International Journal of Uncertainty,
Fuzziness and Knowledge-Based Systems
Vol. 25, Suppl. 2 (December 2017) 21–36

DOI: 10.1142/S0218488517400098

Committee-Based Profiles for Politician Finding

Luis M. De Campos*, Juan M. Fernández-Luna† and Juan F. Huete‡

Computer Science and A.I Department – ETSIIT,
Universidad de Granada, 18014, Granada, Spain
**lci@decsai.ugr.es*
†jmfluna@decsai.ugr.es
‡jhg@decsai.ugr.es

Received 21 February 2017
Revised 24 July 2017

One step towards breaking down barriers between citizens and politicians is to help people identify those politicians who share their concerns. This paper is set in the field of expert finding and is based on the automatic construction of politicians' profiles from their speeches on parliamentary committees. These committee-based profiles are treated as documents and are indexed by an information retrieval system. Given a query representing a citizen's concern, a profile ranking is then obtained. In the final step, the different results for each candidate are combined in order to obtain the final politician ranking. We explore the use of classic combination strategies for this purpose and present a new approach that improves state-of-the-art performance and which is more stable under different conditions. We also introduce a two-stage model where the identification of a broader concept (such as the committee) is used to improve the final politician ranking.

Keywords: User-profile; politician recommendation; expert finding; information retrieval.

1. Introduction

A challenge for any government is to find ways to strengthen and improve relationships between politicians and society (i.e. citizens and/or stakeholders) and this is a growing demand from both sides. A possible situation in this context might be where a user (citizen, businessman, journalist, etc.) wishes to meet a politician to discuss a certain problem. In such a situation, a common mistake is to arrange a meeting with the wrong politician (with the associated costs in terms of time and/or money). There are two reasons for this mistake: firstly, citizens believe that any politician is an expert in every field, a premise which is not necessarily true; and secondly, it is difficult for citizens to know what the real interests of a politician are. In view of this, our main research question is whether we can *effectively* and *efficiently* design an intelligent information system to search for those politicians who are truly interested in a specific matter.

Although the framework of our research is the field of e-Government, it also falls within the broader area of "expert finding"[1,2] where the objective is to find an expert in a given topic (science researchers, professionals, journalists, bloggers, physicians, etc.). In our case, since Members of Parliament (MPs) are the main actors (i.e. the target), this can be viewed as an MP-finding or MP-recommendation task. To the best of our knowledge, this problem has never before been studied in the political context where this research is set and which displays certain peculiarities that must be considered.

In order to understand such peculiarities, let us briefly explain how political activities are organized in a parliamentary context: the key element is the initiative, which represents a plan for discussion on a specific issue, law or decree. If the initiative is admitted, it follows the classic legislative procedure: it is passed by the different commissions or committees and is submitted, if necessary, to the plenary sessions. Thus, in plenary sessions, MPs discuss general topics and in committee sessions they deal with more specific topics, such as agriculture, education, economy, etc. In both types of sessions, usually only one MP for each party (the representative or spokesperson) participates in the discussion of a given initiative. It is important to mention that records of all parliamentary proceedings (the Official Gazette) and full transcriptions of every speech are freely available to the public.

Returning to the MP-finding task, it is necessary for the information associated to MPs (i.e. the underlying topics they know about) to be organized in a sort of container which is called the profile. These profiles might be created manually (where the experts themselves, MPs or their assistants, express their preferences[3,4]) or automatically (by analyzing documents created in the course of routine work[4,5]). There are various reasons why we prefer to use automatic profiling in a political context. Firstly, politicians (and people in general, as Berends *et al.*[6] highlighted) dislike having to complete lengthy forms about their interests. Even then, inaccurate profiles may be obtained since politicians have a wide range of interests and they must select their preferences from a large pool of possibilities and it is also common that they slant the selection process towards a subset of relevant topics. Another reason in favor of automatic profiles is that manually created ones are usually static since people rarely update their preferences. This will represent a problem in our context because it is quite usual that an MP changes their interests from one field to another, maybe due to political issues.

In de Campos *et al.*,[7] we focused on the automatic creation of MP profiles and their effects on the performance of an MP-finding system. In order to tackle the problem, we considered their parliamentary interventions/speeches as the information source. The paper considered two approaches. The first is to use a *single profile* that represents a monolithic view of MP preferences. In this case, the terms comprising it will come from different domains of interest which might cause the profile to only reflect a few interest categories or for non-dominant categories to become underweighted or even unnoticed. For the second approach, we might consider that an MP normally sits on more than one committee and so a *structured profile* might be built. This type of profile consists of various subprofiles to represent the topics of the different committees with which the MP is involved. For example, if a representative has participated in education, culture, environment and economy, four subprofiles could be constructed.

