

CONTEXT-AWARE INFORMATION RETRIEVAL BASED ON USER PROFILES

Santtu Toivonen

VTT Technical Research Centre of Finland, P.O.Box 1000, FIN-02044 VTT, Finland

Heikki Helin

TeliaSonera Finland Oyj, P.O.Box 970, FIN-00051 Helsinki, Finland

Keywords: Information retrieval, context-awareness, user-generated content.

Abstract: In mobile Web usage scenarios, taking advantage of user context in information retrieval (IR) and filtering becomes evident. There are multiple ways to approach this, of which we present some alternatives and discuss their performance independently and in combination. The investigation is restricted to Semantic Web like structured content consisting of statements. The discussed approaches address relative importances of statements constituting the content, dependencies between statements, and close matches. Simulated results of the various approaches are provided.

1 INTRODUCTION

Information overload, a phenomenon envisaged as early as in the 1970s (Toffler, 1970), was in essence finally confronted in the 1990s. This was mainly due to the birth and expansive growth of the Web. Since the Web is an open environment where anyone can create content, people accessing it continuously come across material created by someone unknown to them. Of this information people do not typically know a priori whether it is useful for them or not.

Two current and emerging phenomena bring this information overload problem even further. The first of them is called *Web 2.0*, which “refers to perceived or proposed second generation of Internet-based services—such as social networking sites, wikis, communication tools, and folksonomies—that emphasize online collaboration and sharing among users¹.” In *Web 2.0*, it is even easier than before for anyone to create and share content. The second phenomenon is related to the advent of smart phones and *mobile access* to Web content. It poses unique new requirements for information retrieval and filtering. As opposed to desktop PCs, mobile phones are often used while on the move. This, implying the fact that the

user has to simultaneously pay attention to occurrences taking place around her, calls for effective content retrieval capabilities. Otherwise the user would be swamped with useless information, through which she had no time to go. We recognize these phenomena and turn them into a design principle. By doing this we contribute to the rise of *mobile Web 2.0* (Jaokar and Fish, 2005).

The fundamental requirement we are addressing is: Access to Web content should respect that mobile users’ cognitive tasks and attentional resources are discontinuous and short (Miller, 1956; Cowan, 2001), around four seconds by some studies (Oulasvirta et al., 2005). In practice, this principle shows in our work by putting emphasis on information retrieval in general, and investigating the retrieval of relatively small and compact pieces of information in particular. We are harnessing the backbone of *Web 2.0*, namely user-generated content, to assist mobile users in their cognitive tasks. Such user-generated content is subsequently referred to as *Semantic Notes* (Toivonen et al., 2005; Toivonen and Riva, 2006).

Our research focus is in the area of context-aware retrieval (CAR); see, for example (Jones and Brown, 2000; Brown and Jones, 2001; Rhodes, 1997). We acknowledge that the current activity of a user is often the most important context attribute to be recognized in CAR. However, since the focus is on mobile users,

¹Source (retrieval date: 9.1.2007):
http://en.wikipedia.org/wiki/Web_2

our research includes also other context attributes, and therefore extends many CAR research efforts such as (Rhodes, 2000; Lieberman, 1995; Baeza-Yates and Ribeiro-Neto, 1999). We try to *a priori* specify these relevant context attributes of a mobile user, such as her location and social surroundings. In this respect our methodology departs from the majority of CAR, which can more easily apply techniques of case-based reasoning (Schank, 1983) for trying to determine a stationary user's browsing context.

The rest of the paper is organized as follows. In the next section, we summarize the baseline information retrieval approach, which will later be modified in other approaches. Section 3 outlines the three proposed alternative approaches and is followed by Section 4 discussing their evaluation. Finally, Section 5 concludes the paper and outlines some future work.

