A Logic for Qualified Syllogisms

Daniel G. Schwartz

Florida State University
Tallahassee, Florida, USA

schwartz@cs.fsu.edu

Being a thumbnail sketch of ideas published in:

D.G. Schwartz, Dynamic reasoning with qualified syllogisms, *Artificial Intelligence*, **93**, 1-2 (1997) 103–167.

Presented here to give it visibility within the fuzzy community.

Outline

1. The logic **Q** for Qualified Syllogisms
   - Reasoning with Fuzzy Quantification, Usuality, and Likelihood
2. Application to Nonmonotonic Reasoning
   - Resolving some well-known "puzzles"

Objective of Q

To develop a formal logic that captures the reasoning in syllogisms such as:

Most birds can fly.
Tweety is a bird.

It is *likely* that Tweety can fly.

and

Usually if something is a bird, it can fly.
Tweety is a bird.

It is *likely* that Tweety can fly.

Inuititive Motivation, Part 1

There is a natural connection between

Fuzzy Quantification (*all, most, few*, etc.)

and

Fuzzy Likelihood (*certain, likely, unlikely*, etc.)

in the sense that the statement

MOST birds can fly.

may be regarded as equivalent with

If x is a bird, then it is LIKELY that x can fly.

The underlying connection is provided by the concept of a *statistical sampling*, to wit:

Given a bird arbitrarily selected from the population of all birds, there is a high probability that it will be able to fly.

Psuedo Formalism

The foregoing equivalence may be expressed by

$$(Mx)(\text{Bird}(x) \rightarrow \text{Can\_Fly}(x)) \leftrightarrow (\text{Bird}(x) \rightarrow (L)\text{Can\_Fly}(x))$$

in which case, the foregoing syllogism reduces to an application of classical Modus Ponens:

$$\frac{\begin{array}{l}\text{Bird}(x) \rightarrow (L)\text{Can\_Fly}(x) \\ \text{Bird(Tweety)}\end{array}}{(L)\text{Can\_Fly(Tweety)}}$$

This suggests that fuzzy quantification and likelihood can be formalized by adjoining classical logic with an appropriate set of operators.

Intuititive Motivation, Part 2

There is a similar connection between the foregoing two concepts and

Fuzzy Usuality (*always, usually, seldom*, etc.)

Based on the same idea of a statistical sampling, one has that

USUALLY, if something is a bird, then it can fly.

or in symbols

$$(Ux)(\text{Bird}(x) \to \text{Can\_Fly}(x))$$

is equivalent with both

$$(Mx)(\text{Bird}(x) \to \text{Can\_Fly}(x))$$

and

$$\text{Bird}(x) \to (L)\text{Can\_Fly}(x)$$

Thus one should be able to also include usuality in an extension of classical logic.

The Distinctions

Usuality expresses the results of *past exprience* with a population,

Quantification expresses knowledge about the *current state* of a population,

Likelihood expresses expectations about the *future*.

In practice, the latter two are derived from the former, i.e., both quantification and likelihood are *pragmatically rooted* in past exprience.

The Interrelations

Quantification	*Usuality*	*Likelihood*
all	always	certainly
most	usually	likely
many/about half	frequently/often	uncertain/about 50-50
few/some	ocasionally/seldom	unlikely
no	never	certainly not

This suggests that all three concepts can be modeled by the same semantics.

Represent these modifiers by:

$$\mathcal{Q}_2,\ \mathcal{Q}_1,\ \mathcal{Q}_0,\ \mathcal{Q}_{-1},\ \mathcal{Q}_{-2},$$
$$\mathcal{U}_2, \mathcal{U}_1, \mathcal{U}_0, \mathcal{U}_{-1}, \mathcal{U}_{-2},$$
$$\mathcal{L}_2,\ \mathcal{L}_1,\ \mathcal{L}_0,\ \mathcal{L}_{-1},\ \mathcal{L}_{-2}.$$

Now define system **Q 1.0** as follows.

Languages

Symbols:

1. an *individual variable*, denoted by x,
2. *individual constants*, denoted generically by $a, b, \ldots$,
3. some unary *relation symbols*, denoted generically by $\alpha, \beta, \ldots$,
4. *logical connectives*, denoted by $\neg, \vee, \wedge, \rightarrow, \dot{\rightarrow}, \ddot{\neg}$, and $\ddot{\vee}$,
5. the foregoing *quantifiers*, $\mathcal{Q}_i$, *usuality modifiers*, $\mathcal{U}_i$, and *likelihood modifiers*, $\mathcal{L}_i$,
6. the *parentheses* and the *comma*.

