Inheritance in Hierarchical Relational Structures

Derek P. LONG
Dept. of Computer Science,
University College London,
Gower Street,
London,
ENGLAND.

Roberto GARIGLIANO
School of Engineering and
Applied Science,
University of Durham,
Durham,
ENGLAND.

Abstract

A brief survey is conducted of the inheritance principle - the conveyance of properties between components within a hierarchical relational structure. The standard form of inheritance is considered, using the subset (is-a) relation and highlighted as an example of downward inheritance. Downward inheritance is extended to specialisation of actions, and cases are presented in which the rule fails.

An alternative and less well-known form of inheritance is introduced - upward inheritance. Several examples in which upward inheritance is valid and others in which it is not valid are given, in a treatment highlighting the analogy with downward inheritance. The validity of the rule, in those case in which it operates, is underlined, to distinguish it from induction.

A brief account is given of the search for the underlying reasons for the validity of inheritance rules and these are then given. The solution turns out to be due to a hidden or implicit quantifier within the relations that are used. The semantical nature of the problem and of its solution are stressed, emphasising the impossibility of a purely syntactic analysis and solution to the problem. Various points of interest arising from the analysis are listed and discussed.

1. Introduction

The use of structures for representing knowledge has long been acknowledged as a vital tool in AI. A consequence of the use of structures is the identification of certain kinds of hierarchically organised sub-structures. These have many useful purposes, but one of particular significance is that of avoiding repetition, by storing information at one place in the hierarchy, and then using the hierarchical structure to infer that the property is inherited by many other components of the structure. In addition to the savings in storage that this inheritance over hierarchies can offer, there is the possibility of using the hierarchy to infer new information, so that the process becomes not just a means of saving space, but also of generating new knowledge.

The form of inheritance most frequently studied is what we call "downward inheritance", because it exploits the passage from the general to the particular. There has been a considerable amount of work done about inheritance between objects via total inclusion relations, for the very good reason that this form is always valid /Findler 1979/. This form is sometimes confused with downward inheritance betwen sets, which, as we briefly discuss below, is much more problematic.

We will concentrate on total hierarchical relations, between objects, classes and actions. The inheritance techniques that we consider are thus completely valid. The use of partial hierarchical relations introduces elements of plausibility which are beyond the present context. Touretzky provides a foundational analysis of partial hierarchies /Touretzky 1986/, while their use as the basic structure for semantical analogy is described in Garigliano and Long /Garigliano and Long 1988/.

The concept that other relations, beyond set membership and inclusion, may have inferential properties has been investigated by Schank /Schank 1977/ and especially Wilensky /Wilensky 1980/. Their inferential rules, however, tend to capture specific aspects of real world interaction: it would thus be difficult to consider them on the same plane of generality and abstractness as the standard hierarchical inheritance rules. Upward inheritance, on the other hand, is formally on the same plane as the more common downward form. Our analysis shows that the explanation for this powerful tool resides in one semantical aspect of a large class of very common relations.

2. Downward inheritance

A typical example of the use of downward inheritance is the argument cats are vertebrates, vertebrates have back-bones, so cats have back-bones. We have also identified other hierarchical relations for which downward inheritance is valid: embroidering is a special form of sewing (or, embroidering "specialises" sewing), sewing requires skill, so embroidering requires skill.

The most common use of inheritance - so common that it is often the only inheritance rule discussed - is that which can be expressed abstractly as: every element in A has property P, B is a subset of A, so all elements of B have the property P. This is the rule exemplified by the cats and vertebrates above.

A third kind of inheritance hierarchy is based on properties of sets, rather than of individuals within the sets. For example, I can count the set of first-division footballers, Liverpool United is a subset of first division footballers, so I can count the set of Liverpool United players. It is clear that this is not the same as the cats and vertebrates example, since the claim is about the set as a whole, rather than about the individual members.

It is important to observe that the downward inheritance rule is not universally valid. To see this, consider the example: the Tory government was elected by the British people, the Scots are a subset of the British, so the Tory government was elected by the Scots. It is clear that this example leads to a false conclusion. We will not attempt to explain all the conditions under which a relation can be expected to have the downward inheritance rule for classes in this work, but we do note that one important part of such conditions can be the "homogeneity" of the class in the first relation with respect to the given property and the related entity /Garigliano and Long 1988/. By homogeneity we refer to a measure of the evenness of distribution of a property within a set.

The downward inheritance rule for actions is also dependent on certain conditions being satisfied by the first relation, as can be seen by the example: rolling pins are used in cooking, frying specialises cooking, so rolling pins are used in frying. This is plainly false, so we may infer that there is some condition which "specialises" satisfies, but is violated by "are used in". The solution to this particular example is actually found in the discussion below. There are other examples for which the solution is not quite so readily identified, but it is possible that the solution is analogous to that for classes, in adopting some measurement of homogeneity.

