

Towards an Ontological Account of Agent-Oriented Goals

Renata S.S. Guizzardi¹, Giancarlo Guizzardi^{2,3},
Anna Perini¹, and John Mylopoulos⁴

¹ ITC-irst, Trento-Povo, Italy
{souza,perini}@itc.it

² Department of Computer Science, UFES, Vitória-ES, Brazil

³ Laboratory of Applied Ontologies (ISTC-CNR), Trento, Italy
guizzardi@loa-cnr.it

⁴ Department of Computer Science, University of Toronto, Canada
jm@cs.toronto.edu

Abstract. The software agent paradigm has received considerable attention recently, both in research and industrial practice. However, adoption of this software paradigm remains elusive in software engineering practice. We claim that part of the adoption problem lies with the fact that mentalistic and social concepts underlying agents are subjective and complex for the average practitioner. Specifically, although there are many efforts related to the topic coming from philosophy, cognitive sciences and computer science, a uniform and well-founded semantic view on these concepts is currently lacking. This work extends an existing upper-level ontology and offers it as a foundation for evaluating and designing agent-oriented modeling languages. In particular, the paper focuses on the concept of goal, aiming at disambiguating its definition, discussing its different manifestations, and clarifying its relation to other important agent-related concepts. For that, we examine how goals are conceived and used according to some relevant literature on agent-orientation. In addition, related work on akin fields, especially philosophy and AI are used as a basis for the proposed ontological extensions.

1 Introduction

The agent paradigm is shaped by developments from several research areas, such as Distributed Computing, Software Engineering (SE), Artificial Intelligence (AI), and Organizational Science [Wooldridge and Jennings, 1995]. An AI perspective of agents focuses on their cognitive (or mentalistic) properties, e.g. beliefs, goals and commitments. On the other hand, an SE perspective emphasizes its potential for designing open, distributed, dynamically reconfigurable software, with only lip service paid to mentalistic or cognitive underpinnings. However, given the potential of using agents both for conceptual modeling and system development, such properties may indeed be central to both domain analysis and system development. For instance, understanding agent goals, perceptions and beliefs leads to a deeper understanding of values and strategies

adopted in an organization, thereby contributing to the conception of effective information systems [Guizzardi, 2006] [Dignum, 2004].

Several agent cognitive models are proposed in the AI literature, the best-known among them being the BDI model [Rao and Georgeff, 1991]. This model focuses on three basic mental components of agents: belief, desire and intention. *Belief* refers to knowledge the agent has about the environment and about other agents with whom she interacts. *Desire* refers to the “will” of an agent towards a specific goal, although she might never actually pursue these goals. Finally, *intention* entails specific plans and commitments to achieve specific goals. A different model characterizes the state of an agent as a combination of mental components such as *beliefs*, *capabilities*, *choices*, and *commitments* [Shoham, 1993]. Besides these well-known models, much work related to AI theory, philosophy and cognitive sciences underlies the definition of such cognitive notions, guiding their practical use for modeling and developing multi-agent systems. Among them is the early work of Bratman [Bratman, 1987] on goals, beliefs, intentions and related mental models, and the contribution of Castelfranchi and colleagues on delegation [Castelfranchi and Falcone, 1998], dependency [Conte and Castelfranchi, 1995] and commitments [Castelfranchi, 1995]. In addition to these, work on conceptual formalization through the use of ontologies also provides valuable contribution in this respect [Guizzardi, 2006] [Bottazzi and Ferrario, 2005] [Masolo et al., 2003].

This work constitutes a follow-up to earlier efforts on defining a uniform conceptualization for agent-oriented systems. We aim at investigating diverse definitions and treatments of the agent mentalistic concepts and - where possible - merge these through amalgamation or compromise. In [Guizzardi, 2006], we propose an ontology of agent and related concepts, based on previous results and guidelines presented in [Guizzardi, 2005]. In this earlier work, we use the proposed ontology to guide the understanding and evaluation of modeling languages adopted in the development of agent-oriented knowledge management systems. Regarding the use of ontologies to support the evaluation and re-design of software engineering modeling languages, the role of the ontology is threefold:

- clarify modeling language concepts;
- evaluate and re-design the notation in order to avoid construct *overload*, *excess*, *redundancy* and *incompleteness* [Guizzardi, 2005];
- in cases where different notations A and B are used, assist in the transformation from one notation to the other, by guiding the mapping between the concepts of language A to those of language B.

In this paper, we specifically focus on the concept of *goal*, aiming at clarifying its meaning and finding out its relations to other basic agent-inspired concepts. Goals are widely used in agent-orientation and related fields, ranging from conceptual goal modeling in Agent Organizations and Requirements Engineering to goal execution in AI Planning and Agent Teamwork. In Agent Organizations, for instance, goals are used to describe the objectives of the organization as a whole, being generally associated with roles, which are then assigned to agents that act on behalf of the organization [Hubner et al., 2002]

[Dignum, 2004] [Esteva et al., 2002]. In a few Requirements Engineering approaches, on the other hand, the concept of goal is the basis for requirements analysis, representing the objectives of different stakeholders, and rationalizing strategic dependencies among these stakeholders [van Lamsweerde, 2000] [Yu, 1995] [Bresciani et al., 2004]. In AI Planning [Ghallab et al., 2004], goal is an essential concept, since this area mainly focuses on computational approaches to the problem of reasoning and deliberating about actions that are intended to fulfill a goal. Finally, the research area of Agent Teamwork generally makes extensive use of Planning techniques to support the cooperation of the team agents in the pursuit of a common goal [Boella et al., 1999] [Yen et al., 2001].

It is important to emphasize that although being the result of careful investigation, this work still represents the first steps in the direction of providing uniform semantics to the concept of goal. The remaining of this article is organized as follows: section 2 focuses on the main motivations behind this research initiative; section 3 describes this work's main contribution, by presenting an excerpt of our agent ontology (named UFO-C) and discussing it in comparison with related work; section 4 presents applications of the use of UFO-C to support agent-oriented software engineering; and section 5 finally concludes the paper.