2 THE DEFAULT INFORMATION RETRIEVAL APPROACH

In this section the concept of Semantic Note is defined and discussed. In addition, the general usefulness determination process of an arbitrary Semantic Note is formalized. This is based on the approach described in (Toivonen et al., 2005; Toivonen and Riva, 2006). Semantic Notes are entities used for knowledge sharing among distributed cognitive processes involving the Web. Although not all Web 2.0 content are tagged with rich semantic descriptions, it is obvious that having such is useful for many tasks (Gruber, 2006; White, 2006; Lawrence and Schraefel, 2006). Therefore, in order to assist further retrieval and usage of Semantic Notes, they are serialized in a structured form with well-defined semantics.

A Semantic Note stores and transmits some meaningful piece of information, such as a definition of a complex concept or instructions for completing a procedure. The domain of information stored in Semantic Notes is unrestricted. As a consequence, Semantic Note is better defined functionally as capturing the state of a cognitive process's subprocess, which is distributed (Hutchins, 1996) to involve the Web.

A Semantic Note (n) can be decomposed into its constituents, namely statements (s). The terms (tm) in a statement can be organized in the subject-predicate-object model of RDF, and conform to concepts in an ontology. This kind of machine-accessibility is especially important for software agents and other decision support systems. Combining the notion of statements and the approach adopted in (Williams and Ren, 2001), an agent can be said to understand a statement found in a Semantic Note, as long as it under-

stands all the terms found in the statement.

The level of understanding a Semantic Note (n) is represented by n_u . Let S_n be the set of statements in n so that $s_1, s_2, \dots, s_k \in n$, where $k = |S_n|$. n_u receives values between 0 and 1 based on the number of understood statements ($s_{u_1}, s_{u_2}, \dots, s_{u_k} \in n$) divided by the number of all statements in the Semantic Note ($|S_n|$) as follows:

$$\begin{aligned} 0 \leq n_u &= \frac{1}{|S_n|} \sum_{i=1}^{|S_n|} s_{u_i} \leq 1 & S_n \neq \emptyset \\ n_u &= 0 & S_n = \emptyset \end{aligned} \quad (1)$$

There are many alternatives for computing the relevance of a Semantic Note, of which our work falls in the category of rule-based approaches. These user-specified rules connect the information content, of which the relevance is to be determined, with the user's context (cf. (Jones and Brown, 2000)). Both the information content—that is, the Semantic Notes—and the user context are realized as sets of statements.

If there exists a term (tm_{ctx}) in a statement found in the user context, as well as a term (tm_n) in a statement found in the Semantic Note so that both of those conform to respective concepts ($\phi_{ctx,n}$) which are navigable from the concepts (ϕ_{r_1,r_2}) found in the rule (r), the rule is said to be applicable (r_a). Navigability means that there exists a network of concepts and relationships, realized as one ontology or several connected ontologies, which enables navigating between the two concepts. A positive match indicates that an applicable rule is found, as well as suitable values to satisfy it. Negative match means that there exists an applicable rule, but that the statements plugged in it do not have suitable values. In order to assign relevance values for the Semantic Notes utilizing the applicable rules, the following function is defined:

$$app(r_a) = r_m = \begin{cases} 1 & \text{positive match} \\ 0 & \text{negative match} \end{cases} \quad (2)$$

The function app is realized as various concrete rules, that determine the relevance assignment (r_m , where m indicates "match"). The applicable rules (r_a) as well as the match value (r_m) are utilized in the relevance equation for Semantic Notes. Let R_a be the set of applicable rules so that $r_{a_1}, r_{a_2}, \dots, r_{a_k}$, where $k = |R_a|$. The Semantic Note relevance (n_{rel}) can receive values between 0 and 1 as the ratio between the sum of the match values ($r_{m_1}, r_{m_2}, \dots, r_{m_k}$) and the number of applicable rules ($|R_a|$):

$$\begin{aligned} 0 \leq n_{rel} &= \frac{1}{|R_a|} \sum_{i=1}^{|R_a|} r_{m_i} \leq 1 & R_a \neq \emptyset \\ n_{rel} &= 0 & R_a = \emptyset \end{aligned} \quad (3)$$

