

# Cluster Generation and Labeling for Web Snippets: A Fast, Accurate Hierarchical Solution

Filippo Geraci, Marco Pellegrini, Marco Maggini, and Fabrizio Sebastiani

**Abstract.** This paper describes Armil, a meta-search engine that groups the web snippets returned by auxiliary search engines into disjoint labeled clusters. The cluster labels generated by Armil provide the user with a compact guide to assessing the relevance of each cluster to his/her information need. Striking the right balance between running time and cluster well-formedness was a key point in the design of our system. Both the clustering and the labeling tasks are performed on the fly by processing only the snippets provided by the auxiliary search engines, and they use no external sources of knowledge. Clustering is performed by means of a fast version of the furthest-point-first algorithm for metric  $k$ -center clustering. Cluster labeling is achieved by combining intra-cluster and inter-cluster term extraction based on a variant of the information gain measure. We have tested the clustering effectiveness of Armil against Vivisimo, the *de facto* industrial standard in web snippet clustering, using as benchmark a comprehensive set of snippets obtained from the Open Directory Project hierarchy. According to two widely accepted “external” metrics of clustering quality, Armil achieves better performance levels by 10%. We also report the results of a thorough user evaluation of both the clustering and the cluster labeling algorithms. On a standard desktop PC (AMD Athlon 1-Ghz Clock with 750 Mbytes RAM), Armil performs clustering and labeling altogether of up to 200 snippets in less than one second.

## 1. Introduction

An effective search interface is a fundamental component in a web search engine. In particular, the quality of presentation of the search results often represents one of the main keys to the success of such systems. Most search engines present the results of a user query as a ranked list of web snippets. Ranking algorithms play a crucial role in this approach, since users usually browse at most the

ten top-ranked items. Snippet quality is also an important issue, since good-quality snippets allow the user to determine whether the referred pages match or not his/her information need. In order to provide a useful hint about the real content of the page, a web snippet includes both the page title and a short text fragment, which often displays the query terms in context. *Meta-search engines* (MSEs) integrate the items obtained from multiple “auxiliary” search engines, with the purpose of increasing the coverage of the results. However, without an accurate design, MSEs might, in principle, worsen the quality of the information-access experience, since the user is typically confronted with an even larger set of results (see [Meng et al. 02] for a recent survey of challenges and techniques related to building meta-search engines). Thus, key issues to be faced by MSEs concern the exploitation of effective algorithms for merging the ranked lists of results retrieved by the different auxiliary search engines (while at the same time removing the duplicates) and the design of advanced user interfaces based on a structured organization of the results, so as to help the user to focus on the most relevant subset of results. This latter aspect is usually implemented by grouping the results into homogeneous groups by means of clustering or categorization algorithms.

In the context of the World Wide Web, clustering is not only useful for meta-searching. For example, regular search engines do clustering to some extent to avoid showing the user too many semantically-equivalent documents in the first page of results. This activity is close to the classical *duplicate* or *near-duplicate* detection in information retrieval, and one can take advantage of the fact that duplicates and near duplicate are easy to detect via multiple hashing or shingling techniques, and these tasks can be carried out, at least in part, in an offline setting (see, e.g., [Haveliwala et al. 00]). Moreover, in the duplicate detection activity labeling is not an issue.

This paper describes the Armil system,<sup>1</sup> a meta-search engine that organizes the web snippets retrieved from auxiliary search engines into disjoint clusters and automatically constructs a title label for each cluster by using only the text excerpts available in the snippets. Our design efforts were directed towards devising a fast clustering algorithm able to yield good-quality homogeneous groups and a distillation technique for selecting appropriate and useful labels for the clusters. The speed of the two algorithms was a key issue in our design, since the system must organize the results on the fly, thus minimizing the latency between the issuing of the query and the presentation of the results. Second-level clustering is also performed automatically at query time (i.e., not on demand)

---

<sup>1</sup>An armillary sphere (also known as a spherical astrolabe, armilla, or *armil*) is a navigation tool in the form of a model of the celestial sphere, invented by Eratosthenes in 255 BC. Renaissance scientists and public figures were often portrayed with one hand on an armil, since it represented the height of wisdom and knowledge (see [http://en.wikipedia.org/wiki/Armillary\\_sphere](http://en.wikipedia.org/wiki/Armillary_sphere)). The Armil query interface can be freely accessed and used at <http://armil.iit.cnr.it/>.

to minimize latency. In Armil, an equally important role is played by the clustering component and by the labeling component. Clustering is accomplished by means of an improved version of the furthest-point-first (FPF) algorithm for  $k$ -center clustering. To the best of our knowledge, this algorithm had never been used in the context of web snippet clustering or text clustering. The generation of the cluster labels is instead accomplished by means of a combination of intra-cluster and inter-cluster term extraction, based on a modified version of the information gain measure. This approach tries to capture the most significant and discriminative words for each cluster.

One key design feature of Armil is that it relies almost only on the information returned by the auxiliary search engines, i.e., the snippets; this means that no substantial external source of information, such as ontologies or morphological and syntactic linguistic resources, is used. We use just stop-word lists, stemmers, and language recognition tools. We thus demonstrate that such a lightweight approach, together with carefully crafted algorithms, is sufficient to provide a useful and successful clustering-plus-labeling service. Obviously, this assumption relies on the hypothesis that the quality of the snippets returned by the auxiliary search engines is satisfactory.

We have tested the clustering effectiveness of Armil against Vivisimo, the *de facto* industrial standard in web snippet clustering, using as benchmark a comprehensive set of snippets obtained from the Open Directory Project hierarchy. According to two metrics of clustering quality that are normalized variants of the Entropy and the Mutual Information [Cover and Thomas 91], Armil achieves better performance levels by 10%. Note that, since the normalization reduces the ranges of these measures in the interval  $[0, 1]$ , an increase of 10% is noteworthy.

We have tested the clustering effectiveness of Armil against variants of the classical  $k$ -means clustering algorithm. Our method outperforms the  $k$ -means clustering algorithm in time-quality trade-offs.

We also report the results of a thorough user evaluation of both the clustering and the cluster labeling algorithms.

### 1.1. Outline of the Clustering Algorithm

Clustering and labeling are both essential operations for a web snippet clustering system. However, each such system previously proposed strikes a different balance between these two aspects. Some systems (e.g., [Ferragina and Gulli 05, Lawrie and Croft 03]) view label extraction as the primary goal and clustering as a by-product of the label extraction procedure. Other systems (e.g., [Kummamuru et al. 04, Zamir et al. 97]) view instead the formation of clusters as the most important step, and the labeling phase is considered strictly dependent on the clusters found. We have followed this latter approach. In order to cluster the snippets in the returned lists, we map them into a vector space endowed with a distance function, which is a metric; then a modified furthest-point-first

(M-FPF) algorithm is applied to generate the clusters. The M-FPF algorithm generates the same clusters of the “standard” FPF algorithm [Gonzalez 85, Hochbaum and Shmoys 85], but it uses filters based on the triangular inequality to speed up the computation. As such, M-FPF inherits a very important property of the FPF algorithm, i.e., it is within a factor 2 of the optimal solution for the  $k$ -center problem [Gonzalez 85, Hochbaum and Shmoys 85]. In other words, for any fixed number  $k$ , M-FPF produces a  $k$ -clustering (i.e., a partition of the items into  $k$  non-overlapping clusters) such that the maximum cluster diameter is at most twice as large as that of the “optimal”  $k$ -clustering (i.e., the one that minimizes such maximum diameter). The competitive property of M-FPF is even stronger: the approximation factor of 2 cannot be improved with any polynomial approximation algorithm, unless  $P = NP$  [Hsu and Nemhauser 79]. The strength of this formal property has been our main motivation for selecting M-FPF as the algorithmic backbone for Armil. The second interesting property of M-FPF is that it does not compute centroids of clusters. Centroids tend to be dense vectors, and, as such, their computation and/or update in high-dimensional space is a computational burden.<sup>2</sup> M-FPF relies instead only on pairwise distance calculations between snippets and as such better exploits the sparsity of the snippet vector representations.

## 1.2. Outline of the Cluster Labeling Algorithm

The cluster-labeling phase aims at extracting from the set of snippets assigned to each cluster a sequence of words highly descriptive of the corresponding group of items. The quality of the label depends on its well-formedness (i.e., whether the text is syntactically and semantically plausible), on its descriptive power (i.e., how well it describes what is contained in the cluster), and on its discriminative power (i.e., how well it differentiates what is contained in the cluster with respect to what is contained in other clusters).<sup>3</sup> The possibility to extract good labels directly from the available snippets is strongly dependent on their quality and, obviously, on the homogeneity of the produced clusters. In order to pursue a good trade-off between descriptive and discriminative power, we select candidate words for each cluster by means of  $IG_m$ , a modified version of the *Information Gain* measure [Cover and Thomas 91]. For each cluster,  $IG_m$  allows the selection of those words that are the most representative of its contents and are the least representative of the contents of the other clusters. Finally, in order to

<sup>2</sup>The clustering literature also discusses the notion of a cluster “medoid”; similarly to a centroid, a cluster medoid plays the role of cluster representative, but typically has a sparse representation. Unfortunately, the methods proposed in the literature for finding high-quality medoids are not compatible with the real-time nature of the envisioned online web service [Zhang et al. 96]. In this paper we will give a fast and effective medoid searching technique.

<sup>3</sup>These requirements are akin to some of the desirable properties advocated in [Kumnamuru et al. 04] for taxonomic trees, in particular, property 3 (*Sibling nodes distinctiveness*) and property 4 (*Node label predictiveness*) [Kumnamuru et al. 04, page 660].

construct plausible labels, rather than simply using the list of the top-scoring words (i.e., the ones that maximize  $IG_m$ ), the system looks within the titles of the returned webpages for the substring that best matches the selected top-scoring words.