Formulas:

$$\mathsf{F}_1 = \{\alpha(x)|\alpha \text{ is a relation symbol}\}$$

$$\mathsf{F}_2 = \mathsf{F}_1 \cup \{\neg P, (P \vee Q), (P \wedge Q)|P, Q \in \mathsf{F}_1 \cup \mathsf{F}_2\}$$

$$\mathsf{F}_3 = \{(P \rightarrow Q)|P, Q \in \mathsf{F}_2\}$$

$$\begin{aligned}\mathrm{F}_4 = \{&\mathcal{L}_2(P\dot{\rightarrow}\mathcal{L}_i Q), \mathcal{L}_2(P\dot{\rightarrow}\mathcal{Q}_i Q), \mathcal{L}_2(P\dot{\rightarrow}\mathcal{U}_i Q),\\ &\mathcal{Q}_2(P\dot{\rightarrow}\mathcal{L}_i Q), \mathcal{Q}_2(P\dot{\rightarrow}\mathcal{Q}_i Q), \mathcal{Q}_2(P\dot{\rightarrow}\mathcal{U}_i Q),\\ &\mathcal{U}_2(P\dot{\rightarrow}\mathcal{L}_i Q),\mathcal{U}_2(P\dot{\rightarrow}\mathcal{Q}_i Q),\mathcal{U}_2(P\dot{\rightarrow}\mathcal{U}_i Q)|\\ &P, Q \in \mathrm{F}_2 \cup \mathrm{F}_3, i = -2, \ldots, 2\}\end{aligned}$$

$$\mathrm{F}_5 = \{\mathcal{L}_i P, \mathcal{Q}_i P, \mathcal{U}_i P, |P, Q \in \mathrm{F}_2 \cup \mathrm{F}_3, i = -2, \ldots, 2\}$$

$$\mathrm{F}_6 = \mathrm{F}_4 \cup \mathrm{F}_5 \cup \{\ddot{\neg} P, (P \ddot{\vee} Q)|P, Q \in \mathrm{F}_4 \cup \mathrm{F}_5 \cup \mathrm{F}_6\}$$

All substitutions of individual constants for the individual variable *x*.

Abbreviations:

$(P \ddot{\wedge} Q)$	$\ddot{\neg}(\ddot{\neg}P \ddot{\vee} \ddot{\neg}Q)$
$(P \ddot{\rightarrow} Q)$	$(\ddot{\neg}P \ddot{\vee} Q)$
$(P \ddot{\leftrightarrow} Q)$	$((P \ddot{\rightarrow} Q) \ddot{\wedge} (Q \ddot{\rightarrow} P))$

The first "Tweety" syllogism can now be expressed in a language employing the individual constant Tweety and the unary relations Bird and CanFly as:

$$\frac{\begin{array}{l}\mathcal{Q}_1(\mathrm{Bird}(x) \rightarrow \mathrm{CanFly}(x)) \\ \mathcal{L}_2\mathrm{Bird}(\mathrm{Tweety})\end{array}}{\mathcal{L}_1\mathrm{CanFly}(\mathrm{Tweety})}$$

In words: For *most* x, if x is a Bird then x CanFly; it is *certain* that Tweety is a Bird; therefore it is *likely* that Tweety CanFly.

And the related equivalence is:

$$\mathcal{Q}_1(\mathrm{Bird}(x) \rightarrow \mathrm{Can\_Fly}(x)) \ddot{\leftrightarrow} (\mathcal{L}_2\mathrm{Bird}(x) \ddot{\rightarrow} \mathcal{L}_1\mathrm{Can\_Fly}(x))$$

The syllogism can be derived by first using the first line and the equivalence to derive

$$\mathcal{L}_2\text{Bird}(x)\ddot{\rightarrow}\mathcal{L}_1\text{Can\_Fly}(x)$$

instantiating x with Tweety, giving

$$\mathcal{L}_2\text{Bird(Tweety)}\ddot{\rightarrow}\mathcal{L}_1\text{Can\_Fly(Tweety)}$$

and then applying classical Modus Ponens using the second line as

$$\begin{array}{l} \mathcal{L}_2\text{Bird(Tweety)}\ddot{\rightarrow}\mathcal{L}_1\text{Can\_Fly(Tweety)} \\ \mathcal{L}_2\text{Bird(Tweety)} \\ \hline \mathcal{L}_1\text{CanFly(Tweety)} \end{array}$$

giving the desired

$$\mathcal{L}_1\text{CanFly(Tweety)}$$

Semantics

Two notions of semantics are considered, based on:

1. Bayesian probability theory
   - Probabilities are assigned to propositions subjectively, without reference to an underlying universe.
2. L.A. Zadeh's notion of Sigma Counts (restricted to crisp predicates)
   - A frequentist approach
   - Probabilities are computed by counting individuals in an underlying universe.

