Automatic Thematic Categorization of Documents Using a Fuzzy Taxonomy and Fuzzy Hierarchical Clustering

Manolis Wallace, Giorgos Akrivas and Giorgos Stamou
Image, Video & Multimedia Systems Laboratory
Department of Electrical & Computer Engineering
National Technical University of Athens
9, Iroon Polytechniou Str.,
157 73 Zographou, Athens, Greece
Email: wallace@image.ntua.gr

***Abstract*— In this paper we formally define the problem of automatic detection of thematic categories in a semantically indexed document, and identify the main obstacles to overcome in this process. Furthermore, we explain how detection of thematic categories can be achieved, with the use of a fuzzy quasi-taxonomic relation. Our approach relies on a fuzzy hierarchical clustering algorithm; this algorithm uses a similarity measure that is based on the notion of context.**

I. Introduction

IT is nowadays widely accepted that information retrieval systems are reaching an upper bound of performance, as far as evolution of traditional techniques is concerned, without having met the goal of successfully providing users with the documents that best match their needs. The new direction of research that promises to give new momentum to this field is *intelligent information retrieval*, which may be briefly summarized as the conceptual handling of users, user requests and documents.

It is rather obvious that techniques from term based retrieval do not suffice for conceptual analysis of documents. Knowledge – based approaches, on the other hand, appear to be much more robust and flexible in handling modern documents. An important step in the direction of knowledge – based information retrieval has already been made, with the definition of the *semantic entity*; this corresponds to what we might call a concept, object or event, and aims to replace terms and keywords. Of great importance is, as well, the construction of ontologies, which constitute an attempt to describe the relations between real life entities, in a conceptual level [4].

On the other hand, diversity of the various types of documents, that a current information retrieval system must support, as well as the explosive growth of their number, calls for new schemes for categorization among them. A categorization scheme classifies documents to one or more categories, according to their content. Thus, documents belonging to the same category can be treated similarly, with respect, e.g. to user profiling, document presentation, relevance feedback etc. Moreover, a categorization can enable users to browse a document repository, instead of querying it.

Several schemes for categorization can be applied, such as categorization according to document media (for example film, picture, printed document) and document type (movie, news, TV show, scientific publication). In this paper, we are concerned with thematic categorization, i.e. a categorization in a conceptual level.

A conceptual thematic categorization uses the semantic entities encountered in a document, in order to classify the latter into classes, such as sports, diplomacy, chemistry and so on. Our approach performs a fuzzy hierarchical clustering of the semantic entities, relying on knowledge that is stored in the form of semantic relations. The notion of context has a central role in this process.

The structure of the paper is as follows: In section II we present a novel quasi – taxonomic semantic relation. Based on this relation, after formally defining the problem of thematic categorization in section III, in section IV we rely on the notion of context in order to detect the thematic categories that are related to a document. In section V we discuss the applications of the proposed method, list areas of related future work and present our concluding remarks.

II. The semantic taxonomic relation

Ontologies are an attempt for modelling real world entities [4]. They may be described as follows:

$$O = \{S, \{R_i\}\}, i = 1 \dots n$$

$$R_i : S \times S \rightarrow \{0, 1\}, i = 1 \dots n$$

where O is an ontology, S the set of semantic entities it describes and R_i the i-th semantic relation amongst the semantic entities. The formal definition of ontologies also supports an inference layer, but herein we omit it for the sake of simplicity.

Although any type of relation may be contained in an ontology, the two main categories are taxonomic (i.e. ordering) and compatibility (i.e. symmetric) relations. Compatibility relations have traditionally been exploited by information retrieval systems for tasks such as query expansion. They are ideal for

TABLE I
THE FUZZY SEMANTIC RELATIONS

Sp	Specialization
Ct	Context
Ins	Instrument
P	Part
Pat	Patient
Loc	Location
Ag	Agent

the description of similarities of various natures, but fail to assist in the determination of the context of a query or a document; the use of ordering relations is necessary for such tasks [1]. Thus, a main challenge of intelligent information retrieval is the meaningful exploitation of information contained in taxonomic relations of an ontology.

