when we need it. If you have a physical rolodex, you flip through it. If you have an electronic rolodex, on-line address list, directory, or database, you search it. In each case, you must take action to find contacts. Even in MS Outlook or relationship management software, the system may look-up contacts on your behalf, but only if you begin typing contacts' names into particular text fields, or if you receive a message from (or to) a contact in your database.

This paper presents a system which seeks to improve our ability to capture, reuse, and share personal and organizational contacts. It goes beyond existing systems by proactively retrieving contacts based on the user's context (e.g., the content of a displayed email message or web page), and provides one-click access to relevant contact information without disrupting ongoing tasks.

2. APPROACH

Our approach involves two main components: a capture component and a recommender component. The capture component captures contact information and stores it in a database. The recommender component analyzes the user's context to proactively recommend potentially relevant personal or shared contacts.

The system supports multiple ways to capture contact information. For example, a video guestbook kiosk that resides in a public location near the main entrance to our organization is used to record visitor information (see Figure 1). The kiosk consists of a conversational character running on a touch screen display (to greet and instruct visitors), a business card scanner (to capture visitors' contact information), and a video camera and microphone (to capture faces and name pronunciations).

By default, the kiosk stores visitor contact information in a shared corporate rolodex. Employees can also use the kiosk to add contacts to their own personal rolodex (or the shared rolodex) by scanning business cards they have collected. Finally, the system provides an API to allow contact information to be imported from various sources such as an existing corporate directory or users' electronic address books. All captured information is stored in a

contact database that can be accessed via a web interface for managing personal and corporate rolodexes, or via a web service for contact retrieval, management and contextual matching. The web service can be utilized by corporate applications that require access to contact information, such as the front end of our contextual contact recommender.



Figure 1: Video Guestbook Kiosk

The recommender front-end runs as a toolbar in MS Outlook and Internet Explorer, and monitors employees' email and web use to recommend relevant contacts at the time of need. To accomplish this, the toolbar extracts the full text of the currently displayed email message or web page and asynchronously requests contact recommendations using our contextual contact retrieval web service. Since all contact matching operations run on the server and the web service is called asynchronously, there is no noticeable performance impact on users' web or email experience.

To protect users' privacy, all interactions with the server are logged anonymously, proactively transmitted text is never persisted, and users can easily opt out of proactive recommendations. The server identifies contacts related to the content of the currently displayed email message or web page, and the toolbar unobtrusively displays matched names of individual contacts or organizations (see Figure 2). For instance, an employee reading an email message regarding a corporate partner could receive contact recommendations regarding recent visitors working for that corporate partner.

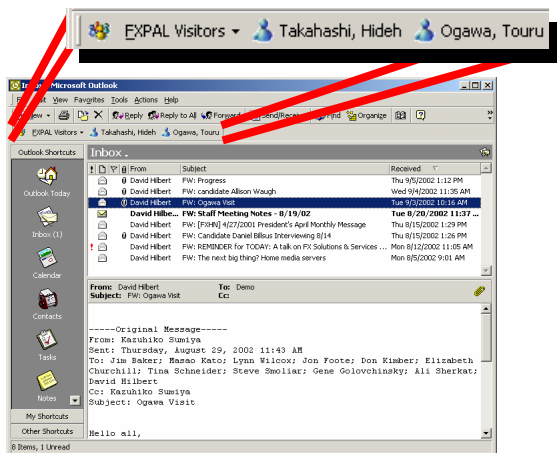


Figure 2: Contextual Contact Bar

By clicking on one of the recommendations displayed in the recommender toolbar, users can access additional information about contacts, such as their business cards or video clips that indicate what they look like, how to pronounce their names, and who they visited at the organization (see Figure 3). Selecting the "Why" link next to their name highlights matched terms in the current email or web page to indicate why the contact was recommended. Selecting a recommended organization shows a list of people from that organization who have visited before. If a user has not met the contact before, he or she can find out who at their organization hosted the visitor or added the contact to the system to learn more.

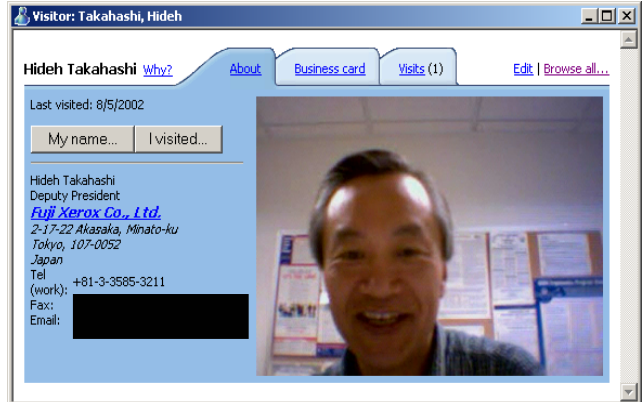


Figure 3: Contact Details

In summary, by providing contextual contact recommendations our system: (a) helps users discover new contacts when they need them, (b) helps users remember details such as what contacts look like, how to pronounce names, or who they know, and (c) leverages existing work practice by presenting recommendations in familiar applications without interrupting ongoing tasks.

3. MATCHING ALGORITHM

The recommender component analyzes transmitted text and uses a matching algorithm to detect explicit occurrences of contact information fragments that match entries in the system's contact database. The algorithm must be sophisticated enough to deal with a wide variety of potential formatting differences of names and contact information.

Our recommender implementation performs a three step fuzzy matching algorithm: (1) first, address information (postal addresses, email addresses, phone numbers) is extracted, matched, and scored; (2) then names and organizations in the database are checked against the document text; (3) Finally, a list of recommendations is constructed and ordered by confidence scores.

