Learning Imprecise Semantic Concepts from Image Databases

D. Sánchez and J. Chamorro-Martínez

Department of Computer Science and Artificial Intelligence
University of Granada, Avda. Andalucia 38, 18071 Granada, Spain
E-mail: {daniel,jesus}@decsai.ugr.es

Abstract

In this paper we introduce a model to represent high-level semantic concepts that can be perceived in images. The concepts are learned and represented by means of a set of association rules that relate the presence of perceptual features to the fulfillment of a concept for a set of images. Since both the set of images where a perceptual feature appears and the set of images fulfilling a given concept are fuzzy, we use in fact *fuzzy association rules* for the learning model. The concepts so acquired are useful in several applications, in particular they provide a new way to formulate imprecise queries in image databases.

Key words: Fuzzy association rule, perceptual feature, image semantics

1 Introduction

The increasing amount of multimedia information stored in databases, digital libraries, and the web, is motivating a big research effort to solve several problems related to the storage and manipulation of this kind of information. Among the areas where this effort becomes apparent are database systems, with much research devoted to geographical information systems and object oriented database systems, and multimedia information retrieval [**?**]. Multimedia data mining is another promising research field that has arisen in recent years.

In this work we focus on image management. One of the first approaches employed to manage images is based on using metadata attached to them, such as captions, tags, and even textual descriptions. But this approach make

Fig. 1. Two images containing vertical information.

little use of images and the generation of such metadata is a hard task that must be performed by humans.

To avoid these drawbacks, much effort is being devoted to the development of management techniques based on the analysis of features appearing in images, using computer vision resources and techniques. However, there is a gap between signal-based descriptions provided by computer vision techniques, and human-generated metadata required to solve many interesting problems, such as conceptual descriptions of image contents, that is far from being filled.

This is specially apparent in the area of image retrieval. Current image retrieval systems are based mainly in representations of images by means of vectors, where each component summarizes information about a basic feature such as color, shape, etc. [**?**,**?**]. The supported queries are images or sketches, and the retrieving process consists of measuring the similarity between the feature vectors of the query and the images in the database. This leads to an important drawback that can be expressed from different points of view:

- The user is forced to use an image or sketch to query the database. But, how can we express by means of an image a query like "give me images containing a few big objects"?
- Images that should be retrieved for certain queries could be not similar in terms of their feature vectors. Indeed, in many occasions the user's idea of similarity cannot be measured in terms of low-level features but on the basis of higher level concepts. As an example, consider the images of figure 1. These images are similar in that they contain *vertical information*. This fact can be perceived by humans. Nevertheless, the usual features of color, shape and the like take very different values for both images, so they are seen as different on the basis of similarity metrics.

Another key point is that both features and concepts can be (and in fact most of the times are) imprecise, in the sense that they could appear in an image to a certain degree. A clear example of imprecise concepts is present in the query "give me images containing a few big objects", that we used before, where the concept *big* is obviously imprecise, and *a few* is an imprecise quantity.

The appearance of features in an image can also be imprecise. For example, the color is one of the most used features in image analysis, but any color is present in an image to a certain degree.

The theory of fuzzy sets provides a suitable tool to represent imprecision. In the first aforementioned example, the presence of *big* objects can be represented as a fuzzy set on the set of images, and *a few* can be seen as a fuzzy relative quantifier, i.e., a fuzzy subset of the real interval $[0,1]$. In a similar way, the presence of a given color in an image may be represented by using a value in $[0,1]$, so that in practice the set of images where a given color appears is a fuzzy set.

In this work we describe a general methodology to learn and represent high-level concepts on the basis of related basic features provided by signal-based analysis of images. Our methodology employs a set of fuzzy association rules to relate the presence of features to the fulfillment of concepts in a given image, taking into account that both features and concepts could be imprecise. We call this set the *model* of the concept.

We obtain the fuzzy association rules that form the model by using the algorithm we introduced in [**?**]. This algorithm works on a dataset that, in this particular case, consists of a description of the basic features that appear in each image of an image database, together with their appearance degree, and user's valuation of the fulfillment degree of several concepts for each image.

The paper is organized as follows. In section 2 we introduce the concept of fuzzy association rule, that is to be used to learn the concepts. The structure of the dataset from which concepts are to be learned and the definition of model of a concept are described in section 3, together with some existing and potential applications of the models learned. Finally, section 4 contains our conclusions and some research directions we shall follow in the future.