Once each cluster has been assigned a set of descriptive and discriminative keywords (we call such set of keywords the cluster *signatures*), all the clusters that share the same signature are merged. This reduces the arbitrariness inherent in the choice of their number  $k$ , that is fixed *a priori* independently of the query.

### 1.3. Outline of the Paper

The paper is organized as follows. In Section 2 we review related work on techniques for the automatic reorganization of search results. Section 3 introduces the data representation adopted within Armil and sketches the properties of the M-FPF clustering algorithm and of the cluster labeling algorithm. In Section 4 the architecture of Armil is described in detail. The results of the system evaluation are reported in Sections 5, 6, 7, and 8. Finally, in Section 9 conclusions and prospective future research are discussed.

## 2. Previous Work

Tools for clustering web snippets have received attention in the research community. In the past, this approach has had both critics [Kural et al. 99, Kural et al. 93] and supporters [Tombros et al. 02], but the proliferation of commercial web services such as Copernic, Dogpile, Groxis, iBoogie, Kartoo, Mooter, and Vivisimo seems to confirm the potential validity of the approach. Academic research prototypes are also available, such as Grouper [Zamir and Etzioni 98, Zamir et al. 97], EigenCluster [Cheng et al. 03], Shoc [Zhang and Dong 04], and SnakeT [Ferragina and Gulli 05]. Generally, details of the algorithms underlying commercial web services are not in the public domain.

Maarek et al. give a precise characterization of the challenges inherent in web snippet clustering and propose an algorithm based on complete-link hierarchical agglomerative clustering that is quadratic in the number  $n$  of snippets. They introduce a technique called “lexical affinity” whereby the co-occurrence of words influences the similarity metric [Maarek et al. 00].

Zeng et al. tackle the problem of detecting good cluster names as preliminary to the formation of the clusters, using a supervised learning approach [Zeng et al. 04]. Note that the methods considered in our paper are instead all unsupervised, thus requiring no labeled data.

The EigenCluster [Cheng et al. 03], Lingo [Osinski and Weiss 04], and Shoc [Zhang and Dong 04] systems all tackle web snippet clustering by performing

a singular value decomposition (SVD) of the term-document incidence matrix<sup>4</sup>; the problem with this approach is that SVD is extremely time-consuming, hence problematic when applied to a large number of snippets. By testing a number of queries on Eigencluster, we have observed that, when operating on many snippets (roughly 400), a reasonable response time (under one second) is attained by limiting the number of generated clusters to a number between five and ten and by avoiding a clustering decision for over 50% of the data. Zamir and Etzioni propose a web snippet clustering mechanism (Suffix Tree Clustering (STC)), based on suffix arrays [Zamir and Etzioni 98, Zamir et al. 97] and experimentally compare STC with algorithms such as  $k$ -means, single-pass  $k$ -means [MacQueen 67], Backshot and Fractionation [Cutting et al. 92], and Group Average Hierarchical Agglomerative Clustering (GAHAC). They test the systems on a benchmark obtained by issuing ten queries to the Metacrawler meta-search engine, retaining the top-ranked 200 snippets for each query, and manually tagging the snippets by relevance to the queries. They then compute the quality of the clustering obtained by the tested systems by ordering the generated clusters according to precision, and by equating the effectiveness of the system with the average precision of the highest-precision clusters that collectively contain 10% of the input documents. This methodology had been advocated [Hearst and Pedersen 96] and is based on the assumption that the users will anyway be able to spot the clusters most relevant to their query. Average precision as computed with this method ranges from 0.2 to 0.4 for all the algorithms tested (STC coming out on top in terms of both effectiveness and efficiency). Interestingly, the authors show that very similar results are attained when full documents are used instead of their snippets, thus validating the snippet-based clustering approach.

Strehl et al. experiment with four similarity measures (cosine similarity, Euclidean distance, Pearson correlation, and extended Jaccard) in combination with five algorithms ( $k$ -means, self-organizing maps, random clustering, min-cut hyper-graph partitioning, and min-cut weighted graph partitioning). Two data sets were used, one formed by 2340 documents pre-classified in 20 categories from the news repository of Yahoo!, the second formed by 966 documents pre-classified in ten categories from the Business Directory of Yahoo!. Overall quality is measured in terms of normalized mutual information. For a specific clustering, also the quality of the single clusters in terms of single cluster purity and single cluster entropy are given. The comparisons are made only in terms of output quality, computing time not being considered relevant in this setting. The highest-quality results are obtained via cosine similarity and Jaccard distance combined with min-cut graph partitioning [Strehl et al. 00]. Our experiments have confirmed that for snippets better results are obtained using variants of Jaccard distance, with respect to standard tf-idf and cosine similarity.

---

<sup>4</sup>At the time of this writing, the Eigencluster system is available online at <http://eigencluster.csail.mit.edu/>.

Lawrie and Croft view the clustering/labeling problem as that of generating multilevel summaries of the set of documents (in this case the web snippets returned by a search engine). The technique is based on first building offline a statistical model of the background language (e.g., the statistical distribution of words in a large corpus of the English language) and on subsequently extracting “topical terms” from the documents, where “topicality” is measured by the contribution of a term to the Kullback-Leibler (KL) divergence score of the document collection relative to the background language. Intuitively, this formula measures how important this term is in measuring the distance of the collection of documents from the distribution of the background language. Additionally, the “predictiveness” of each term is measured. Intuitively, predictiveness measures how close a term appears (within a given window size) to other terms. In the summaries, terms of high topicality and high predictiveness are preferred. The proposed method is shown to be superior (by using the KL-divergence) to a naive summarizer that just selects the terms with highest  $tf \cdot idf$  score in the document set [Lawrie and Croft 03].

Kammamuru et al. propose a classification of web snippet clustering algorithms into *monothetic* (in which the assignment of a snippet to a cluster is based on a single dominant feature) and *polythetic* (in which several features concur in determining the assignment of a snippet to a cluster). The rationale for proposing a monothetic algorithm is that the single discriminating feature is a natural label candidate. The authors propose such an algorithm (a type of greedy cover) in which the snippets are seen as sets of words and the next term is chosen so as to maximize the number of newly covered sets while minimizing the hits with already covered sets. The paper reports empirical evaluations and user studies over two classes of queries, “ambiguous” and “popular.” The users were asked to compare three clustering algorithms over the set of queries and, for each query, were asked to answer six questions of a rather general nature on the generated hierarchy (not on the single clusters) [Kumamuru et al. 04].

Ferragina and Gulli propose a method for hierarchically clustering web snippets and produce a hierarchical labeling based on constructing a sequence of labeled and weighted bipartite graphs representing the individual snippets on one side and a set of labels (and corresponding clusters) on the other side. Data from the Open Directory Project (ODP)<sup>5</sup> is used in an offline and query-independent way to generate predefined weights that are associated online to the words of the snippets returned by the queries. Data is collected from 16 search engines as a result of 77 queries chosen for their popularity among Lycos and Google users in 2004. The snippets are then clustered, and the labels are manually tagged as relevant or not relevant to the cluster to which they have been associated. The clusters are ordered in terms of their weight, and quality is measured in terms of the number of relevant labels among the first  $n$  labels, for  $n \in \{3, 5, 7, 10\}$

---

<sup>5</sup><http://www.dmoz.org/>



[Ferragina and Gulli 05]. Note that in this work the emphasis is on the quality of the labels rather than on that of the clusters (although the two concepts are certainly related), and that the ground truth is defined “a posteriori,” after the queries are processed.

### 3. The Clustering Algorithm and the Labeling Algorithm

We now describe in detail the methods used by Armil for web snippet clustering and cluster labeling.

#### 3.1. The Distance Function

Each snippet is turned into a “bag of words” after removing stop words and performing stemming. In [Geraci et al. 06b] we report experiments using, as a distance function, (a) the cosine distance measure (i.e., the complement to 1 of the cosine similarity function) applied to vectors of terms weighted by  $tf \cdot idf$ ,<sup>6</sup> and (b) a slight modification of the standard Jaccard distance, which we call *weighted Jaccard distance* (WJD); in those experiments, (b) has performed at the same level of accuracy as (a) but has proven much faster to compute.

In this journal version we improve on the preliminary results of [Geraci et al. 06b] by using the generalized Jaccard distance described in [Charikar 02]. Given two “bag-of-words” snippet vectors  $s_1 = (s_1^1, \dots, s_1^h)$  and  $s_2 = (s_2^1, \dots, s_2^h)$ , the *generalized Jaccard distance* (GJD) is

$$D(s_1, s_2) = 1 - \frac{\sum_i \min(s_1^i, s_2^i)}{\sum_i \max(s_1^i, s_2^i)}.$$

The term weights  $s_a^i$  consist of “weighted term frequencies,” obtained as weighted sums of the numbers of occurrences of the term in the snippet, where we take advantage of the structure of the snippets by weighting different parts of the snippet (page title, text fragment, URL) differently; more precisely, weight 3 is assigned to a term occurring in the page title, weight 1 to a term occurring in the text fragment, and weight 0 is assigned to a term occurring in the URL (since, in previous experiments we had run, the text of the URL had proven to give no contribution in terms of cluster quality). Note that, when using unit weights only, the generalized Jaccard distance coincides with the standard Jaccard distance. In [Charikar 02] it is proved that GJD is a metric.