It is well understood that relations among real life entities are always a matter of degree, and are, therefore, best modelled using fuzzy relations. Ontological taxonomies, on the other hand, are crisp in principle. Thus, they fail to fully describe real life concepts, and are limited to α-cuts of the desired relations. This is a very important drawback, that makes such relations insufficient for the services that an intelligent information retrieval system aims to offer.

A. The Fuzzy Quasi – Taxonomic Relation

The authors have proposed fuzzy semantic relations that are most suitable for the modelling of real life information [2]. In this section, we present a few commonly encountered semantic relations that can be modelled as fuzzy ordering relations, and propose their combination for the generation of a meaningful, fuzzy, quasi-taxonomic relation. Based on this relation, in the following sections we will explain how the problem of automatic thematic categorization may be tackled.

The *specialization* relation Sp is a fuzzy partial ordering on the set of semantic entities. $Sp(a, b) > 0$ means that the meaning of a "includes "the meaning of b; the most common form of specialization is sub – classing, i.e. a is a generalization of b. The role of the specialization relation in knowledge – based retrieval is as follows: if a document refers to the meaning of entity b, then it is also related to a, since b is a special case of a. Still, there is no evidence that the opposite also holds; it is obvious that the specialization relation contains important information that can not be modelled in a symmetric relation. The *context* relation Ct is also a fuzzy partial ordering on the set of semantic entities. $Ct(a, b) > 0$ means that b provides the context for a or, in other words, that b is the thematic category that a belongs to. Other relations considered in the following have similar interpretations. Their names and corresponding notations are given in table I.

In this work, fuzziness of the aforementioned relations has the following meaning: High values of $Sp(a, b)$, imply that the meaning of b approaches the meaning of a, in the sense that when a document is related to b, then it is most probably related to a as well. On the other hand, as $Sp(a, b)$ decreases, the meaning of b becomes "narrower" than the meaning of a, in the sense that a document's relation to b will not imply a relation to a as well with a high probability, or to a high degree. Summarizing, the value of $Sp(a, b)$ indicates the degree to which the stored knowledge shows that an occurrence of b in a document implies relation to a. Likewise, the degrees of the other relations can also be interpreted as conditional probabilities or degrees of implied relevance.

The above imply that, for example, $a \neq b \implies Sp(a, b) < 1$ since, if $a \neq b$, then we cannot be sure that both a and b are related to a given document, without first examining the document's context; at this point it is important to remind the reader that a and b are not terms but concepts, which means that $a \neq b$ indicates / ensures a difference in a conceptual level.

A last point to consider is the transitivity of the relations presented above. It is obvious that if b is a specialization of a and c is a specialization of b, then c is a specialization of a. This implies that the specialization relation is transitive. A similar argument can be made for the other relations, as well. Still, the form of transitivity used cannot be $\sup - \min$ transitivity, but one relying on a subidempotent norm. Therefore, we demand that the presented relations are $\sup - t$ transitive, where t is an Archimedean norm.

More formally, the knowledge model presented above may be summarized in the following:

$$O_{\mathcal{F}} = \{S, \{r_i\}\}, i = 1 \dots n$$

$$r_i = \mathcal{F}(R_i) : S \times S \to [0, 1], i = 1 \dots n$$

Based on the relations r_i we construct the following semantic relation:

$$T = Tr^t(\bigcup_i r_i^{p_i}), p_i \in \{-1, 1\}, i \in 1 \dots n$$

where $Tr^t(A)$ is the $\sup - t$ transitive closure of relation A; the transitivity of relation T was not implied by the definition, as the union of transitive relations is not necessarily transitive. In our application we construct the T relation as follows:

$$T = Tr^t(Sp \cup C^{-1} \cup Ins \cup P \cup Pat \cup Loc \cup Ag)$$

Based on the semantics of relations r_i, it is easy to see that T is ideal for the determination of the thematic categories that an entity may be related to, as thematic categories are also semantic entities:

$$TC \subseteq S$$

where $TC = \{tc_i\}, i \in 1 \dots k$ is the set of thematic categories (for example *ball* and *stadium* may be semantic entities, while

football and *sports* are both semantic entities and thematic categories). Unfortunately, the example of the T relation has to be omitted for the sake of space.