2 Fuzzy Association Rules

2.1 Association rules

Association rules [**?**] are one of the best studied data mining techniques. Starting from a set of items $\mathbb{I}$, the data structure from which association rules are extracted is (or can be interpreted as) a set of transactions, where each transaction is a subset of $\mathbb{I}$. We call this structure a T-set.

Association rules that hold in a T-set are "implications" that relate the presence of itemsets (sets of items) in transactions. The classical example are rules extracted from market baskets. Items are things we can buy in a market, and transactions are market basket containing several items. Association rules relate the presence of items in the same basket, for example "every basket that contains bread contains butter", usually noted *bread* $\Rightarrow$ *butter*.

However, it is difficult to find rules that hold with total accuracy in a T-set. Even so, it could happen that the set of transactions where the rule holds is so small that the rule is uninteresting. This leads to the necessity to assess association rules by measuring both their accuracy and their interest.

For this purpose, the ordinary measures proposed in [?] are *confidence* and *support*, both based on the support of an itemset (a set of items). The support of an itemset is the percentage of transactions where the itemset appears. Given a set of items $\mathbb{I}$ and a set of transactions T on $\mathbb{I}$, the support of an itemset $\mathbb{I}_0 \subseteq \mathbb{I}$ is

$$supp(\mathbb{I}_0) = \frac{|\{\tau \in T \mid \mathbb{I}_0 \subseteq \tau\}|}{|T|} \tag{1}$$

i.e., the probability that the itemset appears in a transaction of T. The support of the association rule $A \Rightarrow B$ in T is

$$Supp(A \Rightarrow B) = supp(A \cup B) \tag{2}$$

Support of a rule is indeed the percentage of tuples where the rule holds.

The classical accuracy measure is confidence, defined as

$$Conf(A \Rightarrow B) = \frac{supp(A \cup B)}{supp(A)} = \frac{Supp(A \Rightarrow B)}{supp(A)} \; . \tag{3}$$

However, other accuracy measures have been employed. Shortliffe and Buchanan's certainty factors [?] have recently been proposed as an alternative to confidence, thereby solving the drawbacks which this has [?]. The certainty factor of a rule is defined as

$$CF(A \Rightarrow B) = \begin{cases} \frac{Conf(A \Rightarrow B) - supp(B)}{1 - supp(B)} & \text{when } Conf(A \Rightarrow B) > supp(B) \\ \frac{Conf(A \Rightarrow B) - supp(B)}{supp(B)} & \text{when } Conf(A \Rightarrow B) < supp(B) \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

Certainty factors take values in $[-1, 1]$, indicating the extent to which our belief that the consequent is true varies when the antecedent is also true.

It ranges from 1, meaning maximum increment (i.e., when A is true then B is true) to -1, meaning maximum decrement. Certainty factors also verify the requirements stated by Piatetsky-Shapiro [**?**] for any suitable accuracy measure. In particular, they are able to check statistical independence between A and B, and in that case $CF(A \Rightarrow B) = 0$.

2.2 Fuzzy Transactions

When mining for association rules we assume data comes in the form of (or can be interpreted as) a set of transactions, each one being a subset of items. But this is not always the case. In many occasions, transactions can be interpreted in fact as fuzzy subsets of items.

For example, suppose that we have a set of market baskets (transactions) $B = \{b_1, \ldots, b_u\}$ and we are interested in relating the presence of expensive items only. For that purpose, we could build a new set of baskets B^e by moving the expensive objects of every basket $b_k \in B$ to a basket $b_k^e \in B^e$. Then we can solve the problem by mining for association rules in B^e. However, *expensive* is an imprecise concept, i.e., every item is expensive to a certain degree. Hence, each basket b_k^e is in fact a fuzzy subset of items.

In general, let $\mathbb{I}$ be a set of items, that we assume to be finite.

Definition 2.1 *We call fuzzy transaction to any nonempty fuzzy subset* $\tilde{\tau} \subseteq \mathbb{I}$.

For every $x \in \mathbb{I}$, we note the membership degree of x in a fuzzy transaction $\tilde{\tau}$ as $\tilde{\tau}(x)$. Looking at definition 2.1 it is obvious that ordinary transactions are an special case of fuzzy transactions.

