Representation of user preferences and adaptation to context in multimedia content – based retrieval

Manolis Wallace, Giorgos Akrivas, Giorgos Stamou, and Stefanos Kollias

Image, Video and Multimedia Laboratory,
Department of Electrical and Computer Engineering,
National Technical University of Athens,
15773 Zografou, Greece
{wallace,gakrivas}@image.ntua.gr, {gstam, stefanos}@softlab.ntua.gr
http://www.image.ntua.gr

Abstract. The task of content – based retrieval is to provide users with the multimedia documents that best match their wishes. This process is not free of uncertainty; the role of the user profile is to remove a part of this uncertainty, using the information it contains concerning the user's preferences, and thus improve the precision. For this aim MPEG-7 introduced description schemes representing users' preferences and tools supporting user interaction. In this paper, using a formal representation of user preferences and a semantic knowledge base, we expand the user profile representation introduced in MPEG-7 in order to acquire more semantic information about the user. Furthermore, using the notion of context, we identify the part of the profile that is related to the user's query, thus preventing irrelevant interests from affecting the process of content – based retrieval.

1 Introduction

Digital archiving of multimedia content including video, audio, still images and various types of documents has been recognized by content holding organizations as a mature choice for the preservation, preview and partial distribution of their assets. The advances in computer and data networks along with the success of standardization efforts of MPEG and JPEG boosted the movement of the archives towards the conversion of their fragile and manually indexed material to digital, computer accessible data. By the end of last century the question was not on whether digital archives are technically and economically viable, but rather on how digital archives would be efficient and informative. In this framework, different scientific fields such as, on one hand, development of database management systems, and on the other hand, processing and analysis of multimedia data, as well as artificial and computational intelligence methods, have observed a close cooperation with each other during the last few years. The attempt has been to develop intelligent and efficient human computer interaction systems, enabling the user to access vast amounts of heterogeneous information, stored in different sites and archives.

Database management systems (DBMSs) have been designed that are able to handle such types of access to the stored information. Attaching information bits, called metadata, to the original data is the means for achieving this goal. The focus of technological attempts has been on the analysis of digital video, due to its large amounts of spatiotemporal interrelations, which turns it into the most demanding and complex data structure. Current and evolving international standardization activities, such as of the EBU, MPEG-4 [5],[6], MPEG-7 [9],[10], or JPEG-2000 [13] for still images, deal with aspects related to data structures and metadata. In particular, the new MPEG standards are object-oriented, i.e., adopt video objects as the information units, which is different from the information units used in the current form of video and film, i.e. scenes or shots. Of major importance is the contribution of MPEG-7 and JPEG-2000 to using metadata related to the visual and acoustic content of archived objects. In more detail, MPEG-7 will define a standard for describing multimedia content. The objective is to quickly and efficiently search and retrieve audiovisual material.

To allow interoperability, the standard adopts some normative elements, such as Descriptors (D's), Description Schemes (DS's), the Description Definition Language (DDL) as well as Coding and System Tools. The Descriptors define the syntax and the semantics of the representation of features, while the Description Schemes specify the structure and semantics of the relationships between Descriptors or other Descriptions. Many descriptors have been submitted for MPEG-7, some of which either accepted and included in the eXperimental Model (XM), which is a platform and tool set to evaluate and improve the tools of MPEG-7, or are in the experimentation (Core Experiments, CE) phase. Two parallel levels of descriptors are defined: the syntactic one, which describes the perceptual properties of the content, such as color and motion of spatio-temporal segments and the semantic one, which describes the meaning of content, in terms of semantic objects and events. Syntactic description seems to be well in hand in MPEG-7, but fleshing out the semantic description has not yet received the required attention. It becomes clear among the research community dealing with content – based audio-visual data retrieval and new emerging related standards such as MPEG-7, that the results to be obtained will be ineffective, unless major focus is given to the semantic information level, defining what most users desire to retrieve. Mapping, however, low level, subsymbolic descriptors of a/v archives to high level symbolic ones is in general difficult, even impossible with the current state of technology. It can, however, be tackled when dealing with specific application domains. It seems that the extraction of semantic information from a/v and text related data is tractable taking into account

- the nature of useful queries that users may issue
- the context determined by user profiles

In this paper, we explain that simple models of representation of user profile can provide erroneous interpretation of preferences, in the case where the user has multiple interests, due to possible interference among them, and propose a novel model that overcomes this problem. Moreover, our model introduces negative preferences. Finally, we explain how knowledge and the query context can be used in order to extract a context – sensitive and consistent part of the user profile, which can be subsequently used for personalized content – based retrieval.

