A Theory of Fuzzy Frames

by

Ian Graham and Peter Llewelyn Jones

PART 1

Abstract

The theory of Fuzzy Frames is developed with examples. It is related to existing notions in knowledge representation and fuzzy mathematics, and directions for further research are explored.

KEYWORDS: Fuzzy Set, Frame, Semantic Network, Property Inheritance, Closed World Assumption, Fuzzy Quantifier, Knowledge Representation, Nonmonotonic Logic, Relational Database, Fuzzy Relation, Truth Maintenance, Usuality.

Authors addresses:

Ian Graham,

Logica Financial Systems Ltd.,

64 Newman Street,

London W1A 4SE

Peter Llewelyn Jones,

Creative Logic Ltd.,

Brunel Science Park,

Kingston Lane,

Uxbridge,

Middx UB8 3PQ.

---

1. Introduction

In this paper we introduce a new computational method of representing uncertain and certain knowledge which we have developed as a generalisation of the frame notion introduced by Minsky [8]. Since our generalisation uses, in an intrinsic way, the theory of fuzzy sets due to Zadeh [15], it is natural that we choose to designate our generalised objects 'fuzzy frames'. We hope, in this paper, to lay the basis for the future development of the theory of fuzzy frames, and to show how they may be applied in knowledge engineering applications.

To begin with we review briefly the existing theory of frames (variously called 'schema', 'units', or 'scripts' in the literature [1,3,4,8,11,18,19]) and the, perhaps less familiar, machinery we require from fuzzy set theory. At the end of the paper we explore some of the intriguing questions which our theory raises, some of its problems and suggest topics for further research. In doing this we have cause to compare our approach with the fuzzy quantifiers of Zadeh [16,3,5], truth maintenance systems and nonmonotonic logic [7,10]. We suggest that Fuzzy Frames offer a unified framework for the representation of both certain and uncertain knowledge about objects, and, in a sense to be explained, generalise fuzzy relations and *a fortiori* relations.

2. Representing knowledge about objects

In terms of knowledge engineering, there are many ways to represent knowledge; as productions, by logic, in procedures, etc. These formalisms are usually better at expressing particularly suitable types of knowledge; about causality, relationships, methods, etc. See [2,3] for a more complete treatment. The forms of knowledge representation which seem to best capture knowledge about objects and their properties are those which are generally referred to as semantic networks and frames. In this paper we concentrate on these. Often such knowledge is uncertain, and usually some additional mechanism has to be introduced to model the uncertainty. This can be done by assigning certainty factors or probabilities to the rules or their atomic clauses, or through the use of some truth maintenance procedure, depending on the type of uncertainty involved. Here we concentrate of the kinds of uncertainty which can be readily modelled with fuzzy sets, but, in principle, our arguments should apply to stochastic problems equally well.

2.1 Semantic Networks and Frames

A semantic network consists of a set of nodes and a set of ordered pairs of nodes called 'links', together with an interpretation of the meaning of these. In the terminology of Winston [13] we will restrict ourselves to describing this interpretation using a *descriptive semantics*; that is, a set of statements describing the interpretation. Terminal links are called 'slots' if they represent properties (predicates) rather than objects or classes of objects. A *frame* is a semantic net representing an object (or a stereotype of that object) or a class of objects, and will consist of a number of slots and a number of outbound links. Consider, for clarification, the frame for a toy brick shown in figure 2.1 in the form of a network.

It may also be represented in a tabular form as follows.

<u>Brick-12</u>

IsA: Brick,

Toy

Colour:Red

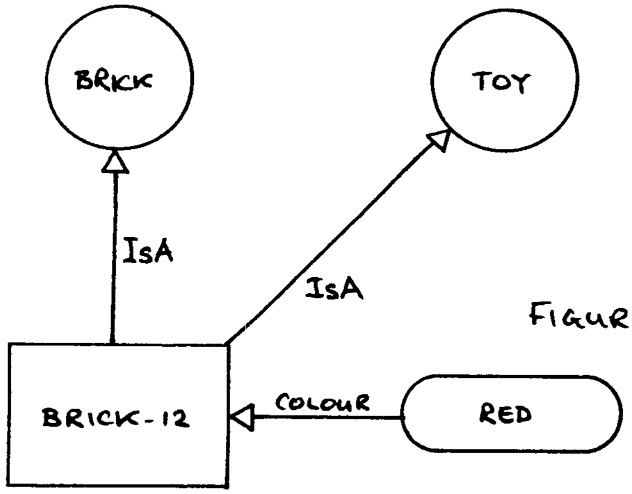

FIGURE 2.1

A collection of frames forming a semantic network will be referred to in this paper as a 'framebase'. In the above example, there are implicitly frames for Brick and Toy.

Figure 2.2 illustrates the inheritance the inheritance of properties into slots as shown also in the form of frames below.