Finally, the usefulness of a Semantic Note for an agent is defined as consisting of both understanding the Semantic Note and considering it relevant. The information usefulness variable (n_{use}) also receives values between 0 and 1, and is formalized as follows:

$$n_{use} = a * n_u + b * n_{rel} \quad (4)$$

where $0 \leq a + b \leq 1$ and $a, b \in \mathbf{R}^+$. Parameters a and b indicate the application-specific weights that are assigned to the understanding (n_u) and relevance (n_{rel}), respectively.

3 ALTERNATIVE APPROACHES FOR INFORMATION RETRIEVAL

This section investigates information retrieval approaches other than the baseline approach presented above. Note that the focus is on information relevance only, that is, information understanding of the baseline approach is left untouched. The motivation for considering other approaches arises in cases where the baseline approach selects too many or too few results, and if those results turn out not to be relevant enough. There are three major subjects in this section, which are:

1. Acknowledging that some statement kinds are more *important* than others.
2. Acknowledging *dependencies* between statements.
3. Acknowledging *close matches* in addition to exact ones.

Next, these three subjects will be separately explained in more detail. In Section 4, their capabilities of finding relevant results among Semantic Notes will be evaluated and contrasted with the baseline approach. In addition to contrasting them separately, a combination of all three approaches will be included.

3.1 Assigning Important Weights for Statement Kinds

In principle the importance weights to various statements can be assigned either based on their semantics, or their position in the Semantic Note. As an example of importance arising from semantics, consider a sailboat traveling in a remote location with very few points of interest. The captain needs to fill up the refrigerator of his boat. In this case location is an important statement kind, since it is useful for the captain to have information about basically all grocery stores relatively close to him. And suppose another case,

where the captain with the same need is in an area equipped with many services, say a busy guest harbor of a big city, but happens to arrive there very late in the evening. Here (opening) time(s) is a much more important statement kind. The user can teach the system about the importances of various statement kinds via a *relevance feedback* loop (Chakrabarti, 2002).

There are alternative approaches to the semantics-based importance assignment, where the user and/or the system has to have a priori insight on which statement kinds are more important than others. The particular approach considered below in more detail is based on the structural position of the statement in the content currently under inspection. The simple rationale behind this approach is that the further deep the statement in the content, the smaller its importance with regard to the big picture. In achieving this, we introduce a new variable d for indicating this depth:

$$0 \leq n_{rel} = \frac{1}{d * |R_a|} \sum_{i=1}^{|R_a|} r_{m_i} \leq 1 \quad R_a \neq \emptyset \quad (5)$$

$$n_{rel} = 0 \quad R_a = \emptyset$$

where $d \in \mathbf{Z}^+$. Parameter d indicates the depth of the statement from the first level. $d = 1$ indicates the first level and $d = 2$ the second, that is:

```
<Note rdf:ID="http://foobar.org/barfoo">
  <FirstLevel>
    This is the first level with d=1
  <SecondLevel>
    This is the second level with d=2
  </SecondLevel>
</FirstLevel>
</Note>
```

Applying the d parameter to the information relevance calculation probably makes little difference if the content structure is very flat. However, in some cases there can be several layers of embedded content. In these situations using d is envisaged to turn out useful, given that the statements closer to the surface can concern larger themes than the ones deeper within the structure, and therefore be considered more important. Besides depth, another metric would be to consider the amount of information "contained" by the statement under inspection. In this model, two statements on the same depth would receive different relevance weights if they would have differing number of sub-statements.

3.2 Recognizing Dependencies Between Statements

A statement in a Semantic Note can be dependent on some other statement. For example, consider a boater

docked in a guest harbor, with the intention of going to a restaurant for a dinner. She has two atomic rules in her profile. The first of them states that she is interested in content created by people she knows. The other one says that she is interested in restaurants rated four stars or higher. Moreover, she has a metarule in her profile stating that while engaged in such activity, she is interested in restaurants which are ranked four stars or more, if the review/rating is created by a friend of hers.

