Exploratory Access to Wikipedia through Faceted Dynamic Taxonomies

Giovanni Maria Sacco

Dipartimento di Informatica, Università di Torino, Corso Svizzera 185, I-10149 Torino, Italy

giovanni.sacco@unito.it

Abstract

Users currently access Wikipedia through two traditional paradigms, text search and hypertext navigation. We believe that user access can be significantly improved by supporting a systematic conceptual exploration of the knowledge base through dynamic taxonomies with a faceted taxonomy organization. This approach allows the easy manipulation of sets of documents and the systematic and intuitive exploration of complex knowledge bases.

Introduction

While the coverage and authoritativeness of Wikipedia have been constantly improving in the past years, its basic knowledge architecture has remained the same and falls short of the potentiality of its knowledge base. In extreme synthesis, the access methods to Wikipedia are basically a simplified text search and a hypermedia navigation. Although a conceptual taxonomy is currently supported, it is a traditional taxonomy that only supports a father-to-son (and son-to-father) navigation by hypertext links and does not take into account the fact that most pages are actually classified under several concepts. In fact, the general architecture of Wikipedia is only slightly different from the traditional encyclopedia *à la* Diderot and D'Alembert.

By using Dynamic Taxonomies (Sacco, 2000), briefly reviewed in the following, we can support an exploratory access to Wikipedia in the following way:

1. the user is presented with the general taxonomy of the encyclopedia,
2. she selects a concept as the focus of interest, e.g. Renaissance,
3. the system automatically prunes from the original taxonomy all those concepts that are not related to Renaissance, giving the user a complete taxonomic summary of the current subset of interest, within the original frame of reference represented by the original taxonomy, which is only modified by pruning irrelevant concepts,
4. the user can then refine the focus of interest by selecting another concept, e.g. Painting, which is combined in AND with Renaissance, and continue exploring.

Among the several advantages of this approach, it is especially important that user navigation is completely free and yet guided. In fact, the user can initially select any concept in order to set her focus of interest, but, subsequently, only concepts that are related to the current interest focus can be selected for refinement.

In the following, we briefly introduce dynamic taxonomies, and discuss our current project, DT-Wikipedia, that aims at applying them to Wikipedia. We discuss our approach and what at present perceive as challenges.

Dynamic Taxonomies

Dynamic taxonomies (Sacco, 2000, later also improperly called *faceted search systems*) are a general knowledge management model based on a multidimensional classification of heterogeneous data items and are used to explore/browse complex information bases in a guided yet unconstrained way through a visual interface. It has been applied to very diverse areas, including electronic commerce (Sacco, 2003), e-government, e-HRM (Berio et al., 2007), multimedia databases with the seamless integration of primitive features (Sacco, 2008), art and museum portals (Yee et al., 2003), and medical diagnosis (Sacco, 2012), among many others. The reader is addressed to Sacco and Tzitzikas, 2009, for the most comprehensive and up-to-date monograph on this model.

The intensional part of a dynamic taxonomy is a taxonomy designed by an expert. It does not require any other relationships in addition to subsumptions (e.g., IS-A and PART-OF relationships).

In the extension, items can be freely classified under several topics at any level of abstraction. This multidimensional classification models common real-life situations. First, items can very often classified under different concepts. Second, items usually have different independent features (e.g. Time, Location, etc.), each of which can be described by an independent taxonomy. These features are often called *perspectives* or *facets*.

In dynamic taxonomies, a concept C is just a label that identifies all the items classified under C. Because of the subsumption relationship between a concept and its descendants, the items classified under C (items(C)) are all those items in the *deep extension* of C, i.e. the set of items identified by C includes the *shallow extension* of C (i.e. all the items directly classified under C) union the deep extension of C's sons. The shallow and the deep extension for a terminal concept are the same, by construction. This set-oriented approach implies that logical operations on concepts can be performed by the corresponding set operations on their extension, and therefore the user is able to restrict the information base (and to create derived concepts) by combining concepts through all the standard logical operations (and, or, not).

