

Applying Category Utility and Partition Utility Functions to Semantic-Based Conceptual Clustering Approach¹

Ivo Marinchev

Institute of Information Technologies, 1113 Sofia

Abstract. *The paper presents an instance of our general semantic-based approach to creating concept hierarchies [7] particularly tuned for better performance on data represented with attribute-value and atomic formulae languages. We achieve this goal by incorporation of category utility and partition utility functions in its distance measure. Finally, we compare the performance of augmented algorithm to that of the COBWEB implementation given in the WEKA 3 ML system.*

Keywords: *lattice theory, distances, category utility, partition utility, hierarchical clustering, least general generalization (lgg).*

1. Introduction

In the context of Machine Learning clustering is an approach to discovering structure in data. Therefore the clusters are usually represented intensionally (in terms of relations, properties, features etc.) and hierarchical clustering techniques are of main interest. The term often used for this area is conceptual clustering. Three basic issues are important in conceptual clustering: the type of the cluster representation, the control strategy used to search the space of possible clusterings and the objective function used to evaluate the quality of clustering. Clusters can be represented by necessary and sufficient conditions for cluster membership (e.g. rules, decision trees) or probabilistically – by specifying the probability distribution of attribute values for the members of each cluster. The control strategy determines the way in which the clustering tree is gener-

¹ The work reported in this paper has been partially supported by Project ИТ-010051 “Advanced methods and tools for knowledge representation and processing”.

ated (e.g. top-down, bottom-up, hierarchical sorting). The objective function can be integrated in the control strategy or used separately. The classical conceptual clustering system COBWEB [4] uses the former approach – the search in the space of possible clusterings is based on an objective function that evaluates the candidate clusterings. The latter approach is taken in another classical system, CLUSTER/2 [8] that generates an initial clustering and then optimizes it iteratively by using an objective function. A more recent approach along these lines is proposed in [3].

Recently a semantic-based approach (a form of distance-based approach) to creating concept hierarchies has been introduced in [7]. It has been described as a general framework for semantic-based inductive learning. It proposes a consistent way to integrate syntactical lggs with semantic evaluation of the hypotheses. For this purpose two different relations on the hypothesis space are used – a constructive one (\leq_1), used to generate lggs and a semantic one (\leq_2) giving the coverage-based evaluation of them. These two relations (when they satisfy certain necessary conditions) jointly implement a semantic distance measure.

In the present paper we represent an instance of our general approach particularly tuned for better performance on data represented with attribute-value and atomic formulae languages. We achieve this goal by incorporating *category utility* and *partition utility* functions in our distance-based algorithm.

2. Basic algorithm

Given a set of examples the task of a hierarchical clustering system is to create a tree structure of clusters² where the children of a cluster partition the examples covered by their parent. The algorithm proposed in [7] accomplishes this task by using a coverage-based (semantic) distance measure and employing a bottom-up control strategy. Using a given set of examples E the algorithm builds a partial lattice structure G , where E is the set of all maximal elements of G . At each step the algorithm selects pairs of examples/clusters to be merged in a single clusters. The selection is based on a distance measure (based on a so-called semantic relation (\leq_1)) and applying a least general generalization operator (lgg) on the selected pair produces the new cluster. The lgg is based on a (so-called) constructive relation (\leq_1) that is used next to remove the examples/hypothesis covered by the newly generated one. The algorithm works iteratively and at each step updates two sets: C – the set of top level clusters in the current hierarchy where the two examples/clusters to be merged are selected from and G – the current hierarchy accumulating all examples and hypotheses (clusters) generated by lgg. Formally the algorithm is as follows:

1. Initialization: $G = E, C = E$
2. If $|C| = 1$ then exit
3. $T = \{h \mid h = \text{lgg}(a, b), (a, b) = \text{argmin}_{a,b \in C} d(a, b)\}$
4. $DC = \{a \mid a \in C \text{ and } \exists h \in T, h \leq_2 a\}$
5. $C = C \setminus DC$
6. $G = G \cup T, C = C \cup T$
7. Go to step 2.

