

Comparing monolithic and committee-based profiles for politician recommendation

Luis M. de Campos
Computer Science and A.I
Department – ETSIIT
Universidad de Granada,
18014, Granada, Spain
lci@decsai.ugr.es

Juan M. Fernández-Luna
Computer Science and A.I
Department – ETSIIT
Universidad de Granada,
18014, Granada, Spain
jmfluna@decsai.ugr.es

Juan F. Huete
Computer Science and A.I
Department – ETSIIT
Universidad de Granada,
18014, Granada, Spain
jhg@decsai.ugr.es

ABSTRACT

In a parliamentary setting, citizen could be interested in knowing those Members of Parliament (MPs) who are working in different areas or involved in the resolution of some people's problems. These topics of interest are usually represented by means of a profile. In this paper, the politicians' profiles are built considering the speeches in parliamentary sessions. However, in most of the cases a single profile is not the best alternative to represent MPs' interests because the specific terms related to a given topic are mixed with others, so that the MPs' preferences are diluted. The alternative is to build different subprofiles containing each one the most representative keywords for each topic, creating in this way a richer representation. We present a first approach to build subprofiles based on the MPs' speeches in different committee and plenary sessions, which will be compared, in terms of performance, to monolithic profiles for an MP content-based recommendation task.

CCS Concepts

•Information systems → Document filtering; Recommender systems; Expert search;

Keywords

Profiles; Subprofiles; Content-based Recommendation; Information Retrieval; Members of Parliaments;

1. INTRODUCTION

Among all the content published in the WWW, we may find more and more frequently information about people's expertise, by means of CVs, scientific papers, blog entries and so on. In many cases we need to find experts in a topic and we search for them using our favorite search engine. Other times, we take a more passive role and we expect that a system recommend them to us. Broadly speaking, this is the so called expert finding problem [8]. Anyhow, and in

order to recommend or find these people, it is required that the information associated to them, i.e. the underlying topics they are interested in, is organized in a kind of container, called profile, so any content-based recommender could easily use it to get a ranking of experts and suggest them. In its more frequent organization, a profile is not more than a set of terms, the most representative ones.

But people may have multiple interests. For example, a Member of a Parliament (MP) could be involved in several subjects: agriculture, education or economy, for example. If a single profile is used to represent them, the terms composing it will come from different domains of interest. This can cause that terms that are important for only a few of the interest categories or for non-dominant categories become underweighted or even unnoticed. Moreover, a single profile composed of terms from different topics will result in a profile pointing halfway between these topics, probably resulting in a less accurate profile. Examples of systems that utilize a single profile for each user are Letizia [11], Amalthea [13], Fab [3], Anatagonomy [10], and more recently [5] and [4].

If a profile is organized in subprofiles, each one representing a topic of interest for the person, the terms that comprise them are closely related among them under the umbrella of a common subject. This is a much more accurate representation and therefore more useful for recommendation purposes. Some examples where subprofiles are used are the following: In [12], a personalized web search assistant that extracts the keyword-based profiles from bookmarked web pages is proposed. Instead of creating a single profile for the user, it uses a set of keyword-based profiles, one per bookmark. Webmate [6] also builds for each user a set of keyword-based profiles, one per user's area of interest, but it automatically learns the interest areas. It uses a fixed number N of different interest areas, and initially each (relevant) document constitutes its own interest area. When there are more than N documents, the two most similar (according to the cosine similarity) are combined into a single profile. The system in [15] utilizes a more sophisticated clustering method to combine the profiles. Alipes [17] creates several profiles for a user and also learns automatically the user's interests. But instead of using a fixed number of interests, it bases the creation of new interests on a similarity threshold. Syskill & Webert [14] learns a separate profile for each topic of each user. The profiles are learned from sets of positive and negative examples of each topic using machine learning algorithms. Another system using multiple profiles is Pro-

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CERI '16, June 14-16, 2016, Granada, Spain

© 2016 ACM. ISBN 978-1-4503-4141-7/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2934732.2934740>

File [2], where each user may define a number of conceptual classes to classify the filtered documents, each class having its own profile. The profiles considered in [1] for each user of a news recommendation system are also divided by topics of interest.

This paper deals with building profiles in a parliamentary context and how to use them in a content-based recommendation task, where the Members of the Parliament are the “objects” to be recommended. MPs’ participation in debates are transcribed in Records of Parliamentary Proceedings. These documents are organized around the concept of an initiative, which are the matters to be discussed in a session. At the same time, these are composed of the MPs’ speeches. The initiative may be discussed in committee or plenary sessions. The former are special sessions where a reduced number of MPs deals with specific topics as agriculture, industry, economy, tourism, among many others. The latter are meetings of all MPs and the topics are more general and politic-oriented, i.e. less technical. It is common that several MPs participate in different committees, so we could say that these politicians are interested in the topics of the committee. Therefore, this would be a direct way of creating MPs’ subprofiles: For a given MP, from her speeches in each committee, including the plenary sessions as a special committee, we could build a subprofile selecting the most relevant terms. For example, if a representative has participated in education, culture, environment and plenary sessions, four subprofiles could be constructed.