A fundamental feature of this model is that dynamic taxonomies can find all the concepts related to a given concept C: these concepts represent the conceptual summary of C. Concept relationships other than subsumptions are inferred on the basis of empirical evidence through the extension only, according to the following *extensional inference rule*: two concepts A and B are related iff there is at least one item *d* in the knowledge base which is classified at the same time under A or under one of A's descendants and under B or under one of B's descendants, or, more formally,

$$A \Leftrightarrow B \text{ iff } items(A) \cap items(B) \neq \emptyset.$$

For example, we can infer an unnamed relationship between *terrorism* and *New York*, if an item classified under *terrorism* and *New York* exists. At the same time, since *New York* is a descendant of *USA*, also a relationship between *terrorism* and *USA* can be inferred.

The extensional inference rule can be easily extended to cover the relationship between a given concept C and a concept expressed by an arbitrary subset S of the universe: C is related to S iff there is at least one item d in S which is also in items(C), or, equivalently,

$$A \Leftrightarrow B \text{ iff } items(C) \cap S \neq \emptyset.$$

Consequently, the extensional inference rule can produce conceptual summaries not only for base concepts, but also for any logical combination of concepts. Moreover, since it is immaterial how S is produced, dynamic taxonomies can summarize sets of items produced by other retrieval methods such as database queries, shape retrieval, etc. and therefore access through dynamic taxonomies can be easily combined with any other retrieval method.

Dynamic taxonomies are defined in terms of conceptual descriptions of items, so that heterogeneous items of any type and format can be managed in a single, coherent framework. Finally, since concept C is just a label that identifies the set of the items classified under C, concepts are language-invariant, and multilingual access can be easily supported by maintaining different language directories, holding language-specific labels for each concept in the taxonomy.

Access through Dynamic Taxonomies

The user is initially presented with a tree representation of the initial taxonomy for the entire infobase. The system can associate with each concept label, a count of all the items classified under it (i.e. the cardinality of items(C) for all C's). This count is an important user feedback in navigation, because when it is sufficiently small, the user usually terminates exploration and inspects the result items.

The initial user focus F is the universe (i.e. all the items in the infobase).

In the simplest case, the user can then select a concept C in the taxonomy and *zoom* over it. The zoom operation changes the current state in two ways. First, concept C is used to refine the current focus F, by intersecting it with items(C); items not in the focus are discarded. Second, the tree representation of the taxonomy is modified in order to summarize the new focus. All and only the concepts related to F are retained and the count for each retained concept C' is updated to reflect the number of items in the focus F that are classified under C'.

The reduced taxonomy is a conceptual summary of the set of documents identified by F, exactly in the same way as the original taxonomy was a conceptual summary of the universe. The term *dynamic taxonomy* is used to indicate that the taxonomy can dynamically adapt to the subset of the universe on which the user is focusing, whereas traditional, static taxonomies can only describe the entire universe.

The exploration process is an iterative thinning of the information base: the user selects a focus, which restricts the information base by discarding all the items not in the current focus. Only the concepts used to classify the items in the focus, and their ancestors, are retained. These concepts, which summarize the current focus, are those and only those concepts that can be used for further refinements. From the human computer interaction point of view, the user is effectively guided to reach his goal, by a clear and consistent listing of all possible alternatives.

Differently from traditional search methods, the exploration process has the goal of reducing the universe to a set of items sufficiently small that they can be manually inspected by the user,

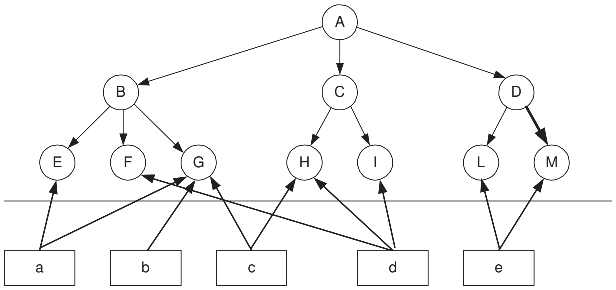

Figure 1 – A dynamic taxonomy: the intension is above, the extension below. Arrows going down denote subsumptions, going up classification

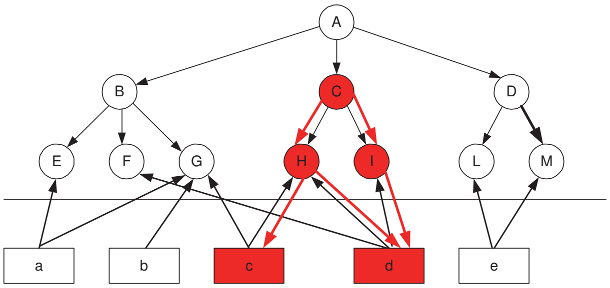

Figure 2 – Focusing on concept C: finding all the items classified under C, i.e. the deep extension of C.