FIGURE 2.2

<u>Brick</u>

IsA: Block,Commodity

Shape: Cuboid

<u>Toy</u>

IsA: Commodity

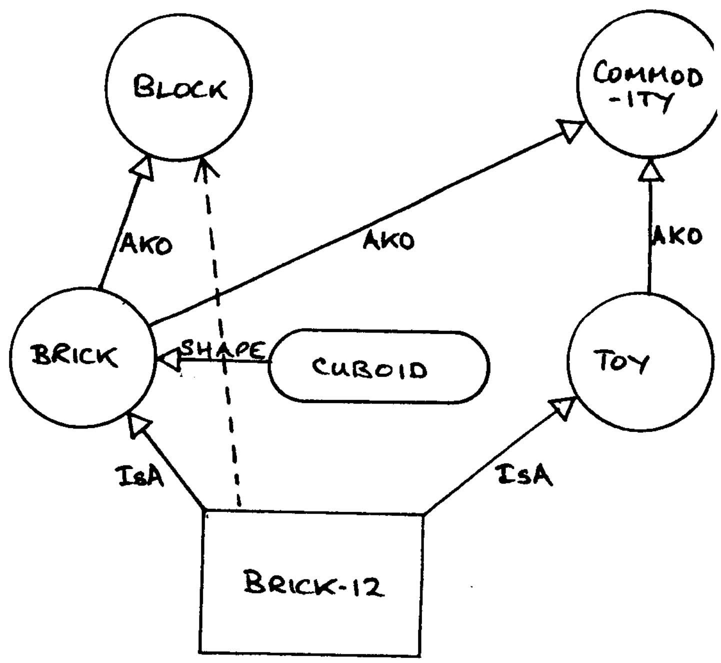

Frames can inherit properties through IsA links, so that Brick-12 inherits the Shape slot's value from Brick in this case, as well as those properties of toys, commodities and blocks which offer no contradiction. Some authors, including Winston, make a distinction between IsA links, which connect individuals to classes, and AKO (a kind of) links, which connect classes. This has relevance for attempts to reduce frame systems to first order logic, such as found in Hayes [4], but need not concern us here. We follow the tradition of FRL [11] and use IsA to stand for both. In any case, Touretzky [12] has argued that such attempts founder when the frame systems permit multiple inheritance (see below).

Winston [13] introduces sub-slots (facets) to permit default values, demons and perspectives. This will be discussed in a forthcoming paper, and only hinted at herein.

2.2 Property Inheritance

Inheritance of the type described provides computer systems with a method of reasoning with implicit facts. Various frame based languages have been implimented; FRL, KRL [11,1,19], ... Most of these systems however suffer from various problems. They usually have no formal semantics, they are not good at reasoning about exceptions (nonmonotonic logic) [12,7,10] and they cannot handle partial inheritance, either in the sense of partially inheriting a property or of inheriting a combination of partially true properties. Other authors have discussed partial inheritance, but from a different point of view. An example is [6]. The programme begun with this paper aims to remedy all these defects within a single unified framework and, additionally, unite the theory with the relational model of data.

Touretzky [12] points out that nonmonotonic or default logics, while possessing a formal semantics, are hopelessly general for practical purposes because they do not have the facility of inheritance systems to reason with implicit data.

We assert here a sort of duality between logics that add extra operators, such as L and N in modal logics and M in nonmonotonic logics, and those that expand the truth space, such as many valued, fuzzy and probability logics. An example of the latter, which we exploit in this paper, is fuzzy logic with Zadeh's fuzzy quantification given the Σ-Count interpretation of his test-score semantics [16,5].

We will show informally later how the mechanisms of Fuzzy Frames can be used to model this logic and compare this approach with that of nonmonotonic logic.

Fuzzy Frames also provide a computationally efficient means of modelling truth maintenance systems or possible worlds, without the introduction of modal operators. We justif_y this claim towards the end of the paper

A certain amount of the machinery of fuzzy logic will be required, and we introduce this as briefly as possible.

3. Basic Fuzzy Set Concepts

The concept of a fuzzy set is due to Zadeh [15] and involves the relaxation of the restriction on a set's characteristic function that it be two valued. In this section we give a very fast summary of the techniques from fuzzy set theory which we require later in this paper. This is merely to fix terminology, and is not intended as a tutorial. Fuller explanations can be obtained from [3] or [5].

3.1 Fuzzy Sets

A fuzzy set is a function λ whose codomain is the unit interval. It may be interpreted as a linguistic value over the variable represented by the domain. For example, if the domain represents wealth (over an arbitrary interval scale) we can introduce fuzzy sets to stand for the imprecise linguistic terms 'rich', 'comfortable' and 'poor' as illustrated in the diagram in figure 3.1. Fuzzy sets are conveniently represented pictorially in this way. They may also be represented as vectors of truth values.

There are several fuzzy logics. In the standard one which we adopt here the operations of the propositional calculus are defined for fuzzy predicates as follows.

λ AND v = $\min(\lambda, v)$

λ OR v = $\max(\lambda, v)$

NOT λ = $1 - \lambda$

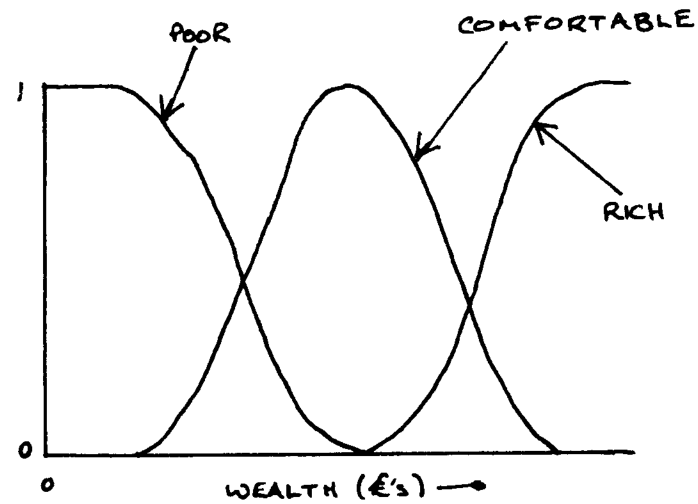

FIGURE 3.1

Implication is then defined in the usual way, by $(\lambda \Rightarrow v) = (\text{NOT } \lambda) \text{ OR } v$

Given a term set of permissible linguistic values, it is possible to extend it using the propositional operators and operations known as hedges. As an example the hedge 'very' is often defined by:

$$\text{VERY } \lambda = \lambda^2$$

Thus, expressions such as 'very rich or not very poor' receive an interpretation as a fuzzy set. This has evident utility in knowledge representation, and this is explored in many other works (cf. [3]).

