# SmarterPhone – Supporting Meetings in a Mobile Context

**Humberto Glória, Tiago Guerreiro, Daniel Gonçalves**
Instituto Superior Técnico – TULisbon / INESC-ID
Rua Alves Redol, 9
1000-029 Lisbon, Portugal
hmggloria@gmail.com, tjvg@vimmi.inesc-id.pt, daniel.goncalves@inesc-id.pt
+351 213100258

## ABSTRACT

We must deal with growing amounts of information, leading to organizational and retrieval problems. This is particularly true in a mobile context. We describe how to proactively present the users with information relevant for a meeting, in a mobile context, based solely on the personal information available in their computers. We performed a study where 100 users were asked about what makes some information important for a given meeting, leading to the creation of SmarterPhone, a mobile application whose interface was crafted to enable users to efficiently access personally relevant information in a particular context. A user study showed that on average 80% of all relevant documents and people are found, demonstrating the validity of our approach and underlying relevance criteria.

## Author Keywords

Personal Information Management, Mobile Applications, Meetings, User Studies.

## ACM Classification Keywords

H.5.2 User Interfaces: User-centered design; H5.2 User Interfaces: Graphical user interfaces; H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

## INTRODUCTION

In recent years, the number and types of electronic devices used by most people has increased dramatically. This is especially true for computers that now take a multitude of shapes and can be used in hitherto unthinkable contexts. Mobile phones in particular are becoming full-fledge computers in their own right. The marked for so-called 'feature phones' is decreasing in favor of 'smartphones'.

If it is true that using a desktop computer we might be able to spend some time and effort looking for a particular piece of information, when using a mobile device we will face countless situations where that is simply not possible (while walking on the street, taking advantage of one last minute before the bus arrives, etc.). Also, the inherent limitations of mobile devices, namely the (comparatively) low processing power and reduced screen real-estate further hinder management and retrieval tasks. We focus on one particular aspect of this problem: timely access to information relevant to a meeting. It would be helpful if, using our mobile phone (that is much more likely to be with us at all times) we were able to efficiently access all personally relevant information for that meeting.

Other works have tried to help users access their personal information. MyLifeBits [3] aims at creating a personal digital archive for the entirety of a user's lifetime. It avoids the typical hierarchical folder organization, as it would scale poorly in this context. One object can easily fit into several categories instead of just one. The Stuff-I've-Seen system [1] also integrates information gathered from the documents and the different applications executed by the users (email, etc.) providing a richer context. Satchel arises from the wish of having access to a document anytime, anywhere [5]. It was developed with the main purpose of retrieving existing documents remotely. Finally, ContextPhone is a true prototyping platform for mobile applications related to the collection of data in the context they're in [6]. Although this project focuses on gathering data and not on Information Retrieval, with an extensive collection of context data it is possible to know more about the user itself and be proactive about it.

All these projects try to help the users by organizing and maintaining their Personal Information, either making it usable in a mobile context or complementing information from the users' computers with that of the context that surrounds them. However, they mostly gather and assemble huge amounts of data, without actively trying to sort through that data, proactively suggesting to the user that which would be more important in a given context. In particular, the meetings scenario is not addressed of any of these and other similar works.

Evidently, just copying all our Personal Information into the mobile phone, or even having a way to remotely access it from it would not suffice. The sheer amount of information we have at our fingertips would make an

efficient retrieval impossible. Hence, we study how it might be possible to automatically infer which information is relevant, so that that information in particular is readily available for the user in a mobile context.

## HOW USERS PREPARE A MEETING

To find *what* information users find to be the most relevant for a particular meeting, we devised a questionnaire where they were asked to classify a set of possible relevance criteria using a 5-point Likert scale. It was available online during a two-week period. We analyzed all questionnaires for quality and consistency. Those found with sub-par data quality (truncated, etc.) were excluded from the analysis. Overall, we had 100 valid respondents (out of 107).

The defining characteristics of a meeting are the *people* that attend it, the *subject* to be discussed by its participants, the *time* when all participants meet, and a location or *place*, either real or virtual. Of these, the less likely to change and more defining of a meeting is *subject*. Most meetings are scheduled to address a particular subject, and only then will the other elements be defined. *People* are tightly coupled with subject, as certain issues will most likely be addressed by certain people. This led us to conclude that we should help the users find *people* and *documents* relevant for a meeting. A document, in this context, should be understood as any discrete piece of information in the user's computer (actual text document, picture, email message, etc.). A list of possible relevance criteria was validated and completed based on the results of preliminary interviews with six users, resulting in those in Tables 1 and 2.