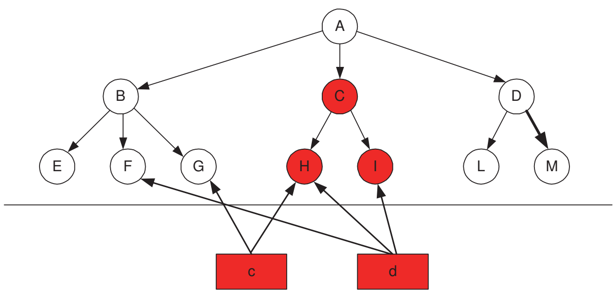

Figure 3 – All the items not classified under C are removed

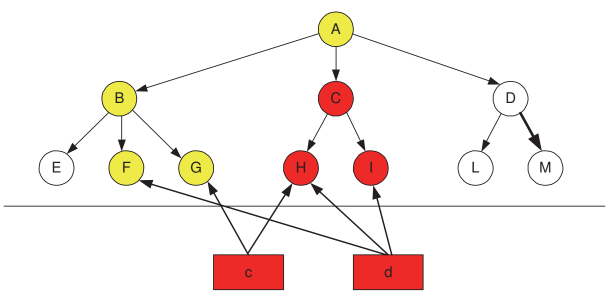

Figure 4 – In yellow, all the concepts under which the items in the focus are classified (and, because of subsumptions) their ancestors are related to C. White nodes are not related to the focus and they will be pruned out.

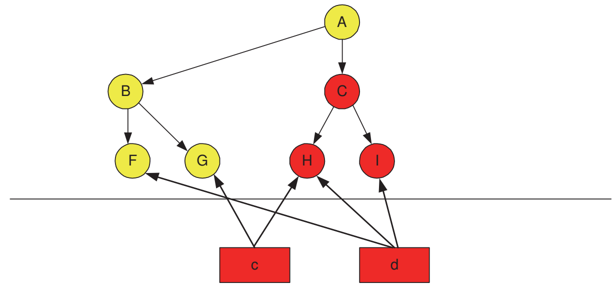

Figure 5 – The reduced taxonomy: all concepts not related to the current focus are pruned.

An Example

Figure 1 shows a dynamic taxonomy: the upper half represents the intensional level where circles represent concepts; the lower half is the extension where rectangles represent items, which are classified according to the concepts in the intension. In the intension, arcs going down represent subsumptions; for instance, the arc from concept A to concept B indicates that B is subsumed by A (is a specialization of A). The classification of items is represented by arcs going up, connecting an item to a concept in the intension. As an example, the arcs from item d to concepts H and I indicate that item d is classified under H and I.

Figures 2 to 5 show how the zoom operation on a concept C works. To simplify the discussion we will assume that the zoom operation is applied to the original taxonomy. In order to compute all the concepts related to C (i.e. to zoom on concept C), we first identify, in figure 2, the user focus. In this case, since we start from the initial taxonomy, the user focus is C or, equivalently, all the items classified under C (that is, the deep extension of C, denoted by items(C)). The deep extension of C is computed by following all the arcs incident to C and to all of its descendants, H and I, and originating from items in the extension. In the example, there is no arc linking items in the extension to C, but there are arcs connecting items c and d to concepts H and I, which are (the only) descendants of C. The current user focus, which is represented by the deep extension of C, items(C), is therefore equal to the set { c, d }.

Once the user focus is computed, all the items in the extension which are not in the user focus can be ignored and logically removed from the extension, as shown in figure 3.