3.2 Rules of Inference

Given these basic definitions it is possible to approach the problem of representing inexact inferences with fuzzy sets. The two kinds of statements we wish to use are assertions of the form 'X is A' and rules of the form 'If X is [not] A [and|or X' is ...] then Y is B'. X, X' and Y stand for objects and A and B stand for fuzzy sets. Simple syllogisms such as

X is A

<u>If X is A then Y is B</u>

Y is B

are handled as follows. The extension in the cartesian product of all linguistic variables appearing as assertions is computed and the intersection E taken. This fuzzy set is interpreted as an elastic constraint on the solution. Next, taking each consequent clause separately, the maximum value of

E is used to determine the truth of the antecedent, so that the fuzzy set B is effectively truncated at this level. The resultant, truncated fuzzy set is the value of Y. When generalised, this method is known as the compositional rule of inference. This rule will be used in this paper, although others have been suggested.

3.3 Defuzzification

In case it is not convenient to work with fuzzy sets as output values, the result may be 'defuzzified' to return a scalar value. There are a number of ways this can be accomplished. We will need to know about two. The 'mean of maxima' (or maximum) method involves returning the scalar in the domain of the resultant fuzzy set which is the arithmetic mean of its maxima. The 'centre of moments' (or moments) method returns the average of all domain values weighted by their truth in the output fuzzy set. The appropriateness of these methods in different applications is discussed extensively in [3].

If a fuzzy set is required as output but some regularity in its form is desirable, then a method known as linguistic approximation may be invoked. This involves the predefinition of an allowed 'term set' of fuzzy sets over the domain. In the case of fuzzy numbers (fuzzy subsets of the real line) an example term set might be {tiny,small,medium,large,huge}. The term set may be extended by fuzzy set operations (e.g. to include 'not very small'). Linguistic approximation returns the term 'closest' to the resultant fuzzy set, according to some stated measure of distance.

3.4 Fuzzy Quantifiers

Fuzzy quantifiers are represented by words such as 'most', 'almost all', 'some' and so on, as opposed to the crisp quantifiers 'for all' and 'there exists'. They often occur implicitly in natural language. Thus, 'birds fly' may be interpreted as 'most birds can fly'. Zadeh [16] introduces rules of inference and a formal semantics for such statements. His test-score semantics interprets quantifiers as elastic constraints on a family of fuzzy relations, which are regarded as entities. This structure is viewed as a test which scores according to the compatibility of the quantifier with the objects in the world (the database). A typical rule of inference for this system is

$$\frac{\begin{array}{l} Q_1 \text{ A are B} \\ Q_2 \text{ (A and B) are C} \end{array}}{Q_1 * Q_2 \text{ A are (B and C)}}$$

where the Q's are fuzzy quantifiers, interpreted as fuzzy numbers, and * stands for the product of fuzzy numbers. For example this justifies the syllogism:

Most (about 90%) birds can fly
<u>Most (about 90%) flying birds have feathers</u>
At least many (about 81%) birds have feathers

Of course there are other rules of inference. The exact meaning of the test scores depends on the measure used for the cardinality of a fuzzy set. The usual one is the Σ-count measure, which is to say the arithmetic sum of the grades of membership (the integral in the continuous case). More details may be found in [5].

This completes the presentation of the minimal set of concepts from fuzzy set theory which we shall require in this paper.

4. Fuzzy Frames

We now come to the main new results of this paper. First, we extend the notion of a frame to that of a fuzzy frame in two ways. First, by allowing slots to contain fuzzy sets as values, in addition to text, list and numeric variables. Second, we allow inheritance through IsA slots to be partial. Later we will generalise in a third dimension by allowing frames to contain more than one 'tuple', or set of slot values, to facilitate the representation of possible worlds or time-variant objects. To explain the advantages and mechanics of fuzzy frames it is preferable to use an example, rather than to develop the formal mathematics. In doing so we will introduce a syntax similar to that used in Leonardo III [18].

Consider the following (toy) problem. You are faced with the problem of estimating the safety implications following on the purchase of various leisure items. Frames give us a way of representing knowledge about and data concerning objects or concepts. The most natural way to analyse (remembering that we are thinking of building a computerised advisor here) the problem is to list the objects involved. Suppose they include a dinghy and a hang-glider among a number of others. These objects are types of more general objects, and have associated with them various properties. Figure 4.1 shows how we can represent our knowledge about some objects interesting in this context using fuzzy frames. We will now consider nine of these frames, in each case annotating them with explanation of the syntactic convention, on the way to a solution of the problem.