These subprofiles could be represented as documents in a collection (each MP could have as many documents as committees she has participated in). An Information Retrieval System (IRS) will index them so given a query formulated by a user, the search engine would compute a ranking of MPs. Therefore, if a user would like to know the politicians who are dealing with the EU Common Agricultural Policy, then she could get a list of MPs sorted by relevance.

The objective of this paper is to determine, in an early stage, if, in the context of MP recommendation, a structured profile based on the committee sessions where the MP has participated, could be worthy in opposition to a single monolithic profile, increasing the performance. This subprofiling approach is the most direct one that may be tried, because the partitions of terms are already given considering the different topics treated in each committee. Our guess is that these so built subprofiles could improve the performance of a single profile, representing much better the topics the MPs are interested in.

The contribution of this paper is a simple method, which takes the most of the distribution of MPs’ interventions across committee sessions, to build subprofiles and use them for recommending MPs. The results obtained in the experiments that we have carried out support the idea of using structured profiles and give us cause for developing more sophisticated and accurated subprofiling methods.

In order to describe these ideas about committee-based subprofiling, this paper is organized as follows: Section 2 shows the way the monolithic profiles and subprofiles are built, presenting the methods to select the terms to be contained in the profiles and how to represent them. Specifically for subprofiles, filtering schemes are introduced in order to get rankings of MPs without repetition. The next section deals with the experimentation. More specifically, Section 3.1 describes how we have designed the experiments to test

our hypothesis and Section 3.2 presents the results and discusses them. This paper concludes with the main outcomes and further works.

2. BUILDING AND USING (SUB)PROFILES FROM MPs’ SPEECHES

In this section we shall describe the procedure of building the MP’s profiles, both monolithic and compound, as introduced in Section 1. First, the source that we have used for this purpose is the speeches of the MPs who have participated in interventions in plenary and committee sessions.

Regarding the profile representation, each MP will be initially represented by a set of documents, where a document is the compilation of her speeches in each single initiative. From this set a profile P_i of the i -th MP in a preliminary stage is considered a bag of n words, where the most important terms, once stop words were removed and terms stemmed (with the Lucene’s SpanishAnalyzer library), are included in it. We also associate weights to terms, which reflect their importance: $P_i = \{(t_{i1}, w_{i1}), \dots, (t_{in}, w_{in})\}$.

Therefore, the building process itself is just a selection of the best n terms according to a weighting scheme. In this research, we have used two measures that take into account the rarity of the term. The first one is the classic and well-known *TfIdf*, while the second one is called Difference (*Diff*), proposed in [16] for search personalization. We have also used other measures in our experiments, but these are the two that offer better performance. Although its details can be found in that paper, broadly speaking, *Diff* measures the extent to which a term is more common inside the i -th MP’s speeches than outside them (the rest of MPs’ speeches). If the first case happens, then that term would be included in the profile.

In order to finally produce a general representation of the profile that could be used by any search engine, instead of including pairs (term, weight) in it, we have opted to replicate the terms according to their original weights. The first approach is called *R – Prop*, based on a linear transformation applied to the weights. Then, that term with the highest value will be replicated n times, while that one with lowest, just once. Consequently, the final profiles are composed of terms with different number of occurrences, allowing in this way that the indexing module of any IRS could deal with them as if they were common documents, making the corresponding arrangements according to the retrieval model used. The second alternative that we have applied to replicate the selected terms is based on the term frequency (*Tf*), noted as *R – Tf*, where a term occurring Tf times is replicated this same number, keeping its original distribution of occurrences. In this case, the previous computed weights of the terms are only useful for selection purposes.

This would be the general process to build monolithic profiles, where the whole collection of interventions of a given MP is considered to build her profile. But the main objective of this paper is to determine whether a profile composed of several subprofiles is better than the full profile, from a performance point of view, i.e. if given a query submitted by a user who is interested in knowing those MPs who work in the query topic, the system recommends (better) those relevant politicians. Taking into account that the interventions of the MPs are carried out in plenary sessions and those committees where she belongs to, a first, direct and intuitive

approach to build subprofiles would be to learn as many of them as different committees they belong to. Each one could specifically reflect the topics these MPs are interested in, maintaining them completely separated and avoiding the mixing of topics that we find when monolithic profiles are used. Let us suppose that a given MP participates in agriculture, education and health committees, as well as in plenary sessions. Then we could build four different subprofiles, one for each type of session. From the subprofile building process point of view, this would be the one already explained in the previous paragraphs, except that instead of considering all the speeches, they would be grouped by committee, so each subprofile would only be built from those speeches belonging to a given type of session.

