

# A Tale of Two Knowledge Servers

Randall Davis

An alien arriving on earth in the late twentieth century was fascinated but confused by what it saw and wanted answers to a few questions. Knowing well that the best place to go for information was an intelligent machine, the alien found itself in the Land of AI, where it encountered two machines claiming to be able to answer questions. The alien approached each in turn.

The first machine said:

*Greetings. I am Spock, a knowledge representation and reasoning service. I provide answers to your queries. Coming to me was the logical thing to do because no matter what the question, I offer two guarantees: (1) if there is an answer, you can be sure I'll supply it and (2) the answer I give you is guaranteed to be logically correct. You can feel free to ask any question you like, in any form that can be expressed precisely. Oh, by the way, before you begin, it would also be logical to read the fine print in my contract.*

Even on Betelgeuse they know about small print, so the alien scanned the contract:

- (a) Questioner agrees to express the question in predicate calculus or any established syntactic variant thereof.*
- (b) Questioner further agrees to wait here for reply, where "wait" is defined as remaining absolutely and completely motionless and absolutely and completely inactive mentally. In effect, questioner agrees to act as if he/she/it has been completely and utterly lobotomized until reply is received.*
- (c) Questioner further acknowledges that a reply might require a substantial period to arrive (the delay may, in fact, best be measured in universe lifetimes).*
- (d) This agreement shall be binding on the questioner, all his/her/its heirs, as well as all future life forms that might evolve therefrom, before reply arrives.*

Being quite long lived but still mortal, the alien decided to bypass this one and try the next. The second machine said:

*Greetings. I am the one called KayEl. I provide answers to your queries and hence I am the right category of entity for your needs. Coming to me was clearly the right thing to do because I offer three guarantees: (1) if there is an answer to your question, you can be sure I'll supply it; (2) the answer I give you is guaranteed to be logically correct; and (3) I will answer you in a very reasonable amount of time, as long as you can phrase it in the appropriate way.*

Then, in a soft voice learned from automobile commercials, the machine added quickly,

*Some additional terms and restrictions apply.*

The alien decided to give it a try, dropped five Universal Galactic Credits into the slot and asked, "I've noticed that there seem to be two disjoint kinds of individuals on this planet. Is it true that every creature is either male or female?"

Whereupon KayEl replied: *Question not askable. See contract.*

So the alien tried again: "Is there any creature that does not belong to either gender?"

But again KayEl replied: *Question not askable. See contract.*

The alien tried yet again: "Is it true that if an individual is not a male, then it is female?"

Alas, once again KayEl replied: *Question not askable*, this time adding, *You might wish to examine the fine print in my contract for help in phrasing your question.*

And so the alien did:

- (a) Questioner agrees to express the question in a restricted subset of predicate calculus or any established syntactic variant thereof.*
- (b) Questioner further agrees that question shall not contain any of the following constructs: (i) disjunction, (ii) negation, (iii) conditionals, (iv) equivalence, (v) recursive definitions, (vi) transitive relations, (vii) functions over ordered sets, (viii) mappings between ordered sets, (ix) functions of sets and functions, (x) binary functions, and (xi)*

*n-ary relations.*

In the face of all these restrictions, the alien shrugged all 16 of its shoulders in an impressive show of despair: The answer from KayEl might be quick in arriving, but it was almost impossible to ask it an interesting question.

This left the alien with quite a conundrum. It seemed as though there were only two options on this planet for getting questions answered: a machine that could answer any question and guaranteed the right answer but might take essentially forever to reply and a machine that could be guaranteed both correct and quick but whose abilities were so narrow as to be almost useless. What a disagreeable situation; what an incredibly primitive culture.

While wandering about trying to decide what to do, the alien realized that it was passing by yet another intelligent machine, one that had been there all the time but had somehow been overlooked. The alien approached the machine and listened to its description:

*Hi there and welcome to the Drop-Off Knowledge Representation and Reasoning Service. We'll do the best we can with whatever tools seem to be appropriate. We often supply the right answer promptly, but we sometimes take a very long time. Our slogan is **We might be fast, but there's no guarantee™.***

*On the other hand, we won't tie you down. You see with us, you can drop your query here and go about the rest of your business. A lot of our customers leave a query here, then spend their time doing other work while we're busy with their query.*

*You can stop back any time. If we're done, terrific. If not, we'll show you what we have so far. But you don't have to stand here idly waiting for us.*

*You should also keep in mind that we have many branch offices here in AI Land, each one specializing in a different approach to answering queries. Because we don't tie you down, a lot of our customers drop off a query here and then go on to several of our other branch offices and drop the query there too. (In fact, one of those offices simply subcontracts out the work to Spock, but we still don't tie you down; you can drop off the query and still go about your business. Of course, there's no guarantee Spock will be finished when you want your answer, but it's often worth a try.)*

*Note also that different levels of service are available at the different branches. If you want an answer guaranteed optimal, it costs more and generally takes longer. We also offer*