A more dramatic example in which downward inheritance fails is the following: cows eat plants, cacti are a subset of plants, so cows eat cacti. The inference is certainly false, but this example is of considerable interest since it appears to follow the same pattern as the cats and vertebrates example. If we follow the direction of Schank and Wilensky, adopting a set of primitives from natural language as the basis of our knowledge representation, we cannot, it appears, identify the inheritance properties of those primitives by a simple syntactic check. This follows from the observation that the syntactic pattern of the cats and vertebrates example was essentially identical to that of the cows and plants, yet the inheritance rule is valid only in the first case.

3. Upward inheritance

We can intuitively understand upward inheritance as a form of inheritance that goes from the specific to the general. An example of such inheritance is the argument: A is smaller than B, B is a subset of C, so A is smaller than C.

Further examples are: a camera creates pictures, pictures are representations, hence a camera creates representations. A pan is for cooking, cooking is a specialization of processing food, hence a pan is for processing food.

Of course, the higher we go up the hierarchy, the less useful the information derived may appear: for example, if we substitute Doing Something for Processing Food, the above inference is still valid, but not very useful. It is important to note, however, that when the relation is an upward one, then the deriving inference is valid, not simply plausible. There is no possibility of this inference being some kind of induction: the explanation for it is to be found elsewhere.

As we mentioned before, just as not all relations allow downward inheritance, so too not all relations allow upward inheritance. Here are some examples when the inference fails:

John is allergic to cats, cats are vertebrates, hence John is allergic to vertebrates.

I can count the size of a football team, a football team is a subset of the world population, hence I can count the size of the world population.

Cats avoid swimming, swimming specifies moving, hence cats avoid moving.

All these examples emphasize the impossibility of using a syntactic check to decide which relations offer a hierarchical inheritance, or, if they offer such an inheritance, which direction it is valid for.

The issue turns around the particular relations used; this clearly calls for an analysis of the underlying structure of these relations.

4. When inheritance is valid

We must now attempt to identify what property of a relation it is that enables it to be used for upward or downward inheritance. First let us explore upward inheritance. There is a strong clue available to us in our search in the following example:

if x is a member of A and A is a subset of B then x is a member of B.

Here we see a set theoretic property which actually obeys the upward inheritance rule. This is a very important example - it is not hard to see why it works. The reason that the property is inherited is that although the relation "is a member" relates an object to a class, at the same time it singles out a very specific part of that class. This part of the class must be carried through to any new class containing the original one. Thus, the significant feature of the relation, "is a member", is that it carries an implicit act of restricting the part of the class to which it refers.

Consider a further example: John owns a cat, cats are animals, so John owns an animal. This is an instance of an upward inheritance rule that works, and the relation is "owns a". If we examine this relation we find that it has the analogous property for objects that "is a member" has for classes -it implicitly restricts that object to which it refers out of all the objects. John does not own all cats, but only a single cat, in the same way that not every element in A is x, but only one. Again, a frying-pan is for frying, frying specialises cooking, so a frying-pan is for cooking is an example of an upward inheritance over actions, using "is for". Here, too, we find that what is being said implicitly is that there is a conceivable instance of frying for which one could use a frying-pan. It does not mean that frying must always be done with a frying-pan - we could use a deep-fat fryer or a wok for example. Thus, when we extend frying to cooking we are actually referring to the same instance of frying in which a frying-pan could be used, and using the fact that this is also an instance of cooking because frying specialises cooking.

So, we have three examples of relations which have the upward inheritance property and seem to have an analogous property - for objects this property is that the relation specifies a particular object of all the possible objects in a class. For classes the property is that the relation highlights a subset of the class (in the "is a member" example this subset is $\{x\}$) and for actions the property is that the relation specifies a particular instance (though possibly hypothetical) of an action in which a certain condition holds.

After a more formal analysis of these examples (considered in detail in /Long and Garigliano 1988/) we find that for a transitive and reflexive hierarchical relation, such as "subset" or "specialises", a second relation has the upward inheritance property if it contains an implicit existential quantifier - that is, if the relation implies a property of a limited part of the class of objects to which it refers, in some sense. This sense has been made more formal in /Long and Garigliano 1988/. Furthermore, this condition is shown to be both necessary and sufficient.

After seeing the pattern for upward inheritance, it is not difficult to find the pattern for downward inheritance. In this case, the inheritance works for reflexive and transitive hierarchies if and only if the inherited property contains an implicit universal quantifier, in a sense which has been made formal. For example, when we say "vertebrates have backbones", we actually mean *all* vertebrates have a backbone.

The analysis indicates that there is a very simple rule that allows us to convert any relation into one for which the upward inheritance rule operates and, conversely, for checking if an inheritance rule will work for a given relation.

For example, suppose we take the relation used above, John is allergic to cats. Now, if we modify "is allergic to" according to the analysis, we may build the new relation, "is allergic to some". Thus, John is allergic to some cats, cats are vertebrates, so John is allergic to some vertebrates. It is quite clear that this new relation has the property of upward inheritance, unlike the original relation.