FIGURE 4.1
Some Fuzzy Frames

```
F1  Commodity
    IsA         :Object
    Uses        :undefined [list]
    Owner       :undefined [text]
    Cost        :undefined [fuzz]
    Necessity   :undefined [fuzz]
    Utility     :high [fuzz]
    Safety      :high [fuzz]

F2  Vehicle
    IsA         :Commodity
    Uses        :(travel,pleasure)
    Keeper      :undefined [text]
    Necessity   :high
    Safety      :high
    Utility     :high
    Cost        :high

F3  Toy
    IsA         :Commodity
    Uses        :(pleasure)
    Keeper      :child
    Necessity   :low
    Safety      :undefined [fuzz]
    Utility     :high
    Cost        :low

F4  Borrowed-object
    IsA         :Object
    Lender      :undefined [text]
    Keeper      :undefined [text];default=finder
    Cost        :mimimal [fuzz]

F5  Dangerous-object
    IsA         :Object
    Safety      :minimal

F6  Dinghy
    IsA         :Vehicle [0.4],
                 Toy [0.6],
                 Dangerous object [0.1]
    Safety      :undefined
    Cost        :undefined
    Draught     :3 [real];IfNeeded=depth-calc

F7  Hang-glider
    IsA         :Vehicle [0.05],
                 Toy [0.7],
                 Dangerous object [0.9]
    Safety      :undefined
    Cost        :undefined

F8  Car
    IsA         :Vehicle [0.9],
                 Toy [0.6],
                 Dangerous object [0.1]
    Safety      :undefined
    Cost        :undefined

F9  Toy-car
    IsA         :Vehicle [0.3],
                 Toy [0.9]
    Cost        :undefined
    Safety      :undefined

F10 Book
    IsA         :Dangerous-object,
                 Toy [0.6]
    Cost        :undefined
    Safety      :undefined

F11 Magazine
    IsA         :Book [0.5],
                 Toy [0.3],
                 Borrowed object [0.5]
    Safety      :undefined
    Cost        :undefined

F12 Dinghy-123
    IsA         :Dinghy
    Draught     :undefined
    Safety      :undefined
    Cost        :undefined

F13 Hang-glider-765
    IsA         :Hang-glider,
                 Borrowed object
    Safety      :undefined
    Cost        :undefined
```

Notice, incidentally, the syntactic provision for defaults and backward chaining demons (IfNeeded procedures). Forward chaining demons would be dealt with similarly.

First, consider the most general concept of a commodity. Our general knowledge about commodities can be summarised in the following structure.

```
F1   Commodity
     IsA        :Object
     Uses       :undefined [list]
     Owner      :undefined [text]
     Cost       :undefined [fuzz]
     Necessity  :undefined [fuzz]
     Utility    :high [fuzz]
     Safety     :high [fuzz]
```

Here the IsA slot points to another frame, in this case the most general one possible. Commodity is a fuzzy frame in two respects. First, the degree of property inheritance from the frame(s) in the IsA slot may be specified as a number between 0 and 1 in square brackets after the name. In this case no value is given and the default value of [1.00] is assumed. Secondly, the other slots may contain fuzzy sets (vectors of truth values) as values. The bracketed expressions indicate the type of the value; either [fuzz], [real], [list] or [text]. The fuzzy sets used in the Commodity frame may be represented as follows:

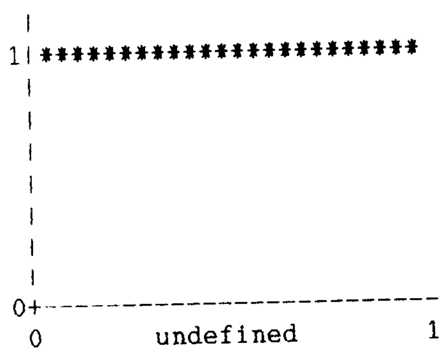

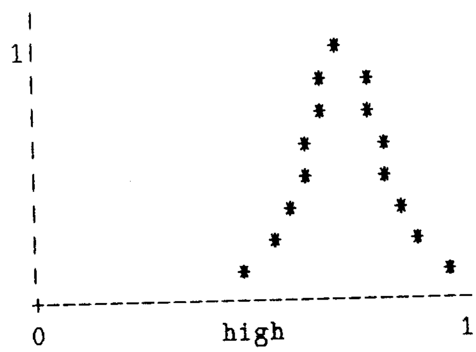

Now for some slightly less general fuzzy frames representing specific types of commodity.

```
F2  Vehicle                              F3  Toy
    IsA        :Commodity                    IsA        :Commodity
    Uses       :(travel,pleasure)            Uses       :(pleasure)
    Keeper     :undefined [text]             Keeper     :child
    Necessity  :high                         Necessity  :low
    Safety     :high                         Safety     :undefined [fuzz]
    Utility    :high                         Utility    :high
    Cost       :high                         Cost       :low
```

The Toy frame will inherit the value 'high' for Safety, since the slot contains 'undefined'; the uniform fuzzy set on the interval scale. The other slots are unaltered. Inheritance occurs based on the immediately superior frame only and then only when the 'child' has an undefined value. In certain applications, such as database ones, it may be preferable to allow inheritance into even those slots which contain defined values. In such a case the inheritance mechanism is modified in such a way that the intersection (minimum) of the fuzzy sets in the parent and child is taken. This corresponds to what we choose to designate the *Fuzzy Closed World Assumption*. That is, if the values assigned to slots represent immutable knowledge about the state of the world and the constraints it imposes, then we would not wish to permit a contradictory reassignment that ignored the influence of the value in a parent. The distinction between the two control strategies is analogous to the one between rules and assertions in fuzzy production systems (see Graham & Jones Chapter 6 [3]). We take this discussion further in the next section.

Two more frames representing general classes of objects must now be discussed. They are:

```
F4   Borrowed-object
     IsA        :Object
     Lender     :undefined [text]
     Keeper     :undefined [text]
     Cost       :minimal
```

```
F5   Dangerous-object
     IsA        :Object
     Safety     :minimal
```

Here we may wish to consider the sad possibility that a borrowed object, such as a book, may pass from *meum* to *tuum* without the transition being too noticable. Thus, in the case of the borrowed magazine in frame F11, only 0.5 of the ownership properties of F4 may be inherited. In particular, the inheritance mechanism attaches a 0.5 certainty factor to the Lender and Keeper values (if they are known) The mechanism for fuzzy slots is that the fuzzy sets (minimal in this case) are truncated at the 0.5 level. Returning to the mainstream of our argument, two new fuzzy sets have been introduced, so we give their definition pictorially, as before.