Various logical connectives can be introduced for expressing different dependency kinds between statements and used in matching (Ranganathan and Campbell, 2003). In the example above, there is an implication relationship between the statements. In other words, the relevance of the rating statement depends on the statement expressing the creator, but not vice versa. There could also be an equivalence, which would entail a mutual dependency.

So, to continue with the example: Suppose there are two restaurants in the neighbourhood with 4-star ratings each. One of the rating annotations is created by someone unknown to the boater, while the other one by someone she knows. Say that there are only these two statements in each annotation (one about the stars, and the other expressing the creator). The relevances for these would be 0 and 1, respectively. In the absence of the implication rule the relevances would be 0.5 and 1.

3.3 Coping with Close Matches in User Profiles

We now introduce two alternative versions of the $app(r_a)$ function. The motivation is to cope with close and partial matches of the information found in user context and in the content to be provided. The first of the functions is applicable in cases when there exists a taxonomy, for example a yellow page like service categorization. The function is based on an algorithm called ‘‘Object Match’’ (OM) (Stojanovic et al., 2001), and is formalized as follows:

$$app_{om}(r_a) = r_m = \frac{|uc(\phi_{ctx}, \phi_{root}) \cap uc(\phi_n, \phi_{root})|}{|uc(\phi_{ctx}, \phi_{root}) \cup uc(\phi_n, \phi_{root})|} \quad (6)$$

where uc refers to the ‘‘upwards cotopy’’ function (Stojanovic et al., 2001). uc returns the distance of the currently analyzed concept from the ontology’s root concept. In the case of services in yellow pages such concept would be SERVICE. More specifically, the $uc(\phi_{ctx}, \phi_{root})$ indicates the distance between the concept found in the user context and the root concept, whereas $uc(\phi_n, \phi_{root})$ means the distance

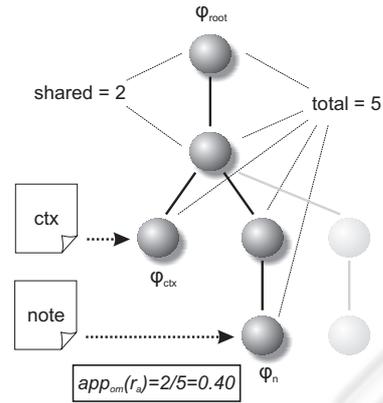


Figure 1: An example hierarchy of concepts.

between the concept found in the Semantic Note and the root concept. Figure 1 depicts an example case of this algorithm.

If the terms in the user context and the Semantic Note have numerical values, the difference between these values can in some cases be used as a measure of relevance. The latter function we present fits these cases. For example, in the case of locations, suppose that the user is interested in content referring to entities close to her, as is often the case. Now, the further away the content in question is from her, the less relevant it can be said to be. The following formula captures this simply by dividing one with the distance of the two numerical values (one in the context and one in the Semantic Note); should their distance be 0, the value would return an exact match as 1:

$$app_{dfr}(r_a) = r_m = \frac{1}{1 + dfr(tm_{ctx}, tm_n)} \quad (7)$$

where the function dfr denotes calculating the difference between the values of tm_{ctx} and tm_n .

4 EVALUATION

The above-mentioned approaches were evaluated and their performances compared with each other, as well as with the baseline approach. Note that the intention is not to find out the single best approach, but instead to show that approaches other than the baseline one can in some cases turn out to be useful alternatives, especially if combined.