All the relations used for the generation of T are partial ordering relations. Still, there is no evidence that their union is also antisymmetric. Quite the contrary, T may vary from being a partial ordering to being an equivalence relation. This is an important observation, as true semantic relations also fit in this range (total symmetricity as well as total antisymmetricity often have to be abandoned when modelling real life). Still, the semantics of the used relations, as well as our experiments, indicate that T is "almost" antisymmetric. Therefore, we categorize to it as quasi – ordering or quasi – taxonomic.

III. Problem definition

Before anything else, let us first present the problem that this work attempts to address, in a more formal manner. The intelligent module presented herein, which we will refer to as module of *Detection of Thematic Categories* (DTC), accepts as input the Semantic Index I. This is in fact a fuzzy relation between documents and semantic entities.

$$I : S \times D \rightarrow [0, 1]$$

The semantic index must be normal for each document, i.e.:

$$\forall d \in D \;\; \exists s \in S \text{ such that } I(s, d) = 1$$

Based on this relation, and the knowledge contained in the available semantic relations R_i, the module of DTC aims to detect the degree to which a given document $d \in D$ is related to a thematic category $tc \in TC$. We will refer to this degree as $R_{\mathrm{TC}}(tc, d)$. In other words, the module of DTC attempts to calculate the relation:

$$R_{\mathrm{TC}} : TC \times D \rightarrow [0, 1]$$

In designing an algorithm that is able to calculate this relation, in a meaningful manner, a series of issues need to be tackled:

1) A semantic entity may be related to multiple, unrelated thematic categories.
2) A document may be related to multiple, unrelated thematic categories.
3) The semantic index may have been created in an automated manner. Thus, existence of random, and therefore misleading semantic entities cannot be excluded. For example, entities that correspond to terms that have been used in a metaphorical sense when annotating a documents may be included in the index.
4) Semantic relations are always a matter of degree. Therefore, correlation between a document and a thematic category is also a matter of degree.

In the following, we provide the principles of the proposed approach to the problem of thematic categorization.

A. Principles of Thematic Categorization

According to issue 1, it is necessary for the algorithm to be able to determine which thematic categories are indeed related to a given document. In order for this task to be performed in a meaningful manner, the common meaning of the remaining entities that index the given document needs to be considered as well.

On the other hand, when a document is related to more than one, unrelated thematic categories, as issue 2 points out, we should not expect all the terms that index it to be related to one another, or to each one of the thematic categories in question. Quite the contrary, we should expect most entities to be related to just one of these thematic categories. Therefore, a clustering of semantic entities, based on the their common meaning, needs to be applied.

In this process, entities that are misleading (eg. entities that resulted from the use of terms in a metaphorical sense) will probably not be found similar with other entities that index a document. Therefore, the cardinality of the clusters may be used to tackle issue 3.

Finally, issue 4 is easily solved by allowing DTC's algorithm to be fuzzy. In the following, we proceed with the presentation of an algorithm for DTC which complies with the above principles.

IV. The algorithm for DTC

The proposed approach may be decomposed into the following steps:

- Perform a fuzzy clustering of semantic entities, using their common meaning as clustering criterion in order to determine the count of distinct topics that a document is related to.
- Find the thematic categories that are related to each cluster.
- Aggregate the findings for each cluster in order to acquire an overall result for the whole document.