² In the current paper clusters are represented intensionally by inductive hypotheses covering their corresponding examples. Thus the terms cluster and hypothesis are interchangeable.




Fig. 1. Visual representation of the algorithm

This algorithm has a big advantage that it is dependent neither on language used for knowledge (examples and hypotheses) representation nor on relations involved. In [7] we present the necessary conditions the two relations (semantic and constructive) have to satisfy in order to be applicable.

Although in the most common form our approach is language independent in this paper we limit our discussion to the two most widely used representations for both examples and hypotheses – the attribute-value language and the language of atomic formulae. All attributes are considered nominal. In the former case every element (example or hypothesis) $e = \{a_i = v_i \mid i = 1, \dots, n\}$ where a_i is the name of the i -th attribute and v_i is its corresponding value. The two relations (\leq_1 and \leq_2) are the set inclusion \subseteq . The lgg (infimum) of two elements (examples and hypotheses) is defined as their intersection, i.e. $\text{lgg}(x, y) = x \cap y$. The distance function is defined as $d(x, y) = 2^{-\text{cov}(x)} + 2^{-\text{cov}(y)} - 2 \times 2^{-\text{cov}(x \cap y)}$, where $\text{cov}(x)$ is the number of examples covered by x in respect to \leq_2 , i.e. $\text{cov}(x) = |\{y \mid y \in E, x \leq_2 y\}|$. In the case of atomic formulae all examples and hypotheses are represented by atomic formulae with the same predicate symbol. The two relations in this setting are θ -subsumption. The definition of distance function is analogous.

Since in most cases the set T created in step 3 is not a single, a restricted version of the algorithm is usually used. The basic idea is instead of the whole set T to use just a single element of it (arbitrarily chosen). This reduces the computational complexity of the algorithm³ and also narrows the created hierarchy.

We illustrate the algorithm by a simple example of hierarchical clustering of a set of 10 animals represented by 6 attributes: covering, milk, homeothermic, habitat, eggs and gills. Table 1 shows⁴ the initial set E and the resulting set G . To display the examples/hypotheses we use positional encoding, i.e. each element of the hierarchy is a list of values or underscores (denoting a missing attribute). In this representation the members of a cluster represented by hypothesis h are all elements of E that h subsumes, i.e. the examples that have the same values at the same positions as h , where an underscore in h matches all values. Obviously because of the random choice in step 3, the clusterings may differ from run to run. Figs. 2a and 2b show two possible hierarchies generated by the algorithm (Fig. 2b corresponds to the set G , shown in Table 1). We are further interested in the following questions: how can we evaluate these hierarchies and can we use this evaluation to improve the quality of clustering? We discuss these questions in the next section.

³ The complexity of the algorithm in this case is $O(n^3)$, where $n = |E|$. More details on this issue can be found in [7].

⁴ In order to minimize the size of tables and figures in the current paper all examples and hypothesis are denoted by ordered lists of their attribute values (positional encoding).

Table 1. Attribute-value representation of examples and hypotheses for the “animals” dataset (t and f stand for true and false respectively)

<i>E</i>	<i>G</i>
[feather, f, t, land, t, f]	[feather, f, t, land, t, f]
[hair, t, t, air, f, f]	[hair, t, t, air, f, f]
[feather, f, t, air, t, f]	[feather, f, t, air, t, f]
[hair, t, t, land, f, f]	[hair, t, t, land, f, f]
[scales, f, f, land, t, f]	[scales, f, f, land, t, f]
[none, f, f, land, t, f]	[none, f, f, land, t, f]
[none, t, t, sea, f, f]	[none, t, t, sea, f, f]
[hair, t, t, sea, t, f]	[hair, t, t, sea, t, f]
[scales, f, f, sea, t, f]	[scales, f, f, sea, t, f]
[scales, f, f, sea, t, t]	[scales, f, f, sea, t, t]
	[scales, f, f, sea, t, _]
	[_, f, f, land, t, f]
	[_, f, f, _, t, _]
	[hair, t, t, _, f, f]
	[feather, f, t, _, t, f]
	[_, f, _, _, t, _]
	[hair, t, t, _, _, f]
	[_, t, t, _, _, f]
	[_, _, _, _, _]