2 Motivation

Concerns with the definition of syntactic and semantic properties of agent-oriented concepts have contributed to the proliferation of research initiatives on metamodels. Many of these works focus on: a) defining organization-centered concepts such as agent, group and roles in order to enable modeling of heterogeneous systems [Odell et al., 2004] [Ferber and Gutknecht, 1998]; b) inter-operating and/or unifying modeling methodologies [Henderson-Sellers et al., 2005] [Perini and Susi, 2005][Bernon et al., 2004]; and c) enabling agent-oriented modeling through the use of CASE tools [Perini and Susi, 2005]. These works have been generally based on a bottom-up strategy, constructing their conceptualizations by abstracting concepts that are present in existing languages, methodologies and formalisms. Modeling Language are sometimes the result of a negotiation process, and commonly incorporate features motivated by reasons other than being truthful to the domain in reality being represented (e.g., increasing computational efficiency, providing compatibility to a computational paradigm, facilitating the translation to a specific implementation environment). Thus, one of the disadvantages of a bottom-up approach such as the ones just mentioned is to incorporate in the produced metamodel many of these improper features.

In contrast, the objective of our research is to employ theories developed in disciplines such as cognitive science, philosophy, as well as social sciences to uncover the kinds of individuals that constitute the social reality as well as to understand the ontological nature of these entities. As a result we aim at producing a Foundational Ontology that explicitly represents these entities.

As argued in [Guizzardi, 2005], the quality of a conceptual modeling language can be systematically evaluated by comparing, on one hand, a metamodel of this

language, and on the other hand, an explicit representation of the subject domain this language is supposed to represent, i.e., a domain ontology. In the ideal case, these two entities are isomorphic and share the same set of logical models. To put it simple terms, in this ideal situation the language is not only able to represent all the relevant concepts of the subject domain at hand, preserving all their properties, but the user of the language can identify in an unambiguous manner what are the domain concepts represented by each of the language's modeling constructs. Thus, if we have a concrete model representing the subject domain, this model can be used for evaluating and (re)designing modeling languages in that domain.

The work described here can then be seen as complementary to the effort of developing metamodels for agent-oriented concepts. First, it can be used to systematically evaluate and perhaps propose modification to these metamodels so that they become isomorphic to this ontology. Second, once the mapping between elements in a metamodel (syntactic elements) and in an ontology are established, the elements of the latter can be used to provide real-world semantics for the elements of the former. In other words, the interpretation mapping from a language construct to a category in an ontology establishes the meaning of that construct in terms of the real-world element represented in that ontology. If the ontology itself is described in a formal language (see [Guizzardi, 2005], this linking also enables the definition of a formal semantics for this language. In this article, however, we do not intend to formally characterize the proposed ontology and, for this reason, the UML diagrams depicting fragments of this ontology are intended here for presentation only. This is mainly due to the fact that this ontology (UFO-C) is still in preliminary stage of development and that we defend the position that we should first concentrate on understanding a certain conceptualization before formally describing it.

3 The UFO Ontology

In this section, we present our conceptualization of goal and related concepts. We base this conceptualization on the *UFO (Unified Foundation Ontology)* defined in [Guizzardi, 2005] [Guizzardi and Wagner, 2005] [Guizzardi, 2006], extending it when necessary.

The UFO ontology is divided into three incrementally layered compliance sets: 1) UFO-A defines the core of UFO, as a comprehensive ontology of *endurants*; 2) UFO-B defines - as an increment to UFO-A - terms related to *perdurants*¹; and 3) UFO-C defines - as an increment to UFO-A and UFO-B - terms related to the spheres of *intentional and social entities*. In this paper, we focus on the UFO-C ontology, referring to the other ontologies only to provide definitions when needed. The ontologies are described here in natural language, and illustrated with the aid of UML class diagrams. Thus, UML is not intended here for formalization purposes but rather for facilitating the visualization of the concepts.

¹ Endurants and perdurants intuitively correspond to *objects* and *events* (respectively) as understood in natural language.

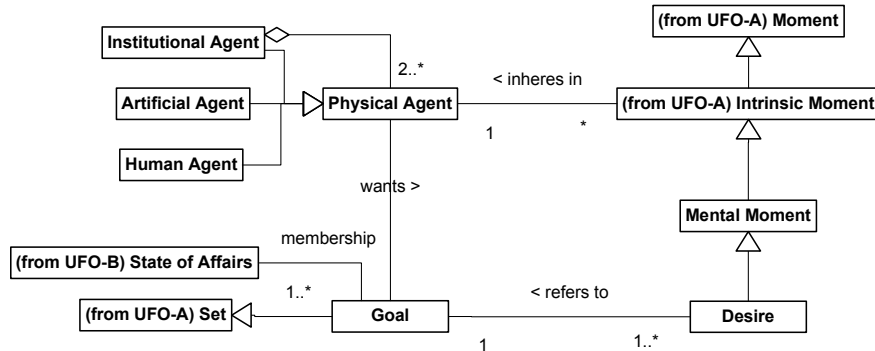


Fig. 1. UML diagram representing a fragment of UFO-C

For an in depth discussion and formal characterization of UFO-A, one should refer to [Guizzardi, 2005]. The formalization of UFO-B and UFO-C is planned as future work, once the semantics of the concepts comprising these ontologies is fully comprehended.

Figure 1 shows an excerpt of UFO-C defining a *goal* in relation to two other important concepts, namely *desire* and *physical agent*. In general, we say that a physical agent has a goal, and this goal is related to the agent's desire. *Desire* here is defined as a *mental moment*, which specializes the concept of intrinsic moment from UFO-A. UFO-A defines a *moment* as an entity whose existence is existentially dependent on another entity. This Husserlian notion of moments is akin to what is termed trope, abstract particular, property instance, or mode in the literature. An intrinsic moment is a special kind of moment that is existentially dependent on one single individual (e.g., the color of an apple depends of the existence of the apple itself). Examples of intrinsic moments of a physical agent are age, height and address. Mental moment is a specialization of intrinsic moment referring to mental components of a physical agent, such as belief, desire, intention, and perception. Summing up, a desire is conceived as a mental moment, which is existentially dependent on a particular agent, being an inseparable part of its mental state.