In the context of MP recommendation, the objective of this paper is to determine whether it would be advantageous to use a structured profile based on the committee sessions in which the MP has participated, rather than a single monolithic profile. For a given query, our objective is to obtain an MP ranking but if we consider structured profiles, an MP might have more than one subprofile in the ranking and so it is necessary to combine the scores of these to obtain the final ranking.

Taking this framework into account, the main contributions of this work can be summarized as follows:

- Although the expert-finding problem has been addressed from different perspectives in academia and industry, we could not find any previous work on expert recommendation in political domains. This work represents a first attempt to tackle this problem from an overall perspective where the use of structured profiles has proved to be a good alternative. We also want to highlight the applied evaluation methodology since it can be used in other political organizations.
- Regarding fusion or combination strategies, this paper explores the performance of a new aggregation strategy, which is a hybrid of both score and ranking-based approaches and based on the discounted cumulative gain evaluation metric. We shall show how this performs well with more stable results than classical approaches.
- Insights into the performance of the structured profiles under a related task (i.e. identification of the committee where user concerns might be discussed) are obtained. Taking this into account, we propose a new two-stage approach for MP-finding which is able to exploit such information.

This paper is organized as follows: the following section presents related work; Section 3 describes the construction of both monolithic profiles and subprofiles and also the fusion strategies used in the case of structured profiles; the next section describes the experimental methodology and results; Section 5 presents the two-stage approach for MP recommendation; and finally, Section 6 outlines the main conclusions and future lines of research.

2. Related Work

The expert-finding problem has received interest (under different formalisms) in fields such as academia[4,6,8–10] and industry[3,11–13] and two different approaches are possible:

- In a first approach,[6] user expertise is represented as a vector of keywords (index terms obtained from a controlled vocabulary such as a thesaurus). A query, which also consists of a set of keywords, is then used to determine an expert for that topic. Nevertheless, such types of queries are useless when considering citizens as users because the keywords are either too general to be useful or too specific and possibly unknown by the citizens.

- A second approach is to tackle this task as an information retrieval problem,[1,3,14] where the profile is considered a "virtual document" based on the documents associated to the candidate or using an approach which first ranks documents in the corpus given a query topic,[1] and then finds the associated candidates from the subset of retrieved documents. Two different alternatives have been proposed: the first, which Balog denoted as Model1,[3] where all the documents are joined into a single document, and Model2, where each document is considered in isolation.

Our approach, which explores the use of profiles for an expert-finding task in the parliamentary domain, is set between both alternatives: we first learn a profile for each candidate from their documents (as in the first approach we obtain a summary of their interests) and this is then transformed into a virtual document (in the same way as the second approach). According to Gauch *et al.*,[15] there are three main representations for user profiles: weighted keywords, semantic networks and weighted concepts. We have considered weighted keyword representation since this is the most common approach. Additionally, these profiles might be used on various IR-related tasks such as personalization (where the original query may be reformulated in order to better adapt the retrieval results to user preferences), content-based recommendation (which helps to filter new documents), clustering (MP segmentation, i.e. grouping MPs with similar preferences), etc.

In the case of using weighted keyword representation, two alternatives can be considered: the first uses a single profile for each user as discussed by Panoptic,[14] Fab,[16] and more recently Cantador *et al.*[17] and Bilenko *et al.*[18] In the case of the user having multiple interests, however, as in the case of an MP being involved in several committees (e.g. agriculture, education and economy), non-dominant categories might be underestimated. The second approach tries to solve this problem by organizing the profile into a set of subprofiles where each represents the person's topic of interest, the terms that comprise them are closely related under the broad umbrella of a common subject. The following examples use subprofiles: Montebello *et al.*[19] design a personalized web search assistant to extract keyword-based profiles from bookmarked web pages; automatic clustering approaches[20–22] are used to learn user interest areas and Syskill and Webert[23] learn a separate profile for each user's topics from positive and negative examples of each topic using machine learning algorithms.

In previous work (de Campos *et al.*[7]), we discussed various strategies to build an MP profile defined as a bag of n words, that contains the most important terms according to a weighting scheme. In that paper, we focused on the impact of weighting criteria, the number of terms in the profile and how to convert a weighted vector of terms into a document. We considered both approaches, i.e. the use of single (monolithic) and compound (structured) profiles.

In this paper, we shall analyze in detail the performance of such profiles for the MP-finding problem and our objective is to obtain an MP ranking in response to the query. Nevertheless, we wish to highlight that profile-based virtual documents have different features from those obtained when considering Balog's Model1 and Model2 approaches, which will be analyzed in greater detail in the experimental section.

