Cooperation = Coordination + Solvability*

Alessandro Agostini
ITC-IRST Trento

February 14, 2003

Abstract

We advance and compare two families of **coalitional paradigms** of solvability.[1] A coalitional paradigm is distinguished from a "noncoalitional" paradigm primarily by its focus on what *groups* of agents can achieve, rather than on what *individual* agents can do—even if cooperating. As a criterion of group formation, our models engage a kind of pairwise, context-dependent coordination between knowledge-based "learning agents," eventually able to communicate the *complete & local* meaning of expressions taken from the literals of a common first-order language. Insights into meaning negotiation and model matching are provided.

Keywords: multyagent systems; knowledge representation; belief revision and update.

1 Introduction

Following [JFL$^+$01], automated negotiation research can be considered to deal with three broad topics, namely: (a) *Negotiation Protocols*: the set of "rules" that govern the interaction among the agents in play. (b) *Negotiation Objects*: the range of issues, or attributes, over which agreement must be reached. (c) *Agents' Decision Making Models*: the decision making system the agents employ to act in line with the negotiation protocol in order to achieve their objectives. Either negotiation protocols togheter with negotiation objects or agents' decision making models is the dominant concern; it depends on the negotiation scenario.

The scope of this article falls primarily into the domain of decision making models, see for instance [JFL$^+$01, Syc89, ZS97] and the references cited there. More precisely, decision making models of choice for coalition formation and teamwork [SLA$^+$99, SK95, Ago00b]. Our main contribution is a framework to study the important topic of group-learning, where coordination for group-formation is explicitely stated. We do this in the tradition of formal learning theory [OdJMW97], that descends from the pioneering studies on inductive inference developed by

---

*Electronic correspondence to: agostini@irst.itc.it.

[1] In this paper we prefer to use "solvability" instead of "learning" under the assumption that "learnable" denotes a wider class of objects than "solvable" does.

Solomonoff, Putman, Gold, L. Blum and M. Blum among others. (See for instance [JORS99] for a survey.)

We provide the framework and some examples of coalition formation in Section 2 and show how *coalitions by coordination* may be eventually used to solve certain classes of structures, in a sense we make precise within a model-theoretic paradigm—we call this paradigm of **group-solvability** (Section 3). Some notions we introduce (e.g., "coordination sentence," "group-solvability") will prove useful to specify the benefits and the limits of our approach for multiagent systems. On the one hand, the technique for group formation we provide suggests new general ways of structuring goal selection and goal decomposition in multiagent systems, and offers a different view of goal management in the spirit, though not in the formal development, of "dynamic coalition formation" (see for instance [KG02]). On the other, we address the solution of a problem of "model matching" [MBDH02] as an example of application.

1.1 What is Coordination?

For the purposes of this paper, we consider coordination as a particular process of negotiation. In other words, negotiation underpins attempts to coordinate and, by starting from coordination, to cooperate. We hope the following equational *slogan* to be of intuitive and immediate help for the reader.

(1) Cooperation = coordination + solvability.

In other words, two agents cooperate whenever they coordinate in solving some goal-problem. Of course, we all have an intuitive sense of what the word "coordination" means. When say Frank and Gertrude pass one another in the hall every day (*The Passing Game*, [Ago01, Ch. 1]), when we watch the Italian volleyball team winning the 2002 Woman's World Championship or, by a counterexample, when we spend hours waiting our best friend in the wrong place or a plane on the "Malpensa" runway because the airline cannot find a gate for it. For many purposes, this intuitive meaning is sufficient. However, in trying to characterize in a formal way the behavior of certain computable agents willing to "coordinate," it is sometimes important to have a more precise idea of what we mean by "coordination." [Dur88, MC94, MO99, RZ94] and, more recently, [CJ02, Sessions 11A and 9B] list a number of definitions that have been suggested for this term. For our purposes here, however, it is useful to begin with the following simple definition.

(2) *Coordination is managing dependencies between activities in a given context.*[1]

As the definition suggests, we believe that it is helpful to use *coordination* in a fairly specialized sense of "coordination in." In this paper we will explain exactly how and why. Moreover, the definition is consistent with the intuition that, nei-

---

[1] This definition is an extension by using the concept of a "context" (see below) to that of [MC94]. We refer the interested reader to the work by Malone and Crowston for further reference that connect our definition to organization theory and coordination science.

ther if there is no interdependance nor if there is no context, there is nothing to coordinate.

2 The Paradigm

We now advance a formal paradigm of "contextual coordination" suitable to represent and analyze coordination processes. An example of a process wherein coordination subsumes a convention over leadership is the following.

(3) EXAMPLE (*A Game of Leadership*) Two players interact extensively by perfect knowledge of the game. Perfect knowledge is required to the agents in order to deal with a "rule of leadership," which is fixed by convention at the beginning of the game. A "leader" chooses an action and a "follower," informed of the leader's choice takes an action as well. In the first part of the game, the players' moves are aimed at making leadership emerge. Then, each player plays according to his role, either of "the leader" or "the follower." Their moves depend on coordination and characterize the game.

