

Learning Non-Unanimous Ontology Concepts to Communicate with Groups of Agents

Mohsen Afsharchi, Behrouz H.Far
Department of Electrical and Computer Engineering
University of Calgary, Calgary, Canada
{mafsharc, far}@ucalgary.ca

Jörg Denzinger
Department of Computer Science
University of Calgary, Calgary, Canada
denzinge@cpsc.ucalgary.ca

Abstract

We present an extension to the definition of a concept in an ontology that allows an agent to simultaneously communicate with a group of agents that might have different understandings of some concepts. We also provide a way to learn such non-unanimous concepts by using a method for learning concepts from a group of teachers. The general idea of non-unanimous concepts is to use the teachers to identify the core of a concept everyone agrees on and what else at least some of the teachers think belongs into the concept. The learning agent also decides what belongs to the concept for itself and whenever it needs to communicate with a group of other agents and needs to be precise it makes use of these three concept aspects by providing additional example objects for what might be misunderstood.

1. Introduction

The ability to communicate with each other is often a necessity for agents in order to achieve both, group and individual, goals. For many purposes the use of a language is a very economic way to convey information between agents. But in order to allow the use of a language, its semantics must be understood by all agents involved, which in itself requires that agents have a common understanding of the concepts that they are communicating about.

The concepts (and their relations) that an agent knows or understands are commonly known as the ontology of this agent. In order to facilitate communication, a common ontology for all agents involved in a task has often been suggested, but this is only a theoretical solution in many cases. Especially in multi-agent systems with agents developed by and representing different groups, there is often no agreement between these groups on a common ontology and even if agreements can be reached, the implemented on-

ologies might still differ, due to different understandings of the agreed-upon concepts by the different groups. To deal with this problem, in the last few years authors suggested to enable agents to learn new concepts for their ontologies and to learn concepts from other agents to create a better basis for communication (see [8], [5], or [2]). While almost all authors have looked at one agent teaching one other agent, in [1] we presented a general framework that allows an agent to learn a concept from a group of teacher agents.

At first glance, learning from a group of agents instead of a single agent only seems to add potential problems, namely that the teachers might not agree on some aspects of a concept to learn so that it is up to the learning agent to decide on these aspects on its own. And this has the consequence that the concept that the learner has learned is some kind of compromise between the concepts the teachers teach. Would it not be better to learn from just one teacher at a time and to make sure that the learning agent learns exactly the concept of this teacher? But what if this learner has to communicate with several other agents? Naturally, the learner could learn the necessary concepts from each of these agents one by one and then it has an ontology that has concepts like "concept X according to agent Y" and "concept X according to agent Z". As pointed out in [2], it is a rather large effort to create such an ontology. But even more, it does not really solve the problem of how to communicate with all the agents at once (in a kind of broadcast or multicast situation).

Being able to address a group of people is a necessity of human communication. In fact, this paper does exactly this. The way we human beings deal with the fact that our listeners (or readers) might have slightly different interpretations of the concepts we address is to have an idea where everyone agrees and where there is potential for misunderstandings. And then we address the potential misunderstandings by providing in more detail our understanding. In this paper, we present the idea of non-unanimous ontology concepts that allow us to express the range of agreements on a particular concept based on what an agent learns from a group of teacher agents. We modify the approach of [1] to

enable us to learn such non-unanimous concepts that represent a whole spectrum of possible definitions for a concept. The basic idea is to let the learning agent query its teachers regarding all positive and negative examples it receives from them and use the positive examples all agree on (and the other examples as negative examples) to learn the *core* of the new concept. It then uses all positive examples received (and the negative examples not contradicting these examples) to learn the *periphery* of the new concept.

The agent itself then chooses a concept definition that encompasses the core and is itself encompassed by the periphery. When communicating about this particular concept, the agent uses its awareness of the difference between its own definition, the core and the periphery to enhance the usage of the concept name with explicit references to objects in the periphery but not in the core that are relevant to the communication. We provide a case study in ontologies about organizational structure and courses of universities and show on the one hand side the potential for misunderstandings between universities and on the other side how these misunderstandings can be at least reduced if we use non-unanimous concepts for communication.

2 Basic Definitions

In this section, we provide some basic definitions around ontologies and agents on which we will build in the following sections.