We represent a set of fuzzy transactions by means of a table. Columns and rows are labelled with identifiers of items and transactions respectively. The cell for item x_k and transaction $\tilde{\tau}_j$ contains a value in $[0, 1]$, the membership degree of x_k in $\tilde{\tau}_j$, noted $\tilde{\tau}_j(x_k)$.

We use the same notation for the membership degree of an itemset $\mathbb{I}_0 \subseteq \mathbb{I}$ to a fuzzy transaction $\tilde{\tau}$, $\tilde{\tau}(\mathbb{I}_0)$. We define

$$\tilde{\tau}(\mathbb{I}_0) = \min_{x \in \mathbb{I}_0} \tilde{\tau}(x)$$

The following example is from [**?**]:

Example 2.2 *Let* $\mathbb{I} = \{x_1, x_2, x_3, x_4\}$ *be a set of items. Table 1 describes three fuzzy transactions defined on items of* $\mathbb{I}$.

	x_1	x_2	x_3	x_4
$\tilde{\tau}_1$	0	0.6	1	1
$\tilde{\tau}_2$	0	1	0	1
$\tilde{\tau}_3$	1	0.4	0.75	0.1

Table 1
Three fuzzy transactions

Here, $\tilde{\tau}_1 = 0.6/x_2 + 1/x_3 + 1/x_4$, $\tilde{\tau}_2 = 1/x_2 + 1/x_4$ *and* $\tilde{\tau}_3 = 1/x_1 + 0.4/x_2 + 0.75/x_3 + 0.1/x_4$. *In particular,* $\tilde{\tau}_2$ *is a crisp transaction,* $\tilde{\tau}_2 = \{x_2, x_4\}$.

We call FT-set a set of fuzzy transactions. Example 2.2 shows an FT-set containing three transactions, $\{\tilde{\tau}_1, \tilde{\tau}_2, \tilde{\tau}_3\}$. Let us remark that an FT-set is a crisp set.

2.3 Association rules in FT-sets

We call *fuzzy association rule* to a rule that hold in an FT-set. This concept was formalized in [**?**] as follows: let $\mathbb{I}$ be a set of items and let T be an FT-set based on $\mathbb{I}$.

Definition 2.3 *A fuzzy association rule is a link of the form* $A \Rightarrow C$ *such that* $A, C \subset \mathbb{I}$ *and* $A \cap C = \emptyset$. A *and* C *are called* antecedent *and* consequent *respectively.*

An association rule is obviously a particular case of fuzzy association rule since, as we remarked before, transactions are a particular case of fuzzy transactions.

The meaning of a fuzzy association rule is the usual for association rules, i.e., C appears in every transaction where A is. The only difference is that the set of transactions where the rule holds is an FT-set, and hence both A and C are in every fuzzy transaction $\tilde{\tau} \in T$ to a certain degree. This leads to the following proposition [**?**]:

Proposition 2.4 *A fuzzy association rule* $A \Rightarrow C$ *holds with total accuracy in* T *when* $\tilde{\tau}(A) \leq \tilde{\tau}(C)\ \forall \tilde{\tau} \in T$.

2.3.1 Assessing fuzzy association rules

To assess fuzzy association rules we extend the measures of support and certainty factors to the case of fuzzy transactions. Certainty factors are obtained from confidence and the support of the consequent, so we extend these measures also.

The calculation of confidence and support in the presence of fuzzy degrees is based on the evaluation of quantified sentences, specifically the method GD [?]. The evaluation of a quantified sentence of the form "Q of H are G" by means of GD, that we note $GD_Q(G/H)$, yields

$$\sum_{\alpha_i \in \Delta(G/H)} (\alpha_i - \alpha_{i+1})\, Q\left(\frac{|(G \cap H)_{\alpha_i}|}{|H_{\alpha_i}|}\right) \tag{5}$$

where $\Delta(G/H) = \Lambda(G \cap H) \cup \Lambda(H)$, $\Lambda(H)$ being the level set of H, and $\Delta(G/H) = \{\alpha_1, \ldots, \alpha_p\}$ with $\alpha_i > \alpha_{i+1}$ for every $i \in \{1, \ldots, p\}$.

The set H is assumed to be normalized. If not, H is normalized and the normalization factor is applied to $G \cap H$. An efficient algorithm to perform evaluation by means of GD is shown in [?].

