Simple Pre and Post-Pruning Techniques for Large Conceptual Clustering Structures

Guy W. Mineau[†], **Akshay Bissoon**[†], **Robert Godin**[‡]

[†] Dept. of Computer Science
Université Laval
Quebec City, Quebec
Canada, G1K 7P4
tel: (418) 656-5189
fax: (418) 656-2324
mineau@ift.ulaval.ca

[‡] Dept. of Computer Science
Université du Québec à Montréal
Montreal, Quebec
Canada, H3C 3P8
tel: (514) 987-3000 x.3088
fax: (514) 987-84776
godin.robert@uqam.ca

Abstract. In (Godin et al., 1995a) we proposed an incremental conceptual clustering algorithm, derived from lattice theory (Godin et al., 1995b), which is fast to compute (Mineau & Godin, 1995). This algorithm is especially useful when dealing with large data or knowledge bases, making classification structures[1] available to large size applications like those found in industrial settings. However, in order to be applicable on large data sets, the analysis component of the algorithm had to be simplified: the thorough comparison of objects normally needed to fully justify the formation of classes had to be cut down. Of course, from less analysis results classes which carry less semantics, or which should not have been formed in the first place. Consequently, some classes are useless in terms of the information needs of the applications that will later on interact with the data. Pruning techniques are thus needed to eliminate these classes and simplify the classification structure. However, since these classification structures are huge, the pruning techniques themselves must be simple so that they can be applied in reasonable time on large classification structures. This paper presents three such techniques: one is based on the definition of constraints over the generalization language, the other two are based on discrimination metrics applied on links between classes or on the classes themselves. Because the first technique is applied before the classification structure is built, it is called a *pre-pruning* technique, while the other two are called *post-pruning* techniques.

Keywords: classification structures, simple pruning techniques, large size applications

Type of submission: long paper

1 INTRODUCTION

Conceptual clustering algorithms form classes of similar objects. In order to do so, objects are compared to one another to detect any similarity that they may convey. Based on this similarity, common generalizations of objects are produced: they justify the formation of classes, they also serve as characteristic descriptions for the classes of objects from which they were inferred.

With our classification algorithm, called MSG (Mineau & Godin, 1995), the generalization language is determined without comparing the objects: it is generated solely based on the description of the individual objects. It must be general enough to include common generalizations among objects. However, if too general, it may dilute the semantics that the classes would otherwise convey. In order to avoid being overly general, a validation of the generalization language could take place. To be thorough, it would imply pairwise comparisons of objects in order to ensure that the generalization language is appropriate to represent the *maximally specific* common generalizations among them. With large data (knowledge) bases, this would bring us back to the complexity problem that prompted us to develop the MSG (summarized in Section 2 below).

---

[1] In this text, even though our clustering method is unsupervised, we use the terms *classification* or (conceptual) *clustering* interchangeably.

Since we can not guarantee that the generalization language as produced with our technique will be appropriate to find maximally specific generalizations among objects, we will therefore need pruning techniques to eliminate classes that do not carry sufficient semantics. Since our goal is to produce efficient classification algorithms for very large datasets, and since pruning techniques must be part of these algorithms, we need to devise pruning techniques which are themselves very simple to apply (otherwise we may lose the computational gains obtained by the MSG).

Section 3 below describes a pruning technique applied on the generalization language itself (the definition of constraints on the generalization process). Section 4 describes two pruning techniques applied on the resulting classification structure: one eliminating *weak links* between classes, the other one eliminating classes which do not show a satisfying distribution of their object set into their children nodes (children classes). Because the first technique is applied before the classification structure is built, it is called a *pre-pruning* technique, while the other two are called *post-pruning* techniques.

2 Conceptual Clustering for Large Size Applications: the MSG

Classification algorithms provide a characterization of the data on which they are applied. This is useful for indexing, retrieval, knowledge discovery and explanation purposes. Industrial applications often deal with datasets that are too large to be considered by traditional approaches. Consequently, some alternative techniques needed to be devised. In (Mineau & Godin, 1995) we introduced one such technique called MSG (Method of Structuring by Generalization) which proposes to divide the classification process into two distinct and successive phases: 1) the formation of *promising* classes, and 2) the subsequent production of characteristic descriptions, but only for the classes identified as relevant by the particular applications which will use the data. Doing so eliminates the need for the class formation process to produce a *complete* characteristic description for each class in order to justify its existence, since it may well be the case that only a portion of the resulting classes will effectively be used by the applications. Consequently, a similarity function that is less time consuming than those that compute maximally specific common generalizations can be used in the class formation phase.

As mentioned in Section 1 above, the MSG first determines a generalization language L_g that represents a subset of all possible generalizations of the objects. Then the objects are mapped onto L_g. The meeting points in L_g will be used to form and describe classes. The idea is to bound L_g so that the determination of common generalizations among objects becomes manageable in terms of complexity. In fact, the complexity of the MSG is directly proportional to the production of L_g and to the number of objects. Naturally, the actual classes found by the MSG are highly dependent on L_g. To explain how the MSG works, Section 2.1 below presents our dataset; while Section 2.2 gives an example of how it would be applied on this dataset.

2.1 Our Application Dataset

As our testbed we used a small knowledge base on impressionist paintings from the National Gallery of Art of Washington D.C. Each painting (from a subset of 100 paintings) was described using conceptual graphs (Sowa, 1984). A part of the description gave bibliographic notes on the painting: author, material, year, era, title, collection, etc.; some other part described the subject matter. Since we aim at developing classification methods for both data and knowledge base environments, we chose to use complex objects from the start. Each object was described by an average of 30 (binary) relations, so altogether the knowledge base contained 3000 (binary) relations.

2.1.1 The Representation of Complex Objects

A conceptual graph (CG) is made of two types of nodes: *concepts* representing objects (either concrete or abstract), and *relations* representing semantic links between concepts. Both nodes are typed; types come from a thesaurus called a *type hierarchy* T. Identifiable individuals are referred to by individual markers. I is the set of all individual markers. All concept nodes in a CG represents an individual of some type, either represented by an *individual marker* (if explicitly known) or by a quantified variable (otherwise). When the variable is existentially quantified, then the * symbol, called the *generic marker*, may be used[2]. An individual represented by a concept node is said to conform to the type of the concept. To that effect, a conformity relation :: is defined between elements of T and $I \cup \{*\}$. We call the elements of $I \cup \{*\}$ *referents*. A concept is thus described using a type t and a referent which conforms to t.

Arcs connect concept nodes to relation nodes in a non-ambiguous way, as described in B, the *canonical basis* of the system. B encodes the *signature* of each relation. For each relation r, B gives the number of parameters of r, their order, and the maximal type of each parameter (which are all concept nodes). By doing so, overgeneralizations are avoided. Canonical formation operators define how concepts and relations may be connected to form new graphs. Consequently, all graphs represented in a CG system are derived from <T,I,::,B>, the *canon* of the system.

2.2 Applying the MSG on our Dataset

With our application database, the MSG must generalize a set of conceptual graphs to produce L_g. Then it must map each object onto L_g to find meeting points in this space of generalizations. These meeting points will be considered as common generalizations. Each common generalization may give rise to the formation of a new class, if the class does not already exist (because of some other common generalization). Also, we may have background information validating the formation of classes (like semantic constraints for instance (Bournaud, 1996)). With our implementation of the MSG, no such background knowledge was used in the formation of the classes.