The structure of the paper is as follows: In section 2, we provide the mathematical background that is needed for the theoretical analysis that will follow, and describe the knowledge that is used by our system. In section 3, we propose a novel profile representation; continuing, in section 4 we tackle the issues of context adaptiveness and preference consistency. Finally, in section 5 we provide simulation examples for the proposed methods and in section 6 we present our final conclusions, as well as possible extensions to our current work.

2 The knowledge that is stored in the system

As mentioned in section 1, our approach relies on the use of knowledge that is stored in the system. Before continuing, we provide a few details on the mathematical notation in general and the formal definition of the system's stored knowledge in particular.

Let S be a crisp set. A fuzzy set F on S is described by a membership function $M_F : S \to [0,1]$. We may describe the fuzzy set F using the sum notation:

$$F = \sum_i s_i/w_i = \{s_1/w_1, s_2/w_2, \ldots, s_n/w_n\}$$

where $i \in \mathbb{N}_n$, $n = |S|$, $w_i = M_F(s_i)$ or, more simply, $F(s_i)$ and $s_i \in S$. The height of the fuzzy set F is defined as $h(F) = \max_i(w_i)$, $i \in \mathbb{N}_{|S|}$; $|S|$ is the cardinality of S.

A fuzzy binary relation on S is a function $R : S^2 \to [0,1]$. Its inverse relation is defined as $R^{-1}(x,y) = R(y,x)$. The intersection, union and sup $-t$ composition of two fuzzy relations P and Q defined on the same set S are defined as $[P \cap Q](x,y) = t(P(x,y), Q(x,y))$, $[P \cup Q](x,y) = u(P(x,y), Q(x,y))$ and $[P \overset{t}{\circ} Q](x,y) = \sup_{z \in S} t(P(x,z), Q(z,y))$ respectively, where t and u are a t-norm and a t co-norm. The standard t-norm and t-conorm are the min and max functions. An Archimedean t-norm also satisfies the properties of continuity and subidempotency, i.e. $t(a,a) < a, \forall a \in (0,1)$. The identity relation R_I is the identity element of the sup $-t$ composition: $R \circ R_I = R_I \circ R = R, \forall R$.

The properties of reflectivity, symmetricity and sup $-t$ transitivity are defined as follows:

- R is called reflective iff $R_I \subseteq R$
- R is called symmetric iff $R = R^{-1}$
- R is called antisymmetric iff $R \cap R^{-1} \subseteq R_I$
- R is called sup-t transitive (or, simply, transitive) iff $R \overset{t}{\circ} R \subseteq R$

A transitive closure of a relation is the smallest transitive relation that contains the original relation. The transitive closure of a relation is given by the formula $Tr(R) = \bigcup_{n=1}^{\infty} R^{(n)}$, where $R^{(n)} = R \circ R^{(n-1)}$ and $R^{(1)} = R$. It is proved that if R is reflective, then its transitive closure is given by $Tr(R) = R^{(n-1)}$, where $n = |S|$ [8].

A fuzzy ordering relation is a fuzzy binary relation that is antisymmetric and transitive. A partial ordering is, additionally, reflective. A fuzzy partial ordering relation R defines, for each element $s \in S$,the fuzzy set of its ancestors (dominating class) $R_{\geq[s]}(x) = R(s,x)$, and its descendants (dominated class) $R_{\leq[s]}(x) = R(x,s)$. For simplicity, we will use the alternative notation $R(s)$ instead of $R_{\leq[s]}$.

2.1 The Semantic Encyclopedia

Content – based retrieval that is based on terms suffers from the problematic mapping of terms to concepts [7]. Specifically, as more than one terms may correspond to the same concept, or more than one concepts may correspond to the same term, the processing of query and index information is not trivial.