4.1 Evaluation Setup

Due to the lack of a common test set, the evaluations were conducted as simulations so that 500 Semantic

Notes, each containing 4-10 statements (based on random assignment), were created. Each statement was randomly assigned to be relevant or not relevant. In addition, each statement was assigned a depth, which is recognized by the “relative” approach. The restriction was that a statement’s depth could not exceed the number of statements in the Semantic Note. Note that assigning a depth to a statement has impact on the depths of the rest of the statements in the respective Semantic Note. The range of the first statement’s d is $1 \leq d_{s_1} \leq (|S_n| + 1)$, the next one’s $1 \leq d_{s_2} \leq (|S_n| + 1 - d_{s_1})$, and so on. Furthermore, each statement was randomly assigned as being dependant on some other statement or not. This has impact on the “dependant” approach.

When generating the relevance values at the Semantic Note level, the basic (baseline) relevance was first calculated by summing up the relevances of the statements, and dividing them by the number of all statements in the respective Note. The relative relevance was calculated in the same manner, but instead of summing up the plain statement relevances, their relative values were instead used. In the case of dependency relevance, it was first checked whether the statement in question is relevant or not. If yes, it was checked whether the statement it was labeled dependant on was relevant or not. If that was true as well, the statement was labeled as “dependency-relevant”. Finally, the “OM relevance” was simulated as follows: If the statement currently under inspection was not relevant, it did not automatically receive a 0 relevance value, but some randomly assigned floating point between 0 and 1. This represents the inclusion of close matches to the calculation.

To complement the above-mentioned relevance kinds, the system randomly—and regardless of the above relevance values based on statements’ relevances—labeled each Semantic Note as “really” relevant or not. This was representing the user’s actual consideration of the Semantic Note, whereas the above-mentioned relevance values represent the decision support capabilities of our system. The difference between the “real” relevance and the statement-based relevance kinds was tested as-is (with 0 correspondence), with 0.5 correspondence, and with 0.9 correspondence. This consideration was justified since it is envisaged that the rules stored in the user profiles have *some* correlation with the actual relevances. That is, if the user creates a rule stating that she is interested in ice cream, it is indeed justified to assume that (all other things equal) she will be more interested in an ice cream parlor than a hot dog stand.

For generating the test set, we modified three things: First, the likelihood of correspondence (Lhc)

was set to be 0, 0.5, or 0.9. Secondly, either half or one quarter of the statements were labeled as “really” relevant (likelihood of relevance, Lhr). Third, the “real” relevances were reassigned based on baseline relevance values, based on combined relevance values, or not at all reassigned. The same value which was used as a threshold for retrieving content throughout the tests, namely 0.5, was also used as a threshold for reassignment. By combining these options, we came up with 18 different test cases, with each of them having 500 generated Semantic Notes.

4.2 Evaluation Results

We now present some results of the simulations. In the following Tables, the approaches are referred to as “Baseline”, “Relative”, “OM”, “Dependant”, and “Combined”. The “Combined” approach is the average of “Relative”, “OM”, and “Dependant” approaches. Naturally, we could have considered other combinations, too. However, contrasting the alternative approaches with “Baseline” separately and as one combination is enough for giving us guidelines on their performance.

Basic instruments of information retrieval, namely precision, recall, and the F-measure, were used in the evaluation. As for relevant documents, we used the “real” relevance, that is, the relevance which was not derived from the number of statements considered relevant. This way we could compare the decision support of the system with the (simulated) true relevance considered by the user. In doing this, precision came to indicate the number of documents which are both retrieved and (“really”) relevant divided by the number of retrieved documents. Recall indicates the number of documents which are both retrieved and (“really”) relevant divided by the number of (“really”) relevant documents. Finally, the F-measure is the harmonic mean of precision and recall, with the formula of $F = 2 * precision * recall / (precision + recall)$.

Table 1 depicts the precision values in the case where none of the “real” relevance values are tampered. There exists no significant variation among the approaches; the average standard deviation (SD) between the approaches in different cases is 0.05. If Table 1 is contrasted with Tables 2 and 3, it is visible that more variation among the approaches emerges. Corresponding SD average for Table 2 is 0.12 and for Table 3 it is 0.15. Naturally, for the first two rows, where the likelihood of correspondence (Lhc) is 0, this does not hold. But once the likelihood grows to 0.5 and especially 0.9, differences start to show. This is especially true in the case where the rearrangement of the “real” relevance values is done based on the combined

Table 1: Comparing the precision values of the various approaches when no “real” relevance values are tampered.