A primary vehicle to extend individual solvability to a cooperative, multi-agent setting is to identify and study the basic processes involved in coordination. Are there fundamental coordination processes ("strategies") that occur in all coordinated systems? If so, how can we represent and analyze these processes? One of the advantages of the definition (2) we have used for coordination is that it suggests a direction for addressing these and related questions.

2.1 A First-Order Framework

The paradigm of coordination on presentation is a model-theoretic paradigm. Our approach follows a first-order perspective, "sequences" and "language" are basic ingredients.

Sequences. We denote the set $\{0, 1, 2, ...\}$ of natural numbers by N. We denote the usual linearly ordered structure with domain N by ω. Let η be an infinite sequence.For $i \in N$, we write $\eta|_i$ for the proper initial sequence of length i in η. We write $\mathit{length}(\sigma)$ for the length of a finite sequence, $\emptyset$ for the finite sequence of length zero, σ_i or also $(\sigma)_i$ for the ith element of σ, $0 \leq i < \mathit{length}(\sigma)$, and $\mathit{last}(\sigma)$ for the last element in σ. The set of elements in a (finite, infinite) sequence τ is denoted by $\mathit{range}(\tau)$.

Language. We write $\mathcal{L}_{form}$ to denote a first-order language with equality built up from a (countable, decidable) vocabulary $\mathcal{L}$ consisting of predicates and function symbols of various arities, along with constants symbols and countably many variables $\mathit{Var} = \{v_i \,|\, i \in N\}$. We write $\mathcal{L}_{sen}$ to denote the subset of $\mathcal{L}_{form}$ containing no free variables (**sentences**). We write $\mathcal{L}_{basic}$ for the set of literals of $\mathcal{L}_{form}$.

Our semantic notions are standard first-order logic. In particular, a variable-mapping, or "assignment" h to a structure $\mathcal{S}$ is complete if h is a mapping **onto** the domain of $\mathcal{S}$.

Let *SEQ* denote the collection of all the *finite* sequences over $\mathcal{L}_{basic}$. We define an **environment** to be any infinite sequence over $\mathcal{L}_{basic}$. To consider *consistent* data-streams, we need to relate them to a structure. We do it in the next definition.

(4) DEFINITION Let structure $\mathcal{S}$ and complete assignment h to $\mathcal{S}$ be given. An environment e is **for $\mathcal{S}$ via h** just in case $range(e) = \{\beta \in \mathcal{L}_{basic} \mid \mathcal{S} \models \beta[h]\}$. An environment e is **for $\mathcal{S}$** just in case e is an environment for $\mathcal{S}$ via some complete assignment.

In other words, an environment for a structure $\mathcal{S}$ via complete assignment h lists the basic diagram of $\mathcal{S}$ under h.

2.2 Components

According to our informal definition of coordination (2), coordination means "managing dependencies between activities in a given context." Therefore, since activities must, in some sense that will be made clearer below, be performed by "actors," the definition implies that

(5) all instances of coordination include *agents* performing *activities* ("actions') that are *interdependent* in a *context.*

Agents are thus the first leading concept of our new paradigm. Once we have defined what is an "agent," all other components will follow in a natural way.

2.2.1 Agents.

An agent in a contextual coordination paradigm, or **learning agent**, is a pair $\langle \Psi , \mathbf{A} \rangle$, where Ψ is any mapping from *SEQ* to $\mathcal{L}_{basic} \times \mathcal{L}_{sen}$, and $\mathbf{A}$ is a nonempty class of structures for $\mathcal{L}$. Thus, $\Psi(\sigma) = \langle (\Psi(\sigma))_0 , (\Psi(\sigma))_1 \rangle$ for all $\sigma \in SEQ$. We say that $(\Psi(\sigma))_0$ is $\langle \Psi , \mathbf{A} \rangle$'s **action** (on σ), and that $(\Psi(\sigma))_1$ is $\langle \Psi , \mathbf{A} \rangle$'s **guess** (on σ). Intuitively, faced with $\sigma \in SEQ$, $\langle \Psi , \mathbf{A} \rangle$ believes $\Psi(\sigma)$ if whenever $\Psi(\sigma)$ is defined, there are some $\mathcal{A} \in \mathbf{A}$ and assignment h to $\mathcal{A}$ such that $\mathcal{A} \models (\Psi(\sigma))_0[h]$ and $\mathcal{A} \models (\Psi(\sigma))_1$. Of the two components of a learning agent, the first is called **communication function** or "ability," while the second component is called **background world**. We say that $\langle \Psi , \mathbf{A} \rangle$ is **based on $\mathbf{A}$**.

(6) REMARK As one might observe, Ψ is rather general: any mapping of the right form is allowed. It would be possible to make agents more "intelligent," for example by letting them have a default (partial) 'revision' mapping on $\mathbf{A}$. In this case, our approach can be comparable to work in update semantics or belief revision, where the "beliefs" of an agent are represented by his background world. Since we take any class of first-order structures as an agent's background word, it follows that beliefs may be inconsistent in principle. In other words, we do not assume that $\mathbf{A}$ is the class of models of any consistent "belief set" [Lev80, Gär88]. We discuss further this point in Section 4 (6.0.2).