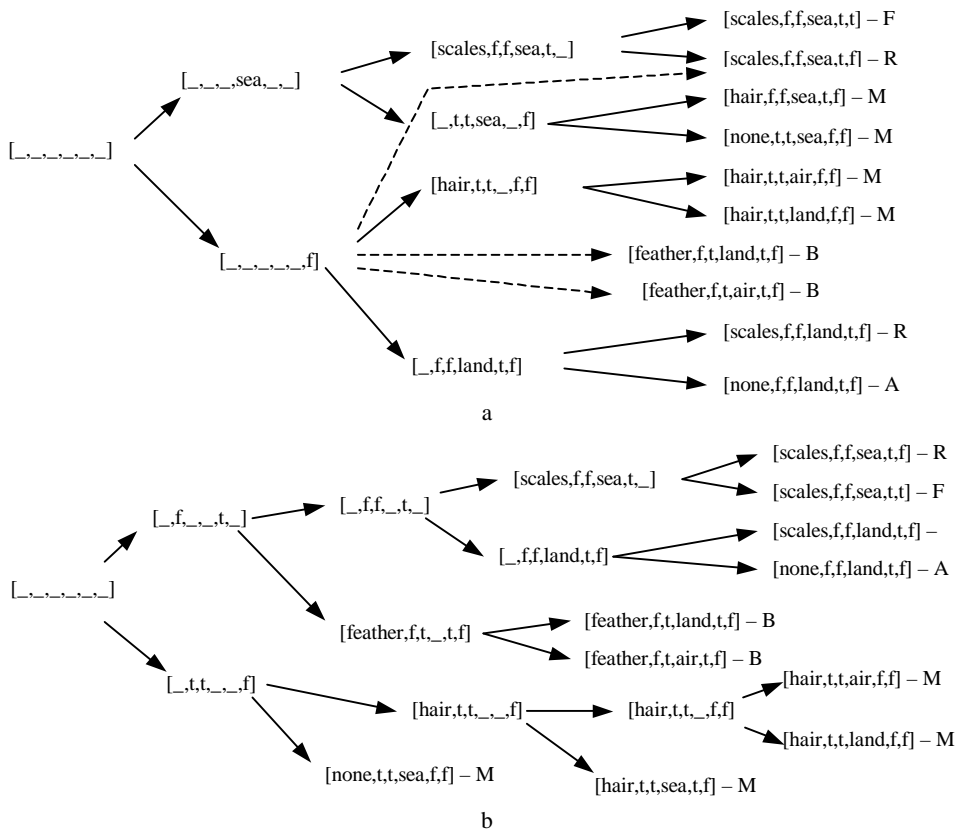


Fig. 2. Concept hierarchies for the “animals” dataset

3. Evaluating clustering quality

A commonly used measure to evaluate clustering quality is the category utility function proposed in [5]. This is the objective function used in COBWEB [4] and in many related systems. Category utility attempts to maximize both the probability that two objects in the same category have attribute values in common and the probability that objects from different categories have different attribute values. The category utility function assigns a value $CU(C_k)$ to each cluster C_k as follows:

$$(1) \quad CU(C_k) = \sum_i \sum_j [P(A_i = V_{ij} | C_k)^2 - P(A_i = V_{ij})^2].$$

The sums are calculated for all attributes A_i and all their values V_{ij} . The probabilities $P(A_i = V_{ij} | C_k)$ and $P(A_i = V_{ij})$ are calculated as relative frequencies of the occurrence of the V_{ij} value for the A_i attribute within the cluster C_k and within the whole set of examples respectively.