Each of the above steps uses the taxonomy relation, in addition to the index. In the following, after discussing the notion of "common meaning", we elaborate on each of these steps.

A. The notion of context

In general, the term *context* refers to whatever is common among a set of elements. In this work, where the elements are semantic entities and documents, the term context may refer to the common meaning of a set of entities, or to the overall topic of a document, respectively.

A document is represented only by its mapping to semantic entities, via the semantic index I. Therefore, the context of a document is again defined via the semantic entities that are related to it. The fact that relation T described in subsection II-A is (almost) an ordering relation allows us to use it in order to define, extract and use the context of a document, or a set of semantic entities in general. Relying on the semantics of the T

relation, we define the *context* $K(s)$ of a semantic entity $s \in S$ as the set of its descendants in relation T:

$$K(s) = T_{\leq}(s)$$

This set also includes the semantic entity in question.

Assuming that a set of entities $S' \subset S$ is crisp, i.e. all considered entities belong to the set with degree one, the context of the group, which is again a set of semantic entities, can be defined simply as the set of their common descendants.

$$K(S') = \bigcap_i K(s_i), s_i \in S'$$

Obviously, as more entities are considered, the context becomes narrower, i.e. it contains less entities and to smaller degrees:

$$S' \supset S'' \implies K(S') \subseteq K(S'')$$

When the definition of context is extended to the case of fuzzy sets of semantic entities, this inequality must still hold. Moreover, we demand that the following are satisfied as well:

- $S'(s) = 0 \implies K(S') = K(S' - \{s\})$, i.e. no narrowing of context.
- $S'(s) = 1 \implies K(S') \subseteq K(s)$, i.e. full narrowing of context.
- $K(S')$ decreases monotonically with respect to $S'(s)$.

Taking these into consideration, we demand that, when S' is fuzzy, the "considered" context $\mathcal{K}(s)$ of s, i.e. the entity's context when taking its degree of participation to the set into account, becomes low when the degrees of taxonomy are low and the degree of participation $S'(s)$ is high. Therefore:

$$cp(\mathcal{K}(s)) \doteq cp(K(s)) \cap (S'(s) \cdot S)$$

where cp is an involutive fuzzy complement, and $\cap$ and $\cup$ correspond to a t-norm and a t-conorm which are dual, with respect to cp. By applying de Morgan's law, we obtain:

$$\mathcal{K}(s) \doteq K(s) \cup cp(S'(s)) \tag{1}$$

Then the set's context is easily calculated as follows:

$$K(S') = \bigcap_i \mathcal{K}(s_i), s_i \in S'$$

Considering the semantics of the T relation and the process of context determination, it is easy to realize that when the entities in a set are highly related to a common meaning, the context will have high degrees of membership for the entities that represent this common meaning. Therefore, the height of the context $h(K(S'))$ may be used as a measure of the semantic correlation of entities in set S'. We will refer to this measure as *intensity* of the context.

B. Hierarchical clustering

Before actually extracting thematic category information from the set of semantic entities $I(d)$ that are related to a document d via the semantic index I, in order to support the possibility of existence of multiple distinct topics in a single document, the support of the document's description, i.e. the set

$$I(d) = \{s \in S : I(s, d) > 0\}$$

of the entities that are related to it needs to be clustered to groups, according to the topics they are related to.

Most clustering methods belong to either of two general methods, partitioning and hierarchical [6]. Partitioning methods create a crisp or fuzzy clustering of a given data set, but require the number of clusters as input. Since the number of topics that exist in a document is not known beforehand, partitioning methods are inapplicable for the task at hand [5]; an hierarchical clustering algorithm needs to be applied.