Using this method, the support of an itemset $\mathbb{I}_0$ is the evaluation of the sentence

$$M \text{ of } T \text{ are } \mathbb{I}_0 \tag{6}$$

where M is a fuzzy relative quantifier [?] defined as $M(x) = x$ on $[0, 1]$. Therefore, the support of a rule $A \Rightarrow C$ in T is the evaluation of

$$M \text{ of } T \text{ are } A \cup C \tag{7}$$

Finally, the confidence of a rule $A \Rightarrow B$ is the evaluation of

$$M \text{ of } A \text{ are } B \tag{8}$$

and the certainty factor is calculated from support and confidence by using (4). One interesting property of this approach is that it generalizes the measures of support, confidence and certainty factor for crisp data. Other properties can be found in [?], together with an algorithm to mine for fuzzy association rules. In fact, the methodology employed in our algorithm can be used to adapt any existing association rules mining algorithm, in order to obtain algorithms to mine for fuzzy association rules.

3 Learning concepts

Our objective is to introduce a methodology to obtain a description of a concept, based on relating the presence of low-level features to the fulfillment of a concept in images. Let us remark that concepts and features are assumed to be related in some way. For example, if we want to learn concepts about orientation, we should use the basic perceptual features for orientation of the human visual system [?].

In our approach, concepts are to be learned from an image database I and user valuations of the images with respect to the concepts. More specifically, for every image in I a set of features F is obtained by using computer vision techniques. These features, together with the valuations obtained from users, form a dataset. We call *representation* of the image the set of features and valuations for a given image. Since we consider that both features and concepts can appear in an image to a certain degree, this representation is in practice a fuzzy transaction.

Starting from a dataset, we shall use fuzzy association rule mining to discover associations between perceptual features and concepts. A set of association rules relating features to a given concept c_k can be interpreted as a description, we call *model*, for c_k. The concept of model has been employed before as a description of approximate dependencies, i.e., functional dependencies with exceptions in relational databases, see [?,?].

3.1 Representation of images

The representation of an image was introduced in [?] as follows: let F be a set of features and let C be a set of concepts.

Definition 3.1 *Given an image $i \in I$, the representation of i, $r(i)$, is a fuzzy subset of $F \cup C$.*

By definition $r(i)$ is a *fuzzy transaction* characterized by a function

$$\mu_{r(i)} : F \cup C \to [0,1] \tag{9}$$

where $\mu_{r(i)}(o)$ is the degree to which the image i verifies o. We shall note $r^F(i) = r(i) \cap F$ and $r^C(i) = r(i) \cap C$.

Definition 3.2 *The representation of an image database I is the set*

$$R(I) = \{r(i) \mid i \in I\} \tag{10}$$

Since $r(i)$ is a fuzzy transaction, $R(I)$ is an FT-set. We note

$$R^F(I) = \{r^F(i) \mid i \in I\} \tag{11}$$

	f_1	f_2	f_3	c_1
$r(i_1)$	0	0.6	1	1
$r(i_2)$	0	1	0	1
$r(i_3)$	1	0.4	0.75	0.1

Table 2
Three fuzzy transactions

$$R^C(I) = \{r^C(i) \mid i \in I\} \tag{12}$$

As we saw in section 2.2, an FT-set can be represented in a table. The following example is from [?]: let $I = \{i_1, i_2, i_3\}$ be a set of three images, and let $F = \{f_1, f_2, f_3\}$ and $C = \{c_1\}$. Table 2 represents an imaginary representation for $R(I)$. The representation of i_1 is the fuzzy set $r(i_1) = 0.6/f_2 + 1/f_3 + 1/c_1$, meaning that the image i_1 verify the features f_2 and f_3 with degrees 0.6 and 1 respectively, and the concept c_1 with degree 1. Also, $r^F(i_1) = 0.6/f_2 + 1/f_3$ and $r^C(i_1) = 1/c_1$. Finally, $R^F(I)$ and $R^C(I)$ are represented by the columns 1-3 and 4 respectively.

In practice, $R^C(I)$ will be provided by an user, while $R^F(I)$ will be obtained by image analysis. In [?] we describe how we analyze features related to the perception of orientation, to be stored in $R^F(I)$, and its application to image retrieval.

3.2 Model of a concept

We are going to model concepts by mining for fuzzy association rules in $R(I)$. As we have described before, items are features and concepts, and the FT-set is $R(I)$, i.e., the set of representations for a given image database I.