<sup>6</sup>The acronym  $tf \cdot idf$  stands for *term frequency* and *inverse document frequency*. It is a standard technique for assigning real vectors to documents belonging to a set of documents (corpus) based on the number of times a term appears in a document and on the number of documents of the corpus containing that term (see, e.g., [Baeza-Yates and Ribeiro-Neto 99, page 29]).



### 3.2. The Clustering Algorithm

We approach the problem of clustering web snippets as that of finding a solution to the classic  $k$ -center problem:

**Problem 3.1.** Given a set  $S$  of points in a metric space  $M$  endowed with a metric distance function  $D$ , and given a desired number  $k$  of resulting clusters, partition  $S$  into non-overlapping clusters  $C_1, \dots, C_k$  and determine their “centers”  $\mu_1, \dots, \mu_k \in M$  so that the radius  $\max_j \max_{x \in C_j} D(x, \mu_j)$  of the widest cluster is minimized.

The  $k$ -center problem can be solved approximately using the furthest-point-first (FPF) algorithm [Gonzalez 85, Hochbaum and Shmoys 85], which we now describe. Given a set  $S$  of  $n$  points, FPF builds a sequence  $T_1 \subset \dots \subset T_k = T$  of  $k$  sets of “centers” (with  $T_i = \{\mu_1, \dots, \mu_i\} \subset S$ ) in the following way:

1. At the end of iteration  $i - 1$ , FPF holds the mapping  $\mu$  defined for every point  $p_j \in S \setminus T_{i-1}$  as

$$\mu(p_j) = \arg \min_{\mu_s} D(p_j, \mu_s),$$

i.e., the center in  $T_{i-1}$  closest to  $p_j$ ;  $\mu(p_j)$  is called the *leader* of  $p_j$ . Note that this mapping is established in the first iteration in time  $O(n)$ .

2. At iteration  $i$ , among all points  $p_j$ , FPF picks

$$\mu_i = \arg \max_{p_j} D(p_j, \mu(p_j)),$$

i.e., the point for which the distance to its leader is maximum, and makes it a new center, i.e., adds it to  $T_{i-1}$ , thus obtaining  $T_i$ . This selection costs  $O(n)$ .

3. Compute the distance of  $\mu_i$  to any point in  $S \setminus T_i$ , and update the mapping  $\mu$  if needed. Thus,  $\mu$  is now correct for the beginning of iteration  $i + 1$ . This update phase costs  $O(n)$ .

The final set of centers  $T = \{\mu_1, \dots, \mu_k\}$  defines the resulting  $k$ -clustering, since each center  $\mu_i$  implicitly identifies a cluster  $C_i$  as the set of data points whose leader is  $\mu_i$ . Note that  $T_1$  is initialized to contain a single point chosen at random from  $S$ ; this random choice is due to the fact that, in practice, both the effectiveness and the efficiency of the algorithm can be seen experimentally to be rather insensitive to this choice.

Most of the computation is actually devoted to computing distances and updating the auxiliary mapping  $\mu$ : this takes  $O(n)$  time per iteration, so the total

computational cost of the algorithm is  $O(nk)$ . Here we define an improved version of this algorithm that exploits the triangular inequality in order to filter out useless distance computations. This modified algorithm (M-FPF), which we now describe, works in any metric space, hence in any vector space.<sup>7</sup>

Consider, in the FPF algorithm, any center  $\mu_x \in T_i$  and its associated set of closest points  $N(\mu_x) = \{p_j \in S \setminus T_i \mid \mu(p_j) = \mu_x\}$ . We store  $N(\mu_x)$  as a ranked list, in order of decreasing distance from  $\mu_x$ . When a new center  $\mu_y$  is added to  $T_i$ , in order to identify its associated set of closest points  $N(\mu_y)$  we scan every  $N(\mu_x)$  in decreasing order of distance, and stop scanning when, for a point  $p_j \in N(\mu_x)$ , it is the case that  $D(p_j, \mu_x) \leq \frac{1}{2}D(\mu_y, \mu_x)$ . By the triangular inequality, any point  $p_j$  that satisfies this condition cannot be closer to  $\mu_y$  than to  $\mu_x$ . This rule filters out from the scan points whose leader cannot possibly be  $\mu_y$ , thus significantly speeding up the identification of leaders. Note that all distances between centers in  $T_i$  must be available; this implies an added  $O(k^2)$  cost for computing and maintaining these distances, which is anyhow dominated by the term  $O(nk)$ . The gain is that, in practice, fewer than  $n$  points need to be scanned at each iteration.

### 3.3. Using Medoids

Motivated by the aim of attaining high-quality clusters efficiently, we have found that we can improve upon a direct application of the M-FPF algorithm by adopting the strategy described below. The M-FPF is applied not to the whole set but only to a random sample of size  $\sqrt{nk}$  of the input points (this sample size is suggested in [Indyk 99]). Afterwards, the remaining points are associated to the closest (according to the generalized Jaccard distance) center. We obtain improvements in quality by making an iterative update of the “center” when a new point is associated to a cluster. Within a cluster  $C_i$  we find the point  $a_i$  furthest from  $\mu_i$  and the point  $b_i$  furthest from  $a_i$  (intuitively, this is a good approximation to a diametral pair). The medoid  $m_i$  is the point in  $C_i$  that has the minimum value of the function

$$|D(a_i, x) - D(b_i, x)| + |D(a_i, x) + D(b_i, x)|,$$

over all  $x \in C_i$ .<sup>8</sup> When we add a new point  $p$  to  $C_i$ , we check if the new point should belong to the approximate diametral pair  $(a_i, b_i)$ , that is, if the distances  $D(a_i, p)$  or  $D(p, b_i)$  are larger than  $D(a_i, b_i)$ , and if so we update the approximate diametral pair and  $m_i$  accordingly. The association of the remaining points is done with respect to the medoids rather than the centers. The application of M-FPF on a sample plus the iterative recomputation of medoids gave us a clustering of better quality than simply using M-PFP on the whole input set.

<sup>7</sup>We recall that any vector space is also a metric space, but not vice versa.

<sup>8</sup>This formula mimics in a discrete setting the task of finding the cluster point closest to the median point to the segment  $(a_i, b_i)$  in an Euclidean space.

### 3.4. The Candidate Words Selection Algorithm

We select candidate terms for labeling the generated clusters through a modified version of the *information gain* function [Cover and Thomas 91, Yang and Pedersen 97]. Let  $x$  be a document taken uniformly at random in the set of documents  $D$ ;  $P(t)$  is the probability that  $x$  contains term  $t$ , and  $P(c)$  is the probability that  $x$  is in category  $c$ . The complementary events are denoted  $P(\bar{t}) = 1 - P(t)$  and  $P(\bar{c}) = 1 - P(c)$ ;  $P(t, c)$  is the probability that  $x$  is in category  $c$  and contains term  $t$ ,  $P(\bar{t}, \bar{c})$  is the probability that  $x$  does not contain  $t$  and is not in category  $c$ ,  $P(t, \bar{c})$  is the probability that  $x$  contains  $t$  but is not in category  $c$ , and  $P(\bar{t}, c)$  is the probability that  $x$  does not contain  $t$  and is in category  $c$ . Clearly, these events being mutually disjoint, it holds that  $P(t, c) + P(\bar{t}, \bar{c}) + P(t, \bar{c}) + P(\bar{t}, c) = 1$ . The information gain [Cover and Thomas 91, Yang and Pedersen 97] is the contribution of the four terms:

$$IG(t, c) = \sum_{a \in \{t, \bar{t}\}} \sum_{b \in \{c, \bar{c}\}} P(a, b) \log \frac{P(a, b)}{P(a)P(b)}.$$

Intuitively,  $IG$  measures the amount of information that each argument contains about the other; when  $t$  and  $c$  are independent,  $IG(t, c) = 0$ . This function is often used for feature selection in text classification, where, if  $IG(t, c)$  is high, the presence or absence of a term  $t$  is deemed to be highly indicative of the membership or non-membership in a category  $c$  of the document containing it. In the text classification context, the rationale of including in the sum, aside from the factor that represents the “positive correlation” between the arguments (i.e., the factor  $P(t, c) \log \frac{P(t, c)}{P(t)P(c)} + P(\bar{t}, \bar{c}) \log \frac{P(\bar{t}, \bar{c})}{P(\bar{t})P(\bar{c})}$ ), also the factor that represents their “negative correlation” (i.e., the factor  $P(\bar{t}, c) \log \frac{P(\bar{t}, c)}{P(\bar{t})P(c)} + P(t, \bar{c}) \log \frac{P(t, \bar{c})}{P(t)P(\bar{c})}$ ); is that, if this latter factor has a high value, this means that the absence (respectively presence) of  $t$  is highly indicative of the membership (respectively non-membership) of the document in  $c$ . That is, the term is useful anyway, although in a “negative” sense.

However, in our context we are interested in terms that *positively describe* the contents of a cluster and are thus only interested in positive correlation:

$$IG_m(t, c) = P(t, c) \log \frac{P(t, c)}{P(t)P(c)} + P(\bar{t}, \bar{c}) \log \frac{P(\bar{t}, \bar{c})}{P(\bar{t})P(\bar{c})}.$$

The keywords selected via  $IG_m$  will not be shown directly to the user but will form the input to the label-generation phase described in Section 4.

## 4. Overview of the System

We discuss here in more detail the architecture of Armil. Overall the computation flow is a pipeline consisting of (a) data collection and cleaning, (b) first-level clustering, (c) candidate word extraction for labeling, (d) second-level clustering, and (e) cluster labeling. Let us review these steps in order.