At this point we would like to point out to the reader that the generation of (part of) L_g from an object o, called the *completion* of L_g with regard to o, and the mapping of o to L_g can be done simultaneously. L_g will be complete and all objects will be mapped onto it once they are all scanned. This is an iterative process over the set of all objects. For *n* objects, it is thus a O(*n*) process, as long as the completion of L_g with regard to any object is done independently of any other object. In order to keep the complexity associated with the production of the classes as low as possible, the completion of L_g with regard to any object must also minimize its associated complexity, as discussed in the Section 2.2.1 below.

2.2.1 Generalizing Conceptual Graphs

There are three generalization operators that can be applied onto a single object in order to generalize it (Michalski, 1983): 1) *turning a constant into a variable*, i.e., replacing an individual marker by the generic marker, 2) *generalizing a term*, i.e., generalizing the concept type of a concept or the relation type of a relation, and 3) *dropping a condition*, i.e., deleting a subgraph. Each of these operations preserve the truth-value of a graph, that is the generalized graph is a logical entailment of the graph it is derived from.

In our implementation of the MSG, to generate the completion of L_g with regard to an object o, we apply these three operators in a predefined sequence. First, we use operator #3, producing a set of triplets: $\langle concept_1, relation, concept_2 \rangle$[3]. So, from a single graph having *k* relation nodes, we produce *k* initial generalizations of the graph by selecting alternatively each triplet (or by dropping every other triplet)[4]. After this operation, the cardinality of G_i, the set of generalizations of graph i, written $|G_i|$, is exactly *k*. Then each of these *k* triplets is generalized according to operator #1, which introduces additional generalizations into G_i. Because relations do not carry referents, only

[2] In our test database, all concepts are either generic (existentially quantified) or represent known individuals.

[3] We assume that all relations in a conceptual graph are binary relations since n-ary relations can be transformed into n binary relations and an additional concept.

[4] This step is done first so that, from that point on, the following generalization steps could be applied in cases where instead of having triplets, we would have simple keyword-based descriptors or attribute-value pairs, making our method capable of handling other types of datasets (like those found in traditional machine learning literature).

the two concepts can be generalized using operator #1. This results in having $|G_i|$ four times larger than what it was previously, in the worst case (i.e., if no common generalizations if found at that point among the triplets describing graph i); so $|G_i| = 4k = O(k)$. Finally, operator #2 is applied on all concept types (and on no relation type[5]) of all triplets in G_i. In the original version of the MSG, in order to keep the complexity of the completion of L_g as low as possible, we used only a single generalization term: the ? wild card symbol, which symbolizes *something*. Doing so would ensure that the generalized triplets would cover triplets of other graphs as well, at a minimal cost. With only two concept types for each concept of each triplet, we ended up multiplying the size of G_i by 4; so after this last operation, $|G_i| = 16k$, in the worst case (if no common generalizations exist among the triplets of G_i). Consequently, for all n objects, we have $L_g = \cup_i G_i$ ($\forall i \in [1,n]$); and we have $|L_g| = 16nk = O(nk)$, in the worst case. It was expected and was verified empirically that the hidden constant is greatly reduced with the overlap that exists among the different G_i (see below).

Once the set of classes if found, classes may be linked in such a way as to represent the partial order of inclusion ($\subseteq$) defined over the sets of objects that they represent. The MSG explicitly represents this $\subseteq$ relation among classes in such a way as to minimize storage and computing resources: 1) transitivity is used to avoid the explicit representation of all $\subseteq$ relations among classes, 2) classes are inserted into the classification structure according to the cardinality of the set of objects that they represent (the smaller sets first). Doing so results in a subquadratic overall complexity for our method in terms of computing time (in average), while the storage requirements remain linear in n in all cases, provided that an upper bound on $|G_i|$ is set (which we found to be a reasonable assumption for most practical cases that we encountered). For a more detailed presented of the MSG itself and its related complexity, the reader is referred to (Mineau & Godin, 1995) and to (Godin et al., 1995a).

2.2.2 Overgeneralizations

It should be clear that since operator #3 only uses one general unifier for obvious computational reasons, this favors the formation of overly general classes. For instance, a class formed solely based on the following triplet: <?,r,?>, where r is a relation , represents all pairs of objects being in relation r with one another. This may be useful in occasional circumstances where relations may be used as indexes, but it does not help to characterize the data since the ? symbol may represent very heterogeneous objects. In fact, it does not carry sufficient semantics to be useful to semantic driven applications. However, we could use T, the hierarchical vocabulary of the system, i.e., the type hierarchy, in order to generalize concept types with more specific unifiers than the ? symbol. The idea is to search T from some concept type (the one in the triplet to generalize) towards the top of the type hierarchy. The length of the search path could be bounded by some value in order to avoid including too general unifiers, as usually found near the top of type hierarchies. Of course this additional search has a cost both in terms of storage and computing resources: it would increase the hidden constant mentioned above because the height of the type hierarchy would now play a part in the complexity of our method. However, in practice, it is often the case that type hierarchies are relatively shallow. In any case, by carefully choosing some bound on the number of types used to replace the concept types in the original triplets, we hope to increase the semantics conveyed by the classes, and at a reasonable computational cost. Using intermediate concept types (those appearing in the original triplet and the top of the type hierarchy), we hope to add useful classes into the classification structure. By avoiding too general unifiers, we hope to eliminate classes which carry overly general semantics and which are thus useless to the applications. Section 3 below presents the findings of this enquiry.

[5] We found that in practice, the relation type hierarchy is so shallow that generalizing relations immediately introduces overgeneralizations. Therefore, we decided to generalize concept types only.

3 PRE-PRUNING

This section presents comparative results pertaining to the use of T, the type hierarchy of the domain, in the generalization process of our classification algorithm. In the first set of experiments, T is used as such (without modifications); in the second set of experiments, T is broken down into a set of disjoint type hierarchies T*. In T every pair of concept types may be unified through some common supertype, therefore T is called a *unified vocabulary*. In contrast, T* is called a *fragmented vocabulary*. Figure 1 below shows a subset of T[6]; its fragmented counter-part is produced by deleting the bold types from T: Physical_Object, Artefact and Material. As mentioned in Section 2.2.2, the search for more general terms (supertypes) of a concept c in a graph will consider all supertypes from the concept type of c up a certain number of levels *l* in T. With a unified vocabulary, it is expected that some overly general supertypes may be introduced even with low values of *l*. It all depends on the initial concept type of c, i.e., on how general is c in the first place. Of course, the maximum length path for T will be greater than that of T*, leading to more general classes, and possibly to overly general classes. Consequently, it is expected that the classes produced under T* will be more specific than those produced under T, for a lower computational cost. However, we must point out that T* requires background knowledge on maximal unifiers (generalizations) of concept types in order to be applicable. The acquisition of this information requires human intervention, which may not be possible to have: this information or the human resources to acquire it may not always be available.