Once the (sub)profiles are built, they will be indexed by an IRS. Then, when a user formulates a query to the system, this will generate a ranking of (sub)profiles. In the case of monolithic profiles, as each profile is univocally associated to an MP, we directly obtain a ranking of MPs. However, in the case of subprofiles, the original ranking could contain subprofiles that belong to the same MP although in different committees. Therefore, and in order to get a final ranking of MPs we have to filter the possible occurrences of the same MP in different subprofiles. To carry out this task, we propose two methods:

- *MAX* → Given an MP, her subprofile with the highest RSV is kept in the ranking, removing the rest. This first filtering method tries to capture the best match and it is intended for very specific queries.
- *SUM* → Given an MP, a new score is computed summing the RSVs of her subprofiles in the ranking, removing all the occurrences of that MP in the ranking and placing a new one with the accumulated score. This second approach could be useful when the query is transversal to several subprofiles, so the more number with higher scores, the better.

3. EXPERIMENTATION AND RESULTS

This section will present both the experimental design and the results we have obtained in it, analyzing them and presenting the main findings.

3.1 Experimental settings

The experiments are based on a collection of Records of Parliamentary Proceedings from the Andalusian Parliament in Spain. More specifically, it contains 5258 initiatives from its 8th term of office, with a total of 12633 interventions. We have selected those MPs who have participated in, at least, 10 initiatives. In this term of office, there were 22 different committees.

With respect to the building process of the (sub)profiles, their size is an important parameter to be considered in order to determine if it affects the performance. Then, the numbers of terms composing the profiles that we have considered are 50, 250, 500, 750 and 1000, values ranging from small to big profiles. Regarding the weighting schemes used to select the most important terms, we have applied the two presented in Section 2, i.e. *TfIdf* and *Diff*. Finally, the replication methods *R - Tf* and *R - Prop* are brought into play with the aim of testing the best alternative.

In terms of profile retrieval, Lucene¹ has been the library that we have used for our experiments, and more specifically its implementation of the BM25 model².

The first input to the retrieval system are the 40 different profile collections, corresponding to the combination of the different parameter values (20) when building the profiles times the two types of profiles we are testing, monolithic and subprofiles. The second input are the queries which, in our case, are the initiative titles, representing short queries that a citizen could employ to find relevant MPs. The output is a ranking of MP profiles according to their relevance degree. In the case of subprofiles, we have also experimented with the two proposed methods (*MAX* and *SUM*) of eliminating repetitions of MPs from the ranking.

Finally, and regarding the details of the performance evaluation, this is based on the random partition of 80% of the initiatives for training (the (sub)profiles are learnt with this set) and 20% for test, repeating this process five times and averaging the results for each partition. As ground truth, we assume that the MPs which are relevant for the title of a test initiative are those MPs who participate in it. The evaluation measure selected is Normalized Discounted Cumulative Gain (NDCG@10) [9], which gives information about the ranking quality.

Although we focus on the performance evaluation of these ways of building and using profiles, we have to mention that there are no significant differences in terms of efficiency between building monolithic profiles or subprofiles and later in their use.

3.2 Result analysis

In this section, the results obtained with the experimentation presented in the previous section are included in Table 1. In it, the data are arranged in seven columns as follows:

- The three first correspond to parameters related to the profile learning stage, i.e. the replication method (Rep.), the term selection method (Sel.) and the size (Size), respectively.
- The fourth column (Profile) contains the NDCDG@10 values obtained by the monolithic profiles.
- The fifth column (Filter) stores the two types of filtering method used to get a ranking with no repeated MPs.
- The next one (SubP.) shows the NDCG@10 values for recommending with subprofiles, and
- Finally, the last column is the percentage of change (%C), from the NDCG@10 values in subprofiles with respect to the monolithic ones.

As it can be noticed in Table 1, the NDCG@10 values from the monolithic profiles are repeated for *MAX* and *SUM* filters from subprofiles in order to ease the comparison. We also have boldfaced the best results for profiles and subprofiles.

¹<https://lucene.apache.org/>

²In addition to BM25, several experiments have also been executed with the Lucene implementations of Vector Space Model and Language Model, but the performance of this first model is superior to the rest, so for clarity and space reasons, only its results are presented in this paper.

Table 1: Parameter configuration, NDCG@10 values for profiles and subprofiles and the percentage of change