Next, we compute B(C), i.e. the set of concepts under which at least one item in the current focus C is classified. B(C) is the conceptual summary of the current focus: it contains all and only those concepts which are used to classify items in the current focus. We first compute

B_{imm}(C) by identifying all the concepts immediately related to C by the extensional inference rule, i.e., by following all the arcs leaving each element in the user focus and adding each destination concept to B_{imm}(C); in the example, B_{imm}(C)={ F, G, H, I }. If a concept belongs to B_{imm}(C) it also belongs to B(C). In addition, if a concept belongs to B_{imm}(C) also all of its ancestors belong to B(C) as well, because the extensional inference rule applied to C states that a concept D is related to C if there is at least one item which is classified under D or under one of D's descendants. Consequently, if a concept is related to C, also all of its ancestors are, because of the inclusion constraint implied by subsumption (Sacco, 2000). Therefore, the set B(C) of concepts related to C is given by B_{imm}(C) union all the ancestors of all the concepts in B_{imm}(C), i.e. the set of all concepts related to C is {F, G, H, I, B, C, A}, as shown in figure 4, where yellow nodes denote the concepts in B(C) which are not descendants.

Finally, in figure 5, all the concepts not related to C are logically pruned, thus producing a reduced taxonomy that fully describes all and only the items in the current focus. Further zoom operations can be performed on concepts belonging to this reduced taxonomy. If another concept, e.g. concept I, is chosen for a subsequent zoom, the current user focus will be computed as the intersection of the previous user focus (C, in this example) with the current selected focus (I, in this example). That is, the user focus is given by items(C) ∩ items(I).

The user is, however, unable to select concept M for a subsequent zoom, because zooming on it would produce an empty result and for this reason M is not present in the reduced taxonomy.

Benefits of Dynamic Taxonomies

The advantages of dynamic taxonomies over traditional methods are dramatic in terms of an extremely fast convergence of exploratory patterns and in terms of human factors. Three zoom operations on terminal concepts are sufficient to reduce a 10,000,000-item information base described by a compact taxonomy with 1,000 concepts to an average 10 items (Sacco, 2006). Dynamic taxonomies only require a very light theoretical background: namely, the concept of a taxonomic organization and the zoom operation, which seems to be very quickly understood by end-users.

Dynamic taxonomies cleanly separate the process of classifying documents from the use of the classification information in the browsing system, and considerably simplify the design of the conceptual taxonomy. First, the extensional inference rule actually performs concept association mining: concept associations, which are often quite dynamic in time, need not be forecasted and accounted for in schema design. In addition, the user is presented with associations the schema designer might not even be aware of.

Second, since dynamic taxonomies synthesize compound concepts, these need usually not be represented explicitly, so that we avoid the exponential growth due to the description of all the possible concept combinations, and the resulting taxonomy is significantly more compact and easier to understand. Sacco (Sacco, 2000; Sacco and Tzitzikas, 2009) developed a number of guidelines for taxonomies that are compact and easily understood by users. Some are superficially similar to the basic faceted classification scheme by Ranganathan (Ranganathan, 1965): the taxonomy is organized as a set of independent, “orthogonal” subtaxonomies (facets or perspectives). As an example, a compound Wikipedia concept such as "Musicians from Mobile, Alabama" need not be explicitly accounted for, because it can be synthesized from its component concepts: *Arts>Musicians* and *Location>USA>Alabama>Mobile*, where Arts and Location are facets.

Benefits of Dynamic Taxonomies in the Context of Wikipedia

Current conceptual access to Wikipedia is currently provided via a traditional, static taxonomy implemented by hypertext links. In a static taxonomy,

1. the taxonomy does not adapt to specific subsets of the universe, but statically summarizes the entire universe,
2. once a branch is chosen, the user can only refine her search by selecting a specialization. All the other branches are unavailable, and, of course, a terminal concept cannot be further refined, thus leading to severe scalability issues (Sacco, 2006),
3. the impossibility to combine concepts through boolean operations requires that compound concepts be explicitly represented in the taxonomy resulting either in extremely large taxonomies (as it is the case in Wikipedia) and/or a gross conceptual granularity (as it is again the case for some parts of Wikipedia).

