Editor: Charles Petrie • petrie@stanford.edu

Pragmatic Semantic Unification

Charles Petrie • *Stanford University*

I was late to *IEEE Internet Computing*'s recent editorial board meeting, and so was volunteered to take over as department editor for Peer to Peer. The good news is that they said I could write about whatever I wanted. So you'll get my opinions, as well as those of guest columnists. And since this column is supposed to be peer-to-peer, your feedback/pushback is crucial.

Please send official "letters to the columnist" to *IC*'s lead editor Rebecca Deuel at rdeuel@computer.org, and include reference to the article and issue. Now, imagine that I've said some legal stuff about the authorities' disavowal of authors' opinions, and let's begin with two subjects for your consideration.

A Possibly Iconoclastic Thesis

As you know, many researchers, mostly academics, are devoting great resources to developing the Semantic Web and the corresponding semantic Web services. Usually based on top-down ontologies, the various proposed technology suites — most notably, OWL-S and the Web Service Modeling Ontology (www.wsmo.org) — are suitably complex for the many masters and PhD theses they will spawn. This is all in pursuit of a grand vision with huge potential, and like many others, I support it.

Such ontologies are intended primarily to solve the problem of *semantic unification*: determining whether two given terms share a common meaning. This process is usually expressed as the merging of ontologies. An ontology differs from a taxonomy in that the terms are constrained not only by some class hierarchy but also by formal axioms. Given that these axioms and class properties define the terms' meanings, unifying the terms means relating the axioms and properties, usually through formal approaches. Given some starting points or mappings between ontologies, we can formally infer further term unification.

Likewise, industry spends a lot of money dealing with semantics, particularly with the emerging methodology of Web service-based process integration, which is increasingly performed with Business Process Execution Language (BPEL) rather than other process-integration systems. Currently, developers must do this integration by hand, creating files, or even databases, of hand-coded links between terms in different applications, and testing what works.

I propose that industry has got a major philosophical point right: term unification is meaningful only with respect to applications. A term's meaning, or semantics, is evident in its use — if the application is successful, the terms have been used correctly. Unifying academic ontologies is unification in a vacuum. The only way to know if we've done it correctly is to test whether the applications associated with the terms interoperate successfully — a notion I call *pragmatic semantic unification*.

This doesn't mean that ontologies, with their declarative axioms, aren't the right way to proceed, but it does indicate that there's more to be done. Ontologies should refer to applications, and unification success should be tested with interoperability. Axioms might need to be specific to application context. Operationally, perhaps the right place to begin is with industry files, rather than with what we academics can conceive of via Kantian a priori analyses.

Crossing the Cultural Divide

The Semantic Web is likely to go the way of expert systems, software agents, and relational databases: highly successful after much of the original academic software is thrown away, but widely implemented 20 years later than possible.

One reason for this is the community's tendency toward incremental pragmatic application of theoretical and potential approaches, but another barrier to technology transfer is that industry and academia have different technical cultures. This can make it difficult to even talk to each other. We use the same words, but there is little pragmatic semantic unification.

As an example, I worked with software agent experts from an unnamed company one summer, planning for weeks on how to make our two agent systems interoperate. When we finally got to the implementation stage, they asked what was, to me, a puzzling question: how would we address a message to an agent? My answer, equally puzzling to them, was to use the agent's IP address. As it turned out, they assumed that agents all ran on a platform with a single IP address, whereas my agents each had individual addresses. We couldn't discover

continued on p. 95

 Published by the IEEE Computer Society 1089-7801/05/$20.00

this assumption through talking as we were using the same words. We had to build an application to discover the true semantics of our terms.

In another example, I recently gave a public talk in which I suggested using BPEL partner links to bootstrap semantic unification. An industry expert on BPEL and I explored this notion, but we found it took days for either of us to understand what the other was saying. With neither the time nor the resources to actually build something together, we left it with just some vague notions.

Workshops might be a shortcut to application construction — if structured appropriately. Two complementary approaches are likely to be effective:

- Let everyone present their papers and then run a facilitated session in which attendees say what they didn't understand about the papers, grounding the discussion in examples. I've tried it informally in some small groups, and it has helped uncover hidden assumptions.
- Build a workshop around a technical challenge. If every paper describes how the authors' approach (partially) solves a common technical problem, it can help clarify the terms of discussion and, thus, the issues.

The latter approach has been successfully tested and, indeed, is increasingly popular. I did a workshop like this on distributed design and planning in 1994, for example (http://www-cdr.stanford.edu/~petrie/caia.html), and Dieter Fensel, of the Digital Enterprise Research Institute (DERI; http://www.deri.org), has helped design such workshop challenge problems for elevator design[1] and online product classification.[2] Both problems are often cited and have generated good work. Tim Finin has organized several challenges, such as the recent Web service integration challenge workshop (www.comp.hkbu.edu.hk/~eee05/contest/) and DERI's recently announced challenge workshop on process integration.

Both approaches to running workshops can facilitate understanding of common problems (and can be used together). They work with audiences of academics from various disciplines, as well as mixed groups of industrial and academic researchers. Ultimately, the more diverse the group, the bigger the semantic unification problem — and the more worthwhile the results.

We had to build an application to discover the true semantics of our terms.

Formal, and informal, semantic unification requires reference applications. Machines and humans demonstrably understand each other just when they work together successfully. Otherwise, it's all just talk.

References

1. K. Poeck et al., "Combining KARL and CRLM for Designing Vertical Transportation Systems," *Int'l J. Human-Computer Studies* (IJHCS), vol. 44, nos. 3–4, 1996, pp. 435–467.
2. E. Schulten et al., "The E-Commerce Product Classification Challenge," *IEEE Intelligent Systems*, vol. 16, no. 4, 2001, pp. 86–89.

Charles Petrie is a senior resarch scientist and consulting associate professor at Stanford University. His research interests include concurrent engineering, virtual enterprise management, and collective work. Petrie has a PhD in computer science from the University of Texas at Austin. He is EIC emeritus and a member of *IC*'s editorial board. Contact him at petrie@stanford.edu.