For each concept, a set of fuzzy association rules is obtained such that the concept is in the consequent of the rule. We restrict the items in the antecedent to those in F, and we note $Rul(I, c_k)$ the set of rules holding in $R(I)$ whose consequent is c_k, i.e., any rule in $Rul(I, c_k)$ takes the form

$$\{f_{i_1}, f_{i_2}, \ldots, f_{i_p}\} \Rightarrow c_k \tag{13}$$

with $i_q \in \{1, \ldots, |F|\}$, and $i_q \neq i_r$.

We introduce the model of a concept to be the following:

Definition 3.3 *We call* model *of a concept $c_k \in C$, obtained from an image database I, to a set of rules $Rul(I, c_k)$.*

Let us remark that we are only interested in rules with $CF \in (0, 1]$, since this means that the presence of the set of features $\{f_{i_1}, f_{i_2}, \ldots, f_{i_p}\}$ in an image increase our certainty that the image verifies the concept. Hence, $Rul(I, c_k)$ contains only those rules with $CF > 0$.

If I is assumed to contain both representative positive and negative examples of the concept, we can assume that we have learned the concept c_k. This knowledge is represented by the model $Rul(I, c_k)$, that we shall note $Rul(c_k)$ since it is assumed to be valid for any image and not only for those in I.

However, since it is an user who provides a degree for the compatibility between a concept and an image, the model $Rul(c_k)$ represents the concept c_k *as seen by the user*, i.e., a subjective view of c_k. We note $Rul_P(c_k)$ the model for c_k as seen by user P.

If we want to obtain an objective model for c_k (or at least, as objective as possible) a possible solution is to aggregate the valuations provided for every image by a group of users. This way we will obtain a prototypical valuation for each image $i \in I$, i.e., how well the images in I fit the concept c_k as seen by all users. If the group of users is considered to be representative, the model obtained from the aggregated information can be considered as an objective model for c_k.

In occasions, and by the nature of the concepts, we shall be interested only in rules of the form $f_j \Rightarrow c_k$, with $f_j \in F$ and $c_k \in C$. This is the case of many important features of the human visual system. In particular, Graham and Nachmias [**?**] have showed the independence between the presence of information at any particular size and orientation.

Figure 3 shows a graphical representation, extracted from [**?**], corresponding to the application of our methodology to concepts related to orientation. It shows some of the models we obtained from a database I containing 75 images. Some images in I are shown in figure 2. The models where obtained for three users $\{P_1, P_2, P_3\}$ and a set C_O of four concepts: *horizontal* (H), *vertical* (V), *diagonal around 45 degrees* (D1) and *diagonal around -45 degrees* (D2).

In these models, a feature $f_j \in F_O$ with $j \in \{1, \ldots, 12\}$ represents a degree of activity at a particular angle in an image, that has been obtained using properties of the human visual system [**?**]. Specifically, f_j corresponds to the activity around the angle $(15 \cdot (j - 7))°$.