In order to overcome such problems, we choose to work directly with concepts, rather than terms. Hereon we shall refer to these concepts as semantic entities [12]. Let $S = \{s_1, s_2, \ldots, s_n\}, n = |S|$, denote the set of semantic entities that are known. The definitions of these entities, together with their textual descriptions, i.e. their corresponding terms, reside in the semantic encyclopedia; the semantic encyclopedia is the system's knowledge base.

The encyclopedia additionally contains various relations of a semantic nature amongst the semantic entities; these relations constitute the stored knowledge [1],[2]. In this work, we focus on just two of them, the specialization relation, Sp, and the part relation, P [12]. Using them, we construct a novel taxonomic relation, the inclusion relation I, which, as we will explain, is suitable for the handling of user preferences.

The specialization relation 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 forms of specialization are sub – classing, i.e. a is a generalization of b, and thematic categorization, i.e. a is the thematic category of b. The role of the specialization relation in content – based retrieval is as follows: if the user query contains a, then a multimedia document, hereon called document, containing b will be of interest to the user, since it contains a special case of a. Still, there is no evidence that the inverse also holds.

The part relation is also a fuzzy partial ordering on the set of semantic entities. $P(a,b) > 0$ means that b is a part of a. For example a could be a human body and b could be a hand. The role of P in content − based retrieval is the opposite of that of Sp; if the user query contains b, then a document containing a will probably be of interest, because a contains a part b.

In this work, fuzziness of the Sp and P relations has an important role. High values of $Sp(a,b)$ imply that the meaning of b approaches the meaning of a, in the sense that when the user query contains a, then the user will most probably be satisfied with documents containing b. 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 containing b will probably not be of interest to the user. Summarizing, the value of $Sp(a,b)$ indicates the degree to which the stored knowledge shows that an occurrence of b in a document satisfies a request for a in a query. Likewise, the degrees of the part relation can also be interpreted as probabilities of user satisfaction.

The above imply that $a \neq b \implies Sp(a,b) < 1$ and $P(a,b) < 1$, since, if $a \neq b$, then we cannot be sure that both a and b interest the user equally; 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 a difference in a conceptual level.

A last point to consider is that of the transitivity of the part and specialization relations. It is obvious that if b is a part of a and c is a part of b, then c is a part of a. This implies that the part relation is transitive. A similar argument can be made for the specialization relation. Let us now consider a more practical example. Let a be the concept of "car", b the concept of "wheel" and c the concept of "rubber". The inclusion $a < b < c$ is rather obvious. Still, it is not equally obvious that a user requesting documents related to rubber will be satisfied when faced with documents that are related to cars. By this example we wish to demonstrate that the form of transitivity used cannot be $max - min$ transitivity, but one relying on a subidempotent norm. Therefore, we demand that Sp and P are $sup - t$ transitive, where t is an Archimedean norm.

The inclusion relation I is constructed from Sp and P as follows:

$$I = (Sp \cup P^{-1})^{n-1} \tag{1}$$

where $n = |S|$. This means that I is the transitive closure of $Sp \cup P^{-1}$; since the union of transitive relations is not necessarily transitive, this closure is necessary. Based on the roles that Sp and P have in information retrieval, it is easy to see that (1) combines them in a way that implies that, if the user query contains a, then $I(a,b)$ indicates that documents that contain b will also be of interest. Of course we use the same Archimedean t-norm for the transitive closure of I, as we did with Sp and P.

3 Profile representation

As we have already mentioned in section 1, a user profile stores the user's global preferences, i.e. information concerning the user's interests. In order to process the user profile using the stored knowledge, the representation of the former needs to be compatible with the semantic encyclopedia.

As all the relations that exist in the encyclopedia are defined on the set S of semantic entities, we need to define user preferences on the same set. A simple representation of user preferences, that also allows for degrees of preference, is the usage of single fuzzy set defined on the set of semantic entities. As we explain below, this approach is not adequate.

First of all, let us consider the (not rare) case in which a user has various preferences. When the user poses a query that is related to one of them, then that preference may be used to facilitate the interpretation of the query, as well as the ranking of the selected documents. Usage of preferences that are unrelated to the query may only be viewed as addition of noise, as any proximity between selected documents and these preferences is clearly coincidental, in the given context. In order to limit this inter – preference noise, we need to be able to identify which preferences are related to the user query, and to what extent. Therefore, each positive preference needs to be stored separately; a single fuzzy set is not sufficient for the representation of positive user preferences.