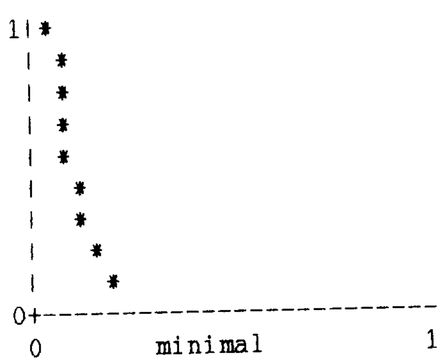

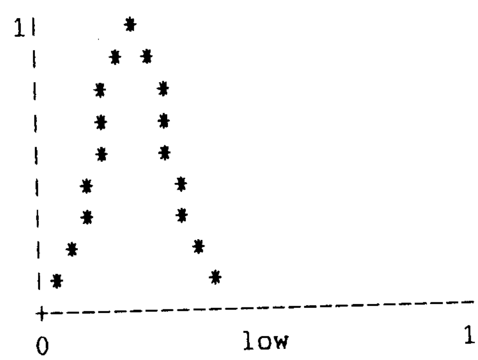

Now we come to the frames describing fairly specific items in the scheme. For example, we have:

```
F6   Dinghy
     IsA        :Vehicle [0.4],
                 Toy [0.6],
                 Dangerous object [0.1]
     Draught    :3 [real]
     Safety     :undefined
     Cost       :undefined
```

```
F7   Hang-glider
     IsA      :Vehicle [0.05],
               Toy [0.7],
               Dangerous object [0.9]

     Safety :undefined
     Cost   :undefined
```

FIGURE 4.2 PARTIAL INHERITANCE NETWORK FOR HANG-GLIDER-765.