In order to conclude this section, we shall briefly review work on data fusion.[24] These techniques are necessary in the case of using structured profiles, where it is possible for a given MP to have more than one subprofile in the final ranking of the retrieval engine. Following the ideas of Refs. 1 and **?**, each (sub)profile can be seen as a vote for the MP represented in it, and it is necessary to combine the different candidate votes to obtain the final ranking. For example, Macdonald and Ounis[25] propose up to twelve different methods which can be categorized into score- and ranking-based, with CombMax, CombSUM and CombMNZ (see Section 3.1) performing best.

3. Using the Profiles for the Expert-Finding Task

In the case of using monolithic profiles, the ranking obtained by the IRS can be used directly for this task (each profile in the ranking represents a single MP) and therefore the top-ranked MPs can be considered "experts" on the subject. However, this does not have to be the case when compound profiles are used. In such a situation, the retrieval engine returns a scored list of subprofiles, $\mathcal{L}_q = \{p_1, p_2, \ldots, p_k\}$, where each p_i is a tuple $p =< mp, c >$ where mp presents the particular MP and c represents a committee. Given the one-to-many relationship between an MP and their subprofiles, it is also possible for a given MP to have more than one listed profile. In order to compute a single score for each MP, we might therefore combine the different results. Once combined, the MPs are sorted in terms of these values and those in the top positions will be returned to the citizen.

3.1. *Fusion strategies*

If we consider each subprofile as a vote for the candidate, it becomes necessary to use some fusion strategy to compute a final score for each MP (all the information presented in their different subprofiles in the ranking might be considered). In this case, however, it is possible that for a given query the number of subprofiles returned by the system could be quite large, with those in the lower positions being barely relevant. Thus, and following the results in Weerkamp *et al.*,[26] we should only focus on those profiles in the top $\mathcal{K}$ positions in the ranking. It appears that taking into account the information provided by these subprofiles (which is marginally relevant to the query) worsens system performance.

We shall now discuss the fusion strategies used in this paper. Although Macdonald and Ounis[25] used twelve fusion strategies, we shall only present four of these which are those that perform best in literature[1] (and also in our experiments).

These can be categorized into rank- and score-based approaches. Two strategies representing a hybridization of both approaches will also be proposed.

We will consider some notation in this context: given a profile p_i in the ranking, $I(p_i)$ represents the MP identity and $s(p_i)$ denotes its score value (similarity between the profile and the query q). For a given MP m, we will also use $\mathcal{P}_{\mathcal{K}}(m,q)$ to denote the set profiles in $\mathcal{L}_q$ including her subprofile with the greatest score in $\mathcal{L}_q$, p_m, but also those among the top-$\mathcal{K}$ which are associated to the MP, i.e.

$$\mathcal{P}_{\mathcal{K}}(m,q) = \{p_m\} \cup \{p_j \in \mathcal{L}_q | I(p_j) = m \ \& \ j \leq \mathcal{K}\}\,, \tag{1}$$

being $m = \arg\max_i\{s(p_i)|I(p_i) = m\}$.

Scored-based approach: only the scores of the retrieved profiles are used to compute the final ranking.

- $CombMax$: Given an MP, the subprofile with the highest score is selected. This first filtering method attempts to capture the best match and is intended for very specific queries:

$$score(m,q) = \max_{p_i \in \mathcal{P}_{\mathcal{K}}(m,q)} \{s(p_i)\} = s(p_m)\,. \tag{2}$$

- $CombSum$: Given an MP, a new score is computed by adding the different scores of their subprofiles in $\mathcal{P}_{\mathcal{K}}$. This second approach could be useful when the query covers various subprofiles, so the greater the number with higher scores, the better:

$$score(m,q) = \sum_{p_i \in \mathcal{P}_{\mathcal{K}}(m,q)} s(p_i)\,. \tag{3}$$

- $CombMNZ$: This includes an extra factor in the CombSum method which takes into account the number of results where the MP appears ($\mathcal{P}_{\mathcal{K}}$). Thus, CombMNZ favors those MPs who appear several times in the retrieved list (reinforcing the aggregated score obtained in such situations):

$$score(m,q) = |\mathcal{P}_{\mathcal{K}}(m,q)| \times \sum_{p_i \in \mathcal{P}_{\mathcal{K}}(m,q)} s(p_i)\,. \tag{4}$$

Rank-based approach: This takes into account the positions of the subprofiles in the ranking.