Let us now consider the case in which we can infer, by monitoring the user's actions, that he is not interested in documents of a specific subject. In such a case, in addition to positive preferences, special care must be taken for the representation of negative preferences.

The above remarks may be summarized in the following:

- each positive interest needs to be stored separately.
- special care must be taken for the representation of negative preferences.

When applying the above in designing a user preference data model, we need to keep in mind that a user is interested in each subject to some degree. Therefore, interests need to be characterized by a degree of intensity. Moreover, the process of mining the interests is not free of uncertainty. Therefore, mined interests also need to be accompanied by a degree of confidence. This leads to the use of two distinct, seemingly independent degrees, for each preference.

Still, it is easy to see that an intense interest will probably be mined with a greater degree of confidence, than one that is not as intense. It is this observation that allows us to suppose that the degree of confidence and the degree of intensity are not independent. In other words, although two distinct degrees are related to each interest, an intensity degree and a confidence degree, a single degree is sufficient for their representation.

In compliance with the principles presented above, we propose the following formal representation of user preferences P in a user profile:

$$P = \{U^+, U^-\}$$

where U^+ refers to the set of positive preferences and U^- refers to the negative preferences,

$$U^- = \sum_i s_i / u_i^-, \quad i \in \mathbb{N}_n$$

where u_i^- is the degree of participation of entity s_i in U^-,

$$U^+ = \{U_i^+\}, \quad i \in \mathbb{N}_k$$

where k is the count of distinct positive interests that are contained in the user profile, and

$$U_i^+ = \sum_j s_j / u_{ij}^+, \quad i \in \mathbb{N}_k, j \in \mathbb{N}_n$$

where u_{ij}^+ is the degree of participation of entity s_j in U_i^+.

It is easy to see that this definition allows the overlapping of positive and negative preferences. Moreover, it allows the participation of a single semantic entity in different interests, and to different degrees.

4 Context adaptiveness and profile consistency

4.1 Context detection

As we have already explained in section 2.1, the relations Sp, P and I have a special meaning in content – based retrieval. Specifically, they indicate the degree to which the presence of a specific semantic entity in a query (in its textual form) implies local preference for other semantic entities as well. In this section we will explain how, relying on this notion, we may define and detect the context of a query. In the following we present a method for the extraction of the common meaning of a set of entities. We refer to this common meaning as the context of the set.

Let us consider as context of a single semantic entity s the set of semantic entities that it implies. Using the inclusion relation I it is easy to see that this set is the set of the descendants $I(s)$. As context $K(U)$ of a set of semantic entities U we should consider the common meaning of the entities that comprise it; a good way to estimate this common meaning is to use the common context of the entities. In this estimation the following guidelines need to be respected.

- If semantic entity s participates in U to a small degree, then it should not greatly affect the context. More formally: $U(s) \to 0 \implies K(U) \to K(U \cap (S - \{s\}))$
- if semantic entity s participates in U to a great degree , then it the context $I(s)$ should be considered an upper bound for the context $K(U)$. More formally: $U(s) \to 1 \implies K(U) \to I(s) \cap K(U \cap (S - \{s\}))$ This simply means that the context $K(U)$ must be shared by all important members of U.
- Finally, the degree to which s affects the context $K(U)$ increases monotonically, with respect to its degree of participation in U.

The above guidelines are obviously satisfied by the following:

$$K(U) = \bigcap_i K(s_i), \quad i \in \mathbb{N}_n \tag{2}$$

where

$$K(s_i)_j = u(I(s_i, s_j), c(U(s_i))) \tag{3}$$

and u and c are a fuzzy co–norm and a fuzzy complement respectively.

When the entities in U are semantically related to a great extend, then they are highly correlated through I. This results in the presence of high values in $K(U)$. Consequently, we may measure the degree to which entities are semantically related, using the height of their common context $h(K(U))$. We shall refer to this as context intensity and denote it as h_U.

4.2 Adaptation to context

As mentioned in section 1, the interference of preferences that are out of context, in the process of information retrieval, may be considered as noise. Therefore, it is necessary to be able to identify and isolate the set of preferences that are related to the context of a specific user query.

