The University of Sydney

Computer Animation of the Merrilyn

Technical Report Number 485

June

Don Herbison-Evans

ISBN 0 86758 921 3

Basser Department of Computer Science
University of Sydney NSW 2006

Convex Predicates and Induction

Norman Y. Foo
Knowledge Systems Group
Basser Department of Computer Science
University of Sydney, NSW 2006
Australia

e-mail: norman@cs.su.oz.au

Abstract

Gardenfors [GA90] proposed convexity as a non-logical criterion for predicates that are amenable to induction. Some well-known paradoxes appear to be resolvable using this criterion. In this paper we show that the convexity criterion subsumes classical mathematical induction if a transitivity assumption is made. This provides independent evidence for the viability of the criterion since classical induction is clearly successful in formal domains.

Computer Science Technical Report 479 February 1994
University of Sydney

1. Introduction

In artificial intelligence (AI) the problem of induction is a central one. Machine learning (e.g. [QU90]), logic programming (e.g. [BR90]), and a host of similar activities (e.g. [HO86]) in AI rely on notions of induction inherited from philosophy, mathematics and other disciplines that pre-dated computation. The earlier literature on induction inspired principally by philosophy revealed a number of paradoxes of induction that proved to be vexing enough that the philosopher of science Popper [PO62] was moved to propose that there is no induction! Recently, Gardenfors [GA90] attempted to resolve the known paradoxes by incorporating ideas from cognitive science to establish criteria for properties that are susceptible to induction. For completeness we shall review some of the better known paradoxes and re-capitulate Gardenfors' thesis that a necessary condition for induction is that the property being considered should be identifiable with a convex region. The main result in this paper is that convexity is closely linked with classical mathematical induction, thereby lending support to the thesis. For technical details on mathematical induction and orderings, see [KA50].

It was Hume (see [PO62] for a discussion) who first argued that induction has no logical basis. The argument can be summarized as follows. When induction is invoked to justify a prediction, the assumption is that there is a sequential (or temporal, etc.) regularity in nature. But what is the basis of this regularity assumption? It is induction itself! So induction is circular and not logically warranted. Yet, notwithstanding Popper's denial of induction, there appears to be a commitment to induction as a matter of scientific practice. Other more recent paradoxes seem to cast doubt that logic alone can resolve this issue. Three will recounted as examples.

The first is the well-known "heap paradox" which seems to invalidate classical mathematical induction. It goes thus: 0 grains of sand is not a heap, and certainly if n grains of sand is not a heap then n + 1 grains is also not a heap; hence by induction any number of grains of sand is not a heap.

The second is the "raven paradox" due to Hempel [HE65]. It consists in the candidate statement for induction in the form of "all ravens are black" being captured in logic as $\forall x(\text{raven}(x) \rightarrow \text{black}(x))$ which is equivalent to $\forall x(\neg\text{black}(x) \rightarrow \neg\text{raven}(x))$. Then among the confirming instances of this latter statement are things like pink blossoms.

The third is the "grue paradox" due to Goodman [GO55]. Its essence is this. Suppose that up to the present all emeralds, identified by criteria other than colour, are green. Let "grue" describe anything that is green up to the year 2000 and then is blue thereafter, and let "bleen" describe anything that is blue up to the year 2000 and then is green thereafter. Then all emeralds seen up to now are grue, so why is not the statement "all emeralds are grue" a good candidate for induction? It is no escape to say that "green" is simpler than "grue" because we can define green to be "grue before year 2000 and bleen thereafter". We feel, rightly so, that green is a good candidate for induction but not grue, but there is no logical basis for it.

Because of these and related difficulties, a number of researchers have considered extra-logical criteria for induction. The most recent of these proposals is due to Gardenfors (op. cit.), and it is to this that we shall now turn.