In figure 3, the graphic for the model $Rul_{P_l}(I, c)$ is a function

$$G_{P_l}^{(I,c)} : F_O \rightarrow [0, 1] \qquad (14)$$

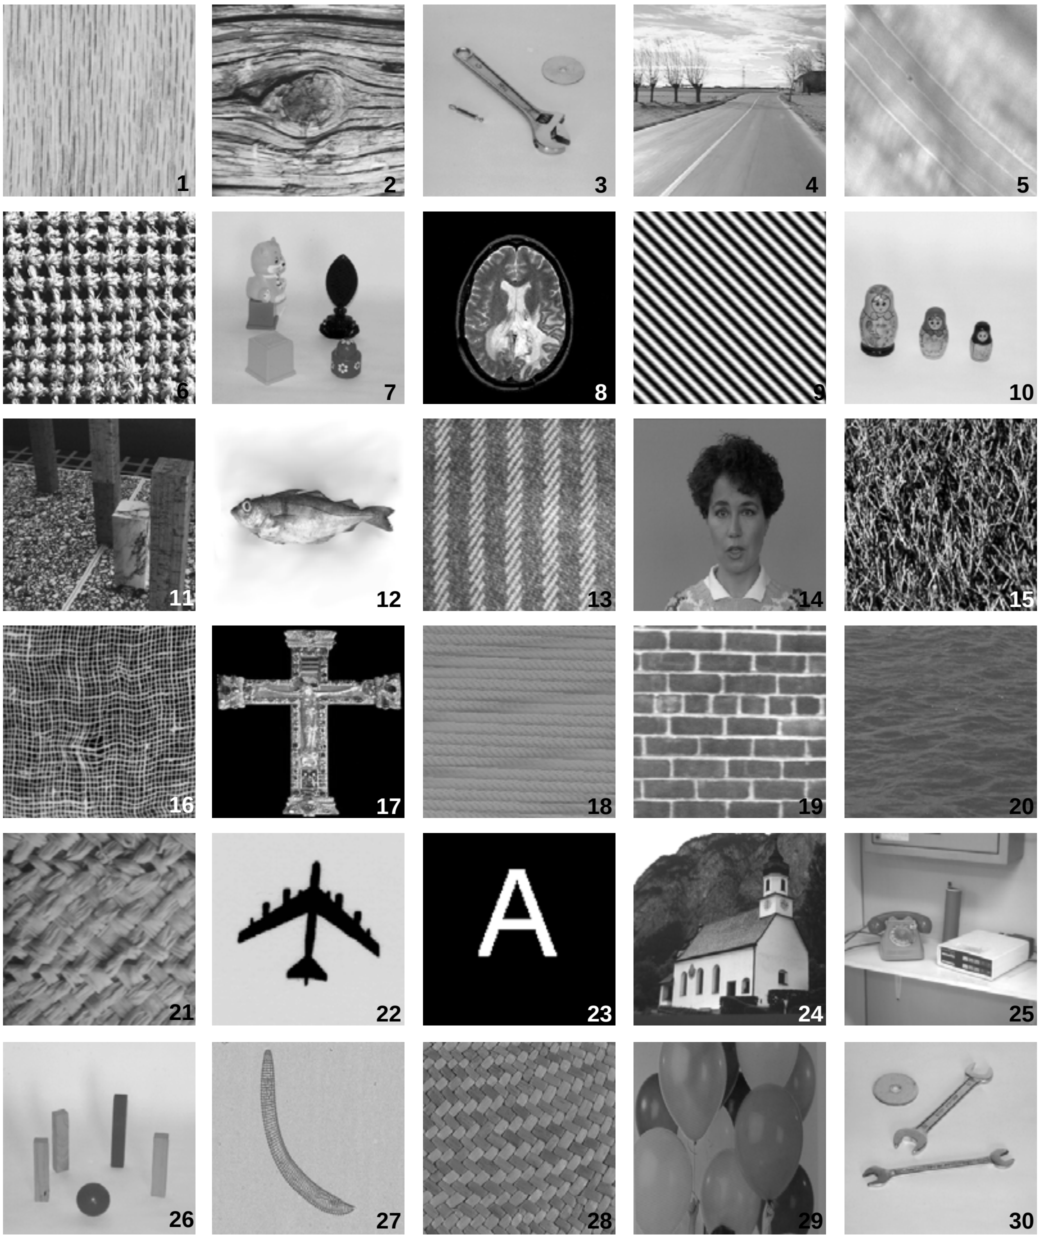

Fig. 2. Some examples of images contained in the database I

such that

$$G_{P_l}^{(I,c)}(fo_k) = \begin{cases} CF(fo_k \Rightarrow c) & \text{when } (fo_k \Rightarrow c) \in Rul_{P_l}(I,c) \\ 0 & \text{otherwise} \end{cases} \tag{15}$$

The points in the graphics have been connected using lines to better appreciate the functions and their differences. Let us remark that, for all the users, all