Hierarchical methods are divided into agglomerative and divisive. Of those, the first are the most widely studied and applied, as well as the most robust. Their general structure, adjusted for the needs of the problem at hand, is as follows:

1) When considering document d, turn each semantic entity $s \in I(d)$ into a singleton, i.e. into a cluster of its own.
2) For each pair of clusters c_1, c_2 calculate a compatibility indicator $CI(c_1, c_2)$. The CI is also referred to as cluster similarity, or dissimilarity, measure.
3) Merge the pair of clusters that have the best CI. Depending on whether this is a similarity or a dissimilarity measure, the best indicator could be the maximum or the minimum operator, respectively.
4) Continue at step 2, until the termination criterion is satisfied. The termination criterion most commonly used is the definition of a threshold for the value of the best compatibility indicator.

The two key points in hierarchical clustering are the identification of the clusters to merge at each step, i.e. the definition of a meaningful measure for CI, and the identification of the optimal terminating step, i.e. the definition of a meaningful termination criterion.

When clustering semantic entities, the ideal similarity measure is one that quantifies their semantic correlation. In subsection IV-A we have defined such a measure; it is the height of their common context. Therefore, the merging of clusters will be based on this measure.

The process of merging should terminate when the entities are clustered into sets that correspond to distinct topics. We may identify such sets by the fact that their common contexts will have low, if not zero, intensity. Therefore, the termination criterion shall be a threshold on the intensity of the common meaning, i.e. a threshold on the selected compatibility measure.

C. Fuzzy clustering

Hierarchical clustering methods are more flexible than their partitioning counterparts, in that they do not need the number of

clusters as an input. Still, they are less robust in other ways:

- They only create crisp clusterings, i.e. they do not support degrees of membership in their output.
- They only create partitions, i.e. they do not allow for overlapping among the detected clusters.

Both of the above are great disadvantages for the problem at hand, as they are not compatible with the task's semantics: in real life, a semantic entity may be related to a topic to a degree other than 1 or 0, and may also be related to more than one distinct topics.

In order to overcome such problems, we describe in the following a method for fuzzyfication of the partitioning. In this way the clusters' cardinalities will be corrected, so that they may be used in subsection IV-D for the meaningful extraction of thematic categories.

Each cluster c is described by the crisp set of semantic entities S_c that belong to it. Using those, we may create a fuzzy classifier, i.e. a function C_c that will measure the degree of correlation of a semantic entity s with the cluster c.

$$C_c : S \rightarrow [0, 1]$$

Obviously, a semantic entity should be considered correlated with c, if it is related to the common meaning of the semantic entities in S_c. Therefore, the quantity

$$Cor_1(c, s) = h(K(S_c \cup \{s\}))$$

where $h(\cdot)$ symbolizes the height of a fuzzy set, is a meaningful measure of correlation. Of course, not all clusters are equally compact; we may measure cluster compactness using the similarity among the entities it contains, i.e. using the intensity of the cluster's context. Therefore, the aforementioned correlation measure needs to be adjusted, to the characteristics of the cluster in question:

$$Cor_2(c, s) = \frac{Cor_1(c, s)}{h(K(c))}$$

It is easy to see that this measure obviously has the following properties:

- $Cor_2(c, s) = 1$ if the semantics of s imply it should belong to c. For example $Cor_2(c, s) = 1, \forall s \in S_c$
- $Cor_2(c, s) = 0$ if the semantics of s imply it should not belong to c.
- $Cor_2(c, s) \in (0, 1)$ if s is neither totally related, nor totally unrelated to c.

These are the properties that we wish for the cluster's fuzzy classifier, so:

$$C_c(s) \doteq Cor_2(c, s)$$

Using such classifiers, we may expand the detected crisp partitions, as to include more semantic entities, as follows: partition c is replaced by cluster

$$c' = \sum_{s \in I(d)} s / C_c(s)$$

Obviously $c' \supseteq c$.

D. Extraction of thematic categories

Thematic categories are semantic entities that have been selected as having a special meaning for the system; more formally:

$$TC \subset S$$

This simplifies the process of automatic thematic categorization: We have already explained that the context of a set of semantic entities is a fuzzy set of semantic entities; this contains the entities that describe the common meaning of the original set. The thematic categories that are contained in the context of a cluster of semantic entities are obviously thematic categories that are related to the whole document. Based on this concept, in the following we present a method for automatic thematic categorization of documents.