- $CombRR$: This is an adaptation of the reciprocal rank[25] method where an MP's rank in the output ranking is determined by the sum of the reciprocal ranks of the related profiles, in such a way that the higher the position in the ranking, the greater its value.

$$score(m,q) = \sum_{p_i \in \mathcal{P}_{\mathcal{K}}(m,q)} \frac{1}{rank(p_i)}\,. \tag{5}$$

Mixed approaches: The value of the score for each subprofile is devalued by taking into account its position in the ranking and these devalued scores are aggregated to compute the final value for a particular MP. Two different devaluation strategies are considered:

- *CombDCS*: This is a mixture of the CombSum and CombRR methods, where the scores are discounted proportionally to the subprofile's position in the ranking.

$$score(m, q) = \sum_{p_i \in \mathcal{P}_{\mathcal{K}}(m,q)} \frac{s(p_i)}{rank(p_i)} . \tag{6}$$

- *CombLgDCS*: This is similar to the previous one, but in this case the scores are reduced logarithmically proportional to the position of the result, as the discounted cumulative gain does for information retrieval evaluation,[27] but as far as we know this approach has not been considered for aggregation purposes. By means of this devaluation, we penalize the occurrence of a subprofile in lower positions of the ranking, but less so than the one used by the reciprocal rank criteria.

$$score(m, q) = \sum_{p_i \in \mathcal{P}_{\mathcal{K}}(m,q)} \frac{s(p_i)}{\log_2(rank(p_i) + 1)} . \tag{7}$$

4. Experimentation and Results

This section will present both the experimental design and the results obtained in addition to analyzing these and presenting the main findings.

4.1. *The data collection*

The experiments are based on a collection of Records of Parliamentary Proceedings from the Andalusian Parliament in Spain and in particular the eighth Term of Office, which has 22 different committees and a total of 132 spokespersons. The collection used in this paper is organized around the initiatives discussed in committee sessions (3001 in total with 7386 interventions). A total of 1063 subprofiles (different MP-Committee pairs) have been learned, where MPs participate on an average of 7.5 committees (standard deviation 4.6, median 7).

Regarding the relevance judgments, we shall use the initiative title as the query (it is worth noting that we do not use this field to build MP profiles) and this enables us to model a typical short query submitted by a citizen. Since the objective is to find MPs who might be familiar with the topic, we have also considered as ground truth that for each query only those MPs who participate in its corresponding initiative are relevant, thereby creating a rather conservative assumption since it is quite reasonable to assume that an initiative will also be relevant to other MPs.[a]

[a]This is equivalent to the common assumption that is made when relevance judgments are considered in the researcher search problem, i.e. an author of a paper is relevant to that paper.

To the best of our knowledge, there is no dataset for expert finding in political frameworks that provides truly relevant judgments for each query-expert pair, but thanks to the methodology used here we can study the feasibility of new approaches for MP finding in other political organizations.

4.2. *Experimental setting*

The set of initiatives is randomly partitioned into a training set (80%) and a test set (20%). As usual, both sets are disjoint. The training set is used to build MP profiles and the test set is used for evaluation purposes. This process is repeated five times, and in this paper the reported results are the average values.

In terms of profile construction, we rely on our previous results in de Campos *et al.*,[7] where we study the impact of the weighting criteria, the number of terms in the profile and also the way to convert a weighted vector of terms into a document. In particular, we only consider the weighting criterion with the best results, i.e. the terms in the profile are those with the greatest *TfIdf* (once stop words have been removed and terms stemmed). Additionally, each selected term is also replicated *Tf* times (where *Tf* is the number of times the term appears in the source initiatives) to build the virtual document. Regarding the number of terms, in this paper we will explore two alternatives: short ($n = 250$) and long profiles ($n = 1000$).

Search Engine. We have used the open source Lucene library using BM25 as the similarity measure. BM25 has two free parameters (k and b) which are set by default to the recommended values of $k = 1.2$ and $b = 0.75$ in the Lucene implementation (additional details on this model and parameter settings can be found in Jones *et al.*[28]). Nevertheless, when considering profile documents these default values do not perform best. Our hypothesis is that profiles are not regular documents for two reasons. Firstly, the selected terms are the most representative and highly related to user preferences. In this situation, a highly frequent term in the document-based profile is not due to either the MP's verbosity or the frequency of the term in language. Consequently, its influence in the ranking must be enhanced in the model. When considering BM25, the parameter k is used to control the impact of term frequency, which asymptotically approaches to a value $k + 1$. In this case, we obtain better results using $k = 15$, which allows us to stretch out the relevance difference between higher and lower term frequency in the profiles. Secondly, since we restrict the number of terms in the profile, it is necessary to dismiss the influence of the document length (it is worth noting that although the profiles have the same number of terms, the length of the document-based profiles can be different since each term appears a different number of times). In BM25 the parameter b, with $0 \leq b \leq 1$, enables adjustment of the influence of the document length: the lower b is, the less it affects the score. In our case, we use $b = 0.25$. These values, which are far from the default values, confirm our hypothesis that profiles are not regular documents.

Table 1. NDCG@10 results obtained for (long and short) Monolithic and Committee-based Profiles. The number of subprofiles used in the different aggregation functions, $\mathcal{K}$, are displayed in the columns.