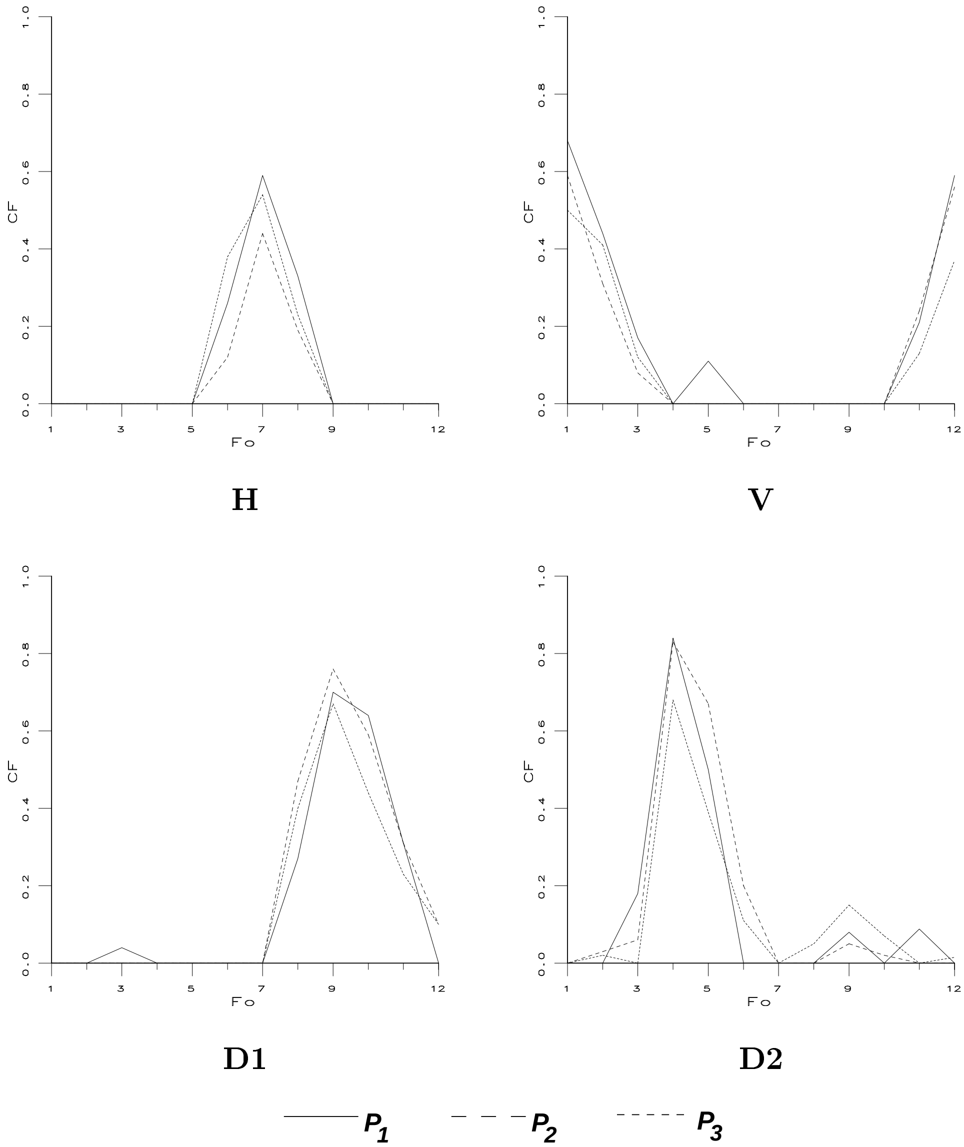

Fig. 3. Graphical representation of the models $Rul_{P_l}(I, c)$ for every user and concept.

the models are coherent in some sense with the corresponding concepts. For example, the certainty that a given image fits the concept *horizontal* increases when the feature f_7 (around 0°) appears. Similar results hold for the rest of concepts in C_O.

3.3 Applications

Now that we have described our methodology to learn and to represent concepts, we shall briefly describe some existing and potential applications.

3.3.1 Image retrieval

In [?] we introduced the application of the learned concepts for image retrieval. The retrieving process we are going to detail now is a direct extension of that described in [?].

Once the model for a certain concept is obtained, we can use it to perform image retrieval in the following way. Suppose we are given an image database I, and we want to retrieve those images in I that verify the concept c_k. We consider that each image verifies c_k with a certainty factor of $Cer(i, c_k) = 0$, meaning ignorance, before the start of the retrieving process. Then we consider $r^F(i)$ and we apply those rules $\{f_{i_1}, \ldots, f_{i_p}\} \Rightarrow c_k$ in $Rul(I, c_k)$ such that

$$\left(\min_{k \in \{1,\ldots,p\}} \mu_{r^F(i)}(f_{i_k}) \right) > 0 \tag{16}$$

i.e., those rules whose antecedent is a set of features appearing in i.

The application of the rule $\{f_{i_1}, \ldots, f_{i_p}\} \Rightarrow c_k$ consists of the following: let α be the certainty that i verifies c_k *before* the rule is applied, and let

$$\delta = \left(\min_{k \in \{1,\ldots,p\}} \mu_{r^F(i)}(f_{i_k}) \right) CF(\{f_{i_1}, \ldots, f_{i_p}\} \Rightarrow c_k) \tag{17}$$

The certainty that i verifies c_k *after* the rule is applied is [?]

$$Cer(i, c_k) = \alpha + \delta(1 - \alpha) \tag{18}$$

At the end of this process, we have a set of values $Cer(i, c_k)$ for every $i \in I$. The final set of retrieved images is obtained by one of these procedures:

- Using a threshold minCer provided by the user. This way, every image i with $Cer(i, c_k) \geq minCer$ is retrieved.
- Performing a clustering of the images in I in terms of the values $Cer(i, c_k)$. The idea is to obtain two clusters, corresponding to retrieved and not retrieved images. This procedure has the advantage that the threshold is computed in an automatic way.