Rep.	Sel.	Size	Profile	Filter	SubP.	%C
R-Prop	Diff	50	0.3415	MAX	0.2798	-18.07
R-Prop	Diff	250	0.3530	MAX	0.3138	-11.11
R-Prop	Diff	500	0.3471	MAX	0.3095	-10.85
R-Prop	Diff	750	0.3452	MAX	0.3114	-9.79
R-Prop	Diff	1000	0.3458	MAX	0.3060	-11.53
R-Prop	TfIdf	50	0.3654	MAX	0.3139	-14.10
R-Prop	TfIdf	250	0.3643	MAX	0.3293	-9.59
R-Prop	TfIdf	500	0.3477	MAX	0.3074	-11.61
R-Prop	TfIdf	750	0.3340	MAX	0.3021	-9.55
R-Prop	TfIdf	1000	0.3227	MAX	0.3041	-5.77
R-TF	Diff	50	0.3325	MAX	0.2819	-15.23
R-TF	Diff	250	0.3379	MAX	0.3357	-0.66
R-TF	Diff	500	0.3443	MAX	0.3394	-1.42
R-TF	Diff	750	0.3469	MAX	0.3461	-0.23
R-TF	Diff	1000	0.3496	MAX	0.3512	<u>0.47</u>
R-TF	TfIdf	50	0.3372	MAX	0.3007	-10.81
R-TF	TfIdf	250	0.3483	MAX	0.3511	<u>0.80</u>
R-TF	TfIdf	500	0.3432	MAX	0.3397	-1.02
R-TF	TfIdf	750	0.3385	MAX	0.3469	<u>2.47</u>
R-TF	TfIdf	1000	0.3314	MAX	0.3532	<u>6.59</u>
R-Prop	Diff	50	0.3415	SUM	0.3131	-8.31
R-Prop	Diff	250	0.3530	SUM	0.3432	-2.76
R-Prop	Diff	500	0.3471	SUM	0.3369	-2.95
R-Prop	Diff	750	0.3452	SUM	0.3329	-3.55
R-Prop	Diff	1000	0.3458	SUM	0.3376	-2.38
R-Prop	TfIdf	50	0.3654	SUM	0.3443	-5.77
R-Prop	TfIdf	250	0.3643	SUM	0.3613	-0.82
R-Prop	TfIdf	500	0.3477	SUM	0.3381	-2.76
R-Prop	TfIdf	750	0.3340	SUM	0.3334	-0.17
R-Prop	TfIdf	1000	0.3227	SUM	0.3304	<u>2.37</u>
R-TF	Diff	50	0.3325	SUM	0.3115	-6.34
R-TF	Diff	250	0.3379	SUM	0.3584	<u>6.05</u>
R-TF	Diff	500	0.3443	SUM	0.3631	<u>5.46</u>
R-TF	Diff	750	0.3469	SUM	0.3653	<u>5.32</u>
R-TF	Diff	1000	0.3496	SUM	0.3697	<u>5.77</u>
R-TF	TfIdf	50	0.3372	SUM	0.3190	-5.38
R-TF	TfIdf	250	0.3483	SUM	0.3721	<u>6.84</u>
R-TF	TfIdf	500	0.3432	SUM	0.3675	<u>7.08</u>
R-TF	TfIdf	750	0.3385	SUM	0.3699	<u>9.28</u>
R-TF	TfIdf	1000	0.3314	SUM	0.3729	<u>12.52</u>

First of all, let us focus on the behavior of monolithic profiles, drawing the following fact-based conclusions from the data:

- $R - Prop$ replication is generally better than $R - TF$, although they are very similar.
- For $R - Prop$, smaller profiles are better than bigger, while for $R - TF$ larger of medium profiles are preferable.
- $Diff$ usually gets better NDCG@10 values than $TfIdf$ in larger profiles and worse in smaller ones.
- The best absolute value, i.e. 0.3654, is found in a profile of size 50, using $R - Prop$ and $TfIdf$.

Regarding the subprofiles, the following findings from the data are drawn:

- SUM filtering is systematically better than MAX . The reason could be that non-specialized queries could occur in the query battery, so in this case, several and different subprofiles could be high in the original ranking, making MP's final score to increase.
- In this case, and in opposition to the results from monolithic profiles, $R - TF$ is better than $R - Prop$.
- $TfIdf$ tends to be better than $Diff$.
- The best absolute value is given by the configuration $R - TF$, $TfIdf$, 1000 and SUM : 0.3729.

In any case, the differences in performance between some parameter values and others, and even between profiles and subprofiles are not bigger enough to tip the scales in favor of ones or the others. But when comparing both knowledge representation techniques, we find that only in 13 experiments out 40 (values underlined in Table 1), the subprofiles perform better than the monolithic profiles, most of them when $R - TF$ and SUM are present in the configurations. The differences in terms of performance, anyway, are very low, i.e. the NDCG@10 values are really similar.

But what is the explanation for this behavior when specialized literature maintains that the application of subprofiles is better than using monolithic profiles? [7]. One possible answer could be that a rather large number of MPs participates a relatively low number of times. Their speeches would not be enough to learn different quality subprofiles. Then our supposition is that for those MPs who do not participate too much in committees or plenary sessions it is better to keep their speeches in only one profile. But, at the same time, those who intervene a lot could have several subprofiles with a minimum of quality to be representative of their interest topics.

A second explanation could be the different meaning of the profile size with respect to the subprofile sizes. When the former is built, for example, with fifty terms, we could say that the best fifty of them are included in it. Meanwhile, when the different subprofiles are learnt, if an MP has participated in six committees, for example, she would have associated six subprofiles with fifty terms each (the best fifty in each committee), i.e. we are using 300 terms in total (although some terms are probably common for all the committees). In the profile, we could say that we have a compact representation of her speeches, including the most representative terms from all the committees, while in the subprofiles, in addition to these terms, we are adding new terms that could not be very useful and closer to noise. Then, specially for a low number of interventions, in this case, subprofiles could decrease the performance of retrieval.