Fig. 1 also defines goal as a set of states of affairs (i.e. a set of world states). This choice has some important implications that deserve debate. We noted two main views on goals in the AI and agent-orientation literature. On one hand, a goal may be seen as a specialization of the concept of mental moment. On the other hand, a goal may be treated as a state of affairs (or set of state of affairs). However, in agent-orientation, both views are possible. In fact, it is common to find works that treat them interchangeably [Conte and Castelfranchi, 1995] [Rao and Georgeff, 1991]. We believe that the reason behind this confusion is the fact that in artificial systems, both the mental states of the agents composing the system and the state of the world are explicit and sometimes treated as the same thing. This approach is illustrated in the context of the CAST architecture supporting Agent Teamwork, where the authors affirm that the team agents

develop an “overlapping shared mental model, which is the source for team members to reason about the states and the needs of others” [Yen et al., 2001]. However, when we consider hybrid systems involving artificial and human agents, we cannot assume anymore the explication of mental moments. Instead, beliefs, intentions and perceptions remain inside the human agent’s mind. With this discussion, however, we do not intend to say that mental moments cannot be considered and represented in an agent-oriented model. What we find important is the realization that there are two distinct concepts involved here: one external and another one internal to the agent. The external concept regards a state of affairs desired by an agent (here called goal), and the internal one is the desire itself, which is part of the agent’s mental state.

In this work, we commit to the definition of goal as a set of states of affairs because we find it more flexible from several different perspectives. For instance, it allows a more flexible view of organizational goals. For now, UFO-C views an organization as an institutional agent constituted by a number of other (physical, artificial or institutional) agents (refer to Fig. 1). Thus, a goal could be seen as a mental moment associated with a sort of collective mind, in the sense of Searle. Nevertheless, [Bottazzi and Ferrario, 2005] see an organization as an abstract social concept, which is separate from the collective body of agents that composes it. Taking this approach leads to the impossibility of considering a goal as a mental moment, since an organization here cannot be conceived as having a mind. Defining goal as a set of states of affairs accommodates both views, i.e. it is always possible to say that an organization (or institutional agent) has a goal. Since our account for organization and related concepts is still preliminary, we prefer to take this more flexible approach².

Another reason for this choice comes from the fact that some ontological theories do admit part-of relations applied to states of affairs but not to moments. Thus, having goal as a mental moment would disallow *goal decomposition* (defined in to Figure 2). However, several approaches foresee the need to refine goals by decomposing it into sub-goals. This is applied, for instance, by some Agent Organization methodologies (e.g. MOISE+ [Hubner et al., 2002] and OperA [Dignum, 2004]) to understand the goals of particular roles by refining general organizational goals. Moreover, this is also common practice for some Requirements Engineering approaches, which use goal decomposition to analyze objectives of particular stakeholders and/or to derive the requirements of supporting information systems [van Lamsweerde, 2000] [Bresciani et al., 2004] [Yu, 1995].

Fig. 2 shows that according to UFO-C a goal decomposition is a kind of *basic formal relation* (from UFO-A) between goals, which is defined in terms of a binary mereological (part-of) relation between these goals. A Goal decomposition groups several sub-goals related to the same super-goal. In other words, suppose

² We do not include here an in depth discussion on organizational goals. In order to be complete, the concepts of roles, commitments/claims and norms would have to be considered. [Guizzardi, 2006] presents our initial views on this topic. However, more remains to be done in the future and is out of the scope of this paper.

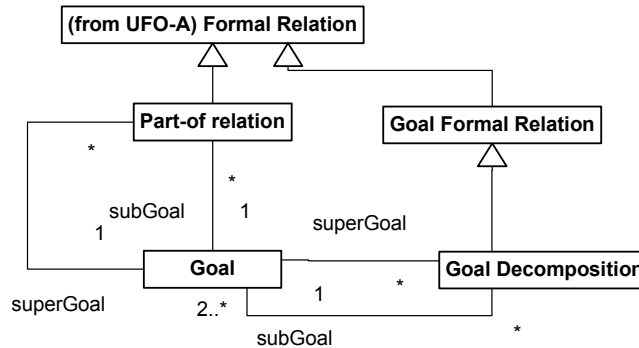


Fig. 2. Goal decomposition

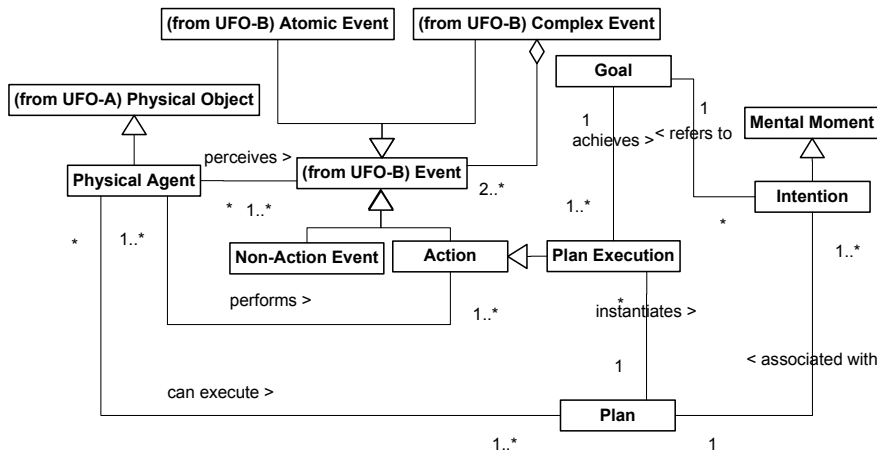


Fig. 3. Differentiating between Goal and Plan

that goals G1 and G2 are *parts of* the super-goal G. Thus, we can say that there is a *goal decomposition relation* between G (as a super-goal) and G1 and G2 (as sub-goals).

Figure 3 focuses on the relation of goal to the actual plan executed to achieve this goal. This leads us to the distinction made in UFO-B between *action* and *non-action events*. The former refers to events created through the action of a physical agent, while the latter are typically events generated by the environment itself and perceived by the agents living in it.