Monolithic	*Aggregation*						
long	—	0.4647					
short	—	0.4371					
Committee		10	20	30	50	70	100
long	MAX	**0.4848**					
	SUM	**0.4779**	0.4643	0.4513	0.4227	0.3977	0.3789
	MNZ	**0.4768**	0.4553	0.4280	0.3713	0.3112	0.2954
	RR	0.4857	0.4881	0.4899	0.4920	0.4926	**0.4932**
	DCS	0.4858	0.4879	0.4890	0.4901	0.4904	**0.4905**
	LgDCS	0.4849	0.4865	0.4878	0.4892	0.4899	**0.4903**
short	MAX	**0.4584**					
	SUM	**0.4543**	0.4414	0.4271	0.3994	0.3794	0.3597
	MNZ	**0.4526**	0.4292	0.4018	0.3439	0.2936	0.2747
	RR	0.4591	0.4613	0.4607	0.4614	0.4614	**0.4617**
	DCS	0.4608	0.4617	0.4624	0.4625	0.4626	**0.4629**
	LgDCS	0.4609	**0.4613**	0.4604	0.4590	0.4591	0.4578

Evaluation Measures. In order to measure the quality of the ranking we will use[b] the Normalized Discounted Cumulative Gain,[27] which focuses on the top ten retrieval results, i.e. NDCG@10.

4.3. *Results for the expert-finding task*

In this section we shall present the results obtained using the different fusion strategies for both short and long profiles. In Table 1 we present the results obtained by the monolithic and combined profiles. It is worth noting that when using either monolithic profiles or CombMax as the fusion strategy, the relative position of MPs in the final ranking are the same as those obtained when the retrieval engine output ranking is considered. However, this does not have to be the case when the other fusion strategies are used where the number of subprofiles in the aggregation process, i.e. the parameter $\mathcal{K}$, has proved decisive. In order to explore this situation in this paper, we consider that $\mathcal{K}$ will take values in $\{10, 20, 30, 50, 70, 100\}$.

These results enable us to conclude that committee-based profiles outperform monolithic profiles and improvements of more than 5% can be achieved. When considering the efficiency of both approaches, we can also see that committee-based profiles obtain the results 80% quicker whereas the index size is only 4% larger.

Regarding the fusion strategies, we can observe that most display a slight trend towards using large $\mathcal{K}$ values (although the results do not differ significantly),

[b]In our experimentation we have used several metrics such as MAP, R-precision, Precision or Recall but due to length restrictions these are not included in this paper. Nevertheless, we should say that the results are highly correlated with NDCG.

except for CombSUM and CombMNZ where the performance drops dramatically as the number of used subprofiles increases. This was an unexpected result since in the literature[1,24–26] CombSUM and CombMNZ are considered among the best alternatives. One explanation for these results is that when subprofile documents are considered, the output ranking exhibits a different distribution of relevant and non-relevant items in the top positions. For example, if a query is related to an economy topic it would be normal for subprofiles associated with the Committee on Economic Affairs to be found in the top positions (each belonging to a different MP) whereas the presence in the ranking of subprofiles related to other committees might be considered as noise (they could be marginally relevant to the query). Since CombSUM and CombMNZ use the raw scores of these 'noisy' subprofiles in the aggregation step, the results worsen as they increase in number. Interestingly, this effect is reduced in those aggregation strategies that take into account the position of the element in the ranking, either in isolation, such as CombRR, or in our hybrid proposals such as CombLgDCS and CombDCS, where the position is used to penalize the score. Thus, the presence of noisy subprofiles in lower positions does not harm the final ranking.

Finally, if we consider the effect of profile size it is clear that we obtain better results with large subprofiles. Regarding the effect of the size on the aggregation functions, the trends are similar although large subprofiles seem to benefit slightly from using large $\mathcal{K}$ values.

5. A Two-Stage Model for MP Finding

As we have hypothesized, the distribution of subprofiles in the output ranking seems to be skewed towards those committees that are most related to the query. In this section, we will consider how such information can be used to improve the MP ranking. The idea is that when a committee appears in the top positions of the ranking it has more probability of being the committee in which the matter is discussed and this can be used to somehow support those MPs on these committees.

In order to tackle this objective, we can consider two different steps:

(1) To obtain a weight for each committee that represents the degree to which it is related to the user query.
(2) To recompute the scores for each subprofile in $\mathcal{L}_q$, promoting or dismissing them on the basis of their occurrence in a given committee.

5.1. *Computing the weights associated to the committees*

In this case, and taking into account that for each subprofile p we can identify the related MP, m, and the related committee, c, we can build a ranking of committees, R_C, by aggregating the scores of the different subprofiles related to the same committee. In this way, and as we did for the MPs, we are considering a structured profile for a committee that will consist of a set of subprofiles, one for each MP belonging to this committee. It is worth noting that by using this approach the same tuple (subprofile) will be part of two different profiles, one relating to an MP and the other to the committee.