Finally, a third possibility could be that the behavior depends on the type of query: some queries, perhaps the more politic-oriented ones, seem to be better to retrieve a monolithic profile, capturing much better the essence of the MPs' discourses, as opposed to other queries, perhaps more specific, more topic-oriented, which are better to use with subprofiles.

With the aim of trying to find if our first hypothesis could make sense, we have carried out a new set of experiments, where the monolithic profiles have been joined to subprofiles,

creating a new collection, noted as *mixed subprofiles*, where the monolithic profile is a subprofile itself. The objective then is to observe the performance of the experiments and determine if it increases mixing both types of profiles. The results of this new experimentation are included in Table 2, where:

- The first four columns of Table 2 contains the values of replication, selection, size and filter methods, as in Table 1.
- For comparison purposes, we have included the column %C-SubP-P, where the percentages of change of NDCG@10 for subprofiles with respect to monolithic profiles are shown (column %C of Table 1).
- The sixth column contains the NDCG@10 values for the mixed subprofiles (MSubP).
- The last two columns represent the percentage of change of NDCG@10 for mixed subprofiles with respect to monolithic profiles (%C-MSubP-P) and subprofiles (%C-MSubP-SubP), respectively.

The following facts about mixed subprofiles can be observed from the results of Table 2:

- $R - TF$ is generally better than $R - Prop$, but again with only small differences between them.
- *Diff* seems to obtain higher values with relatively small sizes (250) with $R - Prop$ and large sizes (1000) with $R - TF$; *TfIdf* is more steady and always gets the highest values when the subprofiles contain 250 terms, independently on the replication method.
- As happened with the subprofile NDCG@10 values, the *SUM* filter is always better than the *MAX* one.
- The best absolute value is 0.3874, found in the $R - TF$, *TfIdf*, 250 and *SUM* parameter combination. This is almost the same combination than for subprofiles, except the size, which is more reduced.

Comparing the NDCG@10 values and the percentage of change of mixed subprofiles, firstly with respect to monolithic profiles, we can observe how the performance increases with 22 percentages of change greater than 0 (column %C-MSubP-P). A noticeable fact is that except two percentages, when the *SUM* filter is applied, the NDCG@10 with mixed subprofiles is better than the obtained with monolithic ones. The same behavior is present when we compare mixed subprofiles with subprofiles (column %C-MSubP-SubP): 27 percentages of change greater than 0, and mixed subprofiles are always better than subprofiles when using the *SUM* filter. Therefore, we may conclude that the inclusion of the monolithic profile as a subprofile itself is positive for increasing the performance of the MP recommendation.

The results of tables 1 and 2 are NDCG@10 values averaged from the values of the five partitions. An insight of what is happening inside these partitions could give us a more accurate picture of the real situation. In this sense, we try to measure in which way the results depend on the number of training data. For this purpose, we focus on the isolated results obtained by the different MPs. In the graph of Figure 1 it can be found that their number of interventions,

represented in decreasing order, were not homogeneously distributed: an MP has got a mean of 92.4 interventions with a standard deviation of 71.3. Also, focusing on the quartiles of the distribution it can be found that the 25%, 50%, 75% and 100% of the MPs participate less than 40, 75, 120 and 380 times in the legislature, respectively.

So, analyzing the results for each MP we can obtain some information about the importance of the number of training data. Particularly, for this purpose, we use as performance measure the recall values obtained for each MP in the test partition, i.e., the ratio of the number of times that the MP was properly recommended and the number of his/her initiatives in the test set. For illustrative purposes, we plot the results (Figure 2) obtained using the different profiles, i.e., monolithic, subprofiles and mixed approaches³ in one of the five partitions, where the X axis represents the number of interventions of the MP in the training data and the Y axis represents the tendencies of the recall values (obtained by means of a linear regression). The averaged recall for each profile is 0.4309, 0.4644 and 0.4729, respectively. We would like to note that we speak about tendencies because the results are highly user dependent, which by itself it could deserve a deeper analysis from the political science perspective.

By means of this graph some conclusions could be obtained (although for the sake of simplicity we only show the results using one partition, the behavior of the rest of partitions is more or less the same): The first one is that when the number of training data increases, the recall values also tend to increase, which might be obvious since we have more data to learn the MPs preferences. This increment is particularly noticeable for subprofiles.

The trends are not the same for the different profile definitions: Monolithic profiles, comparing with subprofiles, seem to be better when the number of interventions is very low (in this fold approximately 20 initiatives, representing the 10% of the MPs). As the number of training data increases the recall values for subprofiles outperform monolithic profiles values. Then, we can say that when the number of MP's interventions is low, we do not have enough information to build quality subprofiles (each one of these subprofiles is obtained with a very low number of interventions –one or two– and therefore it is a low quality profile, being not good enough for representing MP's interests), so it is much better to use monolithic profiles for her.