(1) **Querying one or more search engines.** The user of Armil issues a query string that is re-directed by Armil to the selected search engines (at the moment the user can select Google and/or Yahoo!). As a result Armil obtains a list (or several lists) of snippets describing webpages that the search engines deem relevant to the query. An important system design issue is deciding the type and number of snippet sources to be used as auxiliary search engines. It is well known that the probability of relevance of a snippet to the user information need quickly decays with the rank of the snippet in the list returned by the search engine. Therefore the need of avoiding low-quality snippets suggests the use of many sources each supplying a low number of high-quality snippets (here, *high-quality snippet* implies both relevance of the snippet to the query and representativeness of the snippet with respect to the associated document). On the other hand, increasing the number of snippet sources raises the pragmatic issue of handling several concurrent threads, the need to detect and handle more duplicates, and the need for a more complex handling of the composite ranking by merging several snippet lists (both globally and within each cluster separately). Since we consider Armil a “proof-of-concept” prototype rather than a full-blown service, we have chosen only two (high-quality) sources, Google and Yahoo!. We limit the total number of snippets to be collected to 200 (80 from Yahoo! and 120 from Google; these numbers optimize the total waiting time). We produce the initial composite ranking of the merged snippet list by a very simple method: by alternatively picking snippets from each source list.

(2) **Cleaning and filtering.** Snippets that are too short or have little informative content (i.e., small number of well-formed words) or are in non-Latin alphabets are filtered out. The input is then filtered by removing non-alphabetic symbols, digits, HTML tags, stop words, and the query terms. The latter are removed since they are likely to be present in every snippet and thus are going to be useless for the purpose of discriminating different contexts. On the other hand query terms are very significant for the user, so they are re-introduced in the label-generation phase (7) described below. We then identify the prevalent language of each snippet, which allows us to choose the appropriate stop word list and stemming algorithm. Currently we use the ccTLD (country code top level domain) of the URL to decide on the prevalent language of a snippet. For the purpose of the experiments we only distinguish between English and Italian. For snippets of English webpages, we use Porter’s stemming algorithm, while for Italian ones we use a simple rule-based stemmer we developed in-house. Currently, no other

languages are supported. More advanced language discovery techniques are well known in the literature, and we plan to apply them in the near future; however, it should be noticed that, given the online requirements, methods too expensive in terms of computational effort should be avoided.

**(3) First-level clustering.** We build a flat  $k$ -clustering representing the first level of the cluster hierarchy, using the M-FPF algorithm with medoids and the generalized Jaccard distance as described in Section 3. An important issue is deciding the number  $k$  of clusters to create. Currently, this number is fixed to 30, but it is clear that the number of clusters should depend on the query and on the number of snippets found. A general rule is difficult to find, also because the “optimal” number of clusters to be displayed is a function of the goals and preferences of the user. Also, while techniques for automatically determining such optimal numbers do exist [Milligan and Cooper 85], their computational cost is incompatible with the real-time nature of our application. Therefore, besides providing a default value, we allow the user to increase or decrease the value of  $k$  to his/her liking. Clusters that contain only one snippet are probably outliers of some sort, and we thus merge them under a single cluster labeled “Other topics.”

**(4) Snippets re-ranking.** A cluster small enough that the list of its snippets fits in the screen does not require a sophisticated order of presentation. However, in general, users are greatly facilitated if the snippets of a cluster are listed in order of their likely importance to the user. Our strategy is to identify a geometric “inner core” of each cluster and geometric “outliers.” In order to achieve this aim, we apply the M-FPF algorithm within each cluster as follows. Since M-FPF is incremental in the parameter  $k$ , we increment  $k$  up to a value for which it happens that the largest obtained cluster has less than half of the point of the input cluster. This specific group forms the “inner core”: all other points are termed “outliers.” The points in the “core” are shown in the listing before the “outliers.” Within core and outliers points we use a relative ranking obtained by a linear combination of the native ranking generated by the auxiliary search engines. Note that this computation is done only to decide the order of presentation of the snippets at the first level. It should not be confused with the second-level clustering described below.

**(5) Candidate words selection.** For each cluster we need to determine a set of candidate words to appear in its label; these will hereafter be referred to as *candidates*. For this purpose, for each word that occurs in the cluster, we sum the weights, as defined in Section 3, of all its occurrences in the cluster and pre-select the ten words with the highest score in each cluster.

We refer to this as the *local candidate selection*, since it is done independently for each cluster. For each of the ten selected terms, we compute  $IG_m$ , as explained in Section 3.4. The three terms in each cluster with the highest score are chosen as candidates. We refer to this as *global candidate selection*, because the

computation of  $IG_m$  for a term in a cluster is dependent also on the contents of the other clusters. Global selection has the purpose of obtaining different labels for different clusters. At the end of this procedure, if two clusters have the same signature, we merge them, since this is an indication that the target number of clusters  $k$  may have been too high for this particular query.<sup>9</sup>

(6) **Second-level clustering.** Although the clustering algorithm could in principle be iterated recursively in each cluster up to an arbitrary number of levels, we limit our algorithm to only two levels, since this is likely to be the maximum depth a user is willing to explore in searching for an answer to his/her information needs.<sup>10</sup> Second-level clustering is applied to first-level clusters of size larger than a predetermined threshold (at the moment this is fixed to ten snippets, excluding duplicates). For second-level clustering we adopt a different approach, since metric-based clustering applied at the second level tends to detect a single large “inner core” cluster and several small “outlier” clusters.<sup>11</sup> The second-level part of the hierarchy is generated based on the candidate words found for each cluster during the first-level candidate word selection. Calling  $K$  the set of three candidate words of a generic cluster, we consider all its subsets as possible signatures for second level clusters. A snippet  $x$  is assigned to a signature  $s$  if and only if all the signature elements are in  $x$  and no candidate in  $K \setminus s$  is in  $x$ . If a signature is assigned too few snippets (e.g., one) it is considered as an outlier and it is not shown to the user in the second level. Also, if most of the snippets at the first level end up associated to a single signature, then the second-level clusters are not shown to the user since the second-level subdivision would not be any more useful than the first-level subdivision.<sup>12</sup>

(7) **Labeling.** Many early clustering systems would use lists of keywords as the output to show to the user. In order to offer a more syntactically pleasant label, we decide to give as output well-formed phrases or parts of phrases as it is becoming standard in more recent systems (e.g., [Ferragina and Gulli 05]). We use the candidate keywords just as a basis for generating well-formed phrases that will be shown to the user as real cluster labels.<sup>13</sup> Given the title of the webpage contained in the snippet, considered as a sequence of words (this time including stop words), we consider all its contiguous subsequences and we give each subsequence a cumulative score as follows: candidates are given a high positive score (its  $IG_m$  score), query words a low positive score (set at 0.1),

<sup>9</sup>More precisely, we consider the two original clusters with the same signature as second-level clusters, and we produce for each a different second-level label based on the non-signature keywords of those clusters.

<sup>10</sup>Vivisimo, for example, uses a two-level hierarchy.

<sup>11</sup>We exploited this phenomenon in the snippet re-ranking step (4).

<sup>12</sup>Since second-level clustering is based on first-level candidate words, here we depart slightly from the overall idea of separating clustering and labeling.

<sup>13</sup>For example, “meaning life” and “meaning of life” hold the same information but the latter is a grammatically well-formed phrase.

and all other words have a negative score (set at  $-0.2$ ). For labeling a cluster, among all its snippets we select the shortest substring of a snippet title among those having the highest score. This computation can be done efficiently using a dynamic programming approach. The choice of positive and negative weights is to ensure the balance of two conflicting goals: include most of the candidates but few of the non-candidate connecting words in a label.

**(8) Exploiting duplicates.** Since Armil collects data from several search engines, it is possible that the same URL (maybe with a different snippet fragment of text) is present in duplicate. We consider this fact as an indication of the importance of the URL. Therefore, duplicates are accounted for in determining weights and distances. Since clustering is based on title and text, it is possible that the same URL ends up in different clusters, for which it is equally relevant. However, if duplicate snippets appear in the same cluster, they are listed only once with an indication of the two sources. Thus, duplicate removal is done just before the presentation to the user.

**(9) User Interface.** The user interface is important for the success of a web-based service. We have adopted a scheme common to many search engines and meta-search engines (e.g., Vivisimo), in which the data are shown in ranked list format in the main frame while the list of cluster labels are presented on the left frame as a navigation tree. The interface also allows the user to select the number of clusters, by increasing or decreasing the default value of 30.

## 5. Anecdotal Evidence

In this section we report an analysis of the output of the system on selected queries. In particular, we highlight some very high quality clusters obtained. These results were obtained in November 2005.

### 5.1. Query: Armstrong

The query “armstrong” returns 174 snippets organized in 28 clusters, with a clustering time of 0.94 seconds (see Figure 1). Cluster 2, labeled “Armstrong Ceilings,” contains 12 snippets, among which five are relative to a company manufacturing ceilings and floors, four to companies manufacturing other hardware equipment (e.g., molds, pumps, tools), and one to companies manufacturing medical equipment. Cluster 3, labeled “Lance Armstrong Foundation,” contains 14 snippets, of which 12 are related to the sport champion. Cluster 4, labeled “Louis Jazz,” contains 20 snippets all related to the well-known jazz musician and are mostly in English, while Cluster 6, with the same label, contains 12 snippets of which nine are relative to the same musician but mostly in Italian and three are relative to other musicians named Armstrong. Cluster 12, labeled



- armstrong (29)
  - armstrong air conditioning and heating (6)
  - ⊕ armstrong ceilings (12)
  - ⊕ lance armstrong foundation (14)
  - ⊕ Louis Jazz (20)
  - steam traps humidification systems (6)
  - ⊕ Louis Jazz (12)
  - ⊕ italia (13)
  - salt lake city bed and (5)
  - armstrong county (4)
  - at law with main law (7)
  - guide to the papers of (6)
  - neil armstrong (7)
  - of the village cheese company (4)
  - armstrong biography (9)
  - george (4)
  - armstrong (2)
  - armstrong williams (8)
  - edwin howard armstrong (6)
  - chain market research and consulting (3)
  - sport primi piani ciclismo (7)
  - armstrong harbourtown records (4)
  - armstrong sul sito (4)
  - grande ciclismo tour de france (4)
  - joe armstrong sics (3)
  - armstrong clan society one hundred (1)
  - activity based costing (2)
  - willkommen bei armstrong (2)
  - tim armstrong (2)
  - Other topics (2)

**Figure 1.** Labels generated by Armil for the query “armstrong.”

“Neil Armstrong,” has seven snippets of which six are related to the well-known astronaut. Cluster 19, labeled “George,” contains four snippets, two of which are related to the life of General George Armstrong Custer. Cluster 18, labeled “Edwin Howard Armstrong,” contains six snippets, three of which are devoted to Edwin Howard Armstrong, the inventor of the frequency modulation radio broadcasting technique.