Not rearranged	Baseline	Relative	OM	Dependant	Combined
Lhc=0, Lhr=0.5	0.50	0.32	0.50	0.41	0.53
Lhc=0, Lhr=0.25	0.32	0.27	0.29	0.29	0.35
Lhc=0.5, Lhr=0.5	0.36	0.38	0.42	0.36	0.48
Lhc=0.5, Lhr=0.25	0.20	0.23	0.21	0.21	0.20
Lhc=0.9, Lhr=0.5	0.59	0.79	0.48	0.58	0.69
Lhc=0.9, Lhr=0.25	0.25	0.24	0.23	0.21	0.22

Table 2: Comparing the precision values of the various approaches when “real” relevance values are set to correspond to the basic (baseline) relevance values.

Basic rearranged	Baseline	Relative	OM	Dependant	Combined
Lhc=0,Lhr=0.5	0.47	0.67	0.37	0.55	0.50
Lhc=0,Lhr=0.25	0.19	0.18	0.22	0.20	0.17
Lhc=0.5,Lhr=0.5	0.83	0.90	0.56	0.82	0.92
Lhc=0.5,Lhr=0.25	0.58	0.56	0.34	0.57	0.56
Lhc=0.9,Lhr=0.5	0.97	1.00	0.54	0.98	0.94
Lhc=0.9,Lhr=0.25	0.89	0.91	0.49	0.89	0.87

relevance values.

The combined approach starts to outperform the other approaches, when the “really” relevant values conform with the combined relevance. This is also visible from the last two rows of Table 3. Note that in the case where the comparison is based on the basic relevance, the combined approach performs almost as well as the baseline approach (last two rows of Table 2). The relative approach performs also well, but lags behind in with regard to recall, as is visible from Table 4.

The reason for OM approach having significantly lower precision values than the other approaches is due to the fact that it retrieves so many Semantic Notes, causing many labeled as “really” irrelevant to be included. Naturally, this has an inverse impact on the recall values, since the more retrieved documents, the more chance for “really-relevants” to get included. Table 4 depicts the numbers for various approaches, as far as recall is concerned. The OM approach has the highest average and median values, meaning that it has absorbed more “really” relevant documents than the other approaches. This could be easily prevented by restricting the search space. In other words, the whole taxonomy would not be examined each time, but instead a set of concepts with a prespecified maximum distance from the concept under inspection.

In order to examine the mutual effect of precision and recall, we used the F-measure. Precision is a more important factor than recall in context-aware informa-

tion retrieval systems. This is because not too many results can be simultaneously provided to the user, and the few results that end up being provided should indeed be relevant. If there were 50 possible Semantic Notes which are at the time are passing the threshold of relevance as reasoned by the system, it is not likely that the user will go through all of them, but instead only a small portion. For that reason, it is important that the precision of “real” relevance among these 50 is as high as possible; recall is less important. This is why in addition to using the harmonic F-measure, we also tested the results with a $F_{0.5}$ -measure, which weights precision twice as much as recall.

We grouped the evaluation sets into the following three segments: (i) The first segment consisted of all the 18 cases as presented the above; (ii) The second segment consisted of the cases where no relevance values were rearranged, as well as all the cases where the likelihood of correspondence (Lhc) was 0; (iii) The third segment represents the cases left out from the second segment, namely the cases where the relevances were rearranged based on either the basic relevance or the combined relevance. In general, the approaches perform a little better in terms of the $F_{0.5}$ -measure than the F-measure, as Table 5 depicts.

