Matching Requests for Agent Services with Differentiated Vocabulary

Peter Weinstein and William P. Birmingham

Artificial Intelligence Laboratory, University of Michigan
1101 Beal Ave., Ann Arbor, MI 48109-2110
{peterw, wpb}@eecs.umich.edu

To enable decentralized development of large societies of agents, agents should be able to selectively team with others based on declarative descriptions of services, rather than *a priori* knowledge.[1] This capability is difficult to achieve because descriptions written by different developers may be terminologically heterogenous—including vocabulary from ontologies that are potentially inconsistent. For example, one agent might describe its service as (a formal equivalent of) "query planning for high-school biology", while another agent wants to "find collections for advanced life sciences". We want the latter agent to recognize that the former might satisfy its request.

We have completed research on two aspects of this problem. Our Service Classifier Agent (SCA) supports selection of agent services in societies that are dynamic and evolving, but whose agents all use the same ontologies [Weinstein and Birmingham 97]. We have also developed an algorithm that identifies maximal similarity between concept definitions that are terminologically heterogenous [Weinstein 95].

The SCA uses description logic to maintain a subsumption taxonomy of available services. Agents define their services at *runtime*, using terms from a set of ontologies associated with the SCA (including the taxonomy of services). To find services, agents query the SCA. Queries describe the ideal service desired, but find the best available. If a new agent meets a request better than was previously possible, the requesting agent may automatically switch to using the new agent. The SCA thus facilitates evolution of the society to meet users' needs. Previously, ontologies used for agent communication have described the task domain, rather than agent services.

To assess similarity despite terminological heterogeneity, we build *rough mappings* between source and target concepts. Mappings are sets of one-to-one correspondences between subgraphs in the source and target concepts. Of many possible mappings between a pair of concepts, the largest and most densely linked are evaluated as the best (these ideas come from research in analogy; see Owen [90] for a lucid overview).

We assume a *differentiated vocabulary:* that unshared terms inherit from terms in shared ontologies. The strength of a single correspondence in a mapping is initially a function of the path length between the paired concepts and their least general subsuming concept in the union of ontologies. Path length is a fragile proxy for semantic distance, however, compared to rough mappings, which use the concept definitions. Hence, correspondence strengths based on rough mappings can improve rough mappings for other concepts.

Current research will match requests for services with advertisements in different communities of agents (a community is defined by use of an SCA and its associated ontologies). Two communities will be prepared for inter-ontology translation by *reclassifying* unshared concepts that are defined with terms that are shared. Also, a selection of rough mappings can be edited by humans to verify and refine initial estimates of concept similarity.

Subsequently, inter-community requests for services will be answered with a rough mapping to a candidate service, and also a *mapping negative* that captures differences between the request and the candidate. In any particular problem context, a single difference or a combination of differences may cause the recommended service to fail. Inductive learning within problem contexts, however, can be used to improve assessments. To provide grist for this reasoning we plan to develop a well-organized elaboration of standard ontological relations [Stephens and Chen 95]. For example, there are many kinds of "kind-of" relations. Distinguishing between them can help identify concept definition differences as complementary, or mutually exclusive.

References

Owen, S. 1990. *Analogy for Automated Reasoning*. Academic Press, San Diego.

Stephens, L. M., Yufeng F. C. 1995. Principles for Organizing Semantic Relations in Large Knowledge Bases, Technical Report ECE-LMS-95-10, University of South Carolina.

Weinstein, P. 1995. Matching Expressions with Differentiated Vocabulary: Technology for Knowledge Sharing in Partially Focused Communities. http://www.umich.edu/~peterw/abstract.vocab.html.

Weinstein, P. and Birmingham, W. 1997. Runtime Classification of Agent Services. Proceedings of the AAAI-97 Spring Symposium on Ontological Engineering.

[1] Copyright © 1997, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

This work is supported by the NSF/ARPA/NASA Digital Library Initiative under grant CERA IRI-9411287.