In both versions, an *interpretation* I for a language L consists of

- a *likelihood mapping* l_I which associates each lower-level formula with a number in $[0, 1]$, and
- a *truth valuation* v_I which associates each upper-level formula with a *truth value*, T or F.

The likelihood mappings satisfy the conditions for a probability assignment. For the truth valuations, the interval $[0, 1]$ is divided into five disjoint intervals that cover $[0, 1]$, such as

$$
\begin{aligned}
\iota_2 &= [1, 1] \quad \text{(singleton 1)} \\
\iota_1 &= [\tfrac{2}{3}, 1) \\
\iota_0 &= (\tfrac{1}{3}, \tfrac{2}{3}) \\
\iota_{-2} &= (0, \tfrac{1}{3}] \\
\iota_{-2} &= [0, 0] \quad \text{(singleton 0)}
\end{aligned}
$$

Then for an upper-level formula of the form $\mathcal{M}_i P$ (so that P is a lower-level formula) set

$$v(\mathcal{M}_i P) = T \text{ iff } I(P) \in \iota_i$$

For example, if $\mathcal{M}_i$ is $\mathcal{L}_1$ (standing for *likely*), then

$$v(\mathcal{L}_1 P) = T \text{ iff } I(P) \in \iota_1$$

This leads to two well-defined semantics that validate the foregoing syllogisms.

A Key Result: At the upper level, both semantics validate the axioms and inference rules of Classical Propositional Calculus.

In addition, both semantics validate all formulas having the forms (for open P and Q):

$$\mathcal{Q}_i(P \to Q) \ddot{\leftrightarrow} \mathcal{Q}_3(P \dot{\to} \mathcal{L}_i Q)$$
$$\mathcal{U}_i(P \to Q) \ddot{\leftrightarrow} \mathcal{U}_3(P \dot{\to} \mathcal{L}_i Q)$$

which express salient aspects of the interrelations between quantification, likelihood, and usuality.

Another general form validated by this semantics is

$$\mathcal{Q}_i(P \to Q) \ddot{\leftrightarrow} (\mathcal{L}_3 P \ddot{\to} \mathcal{L}_i Q)$$

which includes the equivalence regarding Tweety discussed previously.

Nonmonotonic Reasoning

Classical formal logical systems are monotonic in that adding new information (axioms) always increases the set the theorems (derivable from the axioms).

Reasoning is nonmonotonic when adding new information cause one to go back and retract old conclusions.

Example:

Suppose that on Tuesday you are told ”Opus is a bird”. Then by default (i.e., in the absence of any countervailing information) you may conclude ”Opus can fly”.

But suppose that on Thursday you are told ”Opus is a penguin”. Now, based on what you know about penguins, you must retract the above and conclude that ”Opus cannot fly”.

General Problem of NMR

How to represent and manage this type of reasoning.

Well-Known Early Approaches

- Truth (or Reason) Maintenance — Jon Doyle, 1979, 1988
- Circumscription — John McCarthy, 1980
- Default Logic — Raymond Reiter, 1980
- Nonmonotonic Logic — David McDermott and Jon Doyle, 1980

Current Threads

- Belief Revision — The AGM Framework (Alchourrón, Gärdenfors, Makinson)
  - 25 years in development by many contributors
- Answer Set Programming (an extension of Prolog)
  - About 15 years in development by many contributors

Neither have so far yielded computational algorithms.

Applying **Q** to Nonmonotonic Reasoning

Requires four additional components.

First is needed a logic for likelihood combination.

For example, if by one line of reasoning one derives *LikelyP*, and by another derives *UnlikelyP*, then one would like to combine these to obtain *UncertainP*.

	2	1	0	-1	-2
2	2	2	2	2	$*$
1	2	1	1	0	-2
0	2	1	0	-1	-2
-1	2	0	-1	-1	-2
-2	$*$	-2	-2	-2	-2

Table: Rules for likelihood combination.

Second is needed a means for providing such inference rules with a well-defined semantics.

Simultaneously asserting *LikelyP* and *UnlikelyP* requires that *P* have two distinct likelihood values, in which case the likelihood mapping *l* would not be well-defined.

Resolved by means of a *path logic*, which portrays reasoning as an activity that takes place in *time*.

Different occurrences of *P* in the derivation path (i.e., the sequence of derivation steps normally regarded as a *proof*) are labeled with a *time stamp*.

Then the likelihood mapping can be defined on labeled formulas, in which case each differently labeled occurrence of *P* can have its own likelihood value.

Third one needs to distinguish between predicates that represent *kinds* of things and those that represent *properties* of things.

To illustrate, in the “Tweety” syllogism, “Bird” represents a kind, whereas “CanFly” represents a property.