## 5.2. Query: Madonna

After adjusting manually the default to 20 clusters, the query “madonna” returns 175 snippets organized in 19 clusters in 0.88 seconds (see Figure 2). Cluster 4, labeled “santuario della madonna del divino,” contains nine snippets, of which six are relative to holy shrines. Cluster 5, labeled “su louise veronica madonna ciccone,” contains 11 snippets, all of them high-quality sites devoted to the well-known popstar. Cluster 7, labeled “Madonna di Campiglio,” contains 13 snippets, of which 11 are related to the Italian ski resort. Cluster 10, labeled “music,” contains 51 snippets, most of which related to the popstar.

## 5.3. Query: Allergy

The query “allergy” returns 100 snippets organized in 27 clusters in 0.94 seconds (see Figure 3). Cluster 3, labeled “Asthma,” contains eight snippets, of which six

- madonna (19)
  - madonnaweb (1)
  - madonnashots (6)
  - madonna tickets madonna concerts (4)
  - santuario della madonna del divino (9)
  - ⊕ su louise veronica madonna ciccone (11)
  - album (22)
  - ⊕ madonna di campiglio (13)
  - ⊕ madonna online (11)
  - madonna university (2)
  - ⊕ music (51)
  - ⊕ madonna news (23)
  - pictures picture gallery pics picture (8)
  - black madonna shrine (6)
  - pictures biography discography screensavers videos (3)
  - umbria madonna delle grazie farmhouse (4)
  - madonna dreams interpreting dreams of (2)
  - madonna (3)
  - madonna paperdoll heaven (2)
  - Other topics (2)

**Figure 2.** Labels generated by Armil for the query “madonna.”

related to asthma and allergies. Cluster 5, labeled “Allergy Journal,” contains four snippets, all of which relate to periodic medical publications. Cluster 6, labeled “guide to allergies and allergy,” contains seven snippets, of which at least five offer general advise on allergy treatments. Cluster 11, labeled “food allergy,” contains five snippets, all of which related to food allergies.

- allergy (27)
  - aaaai american academy of allergy (3)
  - hon allergy glossary a (2)
  - asthma (8)
  - allergy the basics (4)
  - allergy journal (4)
  - guide to allergies and allergy (7)
  - allergy discussion group (2)
  - information symptoms medication products treatment (5)
  - vaccini allergie da acaro vaccino (2)
  - food allergy (5)
  - allergy and asthma faq (5)
  - current allergy and clinical immunology (5)
  - allergy clinic diagnosis and treatment (5)
  - allergy sinusitis ws tichenor (1)
  - institute of allergy and infectious (2)
  - hour local news allergy forecast (3)
  - allergy hot list (4)
  - latex allergy (3)
  - products allergy bedding austin air (4)
  - society for allergy and clinical (7)
  - allergy magazine (2)
  - allergy society of south africa (2)
  - agency index of food recalls (3)
  - publication allergy (2)
  - allergy and immunology internet(2)
  - food allergy news(2)
  - Other topics (13)

**Figure 3.** Labels generated by Armil for the query “allergy.”

## 6. Experimental Evaluation of the Clustering Algorithm

We have performed experiments aimed at assessing the performance of the clustering algorithm underlying Armil.

Assessing the quality of a clustering with respect to a pre-classified collection of items (ground truth) is a well-known problem in information retrieval and objective testing methodologies based on information-theoretic principles have been established. We will compare Armil against the state-of-the-art commercial system Vivisimo and against the classical clustering algorithm  $k$ -means (several variants thereof).

### 6.1. The Baseline

As a baseline against which to compare the clustering capabilities of Armil, we have chosen Vivisimo.<sup>14</sup> Vivisimo is considered an industrial standard in terms of clustering quality and user satisfaction, and in 2001 and 2002 it won the “best meta-search award” assigned annually by the online magazine SearchEngineWatch.com. Vivisimo thus represents a particularly difficult baseline, and it is not known if its clustering quality only depends on an extremely good clustering algorithm, or rather on the use of external knowledge or custom-developed resources. To the best of our knowledge, this is the first published experiment comparing on an objective basis (see below) the clustering quality of an academic prototype and Vivisimo. Vivisimo’s advanced searching feature allows a restriction of the considered auxiliary search engines to a subset of a range of possible auxiliary search engines. For the purpose of our experiment, we restrict our source of snippets to the ODP directory.

Vivisimo’s clustering search engine has been the focus of research in the field of Human Computer Interface and Information systems (see, e.g., [Spink et al. 05, Koshman et al. 05]), with an emphasis on understanding user behavior by analyzing query logs.

### 6.2. Measuring Clustering Quality

Following a consolidated practice, in this paper we measure the effectiveness of a clustering system by the degree to which it is able to “correctly” re-classify a set of pre-classified snippets into exactly the same categories without knowing the original category assignment. In other words, given a set  $C = \{c_1, \dots, c_k\}$  of categories and a set  $\Theta$  of  $n$  snippets pre-classified under  $C$ , the “ideal” term clustering algorithm is the one that, when asked to cluster  $\Theta$  into  $k$  groups, produces a grouping  $C' = \{c'_1, \dots, c'_k\}$  such that, for each snippet  $s_j \in \Theta$ ,  $s_j \in c_i$  if and only if  $s_j \in c'_i$ . The original labeling is thus viewed as the latent, hidden structure that the clustering system must discover.

---

<sup>14</sup><http://vivisimo.com/>

The measure we use is *normalized mutual information* (see, e.g., [Strehl 02, Cover and Thomas 91, Meila 03, page 110]):

$$NMI(C, C') = \frac{2}{\log |C||C'|} \sum_{c \in C} \sum_{c' \in C'} P(c, c') \cdot \log \frac{P(c, c')}{P(c) \cdot P(c')},$$

where  $P(c)$  represents the probability that a randomly selected snippet  $s_j$  belongs to  $c$  and  $P(c, c')$  represents the probability that a randomly selected snippet  $s_j$  belongs to both  $c$  and  $c'$ . The normalization, achieved by the  $\frac{2}{\log |C||C'|}$  factor, is necessary in order to account for the fact that the cardinalities of  $C$  and  $C'$  are in general different [Cover and Thomas 91].

Higher values of  $NMI$  mean better clustering quality. The clustering produced by Vivisimo has partially overlapping clusters (in our experiments Vivisimo assigned roughly 27% of the snippets to more than one cluster), but  $NMI$  is designed for non-overlapping clustering. Therefore, in measuring  $NMI$  we eliminate from the ground truth, from the clustering produced by Vivisimo, and from that produced by Armil, the snippets that are present in multiple copies.

However, in order to also consider the ability of the two systems to “correctly” duplicate snippets across overlapping clusters, we have also computed the *normalized complementary entropy* [Strehl 02, Cover and Thomas 91, page 108], in which we have changed the normalization factor so as to take overlapping clusters into account. The entropy of a cluster  $c'_l \in C'$  is

$$E(c'_l, C) = \sum_{k=1}^{|C|} -\frac{|c'_l \cap c_k|}{|c_k|} \log \frac{|c'_l \cap c_k|}{|c_k|}.$$

The normalized complementary entropy of  $c'_l$  is

$$NCE(c'_l, C) = 1 - \frac{E(c'_l, C)}{\log |C|}.$$

The values of  $NCE$  range in the interval  $[0, 1]$ , and a greater value implies better quality of  $c'_l$ . The complementary normalized entropy of  $C'$  is the weighted average of the contributions of the single clusters in  $C'$ . Let  $n' = \sum_{l \in 1}^{|C'|} |c'_l|$  be the sum of the cardinalities of the clusters of  $C'$ . Note that when clusters may overlap, it holds that  $n' \geq n$ . Thus,

$$NCE(C', C) = \sum_{l \in 1}^{|C'|} \frac{|c'_l|}{n'} NCE(c'_l, C).$$

The  $NCE$  values reported in Table 1 are thus obtained on the full set of snippets returned by Vivisimo.

### 6.3. Establishing the Ground Truth

The ephemeral nature of the web is amplified by the fact that search engines have at best a partial view of the available pages relevant to a given query. Moreover search engines must produce a ranking of the retrieved relevant pages and display only the pages of highest relevance. Thus, establishing a “ground truth” in a context of the full web is problematic. Following [Haveliwala et al. 02], we have made a series of experiments using as input the snippets resulting from queries issued to the Open Directory Project. The ODP is a searchable web-based directory consisting of a collection of a few million webpages (as of today, ODP claims to index 5.1M webpages) pre-classified into more than 590K categories by a group of volunteer human experts. The classification induced by the ODP labeling scheme gives us an objective “ground truth” against which we can compare the clustering quality of Vivisimo and Armil. In ODP, documents are organized according to a hierarchical ontology. For any snippet we obtain a label for its class by considering only the first two levels of the path on the ODP category tree. For example, if a document belongs to class “Games/Puzzles/Anagrams” and another document belongs to class “Games/Puzzles/Crosswords,” we consider both of them to belong to class “Games/Puzzles.” This coarsification is needed in order to balance the number of classes and the number of snippets returned by a query.