First of all, the process of fuzzy hierarchical clustering has been based on the crisp set $I(d)$, thus ignoring fuzziness in the semantic index. In order to incorporate this information in the clusters of semantic entities considered for the process of thematic categorization, we adjust the degrees of membership for them as follows:

$$c''(s) = t(c'(s), I(s, d)) \tag{2}$$

where t is a fuzzy norm. The semantic nature of this operation demands that t is an Archimedean norm. From each one of those clusters, we may extract the corresponding thematic categories. In the following we shall refer to a random fuzzy cluster c'' and its corresponding fuzzy set of thematic categories $R_{\mathrm{TC}}(c'')$.

Obviously, thematic categories that are not contained in the context of c'' cannot be selected as being related to it. Therefore

$$R_{\mathrm{TC}}(c'') \subseteq R^1_{\mathrm{TC}}(c'') \doteq w(K(c'') \cap TC) \tag{3}$$

where w is a weak *modifier*. Modifiers, which are also met in the literature as *linguistic hedges* [3], are used (in this work) to adjust mathematically computed values so as to match their semantically anticipated counterparts. Specifically, our experiments indicate, for example, that a value of 0.7 for the expression $K(c'') \cap TC$ corresponds to a great degree of relevance, and should, therefore, be adjusted accordingly.

In the case that the semantic entities that index document d are all clustered in a unique cluster c'', then $R_{\mathrm{TC}}(d) = R^1_{\mathrm{TC}}(c'')$ is a meaningful approach. On the other hand, when more than one clusters are detected, then cluster cardinalities have to be considered as well.

Clusters of extremely low cardinality probably only contain misleading entities, and therefore need to be ignored in the estimation of $R_{\mathrm{TC}}(d)$. On the contrary, clusters of high cardinality almost certainly correspond to the distinct topics d is related to, and need to be considered in the estimation of $R_{\mathrm{TC}}(d)$. The notion of "high cardinality" is modelled with the use of a "big" fuzzy number L. $L(a)$ is the truth value of the preposition "the cardinality of a is high".

The set of thematic categories that correspond to a document is computed from the remaining clusters, after adjusting membership degrees according to scalar cardinalities, as follows:

$$R_{\mathrm{TC}}(d) \doteq \underset{c'' \in G}{u}(R_{\mathrm{TC}}(c'')) \quad (4)$$

$$R_{\mathrm{TC}}(c'') = R^1_{\mathrm{TC}}(c'') \cdot L(|c''|) \quad (5)$$

where u is a fuzzy co-norm, G is the set of fuzzy clusters that have been detected in $I(d)$ and have had their membership degrees adjusted according to equation 2, and $|b|$ is the scalar cardinality of set b.

It is easy to see that $R_{\mathrm{TC}}(d, tc)$ will be high if a cluster c'', whose context contains tc, is detected in $I(d)$, and additionally, the cardinality of c is high (i.e. the cluster is most probably not comprised of misleading entities) and the degree of membership of tc in the context of c'' is high.

V. Applications, conclusions and future work

In this paper, we started by presenting a fuzzy, quasi – ordering, semantic relation defined on the set of semantic entities. Continuing, we identified the main obstacles that have to be faced in the process of automatic detection of thematic categories that are related to a semantically indexed document, and explained how this can be achieved, using the notion of context; our approach relies on fuzzy hierarchical clustering of the fuzzy index.