The most evident message that was found emerged by comparing the SD values of the F-measure averages, especially between the second and the third segment. The second segment, where the relevance values had not been tampered, showed a 0.04

Table 3: Comparing the precision values of the various approaches when “real” relevance values are set to correspond to the combined relevance values.

Combined rearranged	Baseline	Relative	OM	Dependant	Combined
Lhc=0,Lhr=0.5	0.45	0.60	0.49	0.45	0.56
Lhc=0,Lhr=0.25	0.27	0.29	0.27	0.24	0.28
Lhc=0.5,Lhr=0.5	0.45	0.59	0.32	0.47	0.61
Lhc=0.5,Lhr=0.25	0.59	0.50	0.21	0.55	0.60
Lhc=0.9,Lhr=0.5	0.49	0.73	0.28	0.58	0.91
Lhc=0.9,Lhr=0.25	0.51	0.90	0.28	0.67	0.97

Table 4: Recall trends for different approaches.

Recall trends	Baseline	Relative	OM	Dependant	Combined
Avg	0.31	0.12	0.38	0.19	0.19
Median	0.16	0.07	0.27	0.09	0.10
SD	0.26	0.12	0.24	0.19	0.19

average on the SD values of the approaches. In the third segment the corresponding value was 0.18, indicating a significantly greater variance. (The corresponding numbers for $F_{0.5}$ -measure were 0.06 and 0.18.) This means that if the “real” relevance values correspond to the relevance values as reasoned by the system, there is greater variance between the approaches and choosing an appropriate one is more important. This is a valuable finding and justifies further work on this subject, since it can be assumed, that there is indeed correspondence between what a rational user states as interests in her profile, and what she really considers interesting.

Finally, we examined more closely the performance of different approaches in the cases grouped to the third segment presented above. In particular, we considered the differences between how the approaches perform with regard to the proportion of “really” relevant Semantic Notes. It is noteworthy that most of the approaches perform better when the proportion of “really” relevant Semantic Notes is smaller (0.25). The only exception for the harmonic F-measure is the OM approach, where in 3 out of 4 cases it performs better when the proportion of “really-relevants” is larger (0.5). This is due to OM approach’s relatively good recall values in these cases. We also noted that even though the precision values of the approaches are smaller in the 0.25 case when rearranging based on basic relevance (see Table 2), their F-measure due to better recall is larger. Moreover, this phenomenon gets amplified when re-

arranging is based on the combined relevance (with the exception of OM, which was explained above).

5 CONCLUSIONS

This paper presented some approaches for context-aware information retrieval. The approaches departed from the so-called baseline approach, which has been presented in our previous work in more detail. The approaches put emphasis in importance weights of statements, interdependencies between statements, and close matches in finding appropriate content. Simulated evaluation results for the performances of these approaches were also presented.

The particular approaches presented in this paper are merely a start for considering intelligent retrieval of semantically described content for mobile Web 2.0. In the future we are going to examine new atomic approaches and consider their performance in various cases. In addition, our future work among the area will concentrate on more intelligent ways of combining various approaches. This paper introduced a rather straightforward way of averaging over the selected approaches, but more advanced ways could be introduced. For example, a correlation between the “relative approach” and the “OM-approach” can be envisaged. A statement’s relative relevance considered in this paper arises from its position in the Semantic Note. It can be assumed that it is somehow also semantically related to the statements close to it

Table 5: Comparing the averages with regard to F-measures.

Averages	Segment i	Segment ii	Segment iii
F-measure	0,28	0,15	0,45
F _{0,5} -measure	0,31	0,18	0,48

in the structure. Now the terms in these neighboring statements have corresponding concepts in the ontology. OM-relevance is based on close matches, and the concepts corresponding to the terms in these neighboring statements could be considered.

ACKNOWLEDGEMENTS

We would like to thank Mikko Laukkanen and Matthieu Molinier for their useful comments.