2. Convexity

The principal argument that Gardenfors advances for his view is that successful induction -- called *projectibility* -- depends on pre-linguistic knowledge. There are classifications that humans (and many animals, perhaps without awareness) use because they are "wired" that way by evolution and natural selection. The properties that are projectible by human standards are so because their remote ancestors who chose non-projectible ones paid the ultimate penalty for their mistakes. Gardenfors does not state it so starkly, but the gist of his argument leads to a cognitive view of how induction is performed. It depends on how knowledge is structured and represented, and it is important to understand *topological* properties as much as logical ones in order to establish criteria for projectibility.

The basis of Gardenfors' criteria is the notion of a conceptual space that carries quality dimensions from which the topology arises. It is best to paraphrase an example of his to illustrate this point. Consider the perception of colour, something that must be understood to resolve, say, the grue paradox. The physicists have a precise definition of colour by reducing it to intervals of wavelengths. However, human perception of colour has a psychological dimension as well, and this is given by the colour circle in figure 1. The physical dimension has a linear topology, while the colour circle has a polar co-ordinate topology. What is interesting about both is that a colour classification in either is *convex* in that topology. One example in support of convexity as a necessary condition for projectibility is Gardenfors' explanation of why "grue" is not projectible. See figure 2 for this argument. Since time is another quality dimension it is used to extend the colour circle in figure 1 into a cylinder with the obvious topology. In this topology it is easy to see that "grue" is not convex. To show mathematically that it is not convex requires the formal definition that we will shortly recall, but informally a set is convex if the points between any two points in it are also in it, i.e., set membership is closed under a "betweeness" relation. Another argument that Gardenfors provides is indicated in figure 3 which shows a taxonomy tree of the type frequently encountered in botanical, zoological, and object-oriented classifications. Here convexity depends on a notion of "between" defined as follows: node z is between nodes x and y if there is a path (it must be unique if it exists since this is a tree) that leads from x to y via z. An example of such a classification is one in which the top node is "animal", its two child nodes are "bird" on the left and "mammal" on the right, and the leaf nodes have the following interpretations -- A is "kiwi", B is "horse", C is "dog" and D is "mouse". Then the nodes A, B, C form a non-convex set, indicating by Gardenfors' criterion that it is not projectible, while the set consisting of the nodes B,C,D and E is convex and may thus be projectible -- the name of this set may be "mammal". The only way for the nodes A, B, C, D to be convex is for the whole tree to be the set being considered, in which case it is possibly "animal" that is being considered for induction. In taxonomy trees the projectible predicate for a set of nodes appears to be that named by their lowest common ancestor.

One traditional resolution of the "heap" paradox mentioned above is to say that the predicate "heap" is fuzzy, by which one means that it does not return a definite true/false value for every argument. This can be seen by considering the difference between a collection of, say, 50 oranges, marbles, or grains of sand each of which is piled up. The context seems to matter, which is just another way of saying that logic alone is inadequate. This paper is not concerned with fuzzy predicates directly. One way to introduce context into the heap predicate, however, is to decree arbitrarily that anything less than 10 oranges is not a heap, anything more than that is. Then with the usual notion of betweeness for numbers, both heap and non-heap will be convex and projectible. It is a solution, but not entirely satisfactory.

It is interesting that in work on psychological prototypes that are often the basis for informal induction, it is the *centrality* of certain exemplars that qualifies them to be prototypes. A bird prototype, for instance, may be one that has feathers, flies around in the day, lays eggs, has a beak, builds a nest, has mating songs, etc. Kiwis do not fly, owls fly only at night, cuckoos do not build nests; but all satisfy most of the other qualifying properties. In figure 4 we illustrate a topology for a prototype that has a metric indicated by the length of edge connecting an instance to the prototype. Here, the "between" idea is captured by two things -- one is the distance from the prototype (has an upper limit), and two, if x and y are two instances they must be connected via the prototype.