$CU(C_k)$ estimates the quality of individual clusters. To measure the quality of a clustering we need a function that evaluates the quality of a partition of data. For this purpose the average category utility of clusters in the partition is used, i.e.

$$(2) \quad PU(\{C_1, C_2, \dots, C_n\}) = \frac{1}{n} \sum_k CU(C_k).$$

In hierarchical clustering the partition utility function PU can be used at any level of the concept hierarchy. Applied to the children of the root PU evaluates the overall partition of the examples. For example the PU scores for the top-level partitions of the two concept hierarchies in Figures 2a and 2b are as follows:

$$PU([_, _, _, sea, _, _], [_, _, _, _, f]) = 0.191,$$

$$PU([_, f, _, _, t, _], [_, t, t, _, _, f]) = 0.612.$$

These PU scores show that the second clustering is better than the first one. This also becomes clear if we look at the class memberships of the examples. The classes that the examples belong to are shown at the leaves of the hierarchies as M, B, R, F and A (denoting mammal, bird, reptile, fish and amphibian respectively). The algorithm does not use these labels, however they may be used to evaluate “externally” the created hierarchy. For the bottom clustering, $[_, t, t, _, _, f]$ represents the cluster of mammals, and the other cluster – $[_, f, _, _, t, _]$, groups all non-mammals. The structure of $[_, f, _, _, t, _]$ reflects the natural way of grouping birds, reptiles, fish and amphibians. This observation is also supported by the PU scores of the lower level partitions. For example, the PU scores of the partitions of the mammal clusters in the two hierarchies are:

$$PU([_, _, _, _, f]) = 0.374,$$

$$PU([_, t, t, _, _, t]) = 0.421.$$

Another reason for these scores is that the mammal and non-mammal clusters in the first hierarchy overlap (they share 3 examples) while in the second hierarchy they are completely disjointed.

The above considerations suggest that there is a kind of propagation of “good” PU scores (i.e. good partitions of data) from the lower to the upper levels of the concept hierarchy. This in turn suggests a way to integrate the PU scores as an objective function into the control strategy of our general algorithm. This algorithm uses a bottom-up

control strategy. At each iteration a new hypothesis (cluster) h is added to the hierarchy. Since at this stage the whole partition at that level of the hierarchy is still unknown, h should be evaluated *individually* with respect to its contribution to the quality of clustering for the partition in which it will be included at a later step. That is, h has to maximize its CU score. To implement this we modify step 4 of the basic algorithm accordingly:

$$4. h_{\min} = \operatorname{argmax}_{h \in T} \operatorname{CU}(h), T = \{h_{\min}\}$$

An important issue in bottom-up clustering is how to determine when to stop merging clusters (generalizing hypotheses). In our algorithm the newly generated hypothesis h replaces the hypotheses/examples that it subsumes. Since the goal is to maximize the quality of clustering in terms of PU scores, it is expected that $\operatorname{CU}(h) > \operatorname{PU}(\{h_1, h_2, \dots, h_n\})$, where h_1, h_2, \dots, h_n are the hypotheses/examples that h subsumes, i.e. $h \leq_2 h_i, i = 1, 2, \dots, n$. This condition can be used to stop the process of generating new hypotheses and thus to cut off the hierarchy at a certain level. For a greater cut-off, a stronger condition can be defined as $\operatorname{CU}(h) > \max\{\operatorname{CU}(h_1), \operatorname{CU}(h_2), \dots, \operatorname{CU}(h_n)\}$. Combining this two conditions, we introduce a cut-off parameter $W \in [0, 1]$, that is used to control this process by gradually switching between them: $\operatorname{CU}(h) > \operatorname{PU}(\{h_1, h_2, \dots, h_n\}) + W * (\max\{\operatorname{CU}(h_1), \operatorname{CU}(h_2), \dots, \operatorname{CU}(h_n)\} - \operatorname{PU}(\{h_1, h_2, \dots, h_n\}))$.

The cut-off parameter W is used in step 4 of the algorithm in the following way. If h satisfies the stopping condition, then its successor with maximal CU score is used as a top-level hypothesis (an immediate successor of the root). This is achieved by excluding the latter and all its successors from the candidates for further merging.