Using the method presented in section 4.1, we may extract the context of the query as a fuzzy set $K(Q) = \sum_i s_i/w_i$ on the set of semantic entities S; the query itself is a fuzzy set $Q = \sum_i s_i/q_i$ on the set of semantic entities S. Based on these definitions, we shall extract the fuzzy set of preferences that are related to the query context:

$$U^K = \sum_i U_i^+/k_i$$

Obviously, in the case that no context can be detected in the query, we would like the whole set of user preferences to be selected. More formally: $h(K(Q)) \to 0 \implies U^K \to U^+$. If, on the other hand, the query context is intense, preferences that do not intersect with the context should not be considered, thus eliminating inter – preference noise. The remaining preferences are considered in proportion to the intensity of their intersection with the context.

A simple formula that complies with the above guidelines is the following:

$$k_i = u(\frac{h(U_i^+ \cap K(Q))}{h_Q}, c(h_Q)) \tag{4}$$

where $h_Q = h(K(Q))$ is the intensity of the query context.

Negative preferences that are out of context do not add noise, and consequently do not have to be filtered with the use of the query context. Therefore, the context – adapted user preferences are:

$$P^K = \{U^K, U^-\}$$

We refer to P^K as the local profile.

4.3 Profile consistency

Within a specific query context we may demand, as a minimum consistency criterion, that the local profile does not contain both positive and negative preferences for the same semantic entities. As we have not imposed such a criterion in the process of constructing or expanding the profile, it is possible that P^K is not in accordance with this criterion.

In order to specify the optimal way of altering the local profile, as to make it compatible with our consistency criterion, we start, once again, by identifying some necessary conditions. First of all, positive preferences are generally extracted with greater confidence. Therefore, positive preferences should be treated more favorably than negative ones.

Obviously, if only positive preferences correspond to a specific semantic entity, then their intensities must not be altered. Likewise, if only a negative preference corresponds to a specific semantic entity, then its intensity must not be altered. In general, the intensities of positive preferences should increase monotonically with respect to their original intensity, and decrease monotonically with respect to the original intensity of the corresponding negative preference, and vice versa.

The above conditions are not particularly strict; there are numerous compliant implementations. In the following, we select a simple linear approach:

A degree of favor $a \in [0, 1]$ is defined. It indicates the degree to which positive preferences are favored, with respect to their negative counterparts. When $a = 1$ there is no distinction, while, as a approaches 0, negative preferences are completely ignored. The decision on whether positive or negative preferences dominate a semantic entity s_i is based on the sign of the expression $\max_j \left(k_j u_{ji}^+ - a u_i^-\right)$. If $k_j u_{ji}^+ - a u_i^-$ is positive for at least one j, then the negative preference cannot dominate, i.e $\hat{u}_i^- = 0$; by $\hat{u}_i^-$ we denote the intensity of the negative preference after the adjustment for consistency. Likewise, if $\max_j \left(k_j u_{ji}^+ - a u_i^-\right) < 0 \ \forall j$, then $\hat{u}_i^+ = 0 \ \forall j$.

The adjusting of dominating values is performed using the following formulas:

$$\hat{u}_{ji}^+ = u_{ji}^+ - \frac{a u_i^-}{k_j}$$

$$\hat{u}_i^- = u_i^- - \frac{\underset{j}{max}(k_j u_{ij}^+)}{a}$$

The above approach produces a valid (i.e. consistent) context – sensitive user profile; this profile may only be used in the process of serving the query whose context was used. Therefore, the adaptation to the query context is an online process.

Table 1. Adjusted negative preference values

	"internal combustion"	"two-stroke"	"four-stroke"	"Diesel"
Original intensity u^-	0.5	0.9	0.45	0.4
Adjusted intensity $\hat{u}^-$	0.308	0.713	0.263	0.23

5 Experimental results

The above representation models and context adaptation methods are being applied and tested on the (in development) prototype of the FAETHON system. FAETHON is also equipped with an experimental semantic encyclopedia which is compatible with the one presented in section 2. It contains 80 semantic entities and definitions of 48 semantic relations, out of which 10 are already populated. They have been used, with encouraging results, not only for the handling of user profiles, but also for query expansion [3] and for automatic thematic categorization of indexed documents.