The benefits of dynamic taxonomies in this context are basically:

- a simpler, more compact taxonomy that user can understand and use. Retrofitting the existing Wikipedia taxonomy is discussed below;
- a free, yet guided, exploration of the knowledge base with a comprehensive summary of all the concepts that are related to the current focus. This type of exploration avoids dead-ends by construction and has a superior scalability for a

growing knowledge base. In addition, it provides an exhaustive conceptual description of the current focus, and the user can have the confidence that all possible aspects are being considered,

- the dynamic computation of related concepts implies a dynamicity of relationships. First, relationships among concepts need not be anticipated at design time, as they are established on the basis of empirical evidence. Second, these relationship adapt to changing situations and provide the potential for the discovery of new, unexpected relationships. Discovery of unexpected relationships can represent a major improvement over traditional access,

Finally a good "faceted" taxonomy design coupled with dynamic taxonomies provides an extremely flexible and symmetric way of exploration, and provides answers to questions such as

- "What happened in 1910 in France?", by intersecting *Time>1910* and *Location>France*, and exploring the reduced taxonomy
- "What is known about Mobile, Alabama?", by focusing on *Location>USA>Alabama>Mobile*
- "What is known about Alabama?", by focusing on *Location>USA>Alabama*

which are impossible to answer, or even to frame, in the current approach. In short, dynamic taxonomies and a faceted schema bring to the surface the information buried inside concept labels, and makes it actionable.

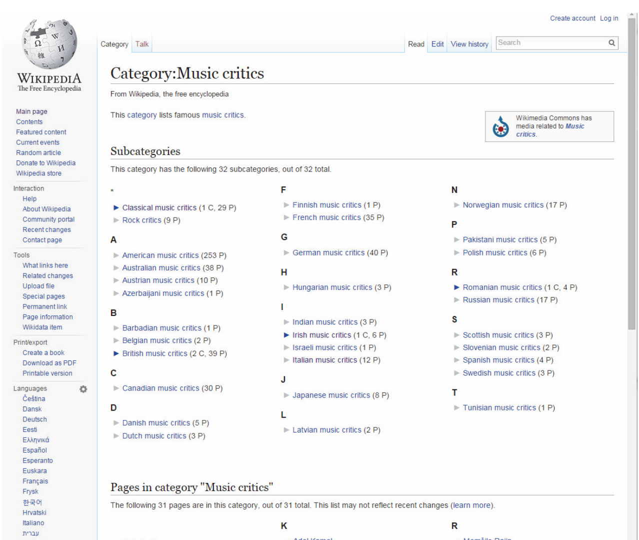

Figure 6 – A category page for music critics, specializing them by nation, and, to some extent, by genre.

As an example, consider the Wikipedia page in figure 6. Here, the representation of a part of the taxonomy as a hyperlinked page makes the information contained in the page non-actionable. What if the user is interested in European music critics, or music critics that are also musicologist, or classical music critics that also write about rock? The information might very well be present in this page or other pages (figure 7 shows a music critic who is a musicologist) but cannot be extracted and used.

DT-Wikipedia

The DT-Wikipedia project is in its initial, exploratory phase. One of the very first, experimental applications of dynamic taxonomies was, several years ago, the electronic version of an Italian encyclopedia. Wikipedia is of course quite different in scope, construction and control, and has peculiar challenges. In providing access through dynamic taxonomies to Wikipedia, we identify two types of problems: system architecture and conceptual architecture.

Figure 7 – A music critic who is a musicologist as well

From the system architecture point of view, dynamic taxonomies cannot be implemented in the extremely simple architecture of Wikipedia. First of all, since the computation of related concepts and reduced taxonomies is done on-the-fly, specialized engines are required in order to provide a real-time interaction.

Second, the interface to conceptual access is obviously more complex than the standard Wikipedia interface and requires a non-trivial amount of work to make it usable over different devices.

However, we believe that the most challenging part of our project is the conceptual architecture.

System Architecture

DT-Wikipedia is based on the Universal Knowledge Processor (UKP), a commercial-grade web implementation that was entirely developed by the author and has been extensively used in our research since 1999.

The engine is designed as a memory-resident specialized architecture with low hardware requirements and support for multimillion item knowledge bases with high user loads. It seamlessly integrates dynamic taxonomy access with full-text retrieval with relevance ranking and quasi-vector-space retrieval, plus db-like features for the management of continuous domains (i.e., dates and numbers) with range queries. These features are especially important because some (less important) concepts can be removed from the taxonomy and their associated documents retrieved through text-retrieval.

The general architecture and an experimental comparison with a memory-resident relational implementation are discussed in (Sacco and Tzitzikas, 2009). The comparison is especially interesting because it shows that the specialized architecture used in UKP is two orders of magnitude faster than the relational architecture which is commonly used in current systems. The same source also discusses a relational implementation, Flexplorer, developed at the University of Crete and also described in (Papadakos, Kopidaki et al., 2009).