Queries are submitted to Vivisimo, asking it to retrieve pages only from ODP. This is done to ensure that Vivisimo and Armil operate on the same set of snippets, hence to ensure full comparability of the results. The resulting set of snippets is parsed and given as input to Armil. Since Vivisimo does not report the ODP category to which a snippet belongs, for each snippet we perform a query to ODP in order to establish its ODP category.

### 6.4. Outcome of the Comparative Experiment

The queries used in this experiment are the last 30 of those reported in Table 5 (the first five have been excluded since too few related snippets are present in ODP). On average, ODP returned 41.2 categories for each query. In Table 1 we report the *NMI* and *NCE* values obtained by Vivisimo and Armil on these data. Vivisimo produced by default about 40 clusters; therefore, we have run Armil with a target of 40 clusters (thus with a choice close to that of Vivisimo and to the actual average number of ODP categories per query) and with 30 (this number is the default used in the user evaluation).

	Vivisimo	Armil(40)	Armil(30)
<i>NCE</i>	0.667	0.735 (+10.1%)	0.683 (+2.3%)
<i>NMI</i>	0.400	0.442 (+10.5%)	0.406 (+1.5%)

**Table 1.** Results of the comparative evaluation.

The experiments indicate a substantial improvement of about 10% in terms of cluster quality of Armil (40) with respect to Vivisimo.<sup>15</sup> This improvement is an important result since, “[T]he scientific literature offers several solutions to the web-snippet clustering problem, but unfortunately the attainable performance is far from the one achieved by Vivisimo” [Ferragina and Gulli 05]. It should be noted moreover that Vivisimo uses a proprietary algorithm, not in the public domain, which might make extensive use of external knowledge. In contrast our algorithm is open and disclosed to the research community.

## 7. Comparison of Armil and $k$ -means

### 7.1. The $k$ -means Algorithm and Its Variants

The set of experiments in this section compare our variant of the furthest-point-first algorithm for  $k$ -center clustering with three recently proposed, fast variants of  $k$ -means.

The  $k$ -means algorithm (see [Forgy 65, Lloyd 82]) can be seen as an iterative cluster quality booster. It takes as input a rough  $k$ -clustering (or, more precisely,  $k$  candidate centroids) and produces as output another  $k$ -clustering, hopefully of better quality. It has been shown [Selim and Ismail 84] that by using the sum of squared Euclidean distances as objective function,<sup>16</sup> the procedure converges to a local minimum for the objective function within a finite number of iterations.

The main building blocks of  $k$ -means are (a) the generation of the initial  $k$  candidate centroids, (b) the main iteration loop, and (c) the termination condition. In the main iteration loop, given a set of  $k$  centroids, each input point is associated to its closest centroid, and the collection of points associated to a centroid is considered a cluster. For each cluster, a new centroid that is a (weighted) linear combination of the points belonging to the cluster is recomputed, and a new iteration starts.<sup>17</sup> Several termination conditions are possible: e.g., the loop can be terminated after a predetermined number of iterations or when the variation that the centroids have undergone in the last iteration is below a predetermined threshold.

<sup>15</sup>For the sake of replicating the experiments, all the search results have been cached and are available upon request.

<sup>16</sup>More precisely, this corresponds to partitioning  $S$ , at every iteration, into non-overlapping clusters  $C_1, \dots, C_k$  and determining their centroids  $\mu_1, \dots, \mu_k \in V$  so that the sum of the squares of the inter-cluster point-to-center distances

$$\sum_j \sum_{x \in C_j} (D(x, \mu_j))^2$$

is minimized.

<sup>17</sup>Note that  $k$ -means is defined on vector spaces but not in general on metric spaces, since in metric spaces linear combinations of points may not be points themselves.

Given the importance of this algorithm, there is a vast literature that discusses its shortcomings and possible improvements to the basic framework. In particular, one well-known shortcoming is that some clusters may become empty during the computation. To overcome this problem, following [Phillips 02] we adopt the “ISODATA” technique that, when a cluster becomes empty, splits one of the “largest” clusters so as to keep the number of clusters unchanged.

It is well known, and our experiments confirm this, that the quality of the initialization (i.e., the choice of the initial  $k$  centroids) has a deep impact on the resulting accuracy. Several methods for initializing  $k$ -means have been compared [Peña et al. 99]. As our baselines we have chosen the three such methods that are most amply cited in the literature while being at the same time relatively simple; we have ruled out more advanced and complex initializations since the possible boost in quality would be paid immediately in terms of computational cost, thus bringing about too slow an algorithm for our intended application.

## 7.2. Algorithms, Baselines, and Variants

The four algorithms we compare in our experiments are the clustering algorithm in Armil and three recently proposed, fast variants of  $k$ -means (RC, RP, MQ) [Phillips 02]. More in detail:

- KC This is our  $k$ -center algorithm with triangular inequality filtering, as described in Section 3.
- RP This is  $k$ -means (as described in Section 7.1) with an initialization based on random perturbation: for each dimension  $d_j$  of the space, the distribution of the projections on  $d_j$  of the data points is computed, along with its mean  $\mu_j$  and its standard deviation  $\sigma_j$ ; the  $k$  initial centroids are obtained through  $k$  perturbations, driven by the  $\mu_j$ ’s and  $\sigma_j$ ’s, of the centroid of all data points [Peña et al. 99].
- MQ This is MacQueen’s variant of  $k$ -means: the initial centroids are randomly chosen among the input points, the remaining points are assigned one at a time to the nearest centroid, and each such assignment causes the immediate recomputation of the centroid involved [MacQueen 67].
- RC This is again  $k$ -means where the initial centroids are randomly chosen among the input points and the remaining points are assigned to the closest centroid [Forgy 65].

## 7.3. Experimental Results

Since web snippet clustering is used as an online support to web browsing, real-time response is a critical parameter. A clustering phase that introduces a delay comparable to the time needed for just downloading the snippets, thus in fact



doubling the user waiting time, is not acceptable for most users. For this reason, instead of using a fixed number of iterations as the termination condition of the  $k$ -means algorithm, we opt for a fixed time deadline (five seconds, in our experiments), and we halt the clustering algorithms at the end of the first iteration that has run over the deadline.

In Table 2 we report clustering time and output quality of several variants of  $k$ -center and  $k$ -means on a sample of 12 queries subdivided, according to the method of [Di Giacomo et al. 05], in three broad families: *ambiguous queries* (armstrong, jaguar, mandrake, java), *generic queries* (health, language, machine, music, clusters), and *specific queries* (mickey mouse, olympic games, steven spielberg). In Table 2, for each query group we indicate averages of the number  $n$  of snippets found, the number  $k$  of ODP classes to which they are associated, and the time  $\tau$  (in seconds) used to build the vectorial representations. For each query group and each algorithm, we indicate the time used for clustering and

Algorithm	NCE	NMI	Time (sec.)
Query group = “ambiguous $\cup$ generic $\cup$ specific” $avg(n) = 157.4$ $avg_{ODP}(k) = 36.9$ $avg(\tau) = 0.729$			
MQ	0.70	0.47	5.59
PTS	0.69	0.45	5.76
RP	0.64	0.41	26.12
KC	0.70	0.47	0.77
Query group = “ambiguous” $avg(n) = 165$ $avg_{ODP}(k) = 38.75$ $avg(\tau) = 0.695$			
MQ	0.70	0.45	5.77
PTS	0.67	0.42	6.0
RP	0.63	0.38	33.3
KC	0.69	0.44	0.70
Query group = “specific” $avg(n) = 107.3$ $avg_{ODP}(k) = 26$ $avg(\tau) = 0.509$			
MQ	0.77	0.52	5.46
PTS	0.77	0.52	5.32
RP	0.71	0.46	18.47
KC	0.76	0.50	0.71
Query group = “generic” $avg(n) = 181.4$ $avg_{ODP}(k) = 42$ $avg(\tau) = 0.889$			
MQ	0.67	0.46	5.51
PTS	0.65	0.44	5.7
RP	0.61	0.40	24.92
KC	0.67	0.47	0.86

**Table 2.** Time/quality results of web snippet clustering for ODP snippets with Armil and variants of  $k$ -means.

the quality of the outcome.<sup>18</sup> The main conclusion of these experiments is that comparable levels of quality are achieved by Armil and  $k$ -means, but Armil is five to ten times faster.

Note that in these experiments we ask our system to partition the snippets into exactly  $k = 30$  clusters, which is the default of Armil. Notice that this number is reasonably close to the average number of classes found in ODP for the test queries.

## 8. User Evaluation of the Cluster Labeling Algorithm

Assessing “objectively” the quality of a cluster labeling method is a difficult problem, for which no established methodology has gained a wide acceptance. For this reason a user study is the standard testing methodology. We have set up a user evaluation of the cluster labeling component of Armil in order to have an independent and measurable assessment of its performance. We performed the study on 22 volunteer master students, doctoral students and post-doctorates in computer science at our departments (University of Siena, University of Pisa, and IIT-CNR). The volunteers all have a working knowledge of the English language.