### Table 1 – Relevance Criteria for People

| | |
|---|---|
| A | Emails about the subject of the meeting were exchanged with that person |
| B | Emails about the subject of the meeting were exchanged with that person *around the time of the meeting.* |
| C | SMSs about the subject of the meeting were exchanged with that person |
| D | A phone call was made to that person shortly after the meeting was scheduled |
| E | A phone call from that person was made shortly before the meeting was scheduled |
| F | Some of the documents relevant to the meeting have that person as an author |
| G | That person was present on previous meetings on the same subject. |

### Table 2 – Relevance Criteria for Documents

| | |
|---|---|
| H | Its title is related to the meeting |
| I | Its content is related to the meeting |
| J | It was received in an email message from a person related to the meeting |
| K | It was sent in an email message from a person related to the meeting |
| L | Its author is someone relevant for the meeting |
| M | It has been relevant on a previous meeting on the same subject. |
| N | It was created around the time the meeting was scheduled. |

Many more criteria might be considered of, but given our goal of finding general criteria applicable to the creation of different applications in different contexts, only the most generally accepted were left in the questionnaire. Even so, we left the users with the choice of indicating "other criteria" they might find relevant.

The 100 respondents of the questionnaire were very varied, indicating that our results can be representative of a wide range of users, and not biased towards a particular subset of them. The users' ages ranged from 22 to 68 years. Ten users were retired from their professions, but still active with other activities that require them to attend meetings on a regular basis. The remainders were professionals in areas ranging from engineering to education or medicine. All attended meetings and used computers regularly.
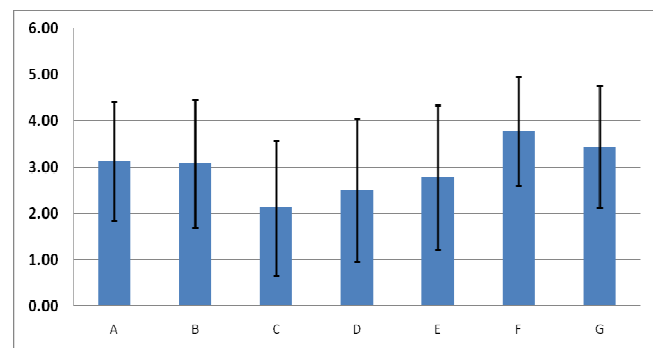


**Figure 1 – Criteria for Relevant People (avg. scores; st.dev.)**

### Relevant People

The most noteworthy results (Figure 1) are that *"F – Have relevant documents authored by the person"* is the most relevant criterion, if only by a small margin, and criterion *"C – SMS are exchanged with that person"* is noticeably less relevant. A 95% significance ANOVA test confirms a statistically significant difference between them and the other criteria, an important design implication.

The users were allowed to suggest other criteria they found of particular relevance. For the most part, those were very specific, tailored for a user's own reality and methods, leading to the conclusion that a system that makes use of these results should account for some user variability and customization. Two comments, however, are applicable to the general case: *a person is relevant to a meeting if this meeting is about that person*, or if *it is noted in the calendar as being a participant*.

We conclude that, regarding which people users find more relevant for a particular meeting, the use of SMS messages is not a good indicator. Existing documents by that person, are the best way to rate its relevance, as is presence in previous meetings, or the fact that the meeting is about that person. Email appears to be a privileged way of arranging meetings in current society, as the rules regarding email as a clue (A and B) are the next best rated.

### Relevant Documents

The results for documents aren't as homogeneous as those for people. One criterion is notably a better indicator of relevance than the others: *"I – The content of the document is related to the meeting"*. This was to be expected but, still, it is important to have verified it (Figure 2).
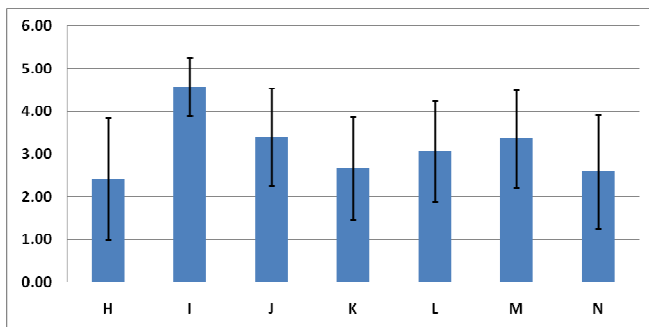


**Figure 2 – Criteria for Relevant Docs. (avg. scores; st.dev.)**

Other notable differences appear between criteria *"J – Got the document by email from someone relevant"* and *"M – The document has been used in a previous related meeting"*. An ANOVA test with 95% confidence showed that indeed there are statistically significant differences between the different criteria.