The existence of this many relations has lead to the divide of the available knowledge amongst them, which in turn results in the need for the utilization of more relations for the generation of an adequate inclusion relation I. For the purpose of handling user preferences we use an inclusion relation that was generated with the use of the following semantic relations:

- Part P, inverted.
- Specialization Sp.
- Example Ex. $Ex(a, b) > 0$ indicates that b is an example of a. For example, a may be "player" and b may be "Jordan".
- Instrument Ins. $Ins(a, b) > 0$ indicates that b is an instrument of a. For example, a may be "music" and b may be "drums".
- Location L. $L(a, b) > 0$ indicates that b is the location of a. For example, a may be "concert" and b may be "stage".
- Patient Pat. $Pat(a, b) > 0$ indicates that b is a patient of a. For example, a may be "course" and b may be "student".
- Property Pr. $Pr(a, b) > 0$ indicates that b is a property of a. For example, a may be "Jordan" and b may be "star".

Of course, since the generation of the I relation did not have a central role in this paper, the methods and formulas it contains are not affected by the fact that more knowledge has been included in it.

The size of the encyclopedia makes it impossible to provide a clear graphical view of the whole I. For this reason we shall restrict ourselves to a small subset, which is adequate for the demonstration of our algorithms. This relation is presented graphically in Figure 1. Values that are implied by transitivity are omitted for the sake of clarity; the archimedean norm we use for the transitive closure is the Yager norm with a parameter $w = 3$.

Let us suppose that a user profile contains the following preferences:
$U_1^+ = \mathrm{ext}/0.81$
$U_2^+ = \mathrm{roc}/0.79 + \mathrm{tur}/0.7 + \mathrm{jet}/0.56 + \mathrm{ext}/0.14 + \mathrm{int}/0.14 + \mathrm{2st}/0.13 + \mathrm{4st}/0.13 + \mathrm{die}/0.12$
$U^- = \mathrm{int}/0.5 + \mathrm{2st}/0.9 + \mathrm{4st}/0.45 + \mathrm{die}/0.4$

Let us now examine how the profile is processed online, for the extraction of the local profile. Let the user query be

$$Q = \mathrm{eng}/1 + \mathrm{air}/0.8$$

By applying the method for context adaptation that is presented in section 4.3, we obtain:

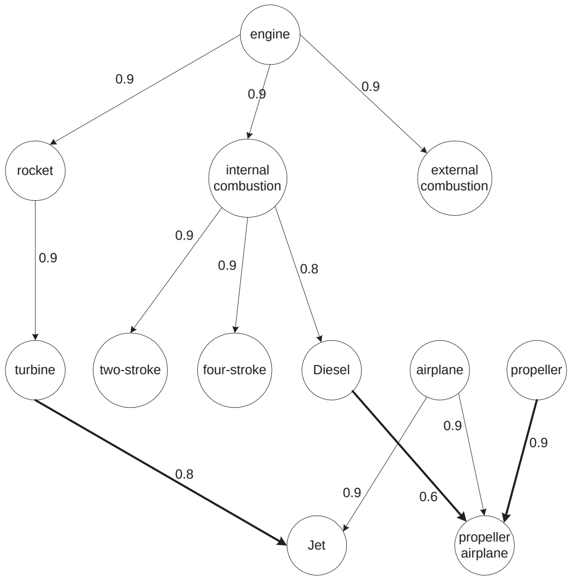

Fig. 1. The Inclusion relation

$$K(Q) = \text{prop-a}/0.582 + \text{jet}/0.785 + \text{turb}/0.2 + \text{roc}/0.2+$$
$$\text{int}/0.2 + \text{2st}/0.22 + \text{4st}/0.2 + \text{die}/0.2 + \text{ext}/0.2 + \text{eng}/0.2$$

and thus:

$$U^K = U_1^+/0.255 + U_2^+/0.713$$

We can observe that the method correctly identified that preference U_1^+ is related to the query context much less than U_2^+. This way, the interference of U_1^+ in subsequent tasks in the process of information retrieval will be reduced.