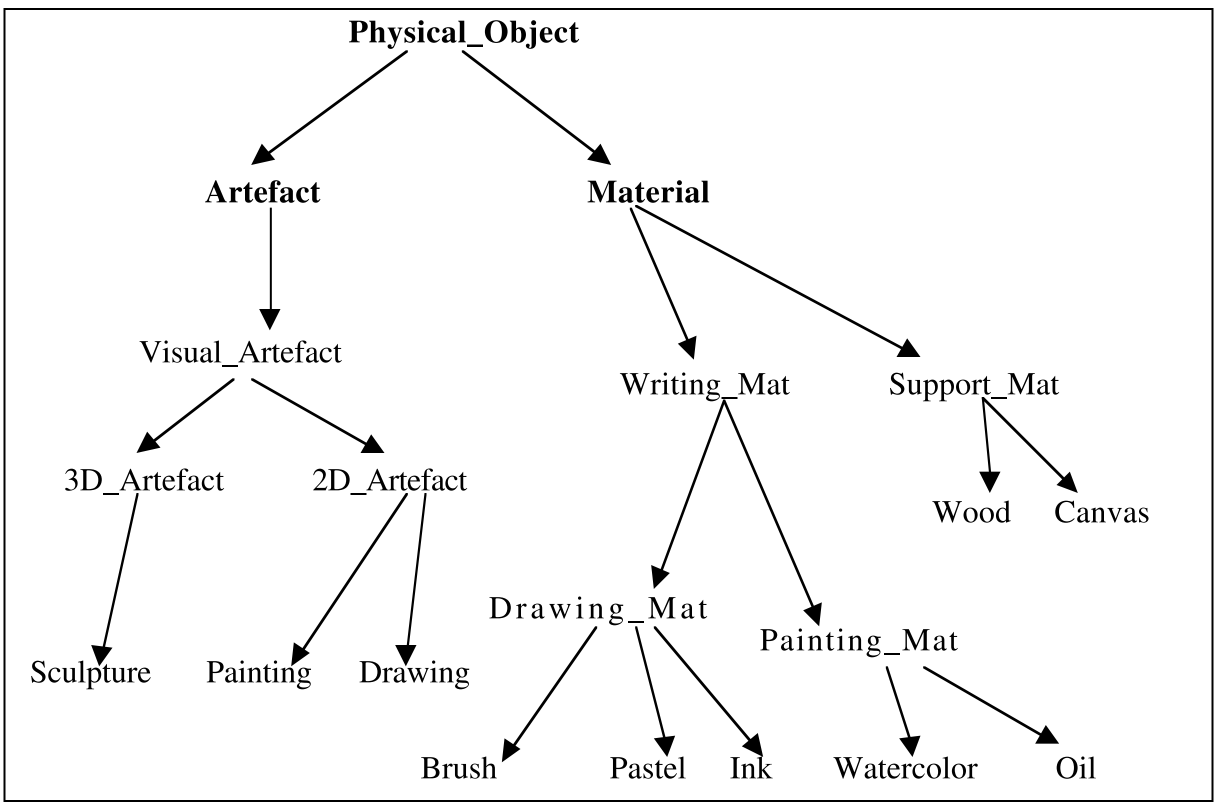

Figure 1. Part of the type hierarchy of the domain.

In the experiments reported below, we increased the maximum path length *l* in order to compare the resulting classification structures in terms of standard performance metrics: number of nodes, of triplets (descriptors), of links, computing time, etc. In the subsequent figures, each classification

[6] Even though in our example the type hierarchy is a tree, in the general case there is no such restriction.

structure is thus identified by h_l. For instance, h_0 is the classification structure produced with no generalization of the concept types; while h_1 implies that each concept type was replaced by its immediate supertypes. With our type hierarchy, the maximum length path is 7 for T and 5 for T*, so h_7 and h_5 are the last classification structures considered by each approach respectively. The tests reported in the following subsections were conducted on a Sun Workstation 20.0 under Solaris 2.5; the algorithms were written in C.

3.1 Number of Nodes

Figure 2 and Figure 3 show the growth in number of nodes (using T and T* respectively) as the maximum path length increases. The reader should notice that in both cases, the number of nodes increases up to some point and then decreases. This was expected, as the introduction of triplets creates new nodes (classes) and eliminates older nodes. Depending on which phenomenon is more important at one point, an increase or a decrease in the total number of nodes is observed. It is normal that new triplets (generated through the generalization of some object) introduce new nodes unless these triplets were already present in the classification structure. In that case, they were previously introduced by the generalization of some other object. Since they now cover more objects, they migrate to more general classes. If the more specific classes from which they migrate become empty, then these nodes are eliminated. We call this phenomenon *total inclusion* of a class description by some other class description. The total number of nodes in a classification structure may decrease as a result of total inclusion.

First, we must say that the number of nodes found by the initial MSG (where only the ? symbol was used as a unique general unifier) was 750. Consequently, with the unified vocabulary we should bound l, the maximum length path, by 3. Otherwise we obtain more classes than with the original version of the MSG. Comparing h_3 of Figure 2 with h_3, h_4 or h_5 of Figure 3 (which are sensibly the same), we see that the fragmented vocabulary improves the number of classes by more than 20%. However, some of the remaining classes may still be too general classes since we observe the same order of improvement with l = 3, 4, or 5 independently. Obviously, other pruning techniques will be needed to further simplify the classification structure (as proposed in Section 4 below).

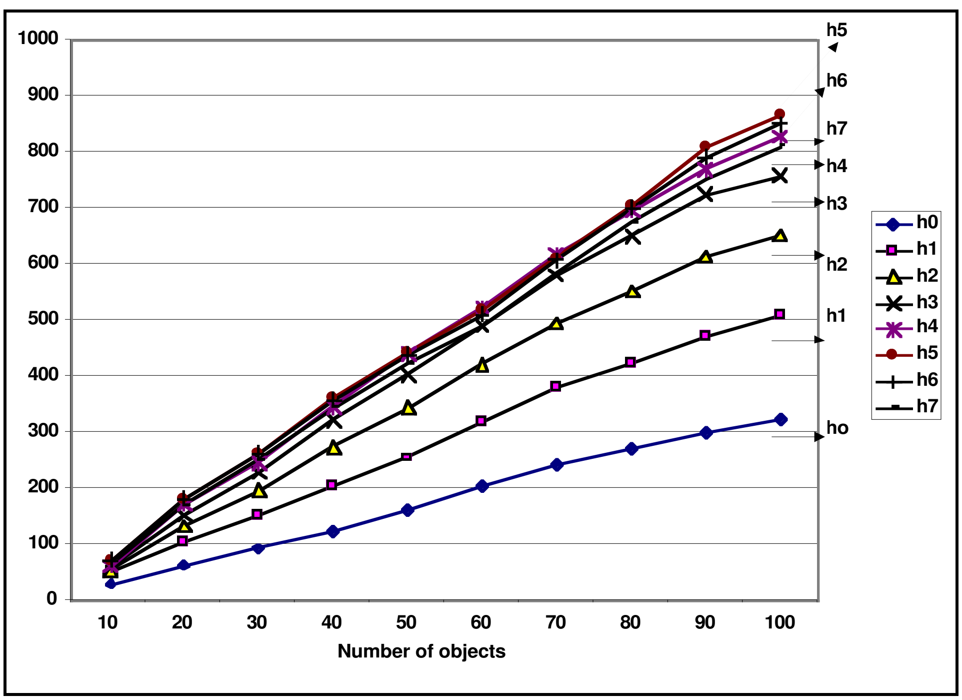

Figure 2. Number of nodes vs number of objects, under T.

