Facilitating Interaction Between Virtual Agents Through Negotiation Over Ontological Representation

E

Fiona McNeill
Heriot-Watt University, UK

INTRODUCTION

It is universally acknowledged that the problem of integration of information across large communities is a difficult and pressing one, particularly when these communities are disparate, widespread and not under centralised control, such as in the Semantic Web (Berners-Lee et al, 2001). The simplest solution to this problem is the enforcement of a single ontology: a single view of the world. However, if the agents interacting are from different organisations or fields, attempting to use a single ontology is usually neither practical nor desirable. Users need to develop a representation that is best suited to their own problems and they need to maintain and update that representation locally. Even if all users do subscribe to a single ontology, integration problems still exist, as changes and updates are made and users tune their ontologies to fit their own needs.

The problem of ontology matching has been widely studied and powerful solutions are available (see (Shvaiko & Euzenat, 2013; Euzenat & Shvaiko, 2013) for a comprehensive survey. However, the ontologies considered are almost always taxonomies, and the problem of ontology matching is concerned with relating a single term in one ontology to one or more terms in another ontology: for example, a term *car* in one ontology may relate to a term *automobile* or a term *carriage* in another ontology. Much less considered is the problem of relating compound terms such as first-order terms or database entries with multiple fields: for example, a term

car(make,model) in one ontology relating to a term *automobile(model,year,brand)* in another. In such situations we still have the problem of relating the single terms contained within these compound terms – e.g., this matching depends on knowing that *car* may be related to *automobile* and that *make* may be related to *brand*. But we must also consider the overall relation of the compound terms, which requires not only semantic but also structural matching.

Another drawback of traditional ontology matching in an online environment is that it tends to assume full knowledge of all relevant ontologies and is generally performed off-line, prior to interaction. These are the assumptions made by the main evaluation processes for Ontology Matching, such as the Ontology Alignment Evaluation Initiative (OAEI[1]). But in large, fast-moving agent communities, or situations where some information may be confidential, we cannot assume that we can have full knowledge of any agent or service we may interact with, nor is it possible to perform the matching off-line if we may not know prior to interaction which agents will need to interact.

In this paper, we introduce our theory of on-the-fly, structured matching and briefly describe the ORS system, which we have developed to implement this theory. Our central hypothesis is that representation – as well as vocabulary and beliefs – must be treated as a fluent and that automated, dynamic matching techniques that can map between structured terms are necessary for full integration of disparate ontologies (Bundy et al, 2006).

BACKGROUND

In a system such as the Semantic Web, where there is no centralised control, we cannot have a complete global overview of the agents and data in the system. Agents may join and leave the system freely and they will all have their own ontologies and data that may be large and complex and may be confidential. We cannot hope for a complete description of the relations between every agent in the system. Our approach is therefore not to consider how such a system can be controlled but how an individual agent can successfully make its way in such a system, interacting with the agents that it needs to interact with, even if these agents are not using the same ontological terms or representations, and even if it is not known in advance of the interaction which agents these will be.

Although many existing ontologies are simple taxonomies, and matching these ontologies is a crucial task, we believe that this kind of matching cannot be sufficient. Agents that are capable of interacting in complex and unpredictable environments need to be able to plan, and planning agents need far richer descriptions of the world: not only taxonomies of classes but also relations and functions between these classes, and planning rules describing how to influence the world. Uniformity of these relations, functions and rules can no more be assumed than uniformity of terms within taxonomies, and therefore matching between these structured objects is just as crucial as the more frequently addressed problem of matching between simple terms within a taxonomy.

Matching large ontologies at run-time, particular ones that contain structured terms, is generally not feasible, but we make this problem tractable by only fixing mismatches when this is demonstrated to be necessary. Since interactions may be frequent and fleeting, there may not be much value in matching the full ontologies, since the interaction that is desired may only require a very small part of the ontologies, and even if inconsistencies exist between the ontologies, these may not lead to communication breakdown between agents during a particular interaction. Our approach is therefore to diagnose mismatches and refine the ontologies accordingly only when these mismatches directly cause communication breakdown.

To this end, we developed ORS (the Ontology Repair System)[2]. This is a tool that an individual agent (which we name PA – the planning agent) can make use of as an aid when communication breaks down. ORS tracks the course of the communication between PA and any agents it may be interacting with (we name these SA – service-providing agents). If communication proceeds successfully then ORS does not need to be utilised. However, if communication breaks down, ORS begins the diagnostic process, analysing the communication so far and prompting PA to ask further questions in order to pin down a specific mismatch between the ontologies of PA and SA, which is then corrected.

The benefit of ORS is therefore that it allows an agent to interact successfully with other agents, even when their ontologies are mismatched in important ways, and even when this mismatch is between complex, structured ontological objects as well as when this mismatch is between simple terms. It works on the fly and fully automatically even when interactions are unpredictable and unforeseen.

ONTOLOGICAL MISMATCH

Planning agents require ontologies that contain three different kinds of objects, which entail three different kinds of mismatches:

1. Purely semantic mismatch, where the mismatch is between words or phrases – for example, *car* is matched to *auto*: this is the problem that is covered by conventional ontology matching.
2. Structural mismatch, where the mismatch is between structured terms (such as relations and functions) – for example, *car(Make,Model)* is matched to *car(Make),*

Related Content

Design of Library Archives Information Management Systems Based on Artificial Intelligence and Multimedia Technology
Ying Li (2023). *International Journal of Information Technologies and Systems Approach (pp. 1-17).*
www.irma-international.org/article/design-of-library-archives-information-management-systems-based-on-artificial-intelligence-and-multimedia-technology/320234

Ecological Performance as a New Metric to Measure Green Supply Chain Practices
June Poh Kim Tamand Yudi Fernando (2018). *Encyclopedia of Information Science and Technology, Fourth Edition (pp. 5357-5366).*
www.irma-international.org/chapter/ecological-performance-as-a-new-metric-to-measure-green-supply-chain-practices/184239

Strategies to Implement Edge Computing in a P2P Pervasive Grid
Luiz Angelo Steffenel, Manuele Kirsch Pinheiro, Lucas Vaz Peresand Damaris Kirsch Pinheiro (2018). *International Journal of Information Technologies and Systems Approach (pp. 1-15).*
www.irma-international.org/article/strategies-to-implement-edge-computing-in-a-p2p-pervasive-grid/193590

Prediction System-Based Community Partition for Tuberculosis Outbreak Spread
Fatima-Zohra Younsiand Djamila Hamdadou (2022). *International Journal of Information Technologies and Systems Approach (pp. 1-20).*
www.irma-international.org/article/prediction-system-based-community-partition-for-tuberculosis-outbreak-spread/289998

Current Status and Future Directions of Blended Learning Models
Michael C. Johnsonand Charles R. Graham (2015). *Encyclopedia of Information Science and Technology, Third Edition (pp. 2470-2480).*
www.irma-international.org/chapter/current-status-and-future-directions-of-blended-learning-models/112663