The user interface of Armil has been modified to show clusters one-by-one and proceed only when the currently shown cluster has been evaluated. The queries are supplied to the evaluators in a round robin fashion from a list of 35 predefined queries (the queries are listed in Table 5). For each query the user must first say whether the query is meaningful to him/her; an evaluator is allowed to evaluate only queries meaningful to him/her. For each cluster we propose three questions

1. Is the label syntactically well-formed?
2. Can you guess the content of the cluster from the label?
3. After inspecting the cluster, do you retrospectively consider the cluster as well described by the label?

The evaluator must choose one of three possible answers (Yes, Sort-of, No), and his/her answer is automatically recorded in a database. Question 1 is aimed at assessing the gracefulness of the label produced. Question 2 is aimed at assessing the quality of the label as an instrument predictive of the cluster content. Question 3 is aimed at assessing the correspondence of the label with the content of the cluster. Note that the user cannot inspect the content of the cluster before answering Questions 1 and 2.

---

<sup>18</sup>For the sake of replicating the experiments, all the search results have been cached and are available at <http://pspl.iit.cnr.it/~mcsoft/fpf/fpf.html>.

### 8.1. Selection of the Queries

Similarly to [Ferragina and Gulli 05, Kummamuru et al. 04], we have randomly selected 35 of the most popular queries submitted to Google in 2004 and 2005.<sup>19</sup> From the selection we have removed queries (such as, e.g., “Spongebob” and “Hilary Duff”) that, referring to someone or something of regional interest only, were unlikely to be meaningful to our evaluators. The selected queries are listed in Table 5.

### 8.2. Discussion of the Results

Each of the 35 queries has been evaluated by two different evaluators, for a total of 70 query evaluations and 1584 cluster evaluations. The results are displayed in Table 3.

	Yes	Sort-of	No
Question 1	60.5%	25.5%	14.0%
Question 2	50.0%	32.0%	18.0%
Question 3	47.0%	38.5%	14.5%

**Table 3.** Results of user evaluation.

Summing the very positive and the mildly positive answers, we can conclude that, in this experiment, 86.0% of the labels are syntactically acceptable, 82.0% of the labels are predictive, and 85.5% of the clusters are sufficiently well described by their label. By checking the percentages of “No” answers, we can notice that sometimes labels considered non-predictive are nonetheless considered well-descriptive of the cluster; we interpret this fact as due to the discovery of meanings of the query string previously unknown to the evaluator.

The correlation matrices in Table 4 show more precisely the correlation between syntax, predictiveness, and representativeness of the labels. Entries in the top part give the percentage over all answers, and entries in the bottom part give percentage over rows.

The data in Table 4 (top) show that there is a strong correlation between syntactic form and predictivity of the labels, as shown by the fact that in a high percentage of cases the same answer was returned to Questions 1 and 2.

The middle and bottom parts of Table 4 confirm that while for the positive or mildly positive answers (Yes, Sort-of) there is a strong correlation between the answers returned to the different questions, it is often the case that a label considered not predictive of the content of the cluster can still be found, after inspection of the cluster, to be representative of the content of the cluster.

<sup>19</sup><http://www.google.com/press/zeitgeist.html>

	2. Yes	2. Sort-of	2. No
1. Yes	42.67%	12.81%	5.11%
1. Sort-of	5.74%	15.27%	4.41%
1. No	1.64%	3.78%	8.52%
1. Yes	70.41%	21.14%	8.43%
1. Sort-of	22.58%	60.04%	17.36%
1. No	11.76%	27.14%	61.08%
	3. Yes	3. Sort-of	3. No
2. Yes	33.52%	12.81%	3.72%
2. Sort-of	11.36%	16.85%	3.66%
2. No	2.14%	8.90%	7.00%
2. Yes	66.96%	25.59%	7.44%
2. Sort-of	35.64%	52.87%	11.48%
2. No	11.88%	49.30%	38.81%
	3. Yes	3. Sort-of	3. No
1. Yes	35.98%	18.93%	5.68%
1. Sort-of	8.64%	12.81%	3.97%
1. No	2.39%	6.81%	4.73%
1. Yes	59.37%	31.25%	9.37%
1. Sort-of	33.99%	50.37%	15.63%
1. No	17.19%	48.86%	33.93%

**Table 4.** Correlation tables of Questions 1 and 2 (top), 2 and 3 (middle), 1 and 3 (bottom).

skype	winmx	nintendo	revolution
pamela anderson	twin towers	wallpaper	firefox
ipod	tsunami	tour de france	weather
matrix	mp3	new orleans	notre dame
games	britney spears	chat	CNN
iraq	james bond	harry potter	simpsons
south park	baseball	ebay	madonna
star wars	tiger	airbus	oscars
london	pink floyd	armstrong	spiderman

**Table 5.** Queries used in the user evaluation.

### 8.3. Running Times

Our system runs on an AMD Athlon (1-Ghz Clock) processor with 750-MB RAM and operating system FreeBSD 4.11 - STABLE. The code was developed in Python V. 2.4.1. Excluding the time needed to download the snippets from the auxiliary search engines, the 35 queries have been clustered and labeled in 0.72 seconds on average; the slowest query took 0.92 seconds.

## 9. Conclusions and Future Work

Why is Armil not “yet another clustering search engine”? The debate on how to improve the performance of search engines is at the core of the current research in the area of web studies, and we believe that so far only the surface of the vein has been uncovered. The main philosophy of the system/experiments we have proposed follows these lines:

1. Principled algorithmic choices are made whenever possible.
2. Clustering is clearly decoupled from labeling.
3. Attention is paid to the trade-off between response time and quality while limiting the response time within limits acceptable by the user.
4. A comparative study of Armil and Vivisimo has been performed in order to assess the quality of Armil’s clustering phase by means of effectiveness measures commonly used in clustering studies.
5. A user study has been set up in order to obtain an indication of user satisfaction with the produced cluster labeling.
6. No use of external sources of knowledge is made.

Further research is needed in two main areas. First, we plan to assess to what extent a modicum of external knowledge can improve the system’s performance without speed penalties. An example of one possible such improvement is the use of extractors of noun phrases in parsing the snippets. Second, it is possible to introduce in the current pipeline (input snippets are clustered, candidates are extracted, then labels are generated) of the architecture a feedback loop by considering the extracted candidates/labels as predefined categories, thus examining which snippets in different clusters are closer to the generated labels. Snippets close to the label of cluster  $C_x$  but in a different cluster  $C_y$  could be shown on the screen as related also to  $C_x$ . This would give the benefits of soft clustering without much computational overload. Finally, methods for detecting automatically the desired number of clusters will be studied.

**Acknowledgments.** A preliminary version of this work appeared in the proceedings of ACM SAC 2006 [Geraci et al. 06b] and SPIRE 2006 [Geraci et al. 06a]. Our work was partially supported by the Italian Registry for the ccTLD “.it.” We wish to thank two anonymous referees, whose comments helped us in improving the style of the presentation, and the participants in the Armil user evaluation experiment.

## References

- [Baeza-Yates and Ribeiro-Neto 99] Ricardo A. Baeza-Yates and Berthier A. Ribeiro-Neto. *Modern Information Retrieval*. New York: ACM Press/Addison-Wesley, 1999.
- [Charikar 02] Moses S. Charikar. “Similarity Estimation Techniques from Rounding Algorithms.” In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, pp. 380–388. New York: ACM Press, 2002.
- [Cheng et al. 03] D. Cheng, R. Kannan, S. Vempala, and G. Wang. “On a Recursive Spectral Algorithm for Clustering from Pairwise Similarities.” Technical Report MIT-LCS-TR-906, Massachusetts Institute of Technology, 2003. Available at <http://cs-www.cs.yale.edu/homes/kannan/Papers/experimentclustering.pdf>.
- [Cover and Thomas 91] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. New York: John Wiley & Sons, 1991.
- [Cutting et al. 92] Douglass R. Cutting, Jan O. Pedersen, David Karger, and John W. Tukey. “Scatter/Gather: A Cluster-Based Approach to Browsing Large Document Collections.” In *Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 318–329. New York: ACM Press, 1992.
- [Di Giacomo et al. 05] Emilio Di Giacomo, Walter Didimo, Luca Grilli, and Giuseppe Liotta. “A Topology-Driven Approach to the Design of Web Meta-Search Clustering Engines.” In *SOFSEM 2005: Theory and Practice of Computer Science, 31st Conference on Current Trends in Theory and Practice of Computer Science, Liptovský Ján, Slovakia, January 17-20, 2005, Proceedings*, pp. 106–116. Berlin: Springer, 2005.
- [Ferragina and Gulli 05] Paolo Ferragina and Antonio Gulli. “A Personalized Search Engine Based on Web-Snippet Hierarchical Clustering.” In *Special Interest Tracks and Posters of the 14th International Conference on World Wide Web*, pp. 801–810. New York: ACM Press, 2005.
- [Forgy 65] E. Forgy. “Cluster Analysis of Multivariate Data: Efficiency vs. Interpretability of Classifications.” *Biometrics* 21:3 (1965), 768–780.
- [Geraci et al. 06a] F. Geraci, M. Pellegrini, F. Sebastiani, and M. Maggini. “Cluster Generation and Cluster Labeling for Web Snippets.” In *String Processing and Information Retrieval: 13th International Conference, SPIRE 2006, Glasgow, UK, October 11–13, 2006, Proceedings*, Lecture Notes in Computer Science 4209, pp. 25–36. Berlin: Springer, 2006.
- [Geraci et al. 06b] Filippo Geraci, Marco Pellegrini, Paolo Pisati, and Fabrizio Sebastiani. “A Scalable Algorithm for High-Quality Clustering of Web Snippets.” In *Proceedings of the 2006 ACM Symposium on Applied Computing*, pp. 1058–1062. New York: ACM Press, 2006.
- [Gonzalez 85] Teofilo F. Gonzalez. “Clustering to Minimize the Maximum Intercluster Distance.” *Theoretical Computer Science* 38:2–3 (1985), 293–306.