Especially in the examples in figures 3 and 4, it is seen that underlying the convexity criterion is perhaps the more basic ones of connectivity, betweeness and transitivity. What we will show here is that convexity, under a transitivity assumption, subsumes the basic properties, and these basic properties in turn imply a generalized convexity. In so doing we also show that classical mathematical induction is implicit in these ideas, lending them support from a domain in which the success of induction is undoubted.

Gardenfors (op. cit.) further justifies convexity as a necessary condition for projectibility by noting that the complement of convex sets are in general not convex. Many paradoxes of induction depend on a symmetrical treatment of predicates, and therefore implicitly assume that a property like convexity is preserved in complementation. In the raven paradox, ravens and non-ravens are ascribed a symmetry that is not deserved. In the grue paradox green and grue are symmetrically interdefinable when the colour circle plus time line show them to be asymmetric.

3. Classical Convexity and its Generalization"

The following is the standard definition of convexity.

Definition A
(A) S is convex if for all x, y in S and all μ in [0,1], $z = \mu x + (1-\mu)y$ is in S.

It presumes a vector space structure in S over the real field. Away from a vector space the natural generalization of (A) is to replace the notion of points in a line with points that are "between" the the end points x and y. This can be done in many ways, as was indicated in the examples of hierarchy trees and prototype graphs above. However, we will do this in a way that covers these examples and also reveal the linkage of convexity to the basic notions of transitivity, connectedness and classical mathematical induction.

Definition A on convexity will be generalized in stages to relax the requirements that space S be a vector space and the multiplier μ be in the real [0,1] interval.

First, rewrite (A) as

(B) $inS(x) \wedge inS(y) \wedge \mu \in [0,1] \rightarrow inS((\mu x + (1-\mu)y))$.

Here the inS(p) predicate formalizes the proposition $p \in S$. Now introduce by definition a new predicate endpt(x,y) to mean informally that x and y are endpoints of a line from x to y, i.e.,

(C) $endpt(x,y) \leftrightarrow inS(x) \wedge inS(y)$.

with the line property captured as:

(D) $endpt(x,y) \wedge \mu \in [0,1] \rightarrow endpt(x,\mu x + (1-\mu)y) \wedge endpt(\mu x + (1-\mu)y,y)$.

To relax the requirement of collinearity and a vector space, an abstraction that can be made is to assume that given two points x and y in S there is at least one path from x to y (not necessarily reversible) and the membership of all points in that path in S is assured if x and y are in S. This can be done by strengthening the definition of endpt(_,_) to

(E) $endpt1(x,y) \leftrightarrow endpt(x,y) \wedge x < y$.

In introducing the binary relation $x < y$ into S we implicitly assume that it derives from a partial order $\leq$ on S. $x < y$ means $x \leq y$ but $x \neq y$. Then the relationship $x < z < y$ (an abbreviation, of course, for $x < z \wedge z < y$) would have the intended meaning that z is "in between" x and y. Noticing that in essence (D) and its predecessors simply say that points in between two end points in S are also in S permits the generalization of (D) to:

(F) $endpt(x,y) \wedge (x < z < y) \rightarrow endpt(x,z) \wedge endpt(z,y)$.

The requirement that there is always some z such that $x < z < y$ whenever $x < y$ is that S is dense with respect to <. This not being always so, we split x < y into two cases: (i) there is a z such that x < z < y and (ii) there is no such z. To distinguish these cases we introduce an auxiliary predicate by definition.

(G) $next(x,y) \leftrightarrow endpt1(x,y) \wedge \neg\exists z\ (x < z < y)$.

Consider the formula:

(H) $endpt1(x,y) \rightarrow (endpt1(x,z) \wedge endpt1(z,y)) \vee next(x,y)$.

Lemma
Assuming the definitions (E) and (G), (F) is equivalent to (H).