Conceptual Architecture

An initial analysis of the current taxonomy in the English version of Wikipedia shows a number of rather severe problems, discussed in the following.

Jumbo Taxonomy

With over 600,000 concepts, the taxonomy of Wikipedia cannot be understood by any user, cannot be shown on a screen, and, unfortunately, cannot be processed manually. This means that ways to process the taxonomy in order to reduce its size, and to make it conform to a facet-like organization are needed. Work in this direction is underway and is described below.

"Cartesian Product" Index Entries

The major problem of the current taxonomy, which is also the major opportunity for a solution, is that most of the entries are caused by the cartesian product of the values of two or more combined facets.

Entries such as "Rugby at the 1900 Summer Olympics" clearly show a cartesian product at work. Here we have three perspectives <sport, time, sport event> that are represented as a single entry, and require the cartesian product of all the values in each perspective, e.g. "Football at the 1900 Summer Olympics", but also, "Rugby at the 1908 Summer Olympics", etc. Once the underlying perspectives are identified, a number of entries equal to the product of the values can be reduced to a number of entries equal to the sum of the values. Each page will then be classified under each applicable perspective value, i.e., the page originally classified under "Rugby at the 1900 Summer Olympics", will be classified instead under *Sport>Rugby*, *Time>1900*, *SportEvent>Summer Olympics*.

One of the basic starting points for automatic processing is to use the two basic facets "Location" and "Time" in order to reduce the number of entries. While we have no data at present, an automatic processor of existing index entries on this basis is under construction and the reduction in the number of entries should be substantial.

It is certainly worth mentioning that the "normalization" of index entries according to perspectives is not only useful in order to reduce the size of the taxonomy, but also, most importantly, in order to allow a better exploration according to perspectives. A page classified under "Rugby at the 1900 Summer Olympics" does not allow any exploration, whereas the normalized index *Sport>Rugby*, *Time>1900*, *SportEvent>Summer Olympics*, allows the user to find which events occurred in 1900, when and where Summer Olympics were held, and where Rugby tournaments were held.

Uneven Coverage

Not surprisingly for an index created without strict enforcement from a supervising authority, the index of Wikipedia is incredibly uneven, going from the extreme detail of "People from Reidsville, North Carolina" to very general headings with no specializations. Additionally, there seems to be a better coverage of certain areas, such as IT, sports, and musical groups. It is obviously outside of the scope of the present project to supply a more even coverage, both for content and for classification. While it is unfeasible to refine the existing classification, it should be rather straightforward to avoid taxonomy branches deeper than necessary by simply pruning specializations that are too deep.

Taxonomy Improvement

In the retrofit of the original Wikipedia taxonomy, we are also considering techniques such as the one implemented in the Wikipedia BitTaxonomy project (Flati, et al., 2014) that can be used in this context, with several modifications, to improve the quality of the resulting taxonomy. Although these techniques are not targeted to dynamic taxonomies, the clean separation in dynamic taxonomies between the schema, the classification of documents, and the navigation system, and the minimal requirements that dynamic taxonomies place on taxonomies make their integration in the present framework viable.

Appropriate Design

While intuitively appealing, we believe that an approach entirely based on the automatic processing of the Wikipedia subject index (e.g., Li et al., 2010) is not the right one.

We still adhere to the principle stated in (Sacco, 2000): the intension of a dynamic taxonomy is designed by an expert. There are a number of reasons for this, and here we will discuss only the two major ones: user orientation and false coordination.

Effective user orientation in a dynamic taxonomy requires an appropriate hierarchical organization of concepts, that can hardly be delegated to some sort of statistical processor. As an example, let us consider *painters*. Should a *painter* be a descendant of *Person*? Or should *painters* be disposed of entirely and instead have *painting* as a descendant of *Art*, and classify persons that are painters under *painting* ? Or should we support both views? If *painters* exists, are they going to be descendants of *Artist*, or *Visual Artist*? And if *Art* exists, should *movies* be descendants of *Art*? What about *tv series*? Most of these practical problems show that the design of a useable dynamic taxonomy requires a large amount of understanding both of data and of user expectations.