As an example of the process of checking whether upward inheritance will apply to a given relation, consider the following. We have the relation "is smaller than", applied between sets, as used in an example above. In order to confirm that upward inheritance can be applied with this relation we must find the construction that yields "is smaller than", given some starting relation. Consider the relation "one up on", defined by:

A is one up on B iff $|A| + 1 = |B|$

Now, "is smaller than" can be defined by:

A is smaller than B iff there is some subset of B which is one up on A.

Thus, "is smaller than" has upward inheritance with respect to the subset hierarchy.

In this case, the relation is a mathematical one, so proving that the inheritance is valid is not difficult, even without the insight we have gained. However, for more general relations, borrowed directly from natural language (in a Schank or Wilensky style for primitives) the insight provides the only tool we have to formally prove an inheritance property.

There is a way of modifying a relation with an implicit existential quantifier so that it becomes a relation with an implicit universal quantifier, in addition to simply adding the universal quantifier on top of the relation - we can restrict the domain of the relation by inserting "the" in front of each group of objects in the domain. For example:

John eats fruit, apples are among fruit so John eats apples

is not valid. Once we modify the domain, however, we obtain:

John eats the fruit, the apples are among the fruit so John eats the apples.

Now, the use of "the" indicates a specification of an exact group of fruit, F say, and of apples, say A, so "the apples are among the fruit" now means:

every member of the set of apples, A, is in the set of fruit, F.

What makes the inference correct is the peculiarity of the use of "the" in conjunction with the relation. It means that the relation is true of every element in the group - so John eats every piece of fruit in F.

For several hierarchies we have studied there are perfectly meaningful and uncontrived relations which have neither the upward or downward inheritance rule. For example, consider the following:

the Tories were elected for by forty per-cent of British, Scots are a subset of British, so the Tories were elected by forty per-cent of Scots.

This is false, so the "elected by forty per-cent of" relation does not have downward inheritance over subsets.

The Tories were elected by forty per-cent of British, British are a subset of Europeans, so the Tories were elected by forty per-cent of Europeans.

Again, this is false.

There are other hierarchies, such as "has part", which holds between an object and each of its parts, for which very few relations have any inheritance properties. In this case it is because there are very few things that can be said of an object which must be true of all its parts, or things which can be said of a part of an object which must be true of the object as well. Positional relations seem to be the only useful relations - all parts of an object must be where the object is, though an object does not have to be where one of its parts is. For example, a tyre must be there where there is a car, but a car does not have to be there where there is a tyre.

Finally, we must highlight a consequence of these findings, which is that those relations for which upward inheritance holds occur in sentential forms with precisely the same type and structure as those for which the rule fails. The only distinguishing feature that we have discovered and that accounts for every example we have considered is the implicit quantifier within the semantic definition of the relation. This leads us to conclude that the power of this rule of inheritance can only be available following a semantical analysis of the relations involved and will not yield to any kind of syntactical or grammatical analysis.

5. Conclusions

In this paper we have described and analysed a technique, the use of inheritance, for transferring properties between components within total hierarchical structures. It is a powerful technique, because it allows valid inferences, as well as being very natural, since it models sentences that could occur in everyday use of language. At the same time, it is widely applicable, using relations which range from the common and simple such as "drink" or "wear", through the fundamental, primitive, relations like "is for" or "causes", to the mathematical relations such as "is a member of" or "smaller than".

We have, in our language, some relations which carry an implicit existential quantifier and others which carry an implicit universal quantifier, but we can identify no particular reason that any of these relations should have developed in that way. The important thing is that we are now able to build relations for which inheritance will work, and also recognise those relations, amongst those we already have, for which the rule applies.

The final point we wish to stress is that this analysis strengthens the argument in favour of the semantical approach to inferences. As we have pointed out before, the distinction between the relations which allow upward inheritance, and those which do not, cannot be identified by purely syntactical means, even if that distinction appears evident to the human ear, and easily recognisable by a semantical analysis.

6. References

Charniak E. and McDermott D. (1985), Introduction to Artificial Intelligence, Addison Wesley.

Findler N.V. (ed.) (1979), Associative Networks: Representation and Use of knowledge by Computers, Academic Press.

Garigliano R. and Long D.P. (1988), A Formal Model through Homogeneity Theory of Adaptive Reasoning, PRG Monograph, Oxford, forthcoming.

Long D.P. and Garigliano R. (1986), Reasoning by Analogy: a Formal Model, in Tenth Annual Conference on Microprocessor Applications, Strathclyde.

Long D.P. and Garigliano R. (1988), Inheritance Hierarchies.

Schank R.C. and Abelson R.P. (1977), Scripts, Plans, Goals, and Understanding, Lawrence Erlbaum Press.

Touretzky D. (1986), The Mathematics of Inheritance Systems, Kaufmann.

Wilensky R. (1980), Understanding Goal-Based Stories, Garland Publishing.