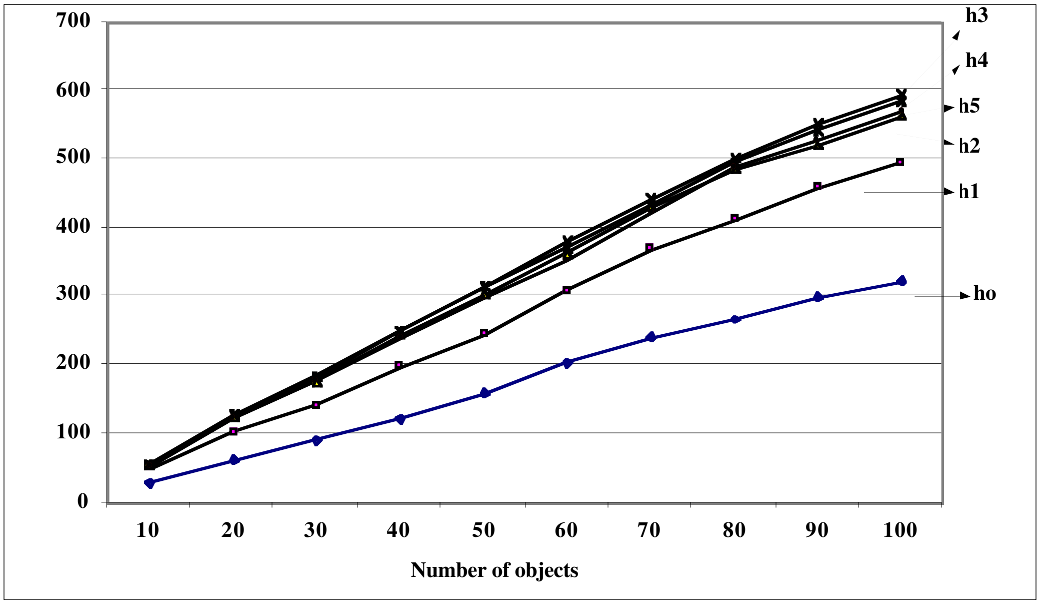

Figure 3. Number of nodes vs number of objects, under T*.

3.2 Number of Links

Figure 4 and Figure 5 show the growth in number of links (using T and T* respectively). We see that the growth in the number of links is directly proportional to the growth in the number of nodes. This was expected because both have a linear growth in *n*, the number of objects. As a matter of fact, since the number of triplets in each G_i is O(*k*) and since only non-empty intersection sets of triplets create classes, both the number of nodes and links are O(*nk*) in the worst case. So the gain in the number of nodes implies an equivalent gain in the number of links, i.e., about 20%. Further reduction in the number of nodes will result in equivalent gains in the number of links, simplifying the classification structure even more. Additional pruning is thus also relevant to improving (reducing) the number of links in the resulting structure.

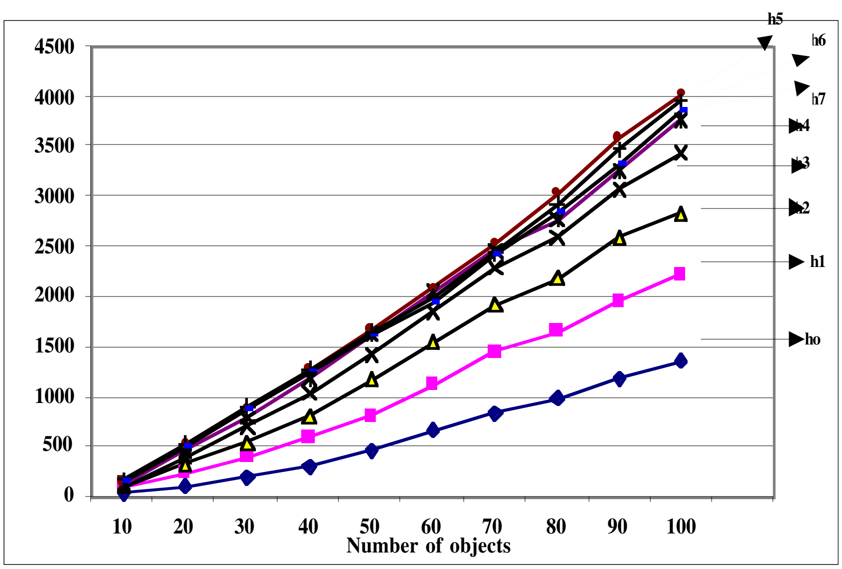

Figure 4. Number of links vs number of objects, under T.

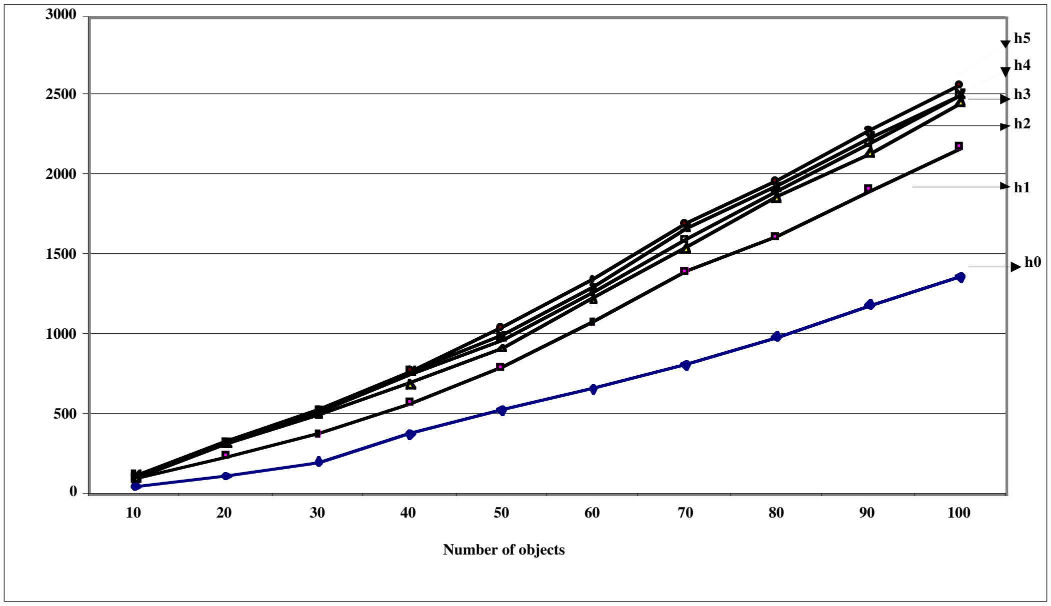

Figure 5. Number of links vs number of objects, under T*.

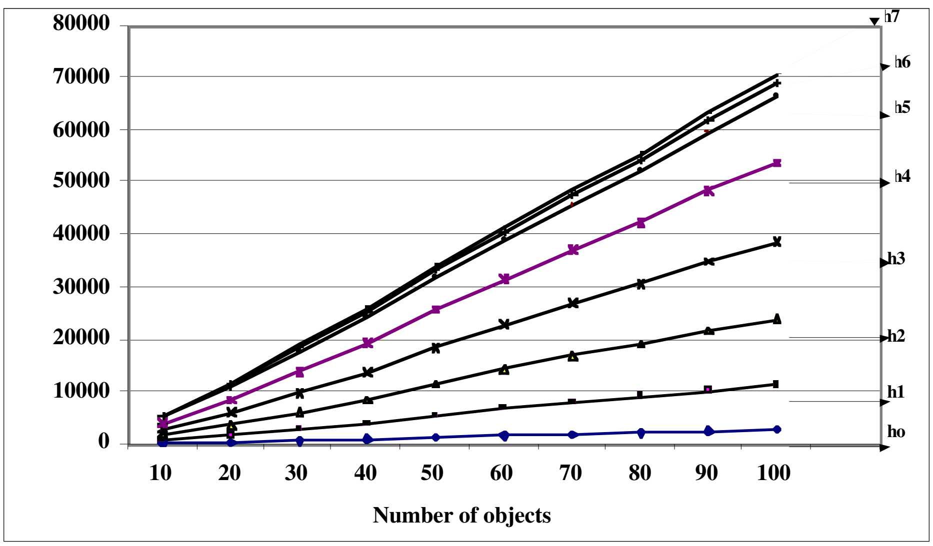

Figure 6. Number of triplets vs number of objects, under T.