A *plan execution* is an intended execution of one or more actions, and is therefore a special kind of action event. In other words, a plan execution may be composed of one or more ordered action events, targeting a particular outcome of interest to the agent. These action events may be triggered by both action and non-action events perceived by the agent. Besides, a plan execution instantiates

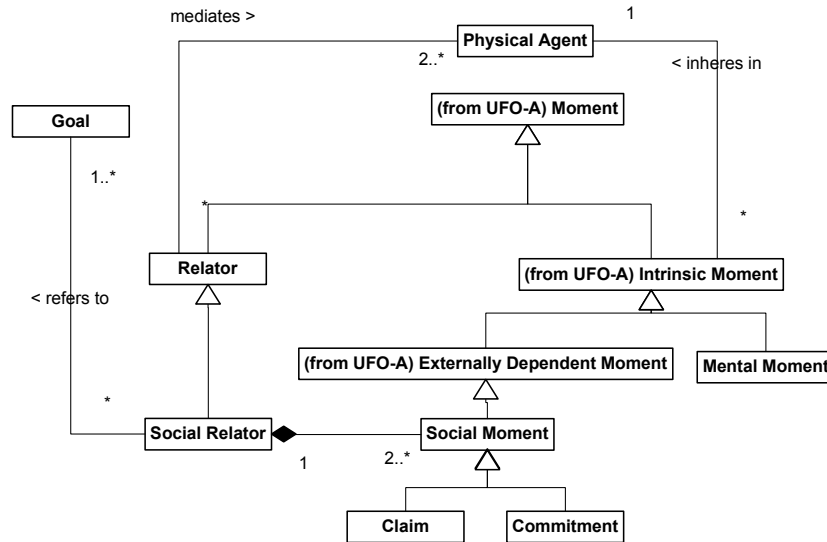


Fig. 4. Commitments and Claims

a plan (or plan type). Thus, when we say that a physical agent executes a plan, we actually mean this agent creates the action events previously specified in the plan. Furthermore, such plan is connected to the agent through a mental moment referred to as *intention*. Agent's intention directly leads to the adoption of certain goals, and is associated with a plan, i.e. a specific way of achieving this specific goal. In fact, the association to a plan is the main differentiation between desire (as in Fig. 1) and intention. To put it differently, while a desire refers to a wish of the agent towards a particular set of state of affairs, an intention actually leads to action towards achieving this goal [Rao and Georgeff, 1991] [Conte and Castelfranchi, 1995] [Boella et al., 1999].

The difference between goal and plan is an important one, not always clear in existing works. For instance, some AI Planning techniques define goals as tasks the system must perform [Ghallab et al., 2004]. MOISE+ [Hubner et al., 2002] also adopts a more operational view on goals as being the tasks performed by the agents of an organization. Examples of work that do make this differentiation include the KAOS [van Lamsweerde, 2000] and i*/Tropos [Yu, 1995] [Bresciani et al., 2004] requirement engineering approaches.

Figure 4 clarifies UFO-C's view on the social concepts of commitment and claim, highly associated with the concept of goal and thus, presenting important contribution to enable the understanding and modeling goal adoption.

First, it is important to have a more detailed view of how UFO-A specializes the concept of moment. Moments can be specialized into intrinsic moments and relators. The former refers to a moment that is existentially dependent on one single individual. In contrast, a relator is a moment that is existentially dependent on more than one individual (e.g., a marriage, an enrollment between a

student and an educational institution). A relator is an individual capable of connecting or mediating entities [Guizzardi, 2005]. For example, we can say that John is married to Mary because there is an individual marriage relator that existentially depends on both John and Mary, thus, mediating the two. Likewise, we can say that Lisa works for the United Nations because there is an employment relator mediating Lisa and the United Nations.

An externally dependent moment is a special kind of intrinsic moment that although inhering in a specific individual, also existentially depends on another one. The employee identifier is an example of externally dependent moment, since although inherent to the employee, is also dependent on the organization where this employee works. The UFO-C notion of *social moment* is a specialization of the concept of externally dependent moment and includes the concepts of *commitment* and *claim*. When two physical agents agree to accomplish goals to one another, a commitment/claim pair is generated between them. These concepts are highly important to regulate the social relations between members of an organization, being related to the deontic notions defined for example in ISLANDER [Esteva et al., 2002] and OperA [Dignum, 2004]. A pair commitment/claim constitutes a *social relator*, which is a particular type of UFO-A *relator*. Fig. 4 also shows that a social relator refers to a goal. When a physical agent A commits to a physical agent B, this means that A adopts a goal of B. Conversely, the social relator created between A and B state that B has the right to claim the accomplishment of this specific goal to A.

Dependency is a common relation explored in Requirements Engineering approaches (e.g. i* [Yu, 1995] and Tropos [Bresciani et al., 2004]) and Agent Organization methodologies (e.g. OperA [Dignum, 2004]). However, the distinction between *dependency* and *delegation* is usually not made. Figure 5 depicts this important distinction. The first difference regards the fact that while a dependency constitutes a *formal relation*, a delegation consists of a *material relation* [Guizzardi, 2005]. This distinction between formal and material relations is elaborated in UFO-A. A formal relation is either an internal relation holding directly between two entities (e.g., instantiation, parthood, inherence), or it is reducible to an internal relation between intrinsic moments of involved relata. Examples of formal relations of the latter type is Lisa ‘is older than’ Mike, and John ‘is taller than’ Mary. In both of cases, these relations are reducible to comparative formal relations between intrinsic moments of the involved relata (individual heights and ages). A material relation, in contrast, cannot be reduced in such a way and has real material content. For a material relation to take place between two or more individuals, something else needs to exist, namely, a relator connecting these entities. The relations ‘married to’ and ‘works for’ aforementioned are examples of material relations founded by relators of type marriage and employment, respectively.