Next, focusing on the tendencies of subprofiles and mixed subprofiles, the number of 105 interventions is approximately the position where both straight lines crossed, in other words, where subprofiles start to outperform the mixed approach (which represents the 35% of the MPs and the 53% of the recommendations). In this case, the average recall starting for those MPs with more than 105 interventions is 0.5136 and 0.4971, respectively⁴. Below this value, representing the 65% of the users, the mixed approach takes advantage from both, subprofiles and monolithic profiles, which collaborate properly in this hybrid approach obtaining in average best results with the mixed profile (the averaged recall values are 0.4473 and 0.4643, respectively⁵). So we can conclude that the subprofiles are better descriptions of the user pref-

³The used configuration was 1000 terms, *Diff*, $R - Tf$ and *SUM*.

⁴For monolithic profiles, this averaged recall is 0.4423.

⁵The averaged recall for monolithic profiles is 0.4269.

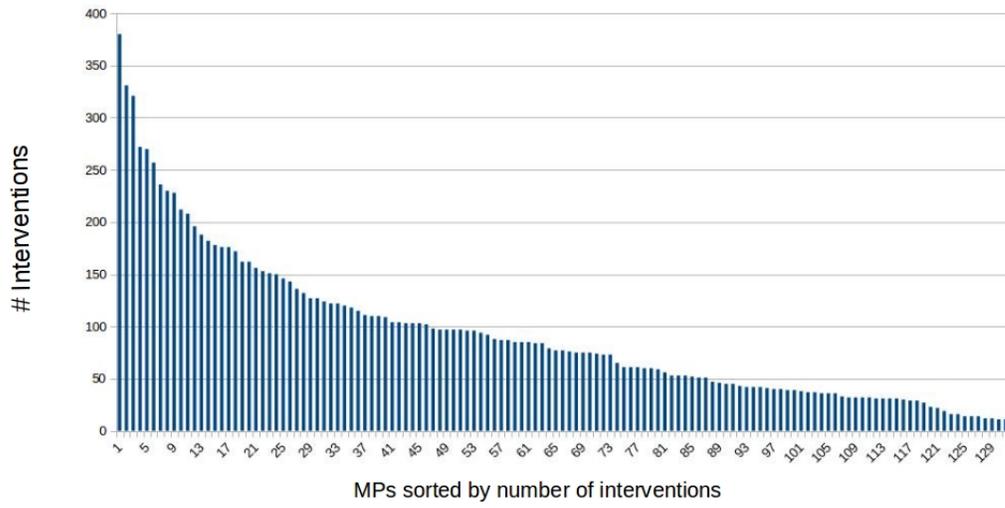


Figure 1: Number of interventions for MPs.

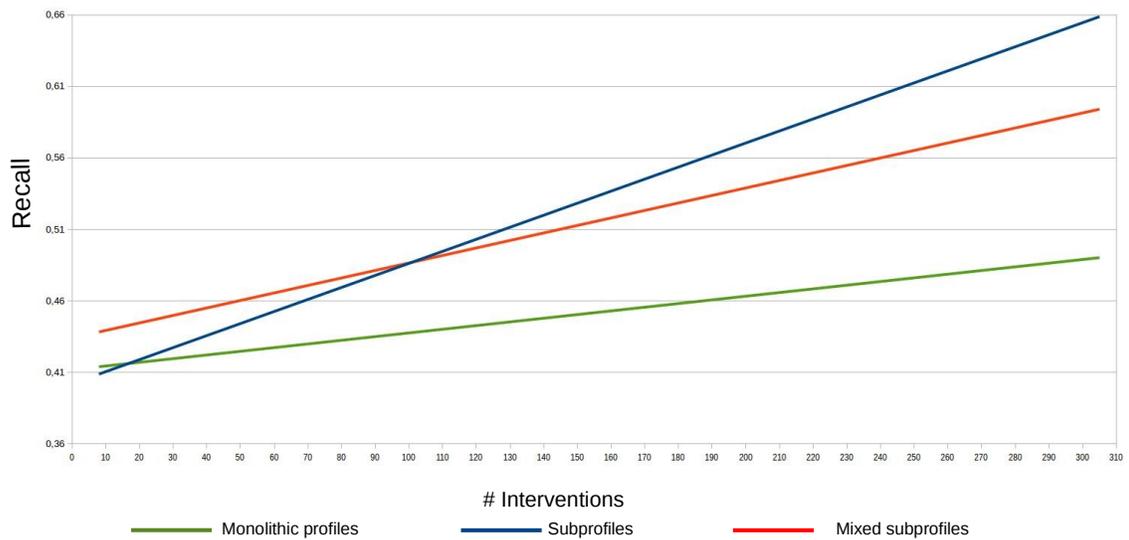


Figure 2: Tendencies of the recall values, as a function of the number of interventions for MPs, for monolithic profiles, subprofiles and mixed profiles, for one of the five partitions.

erences when we have enough data, allowing us to obtain better accuracy in the recommendations.