3.3 Number of Triplets

As expected, the increase in the number of unifying terms resulted in an important overall increase in the size of the classification structure. As theoretically assessed in Section 2.2.1, even though the expected size of the classification structure was O(*nk*) with a hidden constant of 16 in the worst case, the actual evaluation of this constant was around 1.5 (with the original version of the MSG). This was due to the fact that the triplets containing ? symbols were so general that they were shared by many objects, reducing the size of L_g. With the proposed pre-pruning approach, the generalizing terms are more specific than the ? symbol, and the triplets so generated are shared by fewer objects. Consequently, we observed a tremendous increase in the number of triplets. Provided that we compute h_3 using T (see Figure 6), we end up with a classification structure which now carries about 39000 triplets from an initial set of 3000 triplets, which sets the hidden constant at 13. With

T*, the hidden constant would be 5 (see Figure 7), as a total of about 25000 triplets is produced for h_3 thru h_5 independently. Even though the constant is much smaller with T*, our approach still has substantial space requirements; this is the definitely its most constraining effect. However, there is no doubt that the resulting classes will be more specific than those of the classification structure produced by the original version of the MSG. As a matter of fact, we end up with 20% fewer classes while the number of triplets is increased by 3.33. So each class will have a more specific description in terms of the number of triplets composing it, and in terms of the concept types used in these descriptions (because overly general triplets are not part of the structure anymore). This proves that, after pruning, the classes carry more specific semantics than before.

3.4 Computing Time

In either case, as can be seen in Figure 8 and Figure 9, even through there is a slight cost in terms of computing time, the actual platform is fast enough to handle large sets of objects. Extrapolation from either one of these two figures, but especially from Figure 9, tells us that very large classification structures could be produced in precompilation mode in reasonable time. Consequently, computing time is not considered to be a problem.

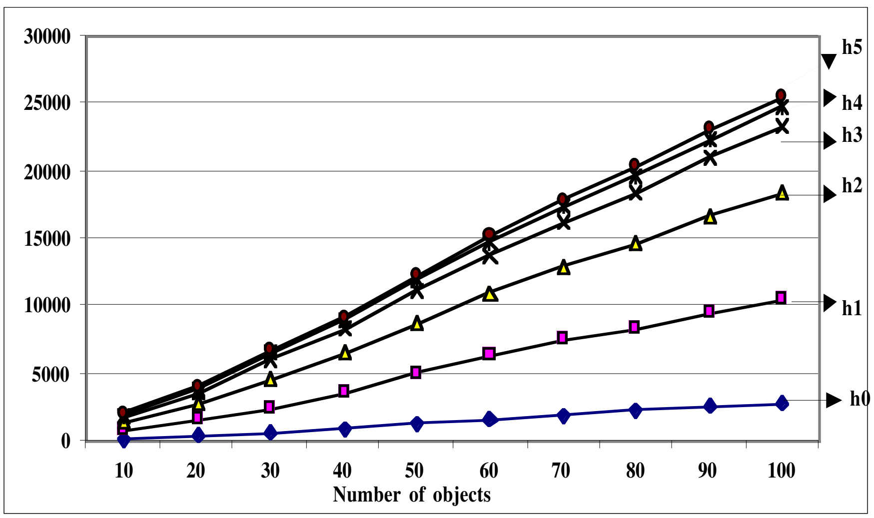

Figure 7. Number of triplets vs number of objects, under T*.

3.5 Summary of Comparison of MSG using T and T*

It is clear from the figures above that all performance metrics indicate that the fragmented approach, when available, should be preferred: it minimizes the added complexity associated with multiple generalizations of concept types. For instance, h_3 in the unified vocabulary is worse than any classification structure built with the fragmented vocabulary (according to *all* performance metrics). Furthermore, it was shown that the classes produced under this fragmented vocabulary approach are more specific than those produced using a unified vocabulary.

Compared with the original version of the MSG, the number of nodes and of links was cut down by 20%. However, the size of the resulting classification structure may be a problem: memory requirements are not negligible as the hidden constant went up from 1.5 to 5 times the size of the original knowledge base. Finally, the increase in computing time was not considered a problem.

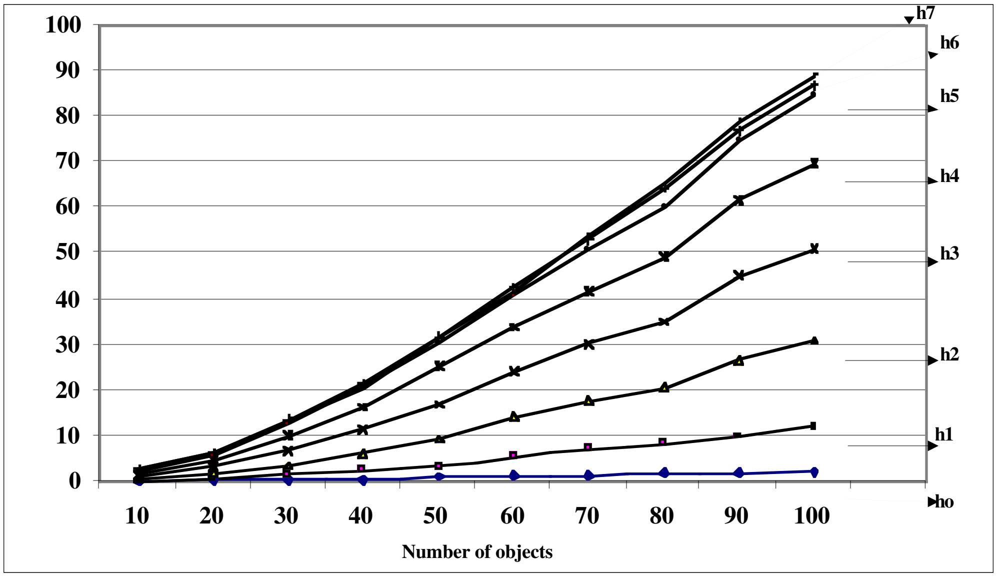

Figure 8. Computing time vs number of objects, under T.