We conclude that a document's subject, inferable from its content, is the best indicator to whether it is relevant for a meeting or not. The relation of those documents to relevant people (authors, senders) is more relevant next. The users' own actions (creating a document or sending it to someone) seem less trustworthy indicators of relevance. Users seem to defer to external authority when determining what documents are relevant. This might stem from the fact that documents they themselves produce and are relevant to a meeting contain information that they are not very concerned with since they intimately know it.

### THE SMARTERPHONE INTERFACE

To validate the effectiveness of the criteria as defined in the interviews we created SmarterPhone, a prototype system that, given a meeting appointment, is able to select potentially relevant people and documents, and make them available to the user in a mobile setting.

To conveniently access all of a user's Personal Information (PI), that information is indexed by an already existing system, Personal Biographer [**Error! Reference source not found.**], outside the scope of this paper. All relevant information, including some coming from the smartphone such as *contacts, agenda, SMS* and *call log*, is stored interconnected in a semantic network where inference rules can be evaluated facilitating high-level reasoning. This is done on a PC, with less memory and computing constraints, and communicated to the smartphone on request, over the Internet.

We tried to keep the interface simple, highlighting relevant information, and following a simplified browsing model. On the first screen, the user is presented with a list of the different meetings known to the system. After choosing one, they are led to a screen where they can decide whether to look at relevant documents or people. Choosing one of the entities will lead to the Entity List screen (Figure 3, left) where all entities of the selected type are displayed. Their description consists not only of details about itself, but also a list of related entities. In the figure, for instance, people related to the "smarterphone.pdf" document are shown (the file was emailed from one to the other). Tapping on any of them will jump to the details screen of that entity. All information made available by SmarterPhone is, thus, interconnected. This makes interacting with it and finding information easier, but serves an additional purpose: related entities function as a justification or explanation for some element. This helps the user understand the significance of the system's choices, increasing trust.
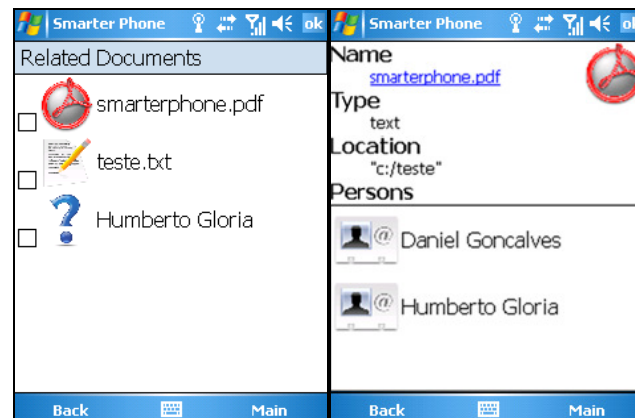


**Figure 3 – SmarterPhone Mobile Interface**

We devised a set of inference rules that embody the criteria found in the aforementioned questionnaire. The knowledge base where all Personal Information is indexed has an expressiveness level equivalent to first order logic, which was enough for our needs. Applying a particular criterion is, thus, just a matter or evaluating the inference rules it gave rise to and gathering the results. The rules determining a document's relevance are highly dependent on relevant persons, and vice-versa. After some relevant documents are found, they can prompt the discovery of relevant people (their authors, for instance). Those persons, in turn, can lead to the discovery of other documents, and so forth. We enhanced the expressiveness of the inference rules to manage this mutually recursive process. SmarterPhone will iteratively evaluate all inference rules for documents, and then those for people, until no new entities are found or a pre-determined maximum number of iterations is reached.

It can happen that the same document or person is identified by more than one rule. This is corroboratory evidence of that entity's relevance. Also, it derives directly from the questionnaires that not all rules equally are important. Each was given a weight based on this fact. All entities

discovered with its help will be given that score. If the entity is found by more than one rule, the scores will be added. The ten (or less, if fewer are found) better ranked documents and persons are the ones displayed in SmarterPhone, sorted according to that rank.