4. Experiments

To evaluate the performance of the algorithm we conducted experiments with some well-known datasets. They include the previously discussed “animals” dataset containing 10 examples with 6 attributes, the “zoo” dataset – 59 examples with 16 attributes and the small soybean dataset – 47 examples with 35 attributes, all obtained from the UCI ML repository [1]. The class attributes were removed from all three datasets and after the concept hierarchies were generated the class information was used as an additional, “external” criterion to evaluate the quality of clustering. The idea was that the top-level partition should reflect the class distribution in the set of examples. As a basis for comparison a version of the COBWEB algorithm provided by the WEKA-3 ML system [11] was used in the experiments. Table 2 shows the PU scores of the top level partition produced by two clustering algorithms – the lattice-based clustering with gain function (denoted as LGG-PU) and COBWEB with the three data sets discussed above. Table 3 shows the performance of the same algorithms on the same datasets measured by the entropy of the class attribute in the clusters (hypotheses) from the top-level partition. The entropy of the class attribute in a cluster C_k is calculated as follows:

$$(3) \quad H(C_k) = \frac{-\sum_i V_i \log_2 \frac{V_i}{|C_k|}}{|C_k|},$$

where V_i are the values of the class attribute C , that occur in the examples belonging to cluster C_k . Then the entropy in a partition $\{C_1, C_2, \dots, C_k\}$, or the partition entropy PH is:

$$(4) \quad \operatorname{PH}(\{C_1, C_2, \dots, C_n\}) = \sum_k \frac{|C_k|}{|E|} H(C_k),$$

where the terms $\frac{|C_k|}{|E|}$ are weight coefficients that account for the cluster size.

Table 2. PU scores for the top-level partition of three datasets. The value of the cut-off parameter for the LGG-PU is $W = 0.5$

Algorithm	LGG-PU	COBWEB
Animals (10 examples, 6 attributes)	0.505	0.505
Zoo (59 examples, 16 attributes)	0.745	0.721
Soybean (47 examples, 35 attributes)	0.750	1.460

Table 3. PH scores/number of clusters for the top-level partition of three datasets. The value of the cut-off parameter for the LGG-PU is $W = 0.5$

Algorithm	LGG-PU	COBWEB
Animals (10 examples, 6 attributes)	0.27/4	0.27/4
Zoo (59 examples, 16 attributes)	1.03/7	0.99/4
Soybean (47 examples, 35 attributes)	0.00/8	0.49/4

In the ideal case each cluster contains examples of a single class and then the *PH* function of this partition is 0. In the general case small PH scores mean that the partition of the data corresponds to the distribution of the class attribute. Of course this also depends on the representativeness of the set of examples (i.e. whether it contains a sufficient number of instances from every class).

The above results are also evidence about the fact that high PU scores and low PH scores do not always mean that the resulting top-level partitioning is “good enough”. It may be too fragmented though. An example of this “feature” can be seen in the Zoo and Soybean results where the LGG-PU algorithm builds much more fragmented partitioning compared to that of the COBWEB. The main reason for these results is the fact that the LGG-PU algorithm performs greedy search in the semi-lattice of all possible clusters, thus like all similar algorithms can fall in the local extremums whereas the COBWEB can “step back” and go out of some of these pitfalls. Removing this “feature” of our algorithm is one direction for future work.

Finally, we present a result from experiments conducted with the relational version of our inductive learning algorithm. Thus we show one advantage of it that stems from the expressive power of this language for knowledge representation.