*the less expensive Good Enough™ service, where the answer might not be optimal, but it will lie within bounded limits of optimality. Most inexpensive (and often fastest) is our Good Guess™ service, where we take our best quick shot at the problem and leave it to you to decide whether our answer will serve your purposes. Most intelligent life forms in the universe find both of these approaches to be very useful in the real world.*

*Because we presume that your supply of Universal Galactic Credits is limited, you might wish to consider how to apportion your Credits among our different branch offices. We suggest you consider the cost of the service at each office (posted on the door) and the likelihood that that office's approach will lead to an answer.*

*That is, after all, the rational thing to do.*

Then it added in a quiet, embarrassed voice:

*Be sure to check my warning label.*

The alien looked around and eventually spotted a small, official looking label stuck to the side of the machine:

**Warning: Use At Your Own Risk.**

This machine is not a certified knowledge representation server: it is not sound, it is not complete, and there is no guarantee that it will produce an answer in a tractable amount of time. There is in fact no precise characterization of its running time at all. How any being in its right mind could use such a service is utterly beyond me.

(signed)

F. Lou deKoop

Knowledge Representation General

The label seemed quite curious to the alien, so it tried asking the machine some questions about it.

"Are there intelligent beings on this planet?" it inquired. Yes, replied the machine.

"Do they ask questions of each other?" Yes, came the answer.

"And do they believe the replies they get?" Sometimes, said the machine.

"And are the beings using reasoning that is known to be sound and complete and whose run-time performance can be specified with complete precision?" *Certainly not*, said the machine, with just a touch of smugness in its voice.

"Well, then," said the alien, "If they are willing to risk asking questions of one another,

why ever do they make you wear that warning label?"

*Yes, that's always been a mystery to me too,* replied the machine, just a bit wistfully.

And with that the alien realized several important lessons about asking questions:

It's useful to know what kinds of questions are guaranteed answerable in a tolerable amount of time. These can provide a reliable foundation for any system.

But the alternative to "guaranteed tolerable" isn't "*guaranteed to be intolerable*," merely "not guaranteed tolerable." The absence of a guarantee is not a guarantee of fatal problems.

Hence, extraordinary efforts to avoid the "not guaranteed tolerable" territory can be self-defeating. For example, such efforts can lead to misguided concerns about speed: Restricting questions to those guaranteed answerable quickly can make it almost impossible to ask anything interesting.

Besides which, there's no reason on Earth (or any other planet) to presume that you have to stand idle while you wait for a query to be answered.

And if you don't have to stand idle, *there's no particular consequence to the possibility that a reply might be a very long time in coming.* That is, undue concern with worst-case analysis can lead to needless paralysis. Rather than standing idle, you can explore in parallel several different approaches to getting the question answered.

Because resources can be large but are inevitably finite, the important question then becomes how to apportion resources among multiple possible approaches. Estimations of cost and benefit might be particularly useful in this decision.

Estimations of *average* cost and benefit are especially appropriate here because it's the expected utility, not the worst case utility that matters.

It might suffice to have only rough estimations of cost and a crude strategy for allocating resources (for example, equipartition) because most of the benefit is in pursuing more than one approach; there might be no need to find a particularly clever allocation strategy.

Fear of the "not guaranteed tolerable" can also lead to undue focus on guarantees:

First, there are many ways to get *good enough* answers to questions if you don't require that the answer be guaranteed correct.

Second, a guaranteed correct answer isn't worth much if you won't live to see it (what is the exact classification of that large green thing

with giant teeth coming toward me?). Conversely, a good guess now (is it likely to mean harm?) can be quite valuable. Good but unsound guesses are vital and demonstrably useful: Every intelligent creature in the universe bets its life every day on inferences that are good but unsound.

Third, if we insist that a method is useless unless its run time can be characterized with mathematical precision, we will indeed find ourselves facing an at times untenable trade-off between expressiveness (Spock) and speed (KayEl). But if the desire for guarantees leads to an untenable situation, consider dropping the insistence on guarantees. And consider that every intelligent creature in the universe bets its life every day on inference processes that lack performance guarantees.

Daily life is formally an intractable problem: Being intelligent isn't predicated on guaranteeing that you can stay out of the complexity; it's a matter of functioning well in the midst of it, in spite of it.

And these were very important lessons indeed.

## Acknowledgments

The story is inspired by "Two Dogmas of Knowledge Representation" by John Doyle and Ramesh Patil, LCS Technical Memo, 387B, Massachusetts Institute of Technology, September 1989. This paper provided the technical result that is summarized briefly (but without exaggeration) in clause b of the second contract. A revised version of this paper appears in *Artificial Intelligence* (volume 3, 1991, pp. 261-297) as "Two Theses of Knowledge Representation: Language Restrictions, Taxonomic Classification, and the Utility of Representation Services." Additional useful ideas and comments were contributed by Howard Shrobe and Peter Szolovits. The satire and any remaining errors are the responsibility of the author.



**Randall Davis** has been wandering about in the Land of AI since 1970, for the most part heading from west to east. He is currently a professor of electrical engineering and computer science at the Massachusetts Institute of Technology (MIT) as well as associate director of the AI Lab at MIT.