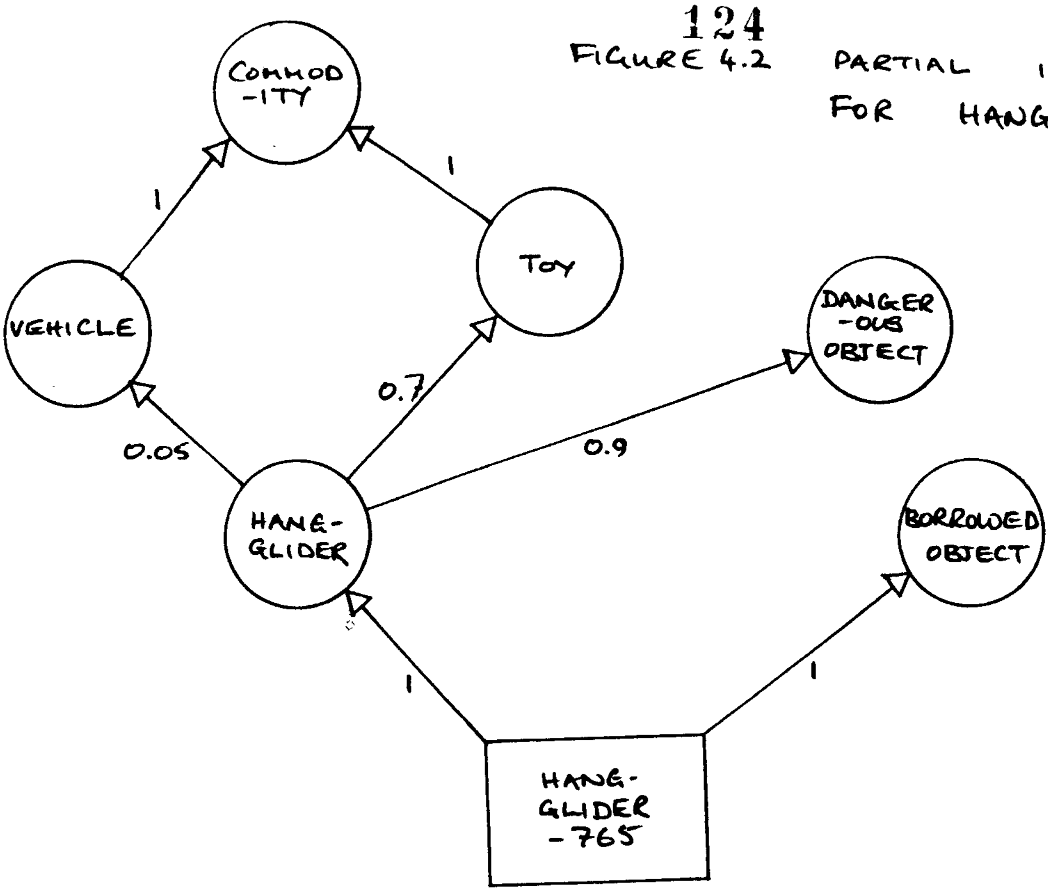

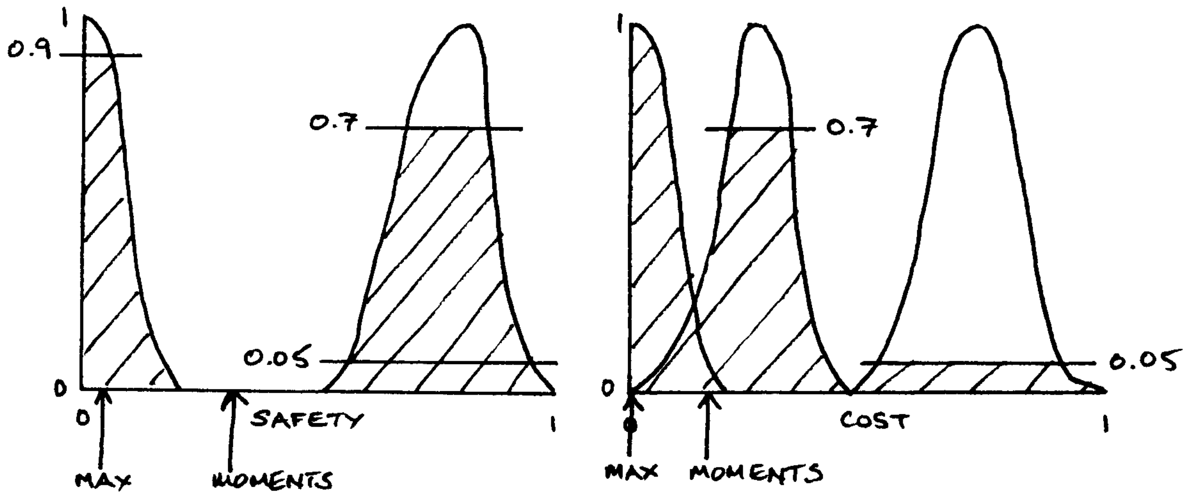

(a) INHERITED FUZZY SETS IN HANG-GLIDER-765

FIGURE 4.3

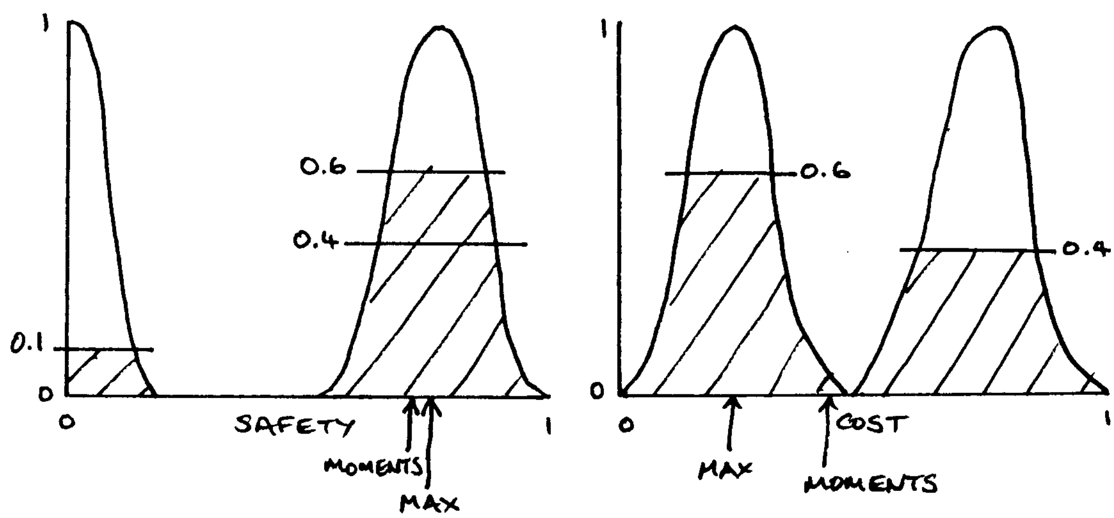

(b) INHERITED FUZZY SETS IN DINGHY-123

We now have to understand how the undefined slots in the lowest level frames (representing individuals) for Dinghy-123 and Hang-glider-765 may be filled. Notice first that we have a non-fuzzy slot for Draught, and multiple inheritance from higher levels. Let us look at the Safety slot of Dinghy first.

Since a dinghy is a vehicle the slot inherits 'high', but as this is only true to the extent 0.4 the fuzzy set is truncated at this level. It also inherits the value 'minimal' from Dangerous-object, but only to degree 0.1. The inheritance path from Commodity via Toy gives the value 'high' in degree 0.6. These fuzzy sets are combined with the union operator as shown in the diagram in Figure 4.3(b). If this were the final result of some reasoning process the resultant fuzzy set would be defuzzified (in this case with the mean-of-maxima operation) to give a truth or possibility value for the term 'safe'. Alternatively linguistic approximation could be applied to return a word corresponding to a normal, convex fuzzy set approximating the returned value. In a different application the moments defuzzification method might be applied. This is a control decision in the same category as the fuzzy closed world assumption, and, we feel, should be left to the discretion of the user or systems designer. The other diagram in Figure 4.3(b) shows the fuzzy set for the dinghy's cost. In the absence of evidence to the contrary, Dinghy-123 inherits both these values.

Cost here is being interpreted as the cost that one might be willing to bear, and thus the cost of a dinghy purchased just for fun ought, normally, to be low. The case of the safety slot of Hang-glider is a little more interesting. The diagrams in Figure 4.3(a) illustrates the text.

Here, the Safety slot inherits the union of the fuzzy set minimal from Dangerous-object [0.9] and high from both vehicle [0.05] and Commodity (via Toy) [0.7]. Applying the operation of union or disjunction to these three fuzzy sets (we of course exclude 'undefined' from this process) to represent the view that IsA attributes are *alternative* viewpoints from which the object may be viewed, we arrive at a resultant fuzzy set. Defuzzification then gives a value close to 0 (or the linguistic approximation 'minimal').

Thus the system is able to deduce correctly that a hang-glider is a very dangerous toy along with the unsurprising conclusion that it doesn't cost much to borrow one in the case of Hang-glider-765.