REFERENCES

- Baeza-Yates, R. and Ribeiro-Neto, B. (1999). *Modern Information Retrieval*. ACM Press / Addison-Wesley, New York, NY.
- Brown, P. and Jones, G. (2001). Context-aware retrieval: Exploring a new environment for information retrieval and information filtering. *Personal and Ubiquitous Computing*, 5(4):253–263.
- Chakrabarti, S. (2002). *Mining the Web: Discovering Knowledge from Hypertext Data*. Morgan-Kaufmann.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, (24):87–114.
- Gruber, T. (2006). Where the social web meets the semantic web. In Cruz, I. et al., editors, *Proceedings of the Fifth International Semantic Web Conference (ISWC 2006)*, page 994, Athens, GA. Springer.
- Hutchins, E. (1996). *Cognition in the Wild*. MIT Press, Cambridge, MA.
- Jaokar, A. and Fish, T. (2005). *Mobile Web 2.0*. Futuretext, London, UK.
- Jones, G. and Brown, P. (2000). Information access for context-aware appliances. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2000)*, pages 382–384, Athens, Greece. ACM Press.
- Lawrence, K. and Schraefel, M. (2006). Bringing communities to the semantic web and the semantic web to communities. In Carr, L. et al., editors, *Proceedings of the 15th international World Wide Web conference (WWW 2006)*, pages 153–162, Edinburgh, Scotland. ACM.
- Lieberman, H. (1995). Letizia: An agent that assists web browsing. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI 1995)*, volume 1, pages 924–929, Montreal, Canada. Morgan Kaufmann.
- Miller, G. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, (63):81–97.
- Oulasvirta, A. et al. (2005). Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile HCI. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI 2005)*, pages 919–928, New York, NY. ACM Press.
- Ranganathan, A. and Campbell, R. (2003). An infrastructure for context-awareness based on first order logic. *Personal and Ubiquitous Computing*, 7(6):353–364.
- Rhodes, B. (1997). The wearable remembrance agent: A system for augmented memory. In *Proceedings of The First International Symposium on Wearable Computers (ISWC '97)*, pages 123–128, Cambridge, MA. IEEE.
- Rhodes, B. (2000). Margin notes: building a contextually aware associative memory. In *Proceedings of the 5th international conference on Intelligent user interfaces (IUI 2000)*, pages 219–224, New York, NY. ACM Press.
- Schank, R. (1983). *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge University Press, New York, NY.
- Stojanovic, N. et al. (2001). Seal: A framework for developing semantic portals. In *Proceedings of the international conference on Knowledge capture (K-CAP 2001)*, pages 155–162, New York, NY. ACM Press.
- Toffler, A. (1970). *Future Shock*. Bantam Books, New York, NY.
- Toivonen, S., Pitkäranta, T., and Riva, O. (2005). Determining information usefulness in the semantic web: A distributed cognition approach. In *Proceedings of the Knowledge Markup and Semantic Annotation Workshop (SemAnnot 2005), held in conjunction with the Fourth International Semantic Web Conference (ISWC 2005)*, pages 113–117, Galway, Ireland.
- Toivonen, S. and Riva, O. (2006). The necessary but not sufficient role of ontologies in applications requiring personalized content filtering. In *Proceedings of the 4th Workshop on Intelligent Techniques for Web Personalization (ITWP 2006), held in conjunction with The 21st National Conference on Artificial Intelligence (AAAI 2006)*, pages 19–30, Boston, MA. AAAI.
- White, B. (2006). The implications of web 2.0 on web information systems. In Cordeiro, J. et al., editors, *Proceedings of the Second International Conference on Web Information Systems and Technologies: Internet Technology / Web Interface and Applications (WEBIST 2006)*, Setúbal, Portugal. INSTICC Press.
- Williams, A. and Ren, Z. (2001). Agents teaching agents to share meaning. In *Proceedings of the fifth international conference on Autonomous agents (AGENTS 2001)*, pages 465–472, Montreal, Quebec, Canada. ACM Press.