Let us examine this difference in further detail. Fig. 5 shows that a dependency connects two physical agents (a depender and a dependee) and a goal (a dependum). An agent A (the depender) depends on an agent B (the dependee) regarding a goal G if G is a goal of agent A, but A cannot accomplish G, and

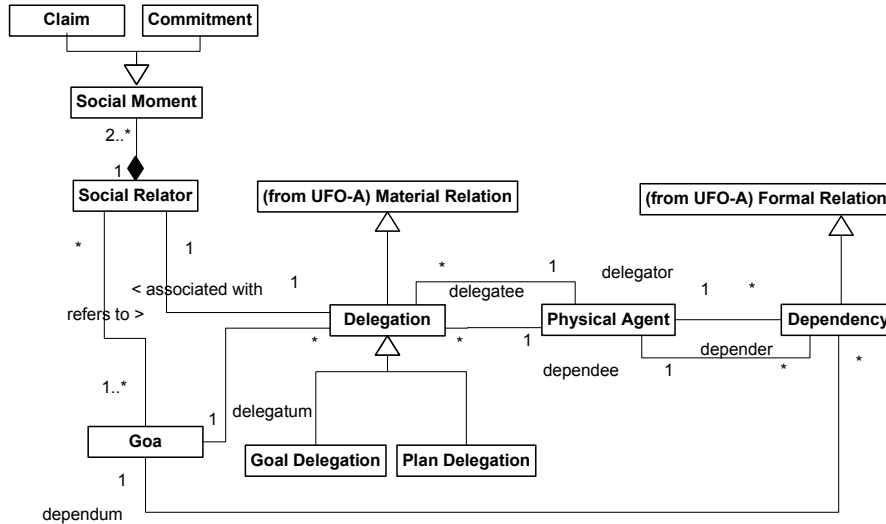


Fig. 5. Goal Delegation and Dependency

agent B can accomplish G. A delegation is associated with a dependency but it is more than that. As a material relation, it is founded on something more than its connected elements. In this case, the connected elements are two physical agents (delegator and delegatee) and a goal (delegatum), and the foundation of this material relation is the social relator (i.e. a commitment/claim pair) established between the two physical agents involved in this delegation. In other words, when agent A delegates a goal G to agent B, besides the fact that A depends on B regarding G, B commits herself to accomplish G on behalf of A, thus adopting the goal of A. Goal and plan delegation refer to what Castelfranchi and Falcone define as open and close delegation [Castelfranchi and Falcone, 1998], meaning that the former leaves the decision regarding the strategy towards goal accomplishment to the dependee. The latter rather prescribes a specific strategy (i.e. a plan) the dependee should adopt towards achieving the delegated goal.

To illustrate the difference between dependency and delegation, consider the following case. Suppose John is a program committee member of a certain conference and that he received from Paul (the conference program chair) an article X to review. Suppose that John cannot review this article by himself, since there are some aspects of the article which are outside his field of competence. Now, suppose that George is a colleague of John who is knowledgeable exactly in those aspects that John needs to review article X. In this case, we could say that John depends on George to review article X. Notice, however, that this relation between John and George can be reduced to relations between the goals and capabilities of these individual agents. Moreover, this relation does not even require that the related agents are aware of this dependence. This is certainly not the case for the relation between Paul and John. As the program committee chair, Paul depends on John to review article X. However, in this case, not only

they are both aware of this dependence but there is the explicit commitment of John to Paul to review article X. In other words, the delegation of Paul to John to review article X cannot be reduced to relations between their intrinsic moments, but it requires the existence of a certain relator (a commitment/claim pair) that founds this relation.

Figure 6 depicts four specializations of the category of goals, namely *depended*, *collaborative*, *shared*, and *conflicting* goals, typical of agent-oriented theoretical and practical works [Boella et al., 1999] [Bresciani et al., 2004] [Yu, 1995] [Conte and Castelfranchi, 1995] [Dignum, 2004] [Yen et al., 2001]. Such distinctions reflect different ways a goal can participate in relations with agents and with other goals, i.e., different roles a goal can play in the scope of certain relations.

Depended goal is the kind already discussed in the context of Fig. 5, i.e. a goal which is a dependum of a dependency relation between two physical agent individuals: the depender and the dependee. In fact, the dependency relation depicted in Fig. 5 is generalized in this model to the category of *Goal Formal Relation involving agents*, which is always a ternary relation between two agents and a goal. A *shared goal* is a set of states of affairs intended at the same time by two different physical agent individuals. In other words, two agents share a goal if they both have individual desires that refer to that same goal. A *collaborative goal* is a special kind of shared goal. A collaborative goal G is the subject of a potential collaboration relation between agents A and B if: (i) G is shared by A and B; (ii) there are at least two sub-goals G1 and G2 of G such that A wants G1 but depends on B to accomplish it, and B wants G2 but depends on A to accomplish it. In other words, a collaborative goal is always composed of at least two depended goals. To illustrate collaborative goals, suppose agents A and B have a shared goal of “taking a heavy table out of the room”. This goal can be decomposed in two sub-goals referring to carrying out each side of the table, which can be respectively adopted by A and B. In this case, one agent depends on the other to accomplish their shared super-goal, thus this goal can only be attained in collaboration. Finally, two goals are *conflicting* if they cannot be achieved at the same time. For instance, taking two conflicting goals G1 and G2, the accomplishment of goal G1 would preclude the achievement of goal G2 and vice-versa. In other words, if we take any two state of affairs S1 and S2, such that S1 satisfies G1 and S2 satisfies G2, we have that S1 and S2 cannot obtain simultaneously (i.e., in the same world or world history).

Note that the definition of these different types of goal also influenced our choice for preferring the definition of goal as a set state of affairs rather than a mental moment. Such definitions are actually facilitated by this choice. For example, a shared goal can be seen as a state of affairs referenced (i.e. intended) at the same time by two physical agents. If it were to be defined as a mental moment, we would have to be careful to talk about shareability, since each agent has its own mental moment and thus, the goals would not be effectively shared. Instead, we would have anyway to assume that these two agents having distinct goals would aim at the same set of state of affairs.