## USER TESTS

The prototype system was used in real-life situations to verify if the relevance criteria are adequate. Fifteen users participated in a user test, conducted over a period of one month. A set of "meeting record sheets" were given to each user. Those sheets had fields where the users could enter relevant information about their meetings, to help them recall and explain them at a later time. The users were requested to record the names of all participants in the meetings and a reference to all documents used or mentioned during them. They were to take a record sheet to each of the meetings they had during the course one month, and to record their observations during the meeting or right after it, to prevent them from forgetting relevant data. At the end of the month, we interviewed all of the users and their personal information was indexed and accessed through SmarterPhone. Its results were compared to the information recorded in the Meeting Record Sheets, allowing us to verify their quality. Due to technical limitations, not all criteria could be implemented, as they were impossible to assess. This might, in some cases, have prevented SmarterPhone from finding relevant results.

Overall, 64 meetings were logged over the one month period, for an average of 4.27 meetings per user (st.dev.=1.39, median=4). SmarterPhone was able to find **78.74%** (st.dev. 16%) of relevant documents and **85.60%** (st.dev. 7.77%) of relevant people (after removing those that didn't exist in electronic format and, thus, would be impossible to find). These are very good results, considering that not all criteria were implemented. Fall-Out, shows how many irrelevant results out of all possible irrelevant values, were returned. For documents, it has a value of 0.095, a very low probability considering that there are thousands of possible documents in the users' hard drives with only very few relevant for a single meeting.

Mean Average Precision gives us an estimate of the quality of our ranking. Ideally, a value of 1 would indicate that all relevant results were the best ranked and appear first in the list. For documents, its value is, on average, **0.25** (stdev=0.09). This shows that the relevant documents aren't the ones best ranked but, rather, appear on the middle of the 10-result list. The ranking function needs to be improved.

Looking at the individual criteria, while some are more important, none are unequivocally able to find the near totality of relevant entities. It is from their use in a synergistic way that good results arise.

## CONCLUSION

Traditional ways of retrieving personal information are becoming increasingly ineffective. The intrinsic limitations of mobile devices make information retrieval tasks even harder, in a mobile setting. This motivates the creation of mechanisms whereby information relevant for a context is efficiently provided to a user.

We have studied how to automatically provide users with all documents and information about persons relevant for a particular meeting, based solely on the information stored in their computers. We were able to find the criteria that users find more relevant when deciding on which what is more relevant in a particular context. Using those insights, we created SmarterPhone, a prototype mobile application that allows users to access the relevant information in a mobile context. User tests showed that, indeed, it is possible to present the users with that information. Even given the existing technical limitations, 80% of all such information was found, validating the criteria that can be used for other applications in this area. Some technical limitations prevented the approach from fulfilling its entire potential.

In the future, we would like to extend the criteria that can be implemented by creating a better index of the users' personal information, resorting to real-time monitoring of their actions. We'll then conduct more thorough tests and evaluate the interface´s usability to discover how more information and scenarios it can be supported.

## ACKNOWLEDGMENTS

## REFERENCES

1. Dumais, S. T. et al. Stuff I´ve Seen: A system for personal information retrieval and re-use. In *Proceedings of SIGIR 2003*.

2. Flynn, M., Pendlebury, D., Jones, C., Eldridge, M., and Lamming, M. The satchel system architecture: mobile access to documents and services. Mob. Netw. Appl. 5, 4 (2000), 243–258.

3. Gemmell, J., Bell, G., Lueder, R., Drucker, S., and Wong, C. Mylifebits: fulfilling the memex vision. In MULTIMEDIA '02: Proceedings of the tenth ACM international conference on Multimedia (New York, NY, USA, 2002), ACM, pp. 235–238.

4. Gonçalves, D., Jorge, J.. In Search of Personal Information: Narrative-Based Interfaces, In *Proceedings ACM IUI 2008*, 13-16 January, Maspalomas, Canary Islands, Spain. ACM Press

5. Lamming, M., Eldridge, M., Flynn, M., Jones, C., and Pendlebury, D. Satchel: providing access to any document, any time, anywhere. ACM Trans. Computer - Human Interact. 7, 3 (2000), 322–352.

6. Raento, M., Oulasvirta, A., Petit, R., and Toivonen, H. Contextphone: A prototyping platform for context-aware mobile applications. IEEE Pervasive Computing 4, 2 (2005), 51–59.