Postal addresses. First, the algorithm extracts zip code candidates by finding 5- or 9-digit numbers adjacent to line and paragraph breaks (since zip codes are the easiest address element to detect and are at the end of addresses). Once a zip code is found, the text before it is scanned for one to three other candidate address lines. If during this search a line is encountered that contains more words than a prescribed limit, the search is aborted and the search for zip codes resumes. The database is queried for organizations with the same zip codes as addresses extracted from the document. All matching entries are given a confidence score,

based upon the number of matching words in the city and street lines, ignoring common words such as “avenue” or “street”.

Email addresses. Regular expressions are used to extract email addresses. If an exact match is found in the database, it is treated as a high-confidence match for the person it relates to. Otherwise, a fuzzy match based on the domain portion of the address is attempted. These are considered to be matches to organizations and are scored based on the ratio of matched domain segments to total segments.

Phone numbers. Regular expressions are used to find 10 digit phone numbers, and for each extracted number, the database is queried for all matching home, work, mobile, and fax numbers. Confidence scores are adjusted depending on the type of phone number matched. For instance, a fax number is given a high score for an organizational match, but a low score for a person match. In contrast, a mobile phone number is scored highly for a person.

Names and organizations. Names and organizations are not extracted from the document as the entities described above. While using a named entity tagger would have been an option, it is likely that explicitly searching for known names leads to more accurate results. Since the number of names in the database is limited, this approach is feasible and efficient enough for real-time matching. Matches of the form [FirstName LastName] or [LastName, Firstname], are scored highly. The score is increased slightly based on the number of times the names are found in the document. Matches of the form [Initial LastName] [LastName, Initial] receive lower scores. An organization match begins when the first word of an organization is found in the document. The score is based on the number of remaining words in the organization name that are found near the first word, ignoring common words, such as “Inc”, “Co”, etc.

Finally, all the matches that relate to the same person or organization are merged into a single entry with a combined score, and entries are sorted according to their scores. Matches that score below a threshold or relate to people who have not been matched by name are deleted. Each match contains a list of strings found in the document. If the user asks for clarification regarding why a person or organization was recommended, each occurrence of the string is highlighted in the original document.

4. EVALUATION

The guestbook kiosk has been deployed at our organization for approximately one year, and is steadily gaining in popularity, i.e. employees frequently encourage visitors to sign in and leave an electronic record of their visit. Numerous visitors and employees have commented that having an internally accessible multimedia database of visitors seems very useful. The contextual contact bar has only recently been made available to employees of our organization. While we have not collected enough usage data for a meaningful statistical evaluation of the recommendation accuracy or service utility, users have reported that the toolbar has suggested relevant contacts, although some have complained that the user interface is so unobtrusive they sometimes do not notice recommendations. The heuristic matching algorithm, in particular, seems to perform exceptionally well—the approach rarely misses

contacts. While our initial subjective impressions are positive, a more formal evaluation will be performed to assess the value of our technology.

5. RELATED WORK

Related work falls into three categories: contact management systems, expertise location systems, and content-based inference engines. Contact management systems include email programs (e.g., Microsoft Outlook), business card scanners and electronic rolodexes (e.g., CardScan and BizCard Reader), and customer relationship management software (e.g., GoldMine, TeleMagic and Act). Expertise location systems include systems such as AnswerGarden2 [2] and ReferralWeb [3]. In contrast to our approach, none of these systems analyzes the user’s context to proactively recommend contacts. Content-based inference engines are commonly used in match-making systems such as Yenta [4], knowledge management systems such as Verity and Autonomy, and just-in-time (JIT) retrieval agents [5, 6, 7]. However, these systems typically recommend documents, not contacts, and use only general textual similarity as opposed to explicit contact information matching.

6. CONCLUSIONS AND FUTURE WORK

The contextual contact retrieval system described in this paper uses a set of simple contact matching heuristics to provide one-click access to relevant contacts, based on currently displayed email messages or web pages. While our initial experience has been encouraging—helping users recall and discover contacts at the time of need—future work will focus on: (1) increasing users’ awareness of recommendations, e.g., by embedding unobtrusive indicators within email and web page text, (2) enriching our retrieval methods with rich content-based models of employees, visitors and contacts; and (3) formally evaluating the utility of this service.

7. REFERENCES

- [1] Nardi, B.A., Whittaker, S., Isaacs, E., Creech, M., Johnson, J., and Hainsworth, J. (2002). Integrating communication and information through ContactMap. *Communications of the ACM*, 45(4).
- [2] Ackerman, M.S. and McDonald, D.W. (1996). Answer Garden 2: Merging Organizational Memory with Collaborative Help. *Proc. of CSCW 1996*: 97-105.
- [3] Kautz, H., Selman, B., & Shah, M. (1997). The Hidden Web. *AI Magazine*, 18(2), 27-36.
- [4] Foner, L. (1997). Yenta: A Multi-Agent, Referral Based Matchmaking System. *Proc. of Agents '97*.
- [5] Rhodes, B. (2000). Just-In-Time Information Retrieval. Ph.D. thesis, MIT Media Laboratory, Cambridge, MA.
- [6] Budzik, J., and Hammond K. (1999). Watson: Anticipating and Contextualizing Information Needs. *Proc. ASIS 1999*.
- [7] P. P. Maglio, R. Barrett, C. S. Campbell, and T. Selker. (2000). SUITOR: An attentive information system. *Proc. of IUI 2000*, pp. 169-176.