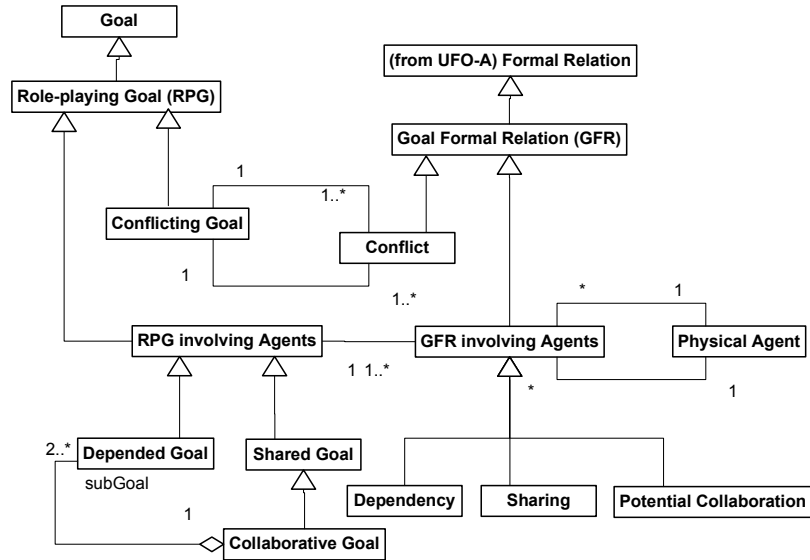


Fig. 6. Different Roles played by Goals in Goal Formal Relations

4 Applications of UFO-C to Support Agent-Oriented Software Engineering

The UFO-C ontology is aimed at providing a consistent understanding of the concepts involved in agent-orientation. In particular, with respect to agent-oriented software engineering, we hope to provide support to: i) clarifying the concepts underlying modeling languages; ii) evaluating and (re)designing modeling languages to make it more consistent and accessible to the user; and iii) interoperating different modeling languages. Figures 7 and 8 present applications of the UFO-C ontology to achieve all these aims.

The Tropos actor diagram of Fig. 7 depicts the main agents and dependencies of a paper review scenario. In the Tropos original language, dependencies and delegations were overloaded in the concept of dependency. In other words, an analysis of this language in light of UFO-C has shown that in many Tropos models, what is called dependency is actually a delegation. In these cases, besides a dependency between agents A and B, the relationship also implies that agent B commits to deliver the dependum (e.g. a goal) to agent A. The diagram of Fig. 7 illustrates this difference. Most relationships shown in the diagram are delegation, for instance, the **PC Chair** depends on the **PC Member** to accomplish the goal of **reviewing papers**. And in this case, the **PC Member** commits herself to this goal. Thus, this is a case of delegation. We can then say that the **PC Chair** delegates the goal of **reviewing papers** to the **PC Member**. On the other hand, the relationship between the **Conference Chair** and the **Paper Author** is an example of dependency. While the former depends on the

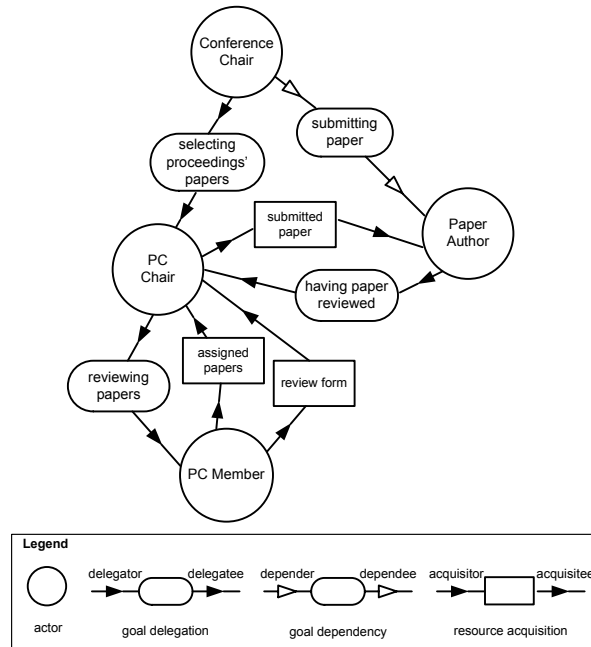


Fig. 7. Tropos actor diagram illustrating a paper review scenario

latter to submit papers in order to guarantee the realization of the conference, she cannot assume that the **Paper Author** will actually do it. In other words, it is possible that no paper is submitted to the conference because there is no commitment from specific paper authors to do so.

Understanding both concepts of dependency and delegation with the aid of UFO-C led to the decision of redesigning Tropos to incorporate both dependency and delegation. This has solved a problem of construct overload, which could prevent the correct understanding of the nature of the relationships while at the same time, has given more expressivity to the language. Benefits gained by considering both concepts are for instance:

- supporting analysts to reason about different *degrees of vulnerability*. In general, a dependency makes the depender more vulnerable than a delegation. This happens because in a delegation, the dependee has an explicit commitment toward the depender in respect to the goal to be accomplished. In a dependency, however, this is not the case. In fact, sometimes, the dependee is not even aware of this dependency (e.g. the dependency between **Conference Chair** and **Paper Author** mentioned above). Consequently, if a goal is depended but not delegated, the depender is less certain of its accomplishment.
- allowing the understanding when the dependee can be *subjected to sanctions*. In the case of a delegation, which assume a commitment from the dependee

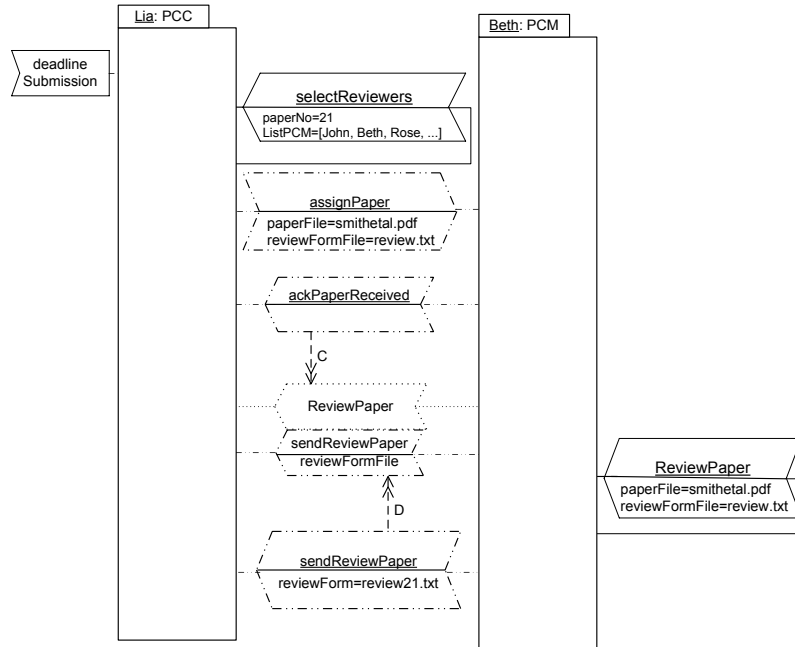


Fig. 8. AORML's interaction sequence diagram

towards the depender, sanctions may be applied in case the dependee fails to accomplish the goal she had committed to.