5.2. *Combining the information*

Thus, for the same query we have two different rankings: the weighted list of committees, R_C, and the profile list $\mathcal{L}_q$. Our objective is to reweight the elements in $\mathcal{L}_q$ by considering the evidence obtained from the rank R_C. In order to do so, we propose the use of the linear combination method, an approach commonly used in data fusion since different weights can be assigned to different systems.[24] Before we present the combination, we should highlight the following two points:

- Taking into account that the scores provided by both lists do not have to be comparable (they have been generated differently), we need to normalize both scores before applying any fusion strategy. One straightforward way to perform such normalization is to divide by the maximal score in each list, i.e. the score of the first committee $w(c_1)$ in R_C and the score of the first profile $s(p_1)$ in $\mathcal{L}_q$.
- We can say that proof of a correct committee is stronger if it is located in the top positions of the ranking (as we shall see, in our experiments the right committee is usually among the top three positions of the ranking, approximately 86% of cases). Therefore, and in order to reward such committees, we propose the use of a quadratic function which greatly penalizes the bottom-ranked committees.

We are now able to describe how the new score of a profile p_i, denoted by $s^*(p_i)$, can be computed as

$$s^*(p_i) = \alpha s(p_i)/s(p_1) + (1-\alpha)(w(c_i)/w(c_1))^2\,, \tag{8}$$

where $0 \leq \alpha \leq 1$ is the weight associated to the original profile list and (1-α) represents the strength of the committee in the final ranking. Thus, α equals to 1 implies that no combination is performed and we use the original scores, whereas α equals 0 implies that we only take into account the value of the committee.[c]

After this process, we can use $s^*(p_i)$ to obtain a new ranking of subprofiles, which can be used as input for the different fusion strategies in order to determine the final MP's ranking. In the next section we will evaluate this approach.

5.3. *Results of the two-stage approach*

The feasibility of this approach depends on the quality of the committee's ranking and so we first analyze this task separately and then evaluate the two-stage approach.

[c]In previous experiments, we have tested a simple convex combination for scoring a profile, but the results have not been good.

Table 2. DCG@10 results for the different aggregation strategies using structured profiles for the committee finding problem.

Committee	*Aggregation*	10	20	30	50	70	100
long	MAX	**0.8002**					
	SUM	0.8110	0.8167	**0.8191**	0.8164	0.8170	0.8184
	MNZ	0.7936	**0.7993**	0.7942	0.7894	0.7838	0.7774
	RR	0.8096	0.8132	0.8149	0.8154	0.8168	**0.8177**
	DCS	0.8162	0.8200	0.8219	0.8234	0.8233	**0.8240**
	LgDCS	0.8146	0.8234	0.8238	0.8277	0.8305	**0.8309**
short	MAX	**0.7801**					
	SUM	0.8089	**0.8115**	0.8089	0.8099	0.8099	0.8105
	MNZ	**0.7951**	0.7949	0.7909	0.7827	0.7797	0.7693
	RR	0.7941	0.7981	0.8000	0.8016	0.8036	**0.8050**
	DCS	0.8058	0.8097	0.8120	0.8142	0.8137	**0.8144**
	LgDCS	0.8093	0.8175	0.8194	0.8239	**0.8254**	0.8248

5.3.1. *Are we able to identify the committee?*

Table 2 presents the performance of the different aggregation functions when focusing on identifying the committee related to the query. From these data, various conclusions can be obtained. Firstly, if we consider the most effective value obtained (an NDCG equal to 0.8309 and with the correct committee in the top positions 67.5% of the time), it seems that it is easier to discover a broader concept as the committee than identify the MP.

Since we are focusing on the quality of the different aggregation strategies, we will also analyze their performance under this particular task. We can see that the trends differ from the MP finding problem: there is a slight deterioration in performance by using SUM and MNZ when aggregating a large number of subprofiles, whereas for the other aggregation criteria it again seems better to use large values of $\mathcal{K}$, although with lower improvement. In this case, since the hybrid LgDCS approach with $\mathcal{K} = 100$ is the method that performs best, this will be the criteria used for this task in the rest of the paper.

Regarding profile size, once again trends are similar to those of the MP-finding problem, but it is interesting to note that performance differences between large and small profiles are much smaller when focusing on committees.

5.3.2. *Results for the combination stage*

Once we have considered the feasibility of our approach for finding the committee, we will analyze its use for the MP-finding task. The first thing that we need to determine is the influence of the parameter α on the linear combination. For this purpose, we used LgDCS for aggregating both the scores of a committee and the scores of the MP, with $\mathcal{K} = 100$. In Table 3 we show how the performance of the system changes with the different values of α. These data have been obtained by

Table 3. NDCG@10 values for the different values of α when combining both committee and profiles ranking. These data have been obtained from Partition A.