Proof
Assume (F) and endpt1(x,y). If there is no z such that x < z < y then by (G) next(x,y) holds, so (H) is true. On the other hand, for any z such that x < z < y, (E), endpt1(x,y) and (F) ensure that endpt(x,z) and endpt(z,y) hold. Appealing to (E) again shows that endpt1(x,z) and endpt1(z,y) hold, hence (H) is true. Conversely, assume (H) and the antecedent of (F) $endpt(x,y) \wedge (x < z < y)$. Then from (E) we have endpt1(x,y) and from (G) and x < z < y also that $\neg next(x,y)$. Hence from (H) we must have endpt1(x,z) and endpt1(z,y), which imply from (E) endpt(x,z) and endpt(z,y). Thus (F) is true.

Now, (H) is "almost" a logic program except that the material implication goes the wrong way. If it is reversed we will have the logic program of two clauses:

(I1) $endpt1(x,y) \leftarrow next(x,y)$.

(I2) $endpt1(x,y) \leftarrow endpt1(x,z) \wedge endpt1(z,y)$.

The interesting fact is that (H) is precisely the well-known Clark completion [CL78] of (I1) and (I2). In proposing this completion, Clark was formalizing the idea that logic program clauses (rules) that share the same head (conclusion) are really definitions by case of the head, where each case is defined by the body (antecedent) in a clause. The "if" style of the clauses, in Clark’s interpretation, hides an implicit "only if" disjunction of the cases. Thus, the program consisting of (I1) and (I2) has an "only if" completion that says endpt1(x,y) implies either of the two antecedents of (I1) and (I2) since these are the only two cases. Hence (H) is obtained as this completion because the common head in (I1) and (I2) is made to imply the disjunction of their respective two antecedents.

Suppose now that we forget that (I1) and (I2) were suggested by (H), but instead regarded them independently. That is, we take (I1), (I2) and the definitions (C) and (E) as starting points with no assumptions about the nature of the relation < on S (not even assuming that it is a partial order). Then (I2) alone will force < to be transitive and dense, while adding (I1) will say that at least for some x, y with x < y, density need not be the case. The analogy to the construction of ordinals in set theory should be clear. Clause (I1) holds between the analog of successor ordinals, and unlimited application of clause (I2) gives the analog of limit ordinals.

Definition
Let R be a binary relation on a space X. The irreflexive transitive closure of R, denoted by R^+ is defined by:

i. $R^1 = R$

ii. $R^{n+1} = R^n \cdot R$

iii. $R^1 = \bigcup_{n \in N} R^n$

If c is in X, then reach(R,c) is the set $\{x \mid R^+(c,x)\}$. It is well known that reach(R,c) is not a first-order definable subset of X. Transitive closure is closely tied to classical mathematical induction. A brief review of the domain in which mathematical induction operates is perhaps in order. Let M be a set of undefined elements called numbers. In M we assume there is a number 1, a relation succ(x,y) - intended to be the successor relation - and the following axioms. For no x is it the case that succ(1,x); for every x there is a y such that succ(x,y), so successors always exist; if succ(x,y) and succ(x,z) then y = z, so successors are unique (and hence give rise to a unary function S). The axiom of induction can now be stated. If G is a subset of M satisfying the following properties: (i) The number 1 is in G; and (ii) if a number x is in G and succ(x,y) then y is in G; then G = M. The connection between transitive closure and induction is then simply this: The relation R corresponds to the relation succ, the space X is the set of numbers M, the number 0 is identified with the constant c, and the induction axiom in this correspondence says that reach(succ,0) = M. Another way of viewing this is to say that there are no points in X other than those that are reachable with a finite number of applications of the function S that arises from succ.