Clearly, the reason we have adopted the view that a hang-glider is only a vehicle to a small extent is that one usually thinks of a vehicle as a safe-ish means of getting from A to B and, indeed, back again. This is not independant from our assumptions about dangerous objects. This warns of a possible design problem for fuzzy framebases. However, there is a further problem. If, as is quite reasonable, a survey showed that most people actually gave a higher value, say 0.95, to the 'toyness' of a hang-glider, then the result would be quite different under the maximum rule of inference, and quite counterintuitive: Hang-gliders would be highly safe. The apparently counter-intuitive nature of this result - which, incidentally only has noticable, serious consequences under the maximum rule - could be due to the incompleteness of our example semantic model. We suggest a way round this problem in the next section, couched in terms of a theory of normal - or non-redundant - forms for fuzzy frames. Currently, the topic of design criteria for fuzzy semantic models is undergoing research, but remains only partially solved. What is required is analogous to the theory of normal forms in database design theory. Another research topic is the design of query languages for fuzzy frame systems.

To put matters right temporarily, let us now explore an example which does conform to intuition more closely than the one chosen above to explicate the syntax and semantics.

```
Immortal                 Mortal           Man
IsA:Category             IsA:Category     IsA:Mortal
Goodness:high                             Goodness:fair
Intellect:omniscient                      Intellect:average

Apollo                   Lucifer          Socrates
IsA:Man[0.4],            IsA:Immortal     IsA:Man,Immortal[0.2]
    Immortal[0.9]        Goodness:low     Intellect:bright
Goodness:undefined                        Goodness:undefined

Philanthropist           JohnpaulgettyIII
IsA:Man,                 IsA:Philanthropist
    Immortal[0.1]        Goodness:undefined
Goodness:high            Fame:undefined
Fame     :high           Intellect:undefined
```

In this case the inhertance mechanism enables us to infer that John Paul Getty III is a nice chap who'll be remembered for quite a while, because philanthropists are usually famous. Apollo, on the other hand, inherits average intelligence as the epitomy of manliness and omniscience from his godliness. We know from our Homer that Apollo was in fact only wise on occasion, and this is reflected in the returned fuzzy set for his intellect, whose linguistic approximation is something like 'bright' if we use the moments rule. Apollo's goodness is also reduced by his manliness. In the absence of a richer structure or, in other words, more knowledge and information, we can only deduce average intelligence for John Paul Getty III, although he might be remembered by posterity as brighter due to the magnifying effect of a degree of immortality (in the sense of living in memory here) on intelligence. Clearly, this kind of frame base has considerable application to computerised models of commonsense reasoning. Zadeh refers to the type of reasoning implied here as that of 'usuality' [17], while Touretzky (erroneously) calls it 'normative' reasoning (he means 'what is normal') [12].

This example, of course, raises many of the usual questions about inheritance that we find in crisp systems. In addition, we are led to ask what would happen in a more complex framebase. In particular, in certain applications it might be necessary to consider that the inheritance of god-like properties by offspring and offspring of offspring should be subject to attentuation but not to complete exception (e.g. Leda, Europa, etc.). In that case we would want to invoke the fuzzy closed world assumption. A classical example of this assumption may be found in the Book of Common Prayer: '... visit the sins of the fathers upon the children unto the third and fourth generation'. The other question raised here, as compared with the previous example, is the evident comparability of the categories represented by the IsA links. This suggests that well designed framebases should evince this property. We have more to say on the soundness of designs later.

For some reason it is apparent that the moments method of defuzzification is the more appropriate one in the example deductions we have discused here. This is because we were dealing with the usuality of properties which are subject to combination in reaching a 'balanced view', rather than ones which contribute to either/or decision making. There could be problems if we had mixed objectives in our use of the framebase. We would at least have to type the IsA links were the two strategies to be required over the same framebase.

We have thus presented, via a couple of very simple examples, the basic theory of fuzzy frames and explained its logic of inheritance. We now justify our efforts by presenting a more practical example.

Consider the problem faced in allocating a marketing budget among the various activities which could lead to higher sales of a product. It is part of the folk-lore of marketing that different types of product will benefit from different allocations. Suppose that the methods at our disposal are:

Advertising	Promotion	Sales Training	Packaging	Direct Mail

Now, suppose that we compare the allocation ideal for breakfast cereals with that for package software. Advertising, promotions and packaging are clearly all useful, but there isn't much point in direct mailing cereal consumers and the degree of training required by the salesmen isn't usually considered to be high. Thus, in the existing situation we might well represent the allocation of resources in the following matrix of percentages.

	Adverts	Promotion	Sales Training	Packaging	Direct Mail
Breakfast Cereal	30	50	5	15	0
Package software	10	5	15	20	50

Warning to marketing executives: These figures are not meant to represent a truly effective strategy!

In the case where knowledge is expressed inexactly we can readily replace these numbers by fuzzy numbers as follows.

	Adverts	Promotion	Sales Training	Packaging	Direct Mail
Breakfast Cereal	about a third	about half	hardly any	a little	none
Package software	less than a little	hardly any	a little	about a fifth	about half

What we have here is two fuzzy relations, for Breakfast-cereal and Package-software. Presumably they can be regarded as part of a larger database of products; in fact product *classes*. Viewing them as classes prompts us to write them down as fuzzy frames and ask about inheritance through IsA links. To see that inheritance (of the properties under consideration) may be partial, consider the fuzzy frame representing the class of commodities called Vending-machine. A vending machine may be viewed as office furniture, catering equipment or even as packaged food depending on the marketing approach taken. Fuzzy Frames give a natural way to build a description of this problem and suggest an implementation which is able to combine evidence and reason with exceptions.