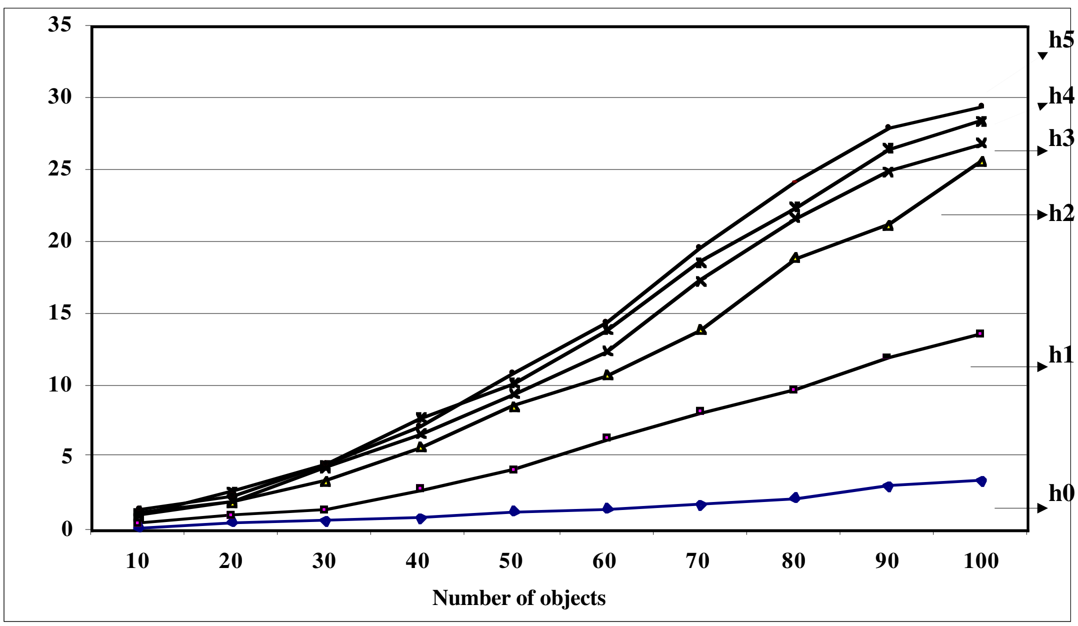

Figure 9. Computing time vs number of objects, under T*.

The evaluation done in this section reveals that a fragmented vocabulary approach should be preferred over a unified vocabulary approach. However, as shown with our test database, the maximum length path *l* could have been set to any number between 3 and 5 and the gain would have been about the same (in terms of the metrics that we used). Consequently, pruning the resulting classification structure may prove to be useful for simplifying it even more, as discussed in Section 4 below.

4 POST-PRUNING

4.1 Pruning Useless Links

A link always connects two distinct classes. Let us define a link k as connecting a class c_1 to some other strictly more specific class c_2 (written $c_1 < c_2$). Since we do not keep transitive links in the structure, c_1 is said to be the *immediate parent* of c_2, and c_2 is said to be the *immediate child* of c_1. In general, let us define the set of immediate children of a class p as C(p) = {class c in the classification structure, for which there is a link between p to c (with p > c)}. Similarly, we can define the set of the immediate parents of a class c as being P(c) = {class p in the classification structure, where $c \in C(p)$ (with p > c)}. Finally, let us define the set of objects represented by a class c as E(c), called the *extension set* of class c.

Either in terms of an entropy metric such as found in traditional decision tree literature (Quilan, 1986), or in terms of the discrimination power of particular objects (Dietterich & Michalski, 1983), or in terms of the optimal height of a classification structure, the most useful links are those which relate a class p to a subset of its children nodes C'(p) such that $\forall c \in C'(p)$, $p \in C(p)$ and $abs((|E(c)|/|E(p)|) - 0.5) < \varepsilon$, where abs() is the absolute value function and ε is a small constant value. Let us define the *selectivity* of link k as sel(k) = (|E(c)|/|E(p)|). Knowing that each time a link is replaced by some other link to a more specific class, the selectivity of the link is less than the selectivity of the deleted link, we now propose a pruning algorithm that deletes links whose selectivity is strictly higher than a certain threshold[7].

1. $M \leftarrow \emptyset$
2. $L \leftarrow$ {all links in the classification structure}
3. Compute sel(k) for all links k in L.
4. For each link k (between p and c, with p > c) such that sel(k) > t, a predefined threshold, the content of c is copied downwards into all its immediate children nodes, i.e., into each class of C(c), so that no information is lost[8]. Then $M \leftarrow M \cup \{p\}$ (for the next iteration).
5. Then each class p in P(c) is connected to each class c in C(c) unless there is already a path between p and c in the classification structure.
6. Finally, c and its incoming and outgoing links can be deleted.
7. If $M \neq \emptyset$ then: $L \leftarrow$ {set of links emanating from the classes in M}, $M \leftarrow \emptyset$, goto step 3.
8. Stop.

So this algorithm is simple enough to be applied on large classification structure. Though we did not assess its asymptotic complexity, since the experiments reported below show that a very small portion of the links were affected by it in our dataset, and since the number of links is $O(nk)$ as mentioned before, it is expected that in average its complexity would probably be bounded by $O(nk)$ as well. Anyhow, it had very limited effect on the classification structure that it was applied on: few links were deleted, and these deletions required substantial additional memory because of redundancy (step 4) and the addition of links (step 5). When applied onto h_5 (built with a fragmented vocabulary) with a threshold of 0.65, the number of links in the [0.40..0.60] selectivity interval increased by only 2%. Furthermore, no gain in the number of nodes was acheived while all other performance metrics showed important drawbacks. This is sufficient to state that the negative side-effects of the method outweigh the slight increase in the selectivity of certain links. Therefore we can not see where such a technique could be applied successfully on large size applications where storage requirements are important.

---

[7] If we deleted links whose selectivity is below some threshold, we would produce links whose selectivity would be even lower than the previous links, i.e., still lower than this threshold, taking us further away from our goal of having sel(k) as near 0.5 as possible.

[8] At this point we would like to remind the reader that our classification structure is an inheritance network.

4.2 Pruning Useless Classes

The other pruning technique that we propose in this section uses a discrimination metric, ent(), based on some entropy evaluation of each class in the classification structure. It is evaluated as: $\mathrm{ent}(p) = (\Sigma_{\forall c \in C(p)}\ \mathbf{abs}(\ |E(c)| - (|E(p)|/|C(p)|)\ )\ )\ /\ |C(p)|$. It favors equal distribution of the set of objects of a class into its immediate children classes (for reasons mentioned in Section 4.1 above).

The classes for which this metric evaluates above some threshold t will be considered useless and will be deleted from the classification structure as was done in the previous algorithm, involving the addition of new links and the copy of triplets (necessary redundancy).

1. $M \leftarrow \varnothing$
2. $L \leftarrow$ {all classes in the classification structure}
3. Compute ent(p) for all classes p in L.
4. For each class p such that ent(p) > t, a predefined threshold, the content of p is copied downwards into all its immediate children nodes, i.e., into each class of C(p). Then $M \leftarrow M \cup P(p)$ (for the next iteration).
5. Then each class q in P(p) is connected to each class c in C(p) unless there is already a path between q and c in the classification structure.
6. Finally, p and its incoming and outgoing links can be deleted.
7. If $M \neq \varnothing$ then: $L \leftarrow M$, $M \leftarrow \varnothing$, goto step 3.
8. Stop.

Again this is a simple algorithm. Its asymptotic complexity has not be assessed, but since the number of nodes in the structure is $O(nk)$ and since the number of parents of a node is $O(k)$, it is expected that it be around $O(nk^2)$. This will be investigated some more in a near future.