As described in [7] this algorithm can handle relational data by using a relational lgg. For our representation such lgg can be implemented by anti-unification. This is an operation that replaces same constants with same variables. For example $\text{lgg}(\{a, b, a\}, \{c, d, c\}) = \{X, Y, X\}$, where X and Y are variables. The main advantage of this representation is the possibility to define explicitly equality of attribute values. To investigate the behavior of our clustering algorithms we used the training sample of the MONK1 dataset [9]. This dataset describes a concept with 6 attributes by 61 positive and 61 negative examples. The propositional description of the target concept is:

{octagon, octagon, _, _, _, _}
 {square, square, _, _, _, _}

{round, round, _, _, _, _}
 {_, _, _, _, red, _}

While the relational one is

{X, X, _, _, _, _}
 {_, _, _, _, red, _}

For clustering we used just the positive examples, because the aim of the MONK1 problem is to build a meaningful description of the target concept and because distance-based inductive algorithms do not need the negative examples to narrow the search or to avoid overgeneralizations. The whole set of examples (including the negatives) was used for “external” evaluation of the top-level clusters. The experiments with the MONK1 data are summarized in Table 4.

Table 4. Results of clustering the MONK1 data. The PU, PH scores, and the number of clusters refer to the top-level partition. PH is calculated by using the whole set of examples (positive and negative)

Algorithm	PU	PH	Number of clusters
COBWEB	0.14	0	7
LGG-PU	0.16	0	2

The main findings reflected in the experiments are the following:

- For propositional data the LGG-PU algorithm performs similarly to COBWEB as long as the PU scores are concerned. In some cases LGG-PU achieves much better PH scores than COBWEB. A cut-off parameter W close to 0 results in high PU scores and low PH scores. With W close to 1, the PU scores are low and the PH scores are high.
- In the relational case both algorithms achieve 0 entropy at the top-level partition. LGG-PU however built two clusters that match exactly the original definition of target concept. COBWEB created very fragmented definition obviously because it does not take into account the relational similarity between examples.
- In the relational example (MONK1) the PU scores are low for both algorithms. This is because the two top level concepts in the original definition of the target concept substantially overlap.

5. Related work

Most of the approaches to hierarchical and conceptual clustering are related to two classical systems – CLUSTER/2 [8] and COBWEB [4]. CLUSTER/2 generates an initial clustering and then optimizes it iteratively by attempting to minimize the cluster overlapping. Clusters are represented as necessary and sufficient conditions for cluster membership (e.g. rules) and are derived by a standard concept learning algorithm. In COBWEB the clusters are represented probabilistically by the probability distribution of the attribute values for the members of each cluster. The control strategy used is based on the PU function. There are many other systems that follow the COBWEB’s approach. Some of them elaborate the objective function by using information-based heuristics (e.g. the information gain used in decision tree induction), Bayesian variants

of PU (used in AUTOCLASS [2]), or the Minimal Message Length (MML) principle (used in SNOB [10]). Others combine the COBWEB and the CLUSTER/2 approaches – they first generate initial clustering and then optimize it by using various objective functions. The approach taken in [3] is based on this idea. It uses hierarchical sorting to induce a clustering and then iteratively simplifies it by applying various techniques as redistribution, reordering etc.

There is third approach to hierarchical clustering that is based on lattice structures. The lattices are a useful mathematical formalism that resembles cluster hierarchies. The problem with the use of lattices for clustering is that the techniques for generating clusters should have some formal properties (usually not present in the heuristic algorithms). The computational complexity of the lattice algorithms is also high. Nevertheless lattice approaches to clustering exist and some of them are successful. For example [6] use Galois lattices to induce hierarchical clusterings for binary attributes. This approach is based on the so-called maximal rectangles and employs some standard algorithms from Galois lattice theory.

Our approach is also based on lattices. Its theoretical basis is the semantic distance measure introduced in [7] that allows evaluating the similarity between examples/clusters as well as applying consistently a generalization operation (*lgg*) to build the lattice structure. The use of a PU-based objective function allows to improve the quality of clustering and to reduce the complexity of the algorithm. The main features of the approach can be summarized as follows:

- LGG-PU represents clusters as propositional or relational rules. Similarly to CLUSTER/2 it allows cluster overlapping.
- It uses a PU-based objective function to evaluate the quality of clustering.
- The control strategy of the algorithm is a bottom-up greedy search based on a consistent integration of a lattice-based distance measure and a PU-based objective function.
- LGG-PU is able to build relational descriptions of the clusters. Although the PU function does not work well on relational data, it is still useful in the control strategy of the algorithm.