The formulas in this paper only include general fuzzy operators, and not specific approaches, such as the standard operators. The reason is that the selection of the "correct" operators is still an ongoing process, which needs thorough theoretical investigation. So far, the best performance has been observed, by trial and error work, when using the operators mentioned below:

- In subsection II-A, the t-norm used for the transitive closure of relation T is Yager's t-norm with parameter 3 [3].
- In equation 1, the union operator u used is the bounded sum, while the complement cp of choice is the standard complement.
- In equation 2, the fuzzy norm used is the product.
- In equation 3, the modifier used is $w(a) = \sqrt{(a)}$
- In equation 4, the standard co-norm is used.

Moreover,

- in section IV-B, the threshold used for the termination criterion of the clustering algorithm is 0.3.
- In equation 5, big fuzzy number L is defined as the triangular fuzzy number $(1.3, 3, \infty)$.

The method presented in this paper has been developed and tested in the experimental (and currently in the process of modules' integration) prototype of the FAETHON multimedia information retrieval system [8]. FAETHON possesses an experimental semantic encyclopedia, as described in section II. It contains definitions for 340 semantic entities, about 20% of which are thematic categories, as well as definitions for 48 semantic relations, out of which 10 are already populated. This encyclopedia has been used, with encouraging results, for the development and testing of algorithms concerning not only automatic detection of thematic categories, but also intelligent, context sensitive query expansion [1] and flexible adaptation to the user [7]. FAETHON also possesses a fuzzy semantic index for numerous documents from the Hellenic Broadcasting Corporation (ERT) and Film Archiv Austria (FAA).

The output of the module of DTC is exploited in numerous ways. As most important we may mention definition and extraction of user preferences at a semantic level, providing of efficient content browsing services to users, timely estimation of the content of relevance feedback based on thematic categorization of documents and automatic suggestion of documents that are related to the document a user is currently viewing.

A major area of future research for this work is the selection of optimal fuzzy operators for most meaningful semantic output. Our findings so far indicate that this selection is not independent from the knowledge itself. In other words, different semantic encyclopedias may perform best for different choices of operators. Thus, the connection between encyclopedia content and operator selection is also an interesting area for research.

Finally, we believe that the definition of context in a semantic manner is an important issue that is not necessarily finalized in this paper. Quite the contrary, we believe that there is a lot of analytic work to be done in the formal representation, extraction and use of context in thematic categorization.

Acknowledgements

This work has been partially funded by the EU IST-1999-20502 FAETHON project.

References

[1] G. Akrivas, M. Wallace, G. Andreou, G. Stamou and S. Kollias, "Context - Sensitive Semantic Query Expansion", Proceedings of the IEEE International Conference on Artificial Intelligence Systems (ICAIS), Divnomorskoe, Russia, September 2002

[2] G. Akrivas G. and G. Stamou, Fuzzy Semantic Association of Audiovisual Document Descriptions, Proc. of Int. Workshop on Very Low Bitrate Video Coding (VLBV), Athens, Greece, Oct. 2001

[3] G. Klir and B. Yuan, Fuzzy Sets and Fuzzy Logic, Theory and Applications, New Jersey, Prentice Hall, 1995.

[4] A. Maedche, B. Motik, N. Silva and R. Volz, MAFRA - An Ontology MApping FRAmework in the Context of the SemanticWeb. Proceedings of the Workshop on Ontology Transformation at ECAI2002, Lyon, France, July 2002.

[5] S. Miyamoto, Fuzzy Sets in Information Retrieval and Cluster Analysis, Kluwer Academic Publishers, Dordrecht / Boston / London, 1990

[6] S. Theodoridis and K. Koutroumbas, Pattern Recognition, Academic Press, 1998.

[7] M. Wallace, G. Akrivas, G. Stamou and S. Kollias, "Representation of user preferences and adaptation to context in multimedia content – based retrieval", Proceedings of the Workshop on Multimedia Semantics at SOFSEM2002, Milovy, Czech Republic, November 2002

[8] IST-1999-20502 Project FAETHON "Unified Intelligent Access to Heterogeneous Audiovisual Content", 2001-2003, http://www.image.ece.ntua.gr/faethon/