- enabling the analyst to find during the analysis, dependencies which can be *opportunities for the establishment of latter delegations*. In other words, if there are dependencies that are critical for the accomplishment of the goals of an agent, then this agent can seek to obtain a commitment from the dependee, lowering her degree of vulnerability. Also in organizational modeling, this analysis can be helpful in the (re)design of the commitments of organizational roles in order for organizational goals to be accomplished more efficiently.

Fig. 8 depicts an AORML (Agent-Object-Relationship Modeling Language) interaction sequence diagram, showing the interactions between **PC Chair** and **PC Member** to accomplish the goal of reviewing papers. This diagram illustrates how UFO-C may assist the interoperation of two notations, namely Tropos and AORML. The delegation between the **PC Chair** and **PC Member** previously analyzed is mapped into an AORML commitment construct during interaction modeling. The **ReviewPaper** commitment is created after the **PC Member** acknowledges that she has received the papers assigned to her for review (view create arrow coming from the **ackPaperReceived** message to the **ReviewPaper** commitment). The **ReviewPaper** commitment has a message attached to it (i.e. a **sendReviewPaper** message), indicating that this commitment is fulfilled if the **PC Member** submits a message of this kind to the

PC Chair. Otherwise, this commitment is broken, giving the **PC Chair** the right to sanction. Fortunately, in this case, the **PC Member** has fulfilled her responsibility (refer to **sendReviewPaper** message which discharges the **ReviewPaper** commitment).

5 Conclusion

This paper presented excerpts of the UFO-C ontology specifically concerned with the concept of *goal*. The UFO-C ontology itself is an extension of the UFO-A and UFO-B ontologies, which together lay out the foundations for domain-independent concepts such as objects, processes, types, properties, state of affairs as well as their relations, such as instantiation, partonomy, participation, inheritance, causality, among many others. In this manner, UFO-C concepts of *agent* and *social moment* can, for instance, be conceived as extension of the UFO-A concepts of *object* and *externally dependent intrinsic moment*, respectively, thus, inhering not only their characterizing ontological meta-properties (e.g., existential (in)dependency, unity), but also the complete formalization of the theories regarding these notions.

We are aware of existing formal addresses of the notion of goal, such as, for example, the logics proposed in [Dastani et al., 2006] [van Riemsdijk et al., 2005] and [Cohen and Levesque, 1990]. Although we also intend in a second stage of this enterprise, to completely formalize the theories put forth here, this work differs from these “logics of goals” in a manner of emphasis. The aim in this particular paper is not to define a formal language that can be used to reason about goals. In contrast, the focus is on the real-world semantics of this concept, i.e., to understand the meaning of the notion of goal by making explicit its ontological meta-properties as well as its relations to other ontological categories (such as state of affairs, mental and social moments, social commitments and claims, objects, processes, etc.) for which a number of formal theories have already been developed in areas such as philosophy and cognitive science.

Several research areas permeating the agent-oriented paradigm make use of the term *goal*. Examples of these areas include Agent Organizations, Requirements Engineering, AI Planning and Agent Teamwork. However, further analysis of these different usages of the term goal in these areas shows that it has been used to represent a number of different and sometimes incompatible notions. In this article, we make use of a comprehensive network of ontological categories to make explicit which ontological elements are referred by these different senses of the term goal used in the literature, as well as the relations they bear to each other.

Finally, although several related works have already been analyzed and discussed, our research agenda for the future includes the study of other works that may provide valuable input to enhance the present conceptualization. In parallel, we aim at extending UFO-C even further, deepening our understanding of other important concepts (for instance, those of action and event, and especially communicative action and communicative event, commitment and claim, etc.). Moreover, we intend to apply UFO-C to evaluate and re-design diverse modeling

languages, proceeding with our previous effort in this direction, while profiting from the advances in the ontology to provide more consistent and semantically uniform languages.