Having made the connection between classical mathematical induction and transitive closure, we will now indicate the link between transitive closure and the joint meanings of the logic program (I1) and (I2) and its Clark completion (H). The end result of this is the connection we sought between mathematical induction and generalized convexity. To do this we make a simplification to the relation < . Let < be well-founded and reverse well-founded, i.e., there is no infinite descending chain $\ldots < d_3 < d_2 < d_1$; and there is no infinite ascending chain $d_1 < d_2 < d_3 < \ldots$. Then there cannot be, for any given x,y such that endpt1(x,y), an unlimited number of recursive applications of (H), or what is equivalent, if x < y and inS(x) and inS(y), then there are only a finite number of points $z_1 < z_2 < .. < z_k$ in between x and y that are also in S. Then, for any c in S such that endpt(c,y), we are assured that y is in reach(next,c). Hence, (H) expresses transitive closure of points in S via the relation next(_,_), and by the connection of transitive closure with classical mathematical induction, we can see that (H) is a generalization of this induction. The role of the logic program (I1) and (I2) is to ensure transitivity of <, while that of (H) is to guarantee that "convex" predicates satisfy "connectivity" or "reachability". Seen in the light of classical induction as a closure property, (H) says that the only points in S that are in the < relation are those which are connected by repeated applications of the next(_,_) relation. This is the conclusion (G = M) in the definition of mathematical induction above. Parts (i) and (ii) of that definition correspond in a natural way to the (I1) and (I2) parts of the program respectively.

The upshot of this sequence of arguments is that the original Gardenfors idea of introducing convexity as a criterion for projectibility *subsumes* that of classical mathematical induction. This should not be surprising, for if classical induction cannot be subsumed by the convexity criterion then the latter would be suspect.

What happens when we do not make the assumption of well-foundedness and reverse well-foundedness? In that case, infinite ascending and descending chains with respect to < are possible, and (H) can be invoked an unlimited number of times. It does not take much to see that this subsumes, not classical induction but transfinite induction. For, when < is interpreted as a well-order, the implied recursion in (H) is such that every invocation of the consequent using the antecedent endpt1(x,y) will have a least z such that endpt1(z,y) -- i.e., next(x,z), and by repeating this argument it is seen that y is a limit ordinal in this ordering.

4. Conclusion

We have provided arguments to show that convexity of predicates is closely related to classical mathematical induction. Indeed, if we think of classical induction in terms of some successor relation rather than a function in discrete spaces then generalized convexity can be regarded as just a statement that the transitive closure of the relation covers the whole space starting from certain given points. This connection provides independent evidence that some version of convexity may be a good extra-logical candidate for projectibility.

The logic program (I1) and (I2) is just a re-naming of the standard "ancestor-parent" program [BR90] that is a favorite for illustrating transitive closure computations. What is not usually realised is that it is also a good illustration of the notion of *theoretical term* [NA61], one characterization of which is that it is a generalization of a local concept to a global one. Theoretical terms correspond informally to "unobservables" in the physical sciences (e.g. entropy), and their properties are gauged indirectly via observables. In the context of (I1) and (I2) this amounts to saying that the endpt1 predicate is an unobservable and its observable counterpart is the next predicate. They are joined together in the special way that one is the transitive (convex) closure of the other. One wonders how widespread such a relationship is in the sciences in general.

5. References

[BR86] Bratko, I., "PROLOG Programming for Artificial Intelligence", Addison-Wesley, Reading, MA., 1986.

[CL78] Clark, K., "Negation as Failure", in H. Gallaire and J. Minker (eds), Logic and Databases, Plenum Press, New York, 1978.

[GA90] Gardenfors, P., "Induction, Conceptual Spaces and AI", Philosophy of Science, 57, (1990), pp 78-95.

[GO55] Goodman, N., "Fact, Fiction and Forecast", Harvard University Press, Cambridge, MA., 1955.

[HE65] Hempel, C.G., "Aspects of Scientific Explanation and Other Essays in the Philosophy of Science", Free Press, New York, 1965.

[HO86] Holland, J., Holyoak, K.J., Nisbett, R.E. and Thagard, P.R., "Induction: Processes of Inference, Learning and Discovery", MIT Press, Cambridge, MA., 1986.

[KA50] Kamke, E., "Theory of Sets", Dover Publications, New York, 1950.