In the experiments reported below, three different simple-to-compute thresholds were used: s1 = $(xm)/100$, s2 = $(x(m/k))/100$, and s3 = x, for $x \in [0,100]$, where $m = |E(c)|$ and c is the class for which ent(c) is computed. All thresholds express an acceptable variability on the entropy of a class. The class will be deleted if ent(c) > t. The two thresholds, s1 and s2, are called *local* thresholds since $|E(c)|$ is part of their evaluation; while s3 is a *global* threshold since it is fixed for all classes. Of course s3 requires no computation, but is harder to assess. Unlike s1 which has more impact on more specific classes, s3 is more effective near the top of the classification structure: no classes c for which $|E(c)| <$ s3 will ever be pruned. Threshold s2 is more drastic than the other thresholds, forcing a partition of extension sets of a class into its immediate children classes. Since our classification structure is not a tree, s2 may overprune the structure. The precise determination of correct values for each of these parameters can be determined based on experience. The following sections present our findings with that regard, and compare the three approaches (thresholds) that were used.

4.2.1 Number of Nodes

Figure **10** shows how the choice of the threshold affects the elimination of nodes in the classification structure. In order to fully compare the different thresholds, the next subsections evaluate some performance metrics on classification structures having the same number of nodes after having been pruned using s1, s2 or s3 as threshold. As expected, s2 produced a drastic pruning right from the start: this may introduce important degeneration of the classification structure.

4.2.2 Redundancy: the Number of Links and Additional Triplets

Figure 11 gives a comparison basis for classification structure h_5 (build using a fragmented vocabulary) and pruned using the different thresholds. For the three thresholds, the number of links decreases gradually. The more drastic gain in the number of links is obtained with s2.

However, **Figure 12** shows the added complexity in terms of triplet redundancy, and s2 is identified as the worse case with regard to that parameter which was earlier identified as a potential problem with large size applications (see Section 3.5). In these case, s1 or s3 could be used as threshold.

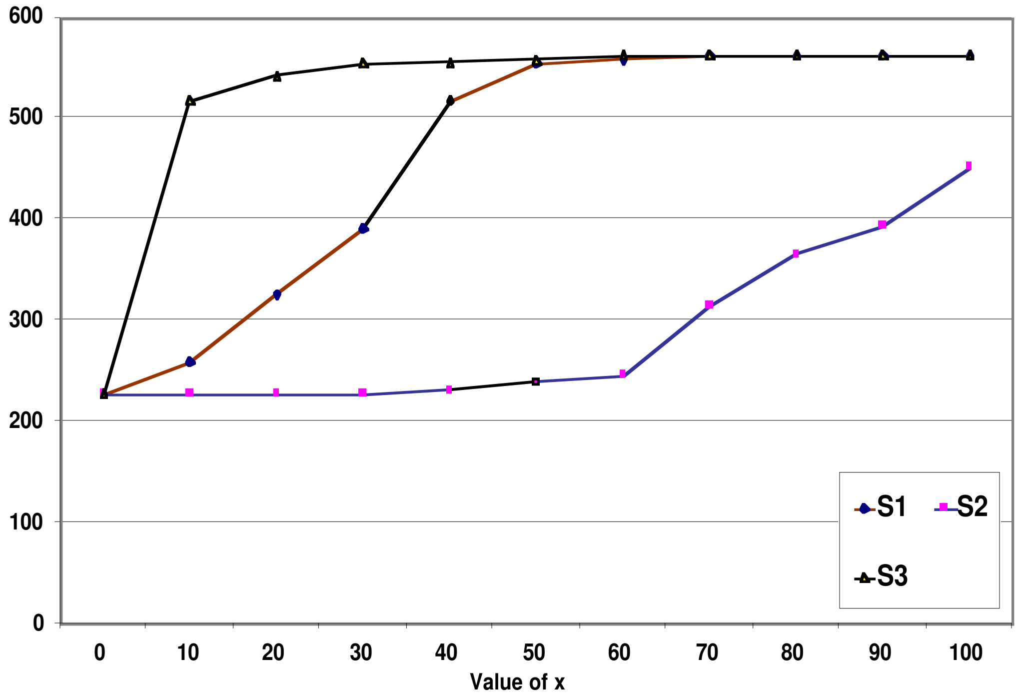

Figure 10. Number of nodes vs *x*.

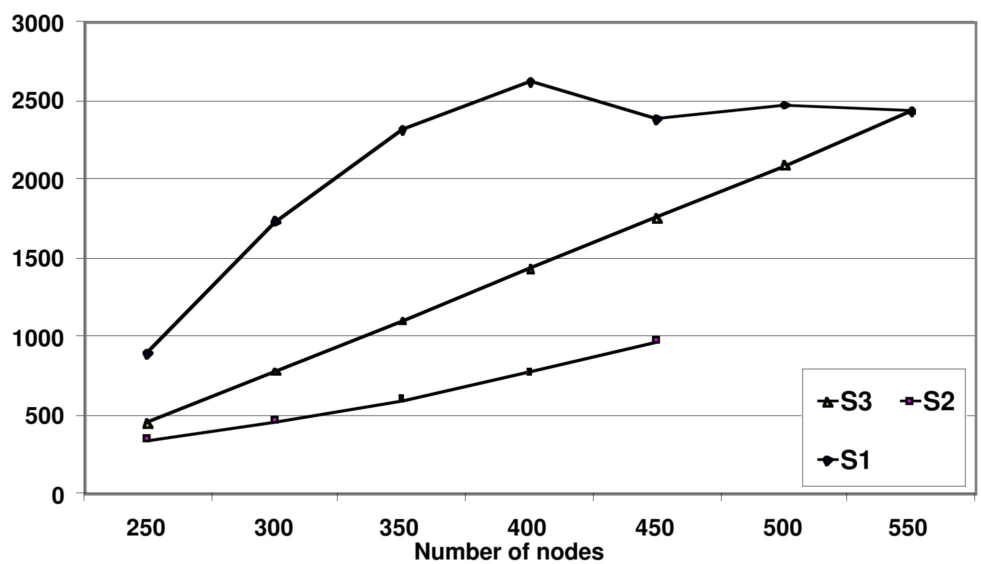

Figure 11. Number of links vs number of nodes.

4.2.3 The Height of the Pruned Classification Structure

As shown in Figure 13 below, the gain in average height is not significant. Once again, s2 is identified as potentially generating too much pruning: a classification with not enough height is a typical sign of a degenerate structure (which would not characterize the domain well). Here also, s1 and s3 seem to have a similar effect on the structure, but on different nodes. However, like with the number of links, s3 performs better than s1.