4. CONCLUSIONS AND FURTHER WORK

In this paper we have presented a comparison between monolithic profiles built from all the initiatives in committee and plenary sessions of the MPs in the Andalusian Parliament and those constructed separately for each type of committee and plenary sessions where the MPs have participated, originating a number of different subprofiles for each MP. These profiles are used as documents in the context of recommending MPs to users who formulate queries.

These structures, which capture the topics where MPs are interested in, are built selecting terms from their interventions by means of *TfIdf* or *Diff* measures, and later replicating the number of terms according with *R - Prop* and *R - TF* techniques. Two ways of filtering subprofiles in rankings have been proposed, *MAX* and *SUM*, in order to obtain non-repeated MPs in the ranking of subprofiles. In the two first types of measures the differences in terms of performance are very low and it is difficult to find a pattern. However, concerning the filters, we may notice clearly better results using *SUM*.

The bare results from the experimentation show that considering interventions in committee and plenary sessions is very similar to grouping all of them and building monolithic profiles. The main problem is that if the number of initiatives where an MP participates is low, then quality subprofiles can not be built, so monolithic ones are preferable. This fact leads us to think that, in a context where subprofiles are taken into account, each MP could have a different scheme depending on the amount of text available for creating profiles.

These approaches could be clearly exported to any other field, outside the parliamentary context. The only requirement is that the document collection is divided in categories from where the subprofiles could be built. Otherwise, if these are not available, then they could be obtained by the application of clustering methods.

A deeper sight to the results shows that the use of subprofiles could be very beneficial for the purposes of recommending MPs. Furthermore, and in this line, we plan to create subprofiles by using clustering algorithms, independently on the committees or plenary sessions where the MPs participate. The main advantage of this approach would be that each MP could have a distinct number of subprofiles, depending on the groups of terms found, and maybe with different sizes.

5. REFERENCES

- [1] Ahn, J., Brusilovsky, P., Grady, J., He, D. and Syn, S.Y. Open user profiles for adaptive news systems: help or harm? In *Proceedings of the 16th international conference on World Wide Web*, pages 11–20, 2007.
- [2] Amati, G., D’Aloisi, D., Gianini, V. and Ubaldini, F. A framework for filtering news and managing distributed data. *Journal of Universal Computer Science* pages 1007-1021, 3, 1997.
- [3] Balabanovic, M. and Shoham, Y. Fab: Content-based collaborative recommendations. *Communications of the ACM*, pages 66–72, 40, 1997.
- [4] Bilenko, M. and Richardson, M. Predictive client-side profiles for personalized advertising. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 413–421, 2011.
- [5] Cantador, I., Bellogín, A. and Vallet, D. Content-based recommendation in social tagging systems. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 237–240, 2010.
- [6] Chen, L. and Sycara, K. A personal agent for browsing and searching. In *Proceedings of the Second International Conference on Autonomous Agents*, pages 132-139, 1998.
- [7] Gauch, S., Speretta M., Chandramouli, A. and Micarelli A. User profiles for personalized information access. In *The Adaptive Web, Lecture Notes in Computer Science*, pages 54–89, 4321, 2007.
- [8] Hofmann K, Balog K, Bogers T and de Rijke M. Contextual factors for finding similar experts. *Journal of the American Society for Information Science and Technology*, pages 994–1014, 61(5), 2010.
- [9] Jarvelin, J. and Kekalainen, J. Cumulative Gain-based Evaluation of IR Techniques. *ACM Transactions on Information Systems*, pages 422–446, 20, 2002.
- [10] Kamba, T., Sakagami, H. and Koseki, Y. ANATAGONOMY: a personalized newspaper on the World Wide Web. *International Journal of Human-Computer Studies*, pages 789–803, 46, 1997.
- [11] Lieberman, H. Letizia: an agent that assists web browsing. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 924–929, 1995.
- [12] Montebello, M., Gray, W. and Hurley, S. A personal evolvable advisor for WWW knowledge-based systems. In *Proceedings of the International Database Engineering and Application Symposium*, pages 224–233, 1998.
- [13] Moukas A. Amalthea: information discovering and filtering using a multiagent evolving ecosystem. *Applied Artificial Intelligence*, pages 437–457, 11, 1997.
- [14] Pazzani M., Muramatsu J. and Billsus D. Syskill & Webert: identifying interesting web sites. In *Proceedings of the 13th National Conference on Artificial Intelligence*, pages 54–61, 1996.
- [15] Somlo, G.L. and Howe, A.E. Incremental clustering for profile maintenance in information gathering web agents In *Proceedings of the 2001 International Conference on Autonomous Agents*, pages 262–269, 2001.
- [16] Vicente-López, E., de Campos, L.M., Fernández-Luna, J.M. and Huete, J.F. Personalization of parliamentary document retrieval using different user profiles. In *Proceedings of the 2nd International Workshop on Personalization in eGovernment Services and Applications (PEGOV2014)*, Volume 1181 of CEUR workshop proceedings, 2014.
- [17] Widyantoro, D.H., Yin, J., El Nasr, M., Yang, L., Zacchi, A. and Yen, J. Alipes: a swift messenger in cyberspace. In *Proceedings of the AAAI Spring Symposium Workshop on Intelligent Agents in Cyberspace*, pages 62–67, 1999.