[NA61] Nagle, E., "The Structure of Science", Routledge and Kegan Paul, London, 1961.

[PO62] Popper, K., "Conjectures and Refutations", Basic Books, New York, 1962.

[QU90] Quinlan, J.R., "Learning Logical Definitions from Relations", Machine Learning, 5, 239-166, 1990.

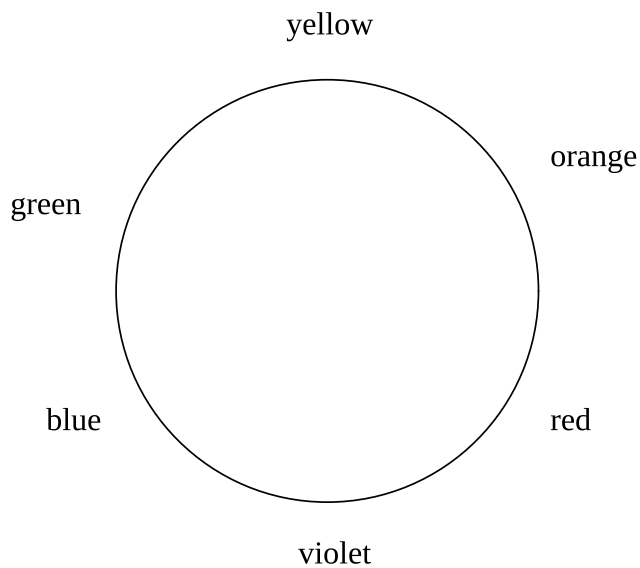

Figure 1. The Colour Circle

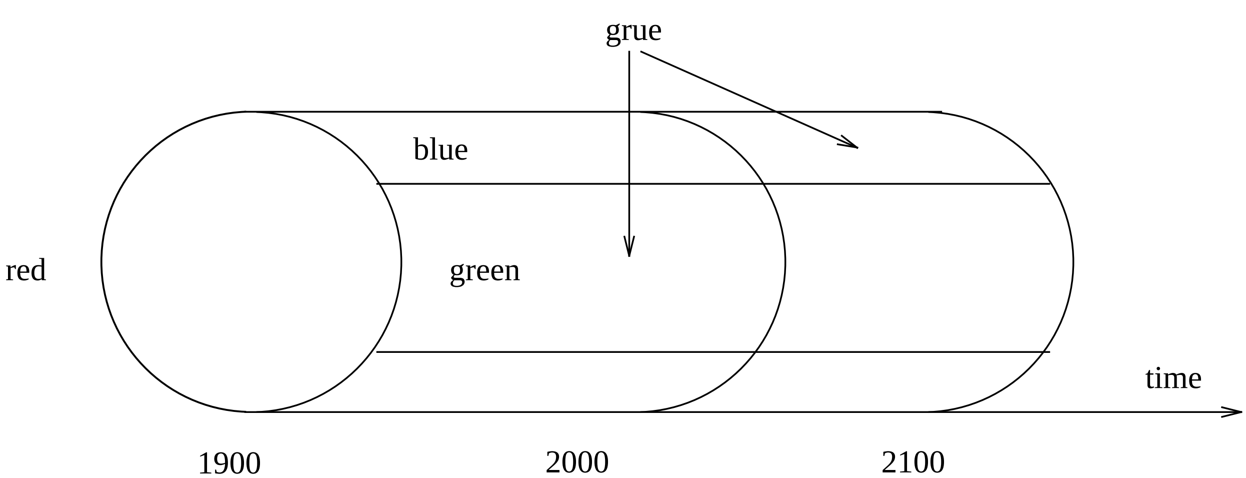

Figure 2. Predicate grue is not convex (after Gardenfors [GA90])