False coordination occurs when a relationship between two concepts A and B is inferred because there is a document d classified under both, when, in fact, the relationship between A and B is not useful. As an example, consider a summary page of news for a certain day in which "Democrats win in Connecticut" and "Hurricane hits Cuba". Let's simplify and assume that the summary page is classified under {Democrats, Connecticut, hurricane, Cuba}. In our framework, there is a relationship between Hurricanes and Connecticut, which can be perceived as a false relationship, in the sense that it does not convey any useful meaning.

False coordination is usually attributed to the post-coordinate approach of Dynamic Taxonomies, in which index entries are combined after the classification rather than before. Pre-coordinate indexing (e.g., {Democrats, Connecticut}, {Hurricane, Cuba}}) establishes the valid relationships when a document is indexed, whereas the post-coordinate approach considers all the permutations valid.

In fact, the problem is not really pre- vs. post-coordination, but is in general caused by the fact that the relationship inferred between two concepts is unnamed, that is: we empirically know that the concepts are related, but we do not know the meaning of their relationships.

In general, however, the user implicitly supplies a name to the relationship, e.g. Hurricane **in** Cuba, and perceives the relationship between Hurricane and Connecticut as false, because Hurricane **in** Connecticut is false (the correct relationship is Hurricane **happens in the same day as some news about** Connecticut). In short, it is a question of expectations because no disambiguation is given.

Sacco (Sacco and Tzitzikas, 2009) discusses a method based on ER modeling for disambiguating the relationship name, when this is useful.

These considerations indicate that taxonomy design requires a lot of thought, and the approach we are pursuing is a computer-assisted design system, rather than an automatic one.

Conclusions

We believe that the conceptual navigation provided by dynamic taxonomies can provide a quantum leap in the usefulness of Wikipedia. Information that is currently present but not available could be exploited. Users would be able to frame complex explorative queries in a free but guided way, taking full advantage of one of the largest knowledge bases in the world.

Our present emphasis on conceptual design is justified by the fact that most of the required infrastructure is already in place and can support heavy user loads and large infobases on inexpensive hardware. In addition to conceptual design, we also plan to investigate techniques to guarantee that the evolution of the taxonomy maintains a high degree of quality.

References

Berio, G., Harzallah M. and Sacco, G. M. 2007. Portals for integrated competence management, in: Encyclopedia of Portal Technology and Applications, A. Tatnall ed., Idea Group Inc

Flati T., Vannella D., Pasini T., and Navigli R. 2014. Two Is Bigger (and Better) Than One: the Wikipedia Bitaxonomy Project, Proc. of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL 2014), Baltimore, USA

Li, C, et al. 2010, Facetedpedia: dynamic generation of query-dependent faceted interfaces for wikipedia, Proceedings of the 19th international conference on World wide web, ACM, 2010.

Papadakos, P., Kopidaki, S.. et al. 2009. Exploratory web searching with dynamic taxonomies and results clustering, Research and Advanced Technology for Digital Libraries. Springer, 106-118.

Ranganathan, S. R. 1965. *The Colon Classification*. Rutgers University Press

Sacco, G. M. 2000. Dynamic Taxonomies: A Model for Large Information Bases. *IEEE Transactions on Knowledge and Data Engineering* 12, 2 , p. 468-479

Sacco, G. M. 2003. The Intelligent E-Sales Clerk: the Basic Ideas, Proc. INTERACT'03 -- Ninth IFIP TC13 International Conference on Human-Computer Interaction, 876-879

Sacco, G. M. 2006. Analysis and Validation of Information Access through Mono, Multidimensional and Dynamic Taxonomies, FQAS 2006, 7th Int. Conf. on Flexible Query Answering Systems, Springer LNAI 4027

Sacco, G. M. 2008. Rosso Tiziano: A System for User-Centered Exploration and Discovery in Large Image Information Bases. DEXA 2008: 297-311

Sacco, G. M. and Tzitzikas, Y. (eds.) 2009. "Dynamic Taxonomies and Faceted Search – Theory, Practice, and Experience", The Information Retrieval Series, vol. 25, Springer

Sacco, G.M. 2012. Global guided interactive diagnosis through dynamic taxonomies. In 25th IEEE International Symposium on Computer-Based Medical Systems (CBMS) , 1-6

Yee, K-P., Swearingen, K. et al.. 2003. Faceted Metadata for Image Search and Browsing, Proc. ACM CHI 2003, 401-408