α	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
	0.443	0.476	0.485	0.488	0.494	0.501	0.506	0.517	0.509	0.499	0.483

Table 4. Results for the different fusion strategies when using the two-step approach for long and short committee-based profiles, $\alpha = 0.7$.

Committee	*Aggregation*	10	20	30	50	70	100
long	MAX	**0.5179**					
	SUM	**0.5179**	0.5129	0.5016	0.4705	0.4422	0.4168
	MNZ	**0.5176**	0.5082	0.4848	0.4097	0.3355	0.3071
	RR	0.5185	0.5199	0.5206	0.5218	0.5219	**0.5231**
	DCS	0.5202	0.5217	0.5219	0.5232	0.5234	**0.5234**
	LgDCS	0.5204	0.5216	0.5236	0.5264	**0.5266**	0.5257
short	MAX	**0.4921**					
	SUM	**0.4923**	0.4873	0.4749	0.4419	0.4162	0.3930
	MNZ	**0.4917**	0.4817	0.4564	0.3793	0.3149	0.2826
	RR	0.4924	0.4935	0.4937	0.4946	0.4950	**0.4951**
	DCS	0.4931	0.4941	0.4943	**0.4952**	0.4950	0.4949
	LgDCS	0.4936	0.4953	0.4951	**0.4968**	0.4960	0.4951

considering only one partition of the data set. It should be noted that if we rank the MPs only according to the committee score ($\alpha = 0$), worse results are obtained.

For the rest of this experimentation, we set $\alpha = 0.7$ (this is the best alternative in our previous experimentation). The objective now is to determine the influence of the committee on the final ranking obtained after applying the different fusion strategies. From the results presented in Table 4 some conclusions can be drawn. The first is that using committee information improves the quality of the subprofile ranking, i.e. those relevant MPs are promoted to top positions in the ranking and this is therefore an approach that should be considered when tackling an MP finding problem (but also in related problems where broader topics should first be determined). This ranking improvement allows us to obtain better results for all the fusion strategies, regardless of profile size and parameter $\mathcal{K}$. In particular, by focusing on the best results, the use of committee information enables us to achieve improvements in the range of 6% to 8% in each configuration. Additionally, and as expected, the trends of the different aggregation strategies are the same as those obtained in Section 4.3.

6. Conclusions and Further Work

In this paper, we explore the application of expert finding in political domains. We believe that this novel application area is worth studying because it is something that society has been demanding and also because there are peculiarities in the domain that deserve to be considered. Our approach is based on the use of MP profiles which are obtained from their speeches and presents a comparison between monolithic (built from all of an MP's initiatives) and structured ones (constructed separately for each type of committee). These profiles are used as documents in the context of recommending MPs to users who formulate queries. We showed that the use of weighted vector-based profiles perform differently when expert and committee finding tasks are considered, mainly due to the distribution of relevant and non-relevant candidates in the ranking. This distribution clearly affects the performance of the different aggregation strategies.

In this sense, we might conclude that our proposal, CombLgDCS, is a strategy that should be considered as a good alternative when considering the expert-finding task. We have seen that the efficiency of the CombMax method relies on the system's ability to find one relevant candidate in the top positions, and that the effectiveness of CombSum and CombMNZ varies greatly with the distribution of relevant and non-relevant elements in the ranking. CombLgDCS, meanwhile, has proved to be a robust, stable alternative under different conditions, and since it benefits from the use of large $\mathcal{K}$ values, it is not necessary to tune this parameter as classical alternatives[1,26] require. Another important contribution of this paper is the two-stage model for expert finding that, in a first step, takes advantage of the feasibility of identifying broader topics (those represented by a committee). This knowledge is used in a second step to re-rank the subprofiles, thereby improving the performance of MP-finding.

These approaches could be clearly exported to any other field. The only requirement is that the candidates are organized into categories (or knowledge areas) from which the subprofiles could be built. Otherwise, if these are not available, they could be obtained using the clustering techniques. These are some of our future lines of work and will include exploring the use of clustering algorithms to automatically create the subprofiles and validating our proposal in other domains.

Acknowledgments

This work has been funded by the Spanish Ministerio de Economía y Competitividad under projects TIN2013-42741-P and TIN2016-77902-C3-2-P, and the European Regional Development Fund (ERDF-FEDER).

References

1. K. Balog, Y. Fang, Maarten de Rijke, P. Serdyukov and L. Si, Expertise retrieval, *Found. Trends Inf. Retr.* **6**(2–3) (2012) 127–256.
2. K. Hofmann, K. Balog. T. Bogers and M. de Rijke, Contextual factors for finding similar experts, *J. Am. Soc. Inform. Sci. Techn.* **61**(5) (2010) 994–1014.
3. K. Balog, People search in the enterprise, Ph.D Thesis (University of Amsterdam, 2008).
4. K. Balog and M. de Rijke, Determining expert profiles (with an application to expert finding), in *Proc. 20th Int. Joint Conf. Artif. Intell.* (*IJCAI07*) (2007), pp. 2657–2662.