An example of this application can be found in [?]. That work focused specifically on concepts related to orientation such as *horizontal* and *vertical*. For that purpose a scheme to represent visual orientation was introduced, based on the perceptual features of the human visual system as described in several physiological studies [?,?,?,?]. Several existing computational vision models are based on these studies [?,?].

We performed some experiments on the image database I containing 75 images mentioned before, that showed that our approach was suitable for the retrieval task, allowing to process queries such as "give images containing diagonal around $-45°$ information" [?]. For this query, the images retrieved among those in figure 2 were $\{3, 5, 9, 27, 28\}$. Notice that this query is not based on using an image or sketch, but in using a previously learned concept.

3.3.2 Image similarity

The problem we face is to measure the similarity of two images. There is a lot of work in this area [?,?,?], but it is based on a direct comparison of low-level data. Our intention is to provide similarity measures on the basis of a given set of concepts

A first approach to the problem is the following:

Definition 3.4 *We introduce the similarity of two images i_1, i_2 on the basis of a single concept c to be*

$$Sim(i_1, i_2, \{c\}) = 1 - \frac{|Cer(i_1, c) - Cer(i_2, c)|}{\max\{Cer(i_1, c), Cer(i_2, c)\}} \tag{19}$$

when $Cer(i_1, c)Cer(i_2, c) > 0$, and 0 otherwise.

Notice that $Sim(i_1, i_2, \{c\}) \in [0, 1]$, 1 meaning maximum similarity. Using this definition we could find that the images in figure 1 are similar in a high degree with respect to the concept *verticality*. On the contrary, similarity measures based on the direct comparison of low-level features won't find them similar.

Once we have defined the similarity based on one concept, the similarity on the basis of a set of concepts $\mathbb{C}$ can be obtained as an aggregation of the similarities $Sim(i_1, i_2, \{c\})$ for every concept $c \in \mathbb{C}$. In particular, quantifier guided aggregation [?] provides a suitable framework for aggregation. Similarity should be measured by evaluating assertions such as "the images are similar for Q concepts", where Q is a quantifier such as *All*, *Many*, *Most*, etc.

For example, for the quantifier *All* we should obtain

$$Sim_{All}(i_1, i_2, \mathbb{C}) = \bigwedge_{c \in \mathbb{C}} Sim(i_1, i_2, \{c\}) \tag{20}$$

where $\wedge$ is a t-norm.

Measuring image similarity is therefore a potential application of the concepts learned that, moreover, can be considered as closely related to the application in image retrieval. In particular, in those cases where a query is formulated by means of an image or sketch, we could retrieve those images similar to the query on the basis of a given set of concepts. These concepts could be fixed, or they could be specified in the query by the user. This way, a query could take the form "give me images that contain the same *vertical* information than this one".

It is our intention to consider this and other approaches to measure image similarity on the basis of semantic concepts; also, to deep in the study of the properties of the resulting similarity measures is an interesting research work that we shall face in the future.

4 Conclusions

We have introduced a general methodology to learn high-level concepts on the basis of basic perceptual features that can appear in images. Our methodology is based on using the concept of fuzzy association rule to describe links between features and concepts. The source of knowledge is a representation of an image database containing both features and user valuations (fuzzy degrees) about the concepts.

Among the applications of the learned models for semantic concepts we have mentioned the development of similarity measures for images, and image retrieval. In the latter area we have developed some retrieval models that have yielded good results.

A valuable feature of our approach is that we can perform a query without providing an image or sketch, but asking for images verifying a certain concept. To our knowledge, this is a possibility not offered by existing approaches. However, the development of similarity measures based on sets of concepts will allow for using the learned concepts together with images or sketches to query an image database.

In the immediate future we plan to apply our general methodology to several perceptual concepts (like size, motion and colour) and to perform experiments in large databases. Other research we will face soon is the study of similarity measures and their properties.

References