References

- [Bernon et al., 2004] Bernon, C., Cossentino, M., Gleizes, M., Turci, P., and Zambonelli, F. (2004). A Study of some Multi-agent Meta-models . In Odell, J., Giorgini, P., and Mller, Jrg, P., editors, *Agent-Oriented Software Engineering V*, volume 3382 of *LNCS*, pages 62–77. Springer-Verlag, Berlin, Germany.
- [Boella et al., 1999] Boella, G., Damiano, R., and Lesmo, L. (1999). A Utility Based Approach to Cooperation among Agents. In *Proceedings of the Workshop on Foundations and applications of collective agent based systems (ESSLLI'99)*, Utrecht, The Netherlands.
- [Bottazzi and Ferrario, 2005] Bottazzi, E. and Ferrario, R. (2005). A Path to an Ontology of Organizations. In *Proceedings of the Workshop on Vocabularies, Ontologies and Rules for The Enterprise (VORTE'05)*, Enschede, The Netherlands. Centre for Telematics and Information Technology (CTIT).
- [Bratman, 1987] Bratman, M. E. (1987). *Intentions, Plans, and Practical Reason*. Harvard University Press.
- [Bresciani et al., 2004] Bresciani, P., Giorgini, P., Giunchiglia, F., Mylopoulos, J., and Perini, A. (2004). Tropos: An Agent-Oriented Software Development Methodology. *International Journal of Autonomous Agents and Multi Agent Systems*, 8(3):203–236.
- [Castelfranchi, 1995] Castelfranchi, C. (1995). Commitments: From Individual Intentions to Groups and Organizations. In *Proceedings of the First International Conference on Multi-Agent Systems*, Cambridge, MA, USA. AAAI-Press and MIT Press.
- [Castelfranchi and Falcone, 1998] Castelfranchi, C. and Falcone, R. (1998). Towards a Theory of Delegation for Agent-Based Systems. *Robotics and Autonomous Systems*, 24(24):141–157.
- [Cohen and Levesque, 1990] Cohen, P. R. and Levesque, H. J. (1990). Intention is Choice with Commitment. *Artificial Intelligence*, 42(3):213–261.
- [Conte and Castelfranchi, 1995] Conte, R. and Castelfranchi, C. (1995). *Cognitive and Social Action*. UCL Press.
- [Dastani et al., 2006] Dastani, M., van Riemsdijk, M. B., and Meyer, J.-J. (2006). Goal Types in Agent Programming. In *Proceedings of the 17th European Conference on Artificial Intelligence*, pages 220–224, Riva del Garda, Italy. IOS Press.
- [Dignum, 2004] Dignum, V. (2004). *A Model for Organizational Interaction: Based on Agents, Founded in Logic*. PhD thesis, Utrecht University, The Netherlands.
- [Esteva et al., 2002] Esteva, M., Padget, J., and Sierra, C. (2002). Formalizing a Language for Institutions and Norms. In Meyer, J.-J. C. and Tambe, M., editors, *Intelligent Agents VIII*, volume 2333 of *LNAI*, page 348 to 366. Springer-Verlag, Berlin, Germany.
- [Ferber and Gutknecht, 1998] Ferber, J. and Gutknecht, O. (1998). A meta-model for the analysis and design of organizations in multi-agent systems. In *ICMAS '98: Proceedings of the 3rd International Conference on Multi Agent Systems*, page 128, Washington, DC, USA. IEEE Computer Society.
- [Ghallab et al., 2004] Ghallab, M., Nau, D., and Traverso, P. (2004). *Automated Planning: Theory and Practice*. Morgan Kaufmann, Sao Mateo, CA, USA.

- [Guizzardi, 2005] Guizzardi, G. (2005). *Ontological Foundations for Structural Conceptual Models*. PhD thesis, University of Twente, The Netherlands.
- [Guizzardi and Wagner, 2005] Guizzardi, G. and Wagner, G. (2005). Some Applications of a Unified Foundational Ontology in Business Modeling. In Rosemann, M. and Green, P., editors, *Ontologies and Business Systems Analysis*, pages 345–367. Idea Group, London, UK.
- [Guizzardi, 2006] Guizzardi, R. S. S. (2006). *Agent-oriented Constructivist Knowledge Management*. PhD thesis, University of Twente, The Netherlands.
- [Henderson-Sellers et al., 2005] Henderson-Sellers, B., Debenham, J., Tran, N., Cossentino, M., and Low, G. (2005). Identification of Reusable Method Fragments from the PASSI Agent-Oriented Methodology. In Kolp, M., Bresciani, P., Henderson-Sellers, B., and Winikoff, M., editors, *Agent-Oriented Information Systems III*, volume 3529 of *LNCS*, pages 95–110. Springer-Verlag, Heidelberg, Germany.
- [Hubner et al., 2002] Hubner, J. F., Sichman, J. S., and Boissier, O. (2002). A Model for the Structural, Functional, and Deontic Specification of Organizations in Multiagent Systems. In Bittencourt, G. and Ramalho, G. L., editors, *Advances in Artificial Intelligence: 16th Brazilian Symposium on Artificial Intelligence (SBIA'02)*, volume 2507 of *LNAI*, pages 118–128. Springer-Verlag, Berlin, Germany.
- [Masolo et al., 2003] Masolo, C., Borgo, S., Gangemi, A., Guarino, N., and Oltramari, A. (2003). Ontology Library, WonderWeb Deliverable. Technical Report D18, LOA-CNR, Trento, Italy.
- [Odell et al., 2004] Odell, J., Nodine, M., and Levy, R. (2004). A Metamodel for Agents, Roles, and Groups. In Odell, J., Giorgini, P., and Mller, Jrg, P., editors, *Agent-Oriented Software Engineering V*, volume 3382 of *LNCS*, pages 78–92. Springer-Verlag, Berlin, Germany.
- [Perini and Susi, 2005] Perini, A. and Susi, A. (2005). Automating Model Transformations in Agent-Oriented modeling. In Mller, J. P. and Zambonelli, F., editors, *Agent Oriented Software Engineering VI*, volume 3950 of *LNCS*, pages 167–178. Springer-Verlag, Berlin, Germany.
- [Rao and Georgeff, 1991] Rao, A. S. and Georgeff, M. P. (1991). Modeling Rational Agents within a BDI-Architecture. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, pages 473–484, Cambridge, MA, USA. Morgan Kaufmann Publishers.
- [Shoham, 1993] Shoham, Y. (1993). Agent-oriented Programming. *Artificial Intelligence*, 60:51 to 92.
- [van Lamsweerde, 2000] van Lamsweerde, A. (2000). Requirements Engineering in the Year 00: A Research Perspective. In *Proceedings 22nd International Conference on Software Engineering*, pages 5–19. ACM Press.
- [van Riemsdijk et al., 2005] van Riemsdijk, M. B., Dastani, M., and Meyer, J.-J. (2005). Semantics of Declarative Goals in Agent Programming. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, Utrecht, The Netherlands.
- [Wooldridge and Jennings, 1995] Wooldridge, M. J. and Jennings, N. (1995). Intelligent Agents: Theory and Practice. *Knowledge Engineering Review*, 10(2):115–152.
- [Yen et al., 2001] Yen, J., Yin, J., Ioerger, T. R., Miller, M. S., Xu, D., and Volz, R. A. (2001). CAST: Collaborative Agents for Simulating Teamwork. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI'01)*, pages 1135–1144, Seattle, WA, USA. Morgan Kaufmann.
- [Yu, 1995] Yu, E. (1995). *Modeling Strategic Relationships for Process Reengineering*. PhD thesis, University of Toronto, Canada.