5. A. Viavacqua and H. Lieberman, Agents to assist in finding help, *Proc. ACM SIGCHI Conf. Human Factors in Comp. Systems* (2000), pp. 65–72.
6. R. Berends, M. De Rijke, K. Balog, T. Bogers and A. Van den Bosch, On the assessment of expertise profiles, *J. Am. Soc. Inform. Sci. Tech.* **64**(10) (2013) 2024–2044.
7. L. M. de Campos, J. M. Fernández-Luna and J. F. Huete, Comparing monolithic and committee-based profiles for politician recommendation, in *Proc. 4th Spanish Conf. Information Retrieval* (ACM, New York, 2016).
8. H. Deng, I. King and M. R. Lyu, Formal models for expert finding on DBLP bibliography data, in *Proc. 8th IEEE Int. Conf. Data Mining* (2008), pp. 163–172.
9. S. D. Gollapalli, P. Mitra and C. L. Giles, Similar researcher search in academic environments, in *Proc. 12th Joint Conf. Digital Libraries* (*JCDL '12*) (ACM, 2012), pp. 167–170.
10. C. Moreira, P. Calado and B. Martins, Learning to rank academic experts in the DBLP dataset, *Expert Systems* **32**(4) (2013) 477–493.
11. Y. Cao, J. Liu, S. Bao, and H. Li, Research on expert search at enterprise track of trec 2005, in *Proc. TREC* (2005).
12. I. Soboroff, A. de Vries, and N. Craswell, Overview of the TREC-2006 enterprise track, in *Proc. TREC* (2006).
13. H.-H. Chen, L. Gou, X. Zhang and C. L. Giles, Collabseer: a search engine for collaboration discovery, in *Proc. 11th Joint Conf. Digital Libraries* (*JCDL '11*) (2011).
14. N. Aswell, D. Hawking, A. M. Vercoustre and P. Wilkins, Panoptic expert: searching for experts not just for documents, in *Ausweb Poster Proc.* (2001).
15. S. Gauch, M. Speretta, A. Chandramouli and A. Micarelli, User profiles for personalized information access, in *The Adaptative Web*, Lecture Notes in Computer Science, Vol. 4321 (2007) pp. 54–89.
16. M. Balabanovic and Y. Shoham, Fab: Content-based collaborative recommendations, *Commun. ACM* **40** (1997) 66–72.
17. I. Cantador, A. Bellogin and D. Vallet, Content-based recommendation in social tagging systems, in *Proc. 4th ACM Conf. Recommender systems* (2010), pp. 237–240.
18. M. Bilenko and M. Richardson, Predictive client-side profiles for personalized advertising, in *Proc. 17th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining* (2011), pp. 413–421.
19. M. Montebello, W. Gray and S. Hurley, A personal evolvable advisor for WWW knowledge-based systems, in *Proc. Int. Database Engineering and Application Symposium* (1998), pp. 224–233.
20. L. Chen and K. Sycara, A personal agent for browsing and searching, in *Proc. 2nd Int. Conf. Autonomous Agents* (1998), pp. 132–139.
21. G. L. Somlo and A. E. Howe, Incremental clustering for profile maintenance in information gathering web agents, in *Proc. Int. Conf. Autonomous Agents* (2001), pp. 262–269.
22. D. H. Widyantoro, J. Yin, M. El Nasr, L. Yang, A. Zacchi and J. Yen, Alipes: a swift messenger in cyberspace, in *Proc. AAAI Spring Symposium on Intelligent Agents in Cyberspace* (1999), pp. 62–67.
23. M. Pazzani, J. Muramatsu and D. Billsus, Syskill and Webert: identifying interesting web sites, in *Proc. 13th National Conf. Artificial Intelligence* (1996), pp. 54–61.
24. S. Wu, *Data Fusion in Information Retrieval* (Springer, 2012).
25. C. Macdonald and I. Ounis, Voting techniques for expert search, *Knowledge and Information Systems* **16**(3) (2008) 259–280.
26. W. Weerkamp, K. Balog, and M. de Rijke, Blog feed search with a post index, *Information Retrieval* **14** (2011) 515–545.

27. J. Jarvelin and J. Kekalainen, Cumulative gain-based evaluation of IR techniques, *ACM Trans. Information Systems* **20** (2002) 422–446.
28. K. S. Jones, S. Walker and S. E. Robertson, A probabilistic model of information retrieval: development and comparative experiments, *Inf. Process. Manage.* **36**(6) (2000) 779–808.