Table 2: Parameter configuration, NDCG@10 values for mixed subprofiles and percentages of change

Rep.	Sel.	Size	Filter	%C-SubP-P	MSubP	%C-MSubP-P	%C-MSubP-SubP
R-Prop	Diff	50	MAX	-18.07	0.2791	-18.26	-0.23
R-Prop	Diff	250	MAX	-11.11	0.3144	-10.94	<u>0.19</u>
R-Prop	Diff	500	MAX	-10.85	0.3096	-10.81	<u>0.05</u>
R-Prop	Diff	750	MAX	-9.79	0.3078	-10.81	-1.14
R-Prop	Diff	1000	MAX	-11.53	0.3145	-9.05	<u>2.80</u>
R-Prop	Tfidf	50	MAX	-14.10	0.3245	-11.19	<u>3.38</u>
R-Prop	Tfidf	250	MAX	-9.59	0.3321	-8.83	<u>0.84</u>
R-Prop	Tfidf	500	MAX	-11.61	0.3074	-11.61	0.00
R-Prop	Tfidf	750	MAX	-9.55	0.2991	-10.45	-0.99
R-Prop	Tfidf	1000	MAX	-5.77	0.3020	-6.42	-0.69
R-TF	Diff	50	MAX	-15.23	0.2815	-15.36	-0.15
R-TF	Diff	250	MAX	-0.66	0.3350	-0.87	-0.21
R-TF	Diff	500	MAX	-1.42	0.3380	-1.83	-0.41
R-TF	Diff	750	MAX	-0.23	0.3447	-0.61	-0.38
R-TF	Diff	1000	MAX	0.47	0.3518	<u>0.65</u>	0.18
R-TF	Tfidf	50	MAX	-10.81	0.3139	-6.92	<u>4.37</u>
R-TF	Tfidf	250	MAX	0.80	0.3509	<u>0.77</u>	-0.03
R-TF	Tfidf	500	MAX	-1.02	0.3365	-1.95	-0.94
R-TF	Tfidf	750	MAX	2.47	0.3396	<u>0.32</u>	-2.10
R-TF	Tfidf	1000	MAX	6.59	0.3445	<u>3.95</u>	-2.48
R-Prop	Diff	50	SUM	-8.31	0.3327	-2.57	6.27
R-Prop	Diff	250	SUM	-2.76	0.3597	<u>1.91</u>	<u>4.80</u>
R-Prop	Diff	500	SUM	-2.95	0.3582	<u>3.18</u>	<u>6.32</u>
R-Prop	Diff	750	SUM	-3.55	0.3539	<u>2.51</u>	<u>6.28</u>
R-Prop	Diff	1000	SUM	-2.38	0.3576	<u>3.40</u>	<u>5.91</u>
R-Prop	Tfidf	50	SUM	-5.77	0.3669	<u>0.41</u>	<u>6.56</u>
R-Prop	Tfidf	250	SUM	-0.82	0.3811	<u>4.63</u>	<u>5.50</u>
R-Prop	Tfidf	500	SUM	-2.76	0.3578	<u>2.90</u>	<u>5.82</u>
R-Prop	Tfidf	750	SUM	-0.17	0.3514	<u>5.21</u>	<u>5.38</u>
R-Prop	Tfidf	1000	SUM	2.37	0.3459	<u>7.19</u>	<u>4.71</u>
R-TF	Diff	50	SUM	-6.34	0.3323	-0.07	6.69
R-TF	Diff	250	SUM	6.05	0.3673	<u>8.68</u>	<u>2.48</u>
R-TF	Diff	500	SUM	5.46	0.3767	<u>9.43</u>	<u>3.77</u>
R-TF	Diff	750	SUM	5.32	0.3796	<u>9.43</u>	<u>3.90</u>
R-TF	Diff	1000	SUM	5.77	0.3811	<u>9.02</u>	<u>3.07</u>
R-TF	Tfidf	50	SUM	-5.38	0.3506	<u>3.97</u>	<u>9.88</u>
R-TF	Tfidf	250	SUM	6.84	0.3874	<u>11.24</u>	<u>4.12</u>
R-TF	Tfidf	500	SUM	7.08	0.3768	<u>9.80</u>	<u>2.54</u>
R-TF	Tfidf	750	SUM	9.28	0.3772	<u>11.42</u>	<u>1.96</u>
R-TF	Tfidf	1000	SUM	12.52	0.3780	<u>14.06</u>	<u>1.37</u>

6. ACKNOWLEDGMENTS

This work has been funded by the Spanish Ministerio de Economía y Competitividad under project TIN2013-42741-P and the European Regional Development Fund (ERDF-FEDER).