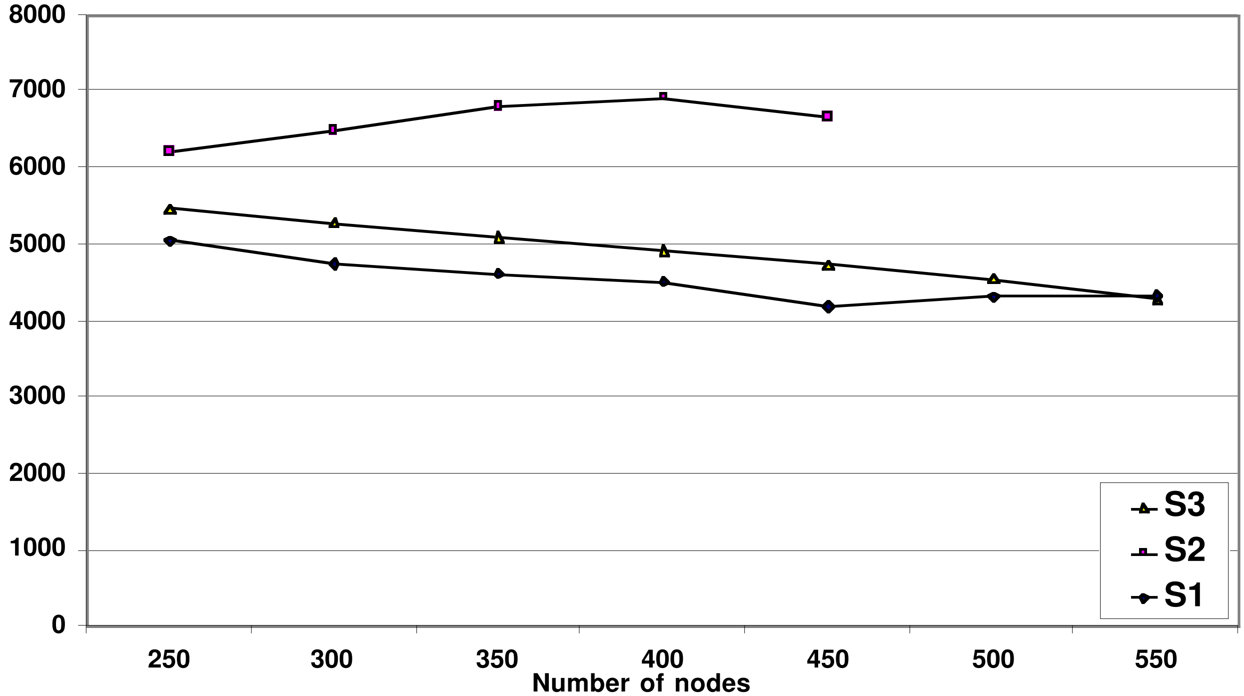

Figure 12. Number of triplets vs number of nodes.

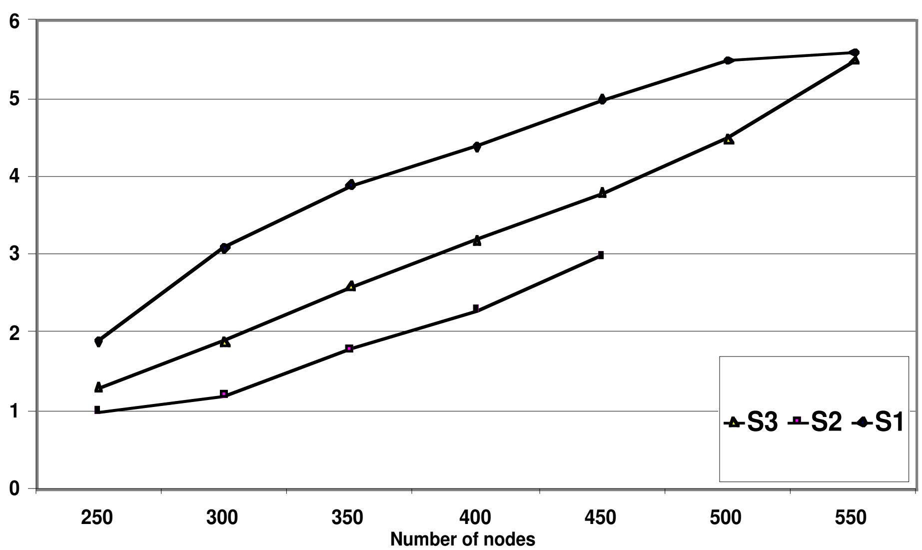

Figure 13. Average height vs number of nodes.

5 Conclusion

In summary, this paper presented simple pre and post-pruning techniques for large classification structures. The use of a fragmented vocabulary in the generalization process of objects clearly improves the standard performance metrics used to assess the quality of such structures. Among other things, the number of nodes and links was reduced by at least 20%. The only drawback is memory requirements (increased by a factor of 3 compared to the original version of the MSG) which may be a serious problem. In any case, it is clear that class descriptions are more specific than with either the original version of the MSG or its unified vocabulary version. It was also clear that additional pruning could reduce even more the size of the resulting classification structure. Therefore, this paper also presented post-pruning techniques.

The first post-pruning technique aimed at deleting weak links in terms of their discrimination power. This experiment failed with our dataset, and seemed to indicate that links are rather incidental with regard to the characterization of a domain, and that therefore, there is not much point in trying to change the topology of a classification structure in that way. Consequently we considered post-pruning methods based on a simple-to-compute entropy metric on classes. Three different thresholds, allowing different variability on the entropy of a class, were used: s1, s2 and s3. Threshold s2 proved to overprune the structure. Threshold s1 adapts to the class for which it is used: it is a local threshold; while threshold s3 is fixed for all classes. Based on a percentage of the extension set of the class, s1 has more effect on specific classes; while s3 has more effect on general classes. Threshold s3 performs better than s1 in all cases, except for a slight increase in the number of triplets. However, depending on the availability of background knowledge, one may be harder to assess than the other. In any case, both techniques proved to be useful. Combined with a pre-pruning approach, additional simplification of a classification structure can be achieved.

Future developments involve the precise assessment of the asymptotic complexity of these methods, and their evaluation on other application domains. Despite these upcoming results, it is clear that the methods proposed in this paper are simple to implement and apply, and it is foreseen that their complexity is manageable. This remains to be fully investigated and will be the subject of a forthcoming paper.

6 References

Bournaud, I., (1996). Regroupement conceptuel pour l'organisation de connaissances. Ph.D. Thesis, Université Paris VI – Institut Blaise Pascal. France.

Dietterich, T.G. & Michalski, R.S., (1983). A Comparative Review of Selected Methods for Learning from Examples. In: Machine Learning: An Artificial Intelligence Approach. R.S. Michalski, J.G. Carbonell & T.M. Mitchell (Eds.). Morgan Kaufmann. 83-129.

Godin, R., Mineau, G.W. & Missaoui, R., (1995a). Incremental structuring of knowledge bases. In: Proceedings of the 1st International Symposium on Knowledge Retrieval, Use, and Storage for Efficiency (KRUSE-95). Santa Cruz, California. Août. G. Ellis, R.A. Levinson, A. Fall & V. Dahl (Eds.). Department of Computer Science, University of California at Santa Cruz. 179-193.

Godin, R., Mineau, G.W., Missaoui, R. & Mili, H., (1995b). Méthodes de classification conceptuelle basées sur les treillis de Galois et applications. In: Revue d'Intelligence Artificielle, volume 9, no 2. 105-137.

Michalski, R.S., (1983). A Theory and Methodology of Inductive Learning. In: Machine Learning: An Artificial Intelligence Approach. R.S. Michalski, J.G. Carbonell & T.M. Mitchell (Eds.). Morgan Kaufmann. 41-81.

Mineau, G.W. & Godin, R., (1995). Automatic Structuring of Knowledge Bases by Conceptual Clustering. In: IEEE Transactions on Knowledge and Data Engineering, volume 7, no 5. 824-829.

Quinlan, J.R., (1986). Induction of Decision Trees. In: Machine Learning Journal, volume 1. Kluwer Academic. 81-106.

Sowa, J.F., 1984. *Conceptual Structures: Information Processing in Mind and Machine*. Addison-Wesley.