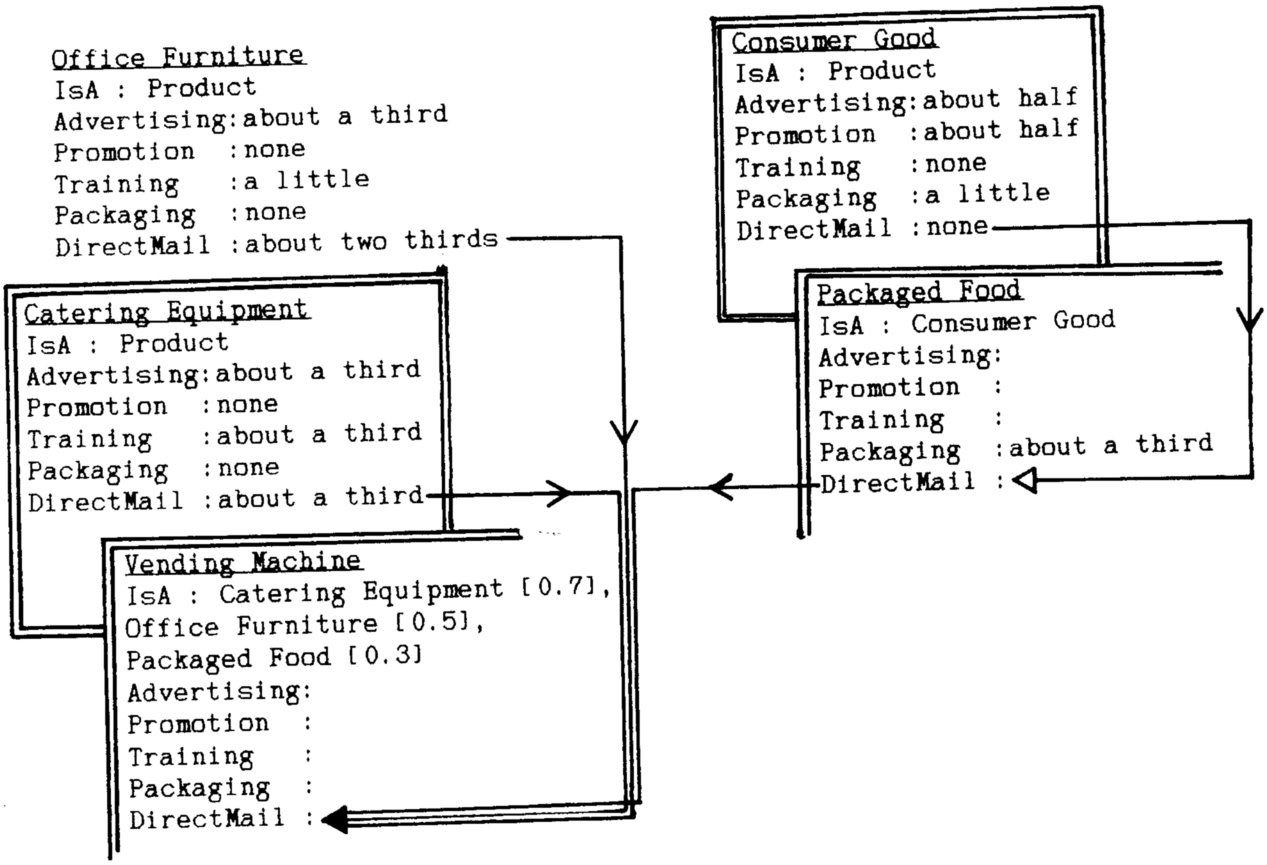

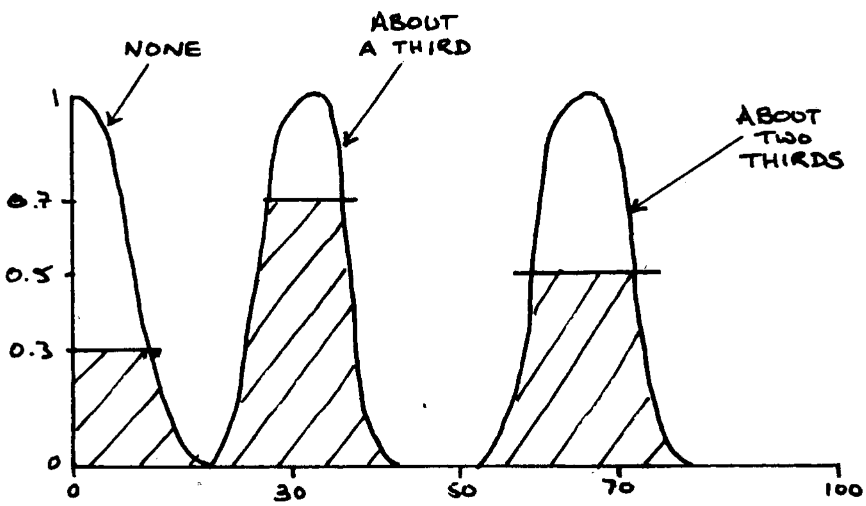

The resultant fuzzy set for Direct Mail

FIGURE 4.4 How to sell vending machines

The figure shows how the combined partial inheritance of fuzzy (i.e. linguistically expressed) allocations from general classes of products may be used to infer an allocation for specific types of product. All slots are fuzzy and undefined slots are left blank.

Another possible practical application is to 'dotted line' relationships in organisations, where the responsibilities of certain specialists to technically related parts of an organisation may override or mingle with those of the formal reporting hierarchy. One application of the framebase is to assist with the decision as to whom should be consulted when the specialist is asked to work overseas for a year. Another concerns the construction of formal models of the sort of loose-tight properties of organisations referred to by Peters and Waterman [9].

It is our belief that there are a tremendous number of opportunities for the application of fuzzy frames. Of course, it may be argued that these applications can be addressed by other technologies, but none that we can think of offer simultaneously the advantages of naturality of expression in a unified representational formalism to the extent that fuzzy frames do.

References are appended to the Part 2 of the paper, to be published in BUSEFAL n°32 (Oct. 1987)