2.1 Ontologies and Concepts

A formal definition for ontology has been presented in [6] in which an *ontology* has been defined as a structure $\mathcal{O} := (C, \leq_C, R, \sigma, \leq_R)$. C and R are two disjoint sets with members of C being called *concept identifiers* and members of R are *relation identifiers*. \leq_C is a partial order on C called *concept hierarchy* or *taxonomy* and \leq_R is a partial order on R , named *relation hierarchy*.

$\sigma : R \rightarrow C^+$ is a function providing the argument concepts for a relation such that $|\sigma(r_1)| = |\sigma(r_2)|$ for every $r_1, r_2 \in R$ with $r_1 \leq_R r_2$ and for every projection π_i ($1 \leq i \leq |\sigma(r_1)|$) of the vectors $\sigma(r_1)$ and $\sigma(r_2)$ we have $\pi_i(\sigma(r_1)) \leq_C \pi_i(\sigma(r_2))$. If $c_1 \leq_C c_2$ for $c_1, c_2 \in C$, then c_1 is called a *subconcept* of c_2 and c_2 is a *superconcept* of c_1 . Obviously, the relation \leq_C is supposed to be connected with how concepts are defined. In the literature, taxonomies are often build using the subset relation, i.e. we have

$$c_i \leq_C c_j \text{ iff for all } o \in c_i \text{ we have } o \in c_j.$$

This definition of \leq_C produces a partial order on C as defined above and we will use this definition in the following for the ontologies that our agents use.

Concepts often are seen as collections of objects that share certain *feature* instantiations. In this work, for an on-

tology \mathcal{O} we assume that we have a set of features $\mathcal{F} = \{f_1, \dots, f_n\}$ and for each feature f_i we have its domain $D_i = \{v_{i1}, \dots, v_{im_i}\}$ that defines the possible values the feature can have. Then an object $o = ([f_1 = v_1], \dots, [f_n = v_n])$ is characterized by its values for each of the features (often one feature is the identifying name of an object and then each object has a unique feature combination). By \mathcal{U} we denote the set of all (possible) objects. In machine learning, often every subset of \mathcal{U} is considered as a concept. In this work we want to be able to characterize a concept by using feature values. Therefore, a *symbolic concept* c_k is denoted by $c_k([f_1 = V_1], \dots, [f_n = V_n])$ where $V_i = \{v'_{i1}, \dots, v'_{ij_i}\} \subseteq D_i$ (if $V_i = D_i$ then we often omit the entry for f_i). An object $o = ([f_1 = v_1], \dots, [f_n = v_n])$ is *covered* by a concept c_k , if for all i we have $v_i \in V_i$. In an ontology according to the definition above, we assign a concept identifier to each symbolic concept that we want to represent in our ontology.

2.2 Agents

A general definition that can be instantiated to most of the views of agents in literature sees an agent $\mathcal{A}g$ as a quadruple $\mathcal{A}g = (Sit, Act, Dat, f_{\mathcal{A}g})$. Sit is a set of situations the agent can be in, the representation of a situation naturally depending on the agent's sensory capabilities, Act is the set of actions that $\mathcal{A}g$ can perform and Dat is the set of possible values that $\mathcal{A}g$'s internal data areas can have. In order to determine its next action, $\mathcal{A}g$ uses $f_{\mathcal{A}g} : Sit \times Dat \rightarrow Act$ applied to the current situation and the current values of its internal data areas.

As we want to focus on the knowledge representation used by agents and how this is used for communication, we have to look more closely at Dat . We assume that every element of Dat of an agent $\mathcal{A}g$ contains an ontology area $\mathcal{O}_{\mathcal{A}g}$ that represents the agent's view and knowledge of concepts. There might be additional data, beyond features, that the agent requires from time to time, about concepts and this data is naturally also represented in Dat . Also, there will be additional data areas representing information about the agent itself, knowledge about other agents and the world that the designer of the agent may want to be represented differently than in $\mathcal{O}_{\mathcal{A}g}$.

3 Learning Concepts from Several Teachers

In this section we provide a brief description of the multi-agent concept learning we presented in [1]. As already stated, we have developed a method that demonstrates how an agent can learn new concepts for its ontology with the help of several other agents. This assumes that not all agents have the same ontology. We additionally assume that there are only some base features $\mathcal{F}_{base} \subseteq \mathcal{F}$ that are known