- [Haveliwala et al. 00] Taher H. Haveliwala, Aristides Gionis, and Piotr Indyk. “Scalable Techniques for Clustering the Web.” In *The World Wide Web and Databases: Third International Workshop WebDB 2000, Dallas, TX, USA, May 18–19, 2000, Selected Papers*, Lecture Notes in Computer Science 1997, pp. 129–134. Berlin: Springer, 2000.
- [Haveliwala et al. 02] Taher H. Haveliwala, Aristides Gionis, Dan Klein, and Piotr Indyk. “Evaluating Strategies for Similarity Search on the Web.” In *Proceedings of the 11th International Conference on World Wide Web*, pp. 432–442. New York: ACM Press, 2002.
- [Hearst and Pedersen 96] Marti A. Hearst and Jan O. Pedersen. “Reexamining the Cluster Hypothesis: Scatter/Gather on Retrieval Results.” In *Proceedings of the 19th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 76–84. New York: ACM Press, 1996.
- [Hochbaum and Shmoys 85] D. S. Hochbaum and D. B. Shmoys. “A Best Possible Approximation Algorithm for the  $k$ -center Problem.” *Mathematics of Operations Research* 10:2 (1985), 180–184.
- [Hsu and Nemhauser 79] W. Hsu. and G. Nemhauser. “Easy and Hard Bottleneck Location Problems.” *Discrete and Applied Mathematics* 1 (1979), 209–216.
- [Indyk 99] Piotr Indyk. “Sublinear Time Algorithms for Metric Space Problems.” In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, pp. 428–434. New York: ACM Press, 1999.
- [Koshman et al. 05] S. Koshman, A. Spink, and B.J. Jansen. “Using Clusters on the Vivisimo Web Search Engine.” In *Proceedings of HCI International 2005*. Mahwah, NJ: Lawrence Erlbaum Associates, 2005.
- [Kummamuru et al. 04] Krishna Kummamuru, Rohit Lotlikar, Shourya Roy, Karan Singal, and Raghu Krishnapuram. “A Hierarchical Monothetic Document Clustering Algorithm for Summarization and Browsing Search Results.” In *Proceedings of the 13th International Conference on World Wide Web*, pp. 658–665. New York: ACM Press, 2004.
- [Kural et al. 93] Yasemin Kural, Stephen Robertson, and Susan Jones. “Deciphering Cluster Representations.” *Information Processing and Management* 37 (1993), 593–601.
- [Kural et al. 99] Yasemin Kural, Stephen Robertson, and Susan Jones. “Clustering Information Retrieval Search Outputs.” In *Proceedings of the 21st Annual BCS IRSG Colloquium on Information Retrieval*, 1999. Available at <http://www.bcs.org/server.php?show=ConWebDoc.4252>.
- [Lawrie and Croft 03] Dawn J. Lawrie and W. Bruce Croft. “Generating Hierarchical Summaries for Web Searches.” In *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 457–458. New York: ACM Press, 2003.
- [Lloyd 82] S.P. Lloyd. “Least Squares Quantization in PCM.” Technical Report, Bell Laboratories, 1957. Reprinted in *IEEE Transactions on Information Theory* IT-28:2 (1982), 29–137.



- [Maarek et al. 00] Yoelle Maarek, Ron Fagin, Israel Ben-Shaul, and Dan Pelleg. “Ephemeral Document Clustering for Web Applications.” Technical Report RJ 10186, IBM, 2000. Available at <http://citeseer.ist.psu.edu/maarek00ephemeral.html>.
- [MacQueen 67] J.B. MacQueen. “Some Methods for Classification and Analysis of Multivariate Observations.” In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, pp. 281–297. Berkeley: University of California Press, 1967.
- [Meila 03] M. Meila. “Comparing Clusterings.” In *Learning Theory and Kernel Machines: 16th Annual Conference on Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24–27, 2003, Proceedings*, Lecture Notes in Computer Science 2777, pp. 173–187. Berlin: Springer, 2003.
- [Meng et al. 02] Weiyi Meng, Clement T. Yu, and King-Lup Liu. “Building Efficient and Effective Metasearch Engines.” *ACM Computing Surveys* 34:1 (2002), 48–89.
- [Milligan and Cooper 85] G. W. Milligan and M. C. Cooper. “An Examination of Procedures for Determining the Number of Clusters in a Data Set.” *Psychometrika* 50:2 (1985), 159–179.
- [Osinski and Weiss 04] Stanislaw Osinski and Dawid Weiss. “Conceptual Clustering Using Lingo Algorithm: Evaluation on Open Directory Project Data.” In *Intelligent Information Processing and Web Mining: Proceedings of the International IIS: IIPWM '04 Conference Held in Zakopane, Poland, May 17–20, 2004*, Advances in Soft Computer 25, pp. 369–377. Berlin: Springer, 2004.
- [Peña et al. 99] José Manuel Peña, Jose Antonio Lozano, and Pedro Larrañaga. “An Empirical Comparison of Four Initialization Methods for the  $k$ -means Algorithm.” *Pattern Recognition Letters* 20:10 (1999), 1027–1040.
- [Phillips 02] Steven J. Phillips. “Acceleration of  $k$ -means and Related Clustering Algorithms.” In *Algorithm Engineering and Experiments: 4th International Workshop, ALENEX 2002, San Francisco, CA, USA, January 4–5, 2002, Revised Papers*, Lecture Notes in Computer Science 2153, pp. 166–177. Berlin: Springer, 2002.
- [Selim and Ismail 84] S.Z. Selim and M.A. Ismail. “ $K$ -Means Type Algorithms: A Generalized Convergence Theorem and Characterization of Local Optimality.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6:1 (1984), 81–87.
- [Spink et al. 05] A. Spink, S. Koshman, B.J. Jansen, M. Park, and C. Field. “Multi-tasking Web Search on Vivisimo.com.” In *Proceedings of International Conference on Information Technology Coding and Computing 2005*, Vol. 2, pp. 486–490. Los Alamitos, CA: IEEE Press, 2005.
- [Strehl 02] Alexander Strehl. “Relationship-Based Clustering and Cluster Ensembles for High-Dimensional Data Mining.” PhD thesis, University of Texas, Austin, 2002.
- [Strehl et al. 00] Alexander Strehl, Joydeep Ghosh, and Raymond J. Mooney. “Impact of Similarity Measures on Web Page Clustering.” In *Proceedings of the AAAI Workshop on AI for Web Search*, pp. 58–64. Menlo Park, CA: AAAI Press, 2000.

- [Tombros et al. 02] Anastasios Tombros, Robert Villa, and Cornelis J. van Rijsbergen. “The Effectiveness of Query-Specific Hierarchic Clustering in Information Retrieval.” *Information Processing and Management* 38:4 (2002), 559–582.
- [Yang and Pedersen 97] Yiming Yang and Jan O. Pedersen. “A Comparative Study on Feature Selection in Text Categorization.” In *Proceedings of the Fourteenth International Conference on Machine Learning*, pp. 412–420. San Francisco: Morgan Kaufmann, 1997.
- [Zamir and Etzioni 98] Oren Zamir and Oren Etzioni. “Web Document Clustering: A Feasibility Demonstration.” In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 46–54. New York: ACM Press, 1998.
- [Zamir et al. 97] Oren Zamir, Oren Etzioni, Omid Madani, and Richard M. Karp. “Fast and Intuitive Clustering of Web Documents.” In *Proceedings of the Third International Conference on Knowledge Discovery and Data Mining*, pp. 287–290. Menlo Park, CA: AAAI Press, 1997.
- [Zeng et al. 04] Hua-Jun Zeng, Qi-Cai He, Zheng Chen, Wei-Ying Ma, and Jinwen Ma. “Learning to Cluster Web Search Results.” In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 210–217. New York: ACM Press, 2004.
- [Zhang and Dong 04] D. Zhang and Y. Dong. “Semantic, Hierarchical, Online Clustering of Web Search Results.” In *Advanced Web Technologies and Applications: 6th Asia-Pacific Web Conference, APWeb 2004, Hangzhou, China, April 14–17, 2004, Proceedings*, Lecture Notes in Computer Science 3007, pp. 69–78. Berlin: Springer, 2004.
- [Zhang et al. 96] Tian Zhang, Raghu Ramakrishnan, and Miron Livny. “BIRCH: An Efficient Data Clustering Method for Very Large Databases.” In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pp. 103–114. New York: ACM Press, 1996.

Filippo Geraci, Istituto di Informatica e Telematica, Consiglio Nazionale delle Ricerche, Via G Moruzzi, 1 – 56124 Pisa, Italy, and Dipartimento di Ingegneria dell’Informazione, Università di Siena, Via Roma, 56 – 53100 Siena, Italy (f.geraci@iit.cnr.it)

Marco Pellegrini, Istituto di Informatica e Telematica, Consiglio Nazionale delle Ricerche, Via G Moruzzi, 1 – 56124 Pisa, Italy (m.pellegrini@iit.cnr.it)

Marco Maggini, Dipartimento di Ingegneria dell’Informazione, Università di Siena, Via Roma, 56 – 53100 Siena, Italy (maggini@dii.unisi.it)

Fabrizio Sebastiani, Istituto di Scienza e Tecnologia dell’Informazione, Consiglio Nazionale delle Ricerche, Via G Moruzzi, 1 – 56124 Pisa, Italy (fabrizio.sebastiani@isti.cnr.it)

Received March 1, 2007; accepted July 5, 2007.