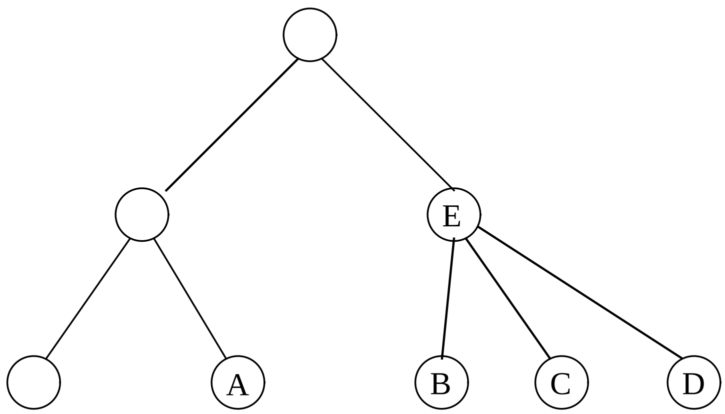

Figure 3. {A,B,C} is not convex; {B,C,D,E} is convex.

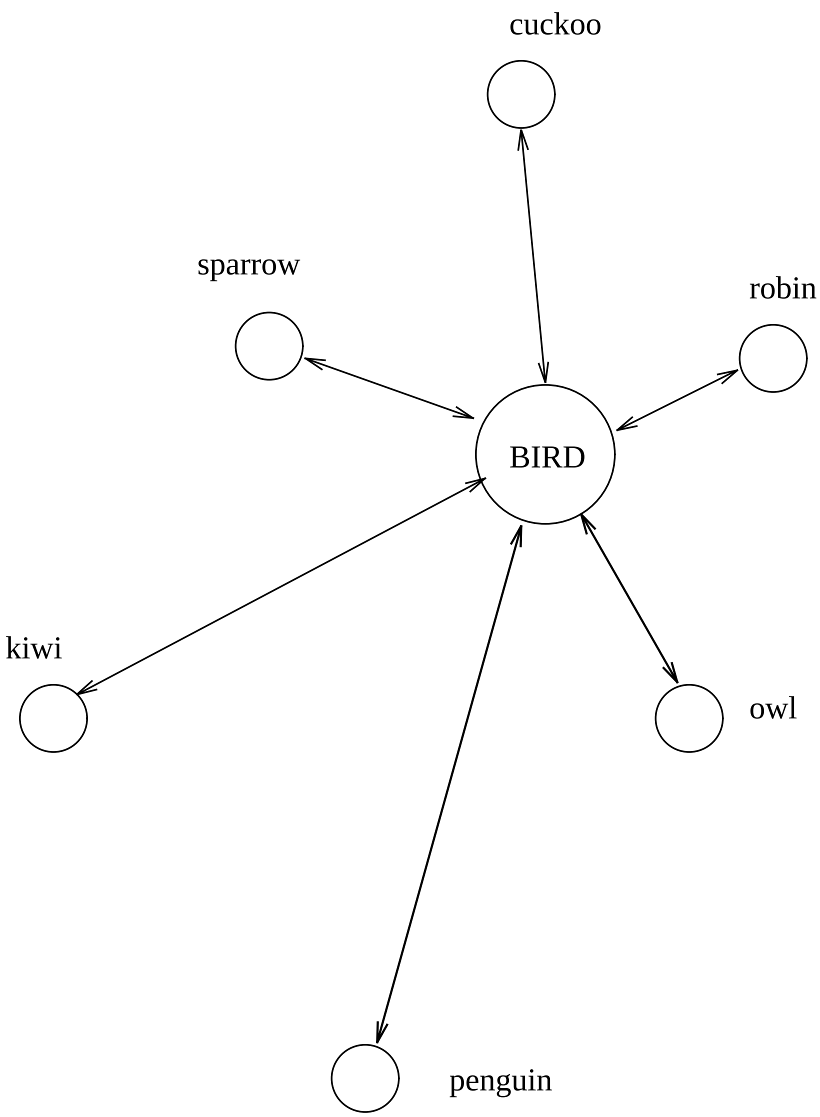

Figure 4. A prototype and some instances