6. Conclusion

We described an approach to hierarchical clustering based on the use of a least general generalization (*lgg*) operator to induce a lattice structure of clusters and a category utility objective function to evaluate the clustering quality. The objective function is integrated with a lattice-based distance measure into a bottom-up control strategy for clustering. The preliminary experiments showed that the approach compares well with other similar approaches and in some cases outperforms them. The future work will address the following issues:

- Extending the algorithm to handle numeric attributes. This can be achieved by extending the partial ordering relation (presently set inclusion) and by defining a proper *lgg* that can generalize numeric attributes too. Then we can use the same distance measure and a PU function that uses probability distributions.
- Although the algorithm evaluates clusters locally, it also maximizes the overall clustering quality since the local evaluation goes through all levels of the hierarchy in a bottom-up fashion. Of course, as this is a greedy search, there is no guarantee

that the global maximum can be found. With this respect other approaches to evaluating the quality of clustering have to be investigated too (e.g. the recursive ones used in COBWEB).

References

1. Blake, C., C. Merz. UCI Repository of Machine Learning Databases. 1998.
2. Cheeseman, P., J. Kelly, M. Self, J. Stutz, W. Taylor, D. Freeman. AutoClass: a Bayesian classification system. – In: Proceedings of the 5th International Workshop on Machine Learning, Ann Arbor, San Mateo, CA. Morgan Kaufmann, 1988, 54-64.
3. Fisher, D. Iterative optimization and simplification of hierarchical clusterings. – Journal of Artificial Intelligence Research, **4**, 1996, 147-179.
4. Gennari, J. H., P. Langley, D. Fisher. Models of incremental concept formation. – Artificial Intelligence, **40**, 1989, 11-61.
5. Gluck, A., J. Corter. Information, uncertainty, and the utility of categories. – In: Proceedings of the 7th Annual Conference of the Cognitive Science Society, Hillsdale. Lawrence Erlbaum, 1985, 283-287.
6. Guenoche, A., I. V. Mechelen. Galois approach to the induction of concepts. – In: Categories and Concepts: Theoretical Views and Inductive Data Analysis, I. V. Mechelen, J. Hampton, R. Michalski, & P. Theuns (Eds.). Academic Press, 1993, 287-308
7. Marinchev, I. Semantic Height Functions and their Application to the Inductive Machine Learning. PhD Thesis, Sofia University, 2001.
8. Michalski, R., R. Stepp. Learning from observation: conceptual clustering. – In: Machine Learning: Artificial Intelligence Approach (R. Michalski, J. Carbonell, T. Mitchell, Eds.). Vol. 1, Tioga, 1983, 331-363.
9. Thrun, S. B., et al. The MONK's Problems – a Performance Comparison of Different Learning Algorithms. Tech. Rep. CS-CMU-91-197, Carnegie Mellon University, 1991.
10. Wallace, C., D. Dowe. Intrinsic classification by mml { the snob program. – In: Proceedings of the 7th Australian Joint Conference on Artificial Intelligence, Armidale, Australia. World Scientific, 1994, 37-44.
11. Witten, I., E. Frank. Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations. Morgan Kaufmann, 1999.

Приложение на функциите за категориална полезност и кластеризационна полезност към семантичния подход за концептуална кластеризация на данни

Иво Маринчев

Институт по информационни технологии , 1113 София, E-mail ivo_m@iinf.bas.bg

(Резюме)

Настоящата статия представя вариант на нашия семантично базиран подход към построяването на концептуални йерархии [7], специално пригоден за ефективна обработка на данни, представени с езика "атрибут–стойност" и езика на атомарните формули. Ние постигаме този резултат чрез прилагане на функциите за категориална полезност и кластеризационна полезност. Накрая сравняваме ефективността на така модифицираната система с тази на реализацията на COBWEB, налична в системата WEKA 3 ML.