For this purpose employ *typed predicate symbols*, indicated by superscripts as in $\text{Bird}^{(k)}$ and $\text{CanFly}^{(p)}$.

Fourth is needed a way of expressing a *specificity relation* between kinds of things, together with an associated *specificity rule*.

For example, if "*All*($\text{Penguin}^{(k)}(x) \to \text{Bird}^{(k)}(x)$" is asserted in the derivation path, asserting in effect that the set of penguins is a subset of the set of birds, then one needs to make an extralogical record that $\text{Penguin}^{(k)}$ is more specific than $\text{Bird}^{(k)}$.

Given this, one can apply the principle that more specific information takes priority over less specific.

Collectively, these components comprise a system for a style of nonmonotonic reasoning known as as default reasoning with exceptions.

The problems associated with formulating this kind of reasoning have been illustrated by a variety of conundrums, the most well-known being the situation of Opus the penguin as illustrated in the following figure.

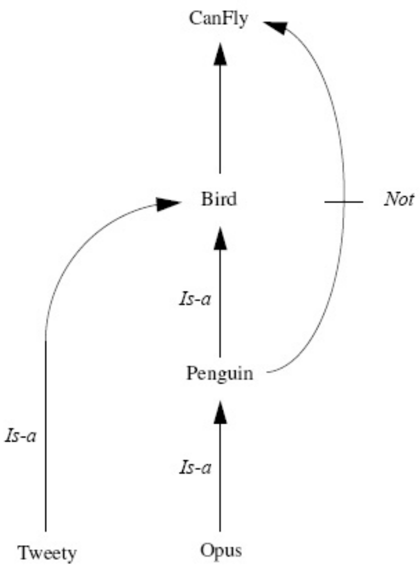

Figure: Tweety can fly, but can Opus?

The puzzle can be resolved in the present system as shown in:

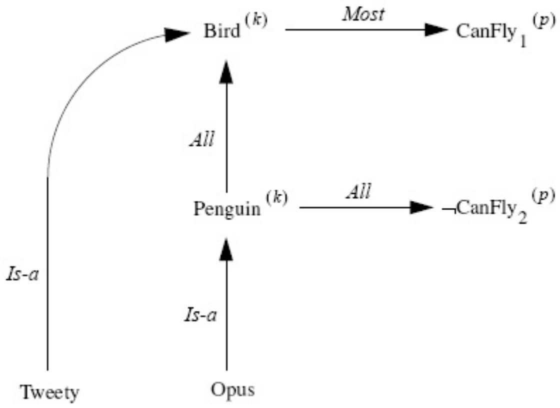

Figure: Tweety *likely* can fly, and Opus *certainly* cannot.

For Opus the inheritance of CanFly_1 is blocked by the more specific information $\neg\text{CanFly}_2$.

Another famous puzzle, the Nixon Diamond.

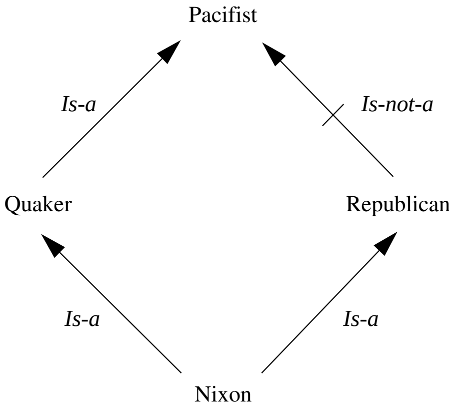

Figure: Is Nixon a pacifist or not?

This can be resolved in the present system with:

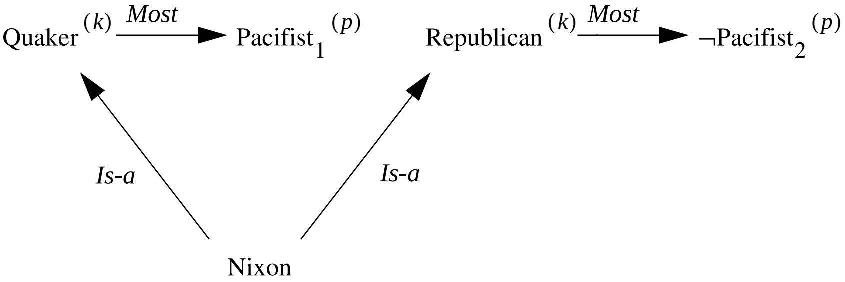

Figure: It is *uncertain* whether Nixon is a pacifist.

Conclusion

- The logic **Q** for qualified syllogisms provides an intuitively plausible method for default reasoning with exceptions.
- Future Reseach
  - Applications
    - Frame-based expert systems
    - Robot motion planning
  - Theory
    - Rules governing typed predicate sytmbols
    - Sound and semantically complete axiomatization