This local profile is obviously not consistent, as it describes both positive and negative preferences for semantic entities "internal combustion", "two–stroke', "four–stroke" and "Diesel". Applying the criterion described in section 4.3, for $a = 0.5$, we can see that negative preferences dominate all these entities. Therefore these entities will not participate in U_1^+ and U_2^+; some may prefer to consider that they will participate to a degree of 0. Table 1 shows how the intensity of the negative preferences change, as the local profile is re–adjusted for consistency.

6 Conclusions and future work

In this paper we studied the role of user profiles in content – based retrieval. We explained the notion of inter – preference noise and provided a novel model for the representation of user preferences; this model does not suffer from inter – preference noise. Furthermore, we explained how knowledge may be used for the extraction of the query context, as well as how this context can be used to extract a context – sensitive and consistent part of the user profile, to be used in the information retrieval process.

The methods in this paper, are quite general and may be implemented in various ways. Still, the selection of a proper implementation appears to be critical and dependant on the available knowledge. The authors believe there is more work to be done, not only in the selection of the fuzzy operators for the algorithmic part, but also in the selection of a proper knowledge representation.

Moreover, the generation of knowledge remains an open issue. Although in this work the discussion is limited to the way the inclusion relation may be produced, given a set of semantic fuzzy relations, the generation of these relations by a human expert is not a trivial task. Still, in parallel to this work, the authors also explore methods for automatic generation of the inclusion relation, either from usage history [4], or with the combination of existing crisp ontologies and relevance feedback.

Finally, it would be very interesting to explore intelligent and user – oriented ways to mine the distinct semantic user preferences automatically [11], while taking advantage of the system's knowledge, as well as ways to automatically identify the extent to which the profile should contribute in the tasks of query interpretation and document ranking.

7 Acknowledgements

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

References

1. Akrivas G., Stamou G., Fuzzy Semantic Association of Audiovisual Document Descriptions, proc. of International Workshop on Very Low Bitrate Video Coding (VLBV), Athens, Greece, Oct. 2001
2. Akrivas, G., Stamou G., Kollias, S., Semantic Association of Multimedia Document Descriptions through Fuzzy Relational Algebra and Fuzzy Reasoning , submitted
3. Akrivas, G., Wallace, M., Andreou, G., Stamou, G., Kollias, S.,Context Sensitive Semantic Query Expansion, proc. of IEEE International Conference on Artificial Intelligence Systems (AIS), Divnomorskoe, Russia, Sept. 2002
4. Akrivas, G., Wallace, M., Stamou, G., Construction of Quasi - Ordering Fuzzy Taxonomies out of Usage History, submitted
5. Battista, S., Casalino, F., Lande C.: MPEG-4: A Multimedia Standard for the Third Millenium, Part 1. IEEE Multimedia 6 (4) (1999) 74-83.
6. Battista, S., Casalino, F., Lande, C.: MPEG-4: A Multimedia Standard for the Third Millenium, Part 2. IEEE Multimedia 7 (1) (2000) 76-84.
7. Kraft D.H., Bordogna G., Passi G., Information Retrieval Systems: Where is the Fuzz?, IEEE International Conference on Fuzzy Systems, Anchorage, Alaska, May 1998
8. Klir G., Bo Yuan, Fuzzy Sets and Fuzzy Logic, Theory and Applications, New Jersey, Prentice Hall, 1995
9. Nack, F., Lindsay, A.: Everything You Wanted to Know About MPEG-7: Part 1. IEEE Multimedia 6 (3) (1999) 65-77.
10. Nack, F., Lindsay, A.: Everything You Wanted to Know About MPEG-7: Part 2. IEEE Multimedia 6 (4) (1999) 64-73.
11. Wallace, M., Stamou, G., "Towards a Context Aware Mining of User Interests for Consumption of Multimedia Documents", Proceedings of the IEEE International Conference on Multimedia and Expo (ICME), Lausanne, Switzerland, August 2002
12. ISO/IEC JTC 1/SC 29 M4242, Text of 15938-5 FDIS Information Technology – Multimedia Content Description Interface – Part 5 Multimedia Description Schemes, October 2001.
13. ISO/IEC JTC1/SC29/WG1 N1646R: JPEG 2000 Part I Final Committee Draft Version 1.0 (2000).