2.2.2 Protocols.

Outputs of the learning agents'communication functions are at the basis of agents' successive interactions, the combination thereof constituting the "rules of the game." We call such rules a communication protocol. For simplicity, the communication protocol we consider in this paper is pairwise, so that the kind of models we discuss is suitable for modeling group situations where communication is not with the whole group, as in an auction, but indeed pairwise, as happens for instance in commercial transactions. Apart from this constraint, we allow any kind of protocol, in order to make our model applicable to a broad range of pairwise communication scenarios. Our presentation may be adapted for a number of different protocols—examples are parallel/sequential moves, n-person communication, explicit guesses.

2.2.3 Contexts.

By a (coordination) **context** we mean a set of sentences, and precisely any subset of $\mathcal{L}_{sen}$. Intuitively, given a set of learning agents, a coordination context expresses the set of "intelligible hypotheses" π each learning agent in the set may eventually use to coordinate. Coordination contexts determine what agents or designers think to be an "interesting" set of goals for coordination.

2.2.4 Success.

Learning agents have to describe by induction a structure representing a part of their beliefs, eventually after a finite number of failures and a finite sequence of moves. Each description is an appropriate response to the behavior of an agent, who acts by playing on the basis of the "behavioral" language $\mathcal{L}_{basic}$. On the other hand, the agents' moves are concerned with success, in particular, with the "coordination sentence" chosen by each agent at any step of the process of coordination. According to the next definition, the agents can restart their mutual interaction finitely many often, but after the last disagreement they must eventually coordinate. (We write $\overline{\Psi}_0$, $\overline{\Psi}_1$, ${}_k|\overline{\Psi}_0$ and ${}_k|\overline{\Psi}_1$ to denote the infinite sequences of responses, according to the chosen protocol and agents' communication functions Ψ, Φ, of the first agent to the second agent, possibly after k erroneous moves. 0, 1 mean response by the first an second component of the functions, respectively. To write $R(\Psi_0, \Phi)$ for $\overline{\Psi}_0$ will be useful elsewhere.)

(7) Definition Let $\pi \subseteq \mathcal{L}_{sen}$ and learning agents $\langle \Psi , \mathbf{A} \rangle$, $\langle \Phi , \mathbf{B} \rangle$ be given. We say that $\langle \Psi , \mathbf{A} \rangle$ **π-coordinates with** $\langle \Phi , \mathbf{B} \rangle$ (written: $\Psi \rightleftharpoons_\pi \Phi$) just in case for some $s, t \in N$, there is $\theta \in \pi$ such that:

(a) ${}_s|\overline{\Psi}_0$ is an environment for some $\mathcal{A} \in \mathbf{A}$;
(b) ${}_t|\overline{\Phi}_0$ is an environment for some $\mathcal{B} \in \mathbf{B}$;
(c) for cofinitely may $n \in N$, $({}_s|\overline{\Psi}_1)_n = ({}_t|\overline{\Phi}_1)_n = \theta$;
(d) $\mathcal{A} \models \theta$ and $\mathcal{B} \models \theta$.

In this case, θ is said to be a **coordination sentence**. If (c) holds, we say that $\langle \Psi , \mathbf{A} \rangle$ ($\langle \Phi , \mathbf{B} \rangle$) **guesses** θ.

Observe that $\rightleftharpoons_\pi$ is symmetric, that is, for all learning agents $\langle \Psi , \mathbf{A} \rangle$, $\langle \Phi , \mathbf{B} \rangle$, if $\Psi \rightleftharpoons_\pi \Phi$ then $\Phi \rightleftharpoons_\pi \Psi$. This seems to be a natural property of coordination. However, $\rightleftharpoons_\pi$ is neither reflexive nor transitive, as it is easy to verify.

According to Definition (7), given a coordination context π learning agents must coordinate twice. First, the agents eventually stabilize on an environment for a structure $\mathcal{S}$ in their own background world. This step is done by the agents through their external, behavioral ability component. Second, the agents eventually stabilize to a "coordination sentence" in π true in $\mathcal{S}$. This second step is done by the agents through their internal, hidden ("cognitive") ability. Both steps are processes in the limit, which eventually start after a finite number of disagreements. π-coordination highlights crucial aspects of communication, sometimes referred to as "meaning negotiation" [BW02].

(8) EXAMPLE (*Meaning Negotiation*) In any system of autonomous and distributed agents willing to coordinate, autonomy is the core condition for agents to make independent assignments of meaning to world objects. (See [Ago02] for a full position on the topic.) In fact, the choice of $\mathcal{A} \in \mathbf{A}$ and $\mathcal{B} \in \mathbf{B}$ in Definition (7) is a local affair. The question π-coordination poses in terms of negotiation of *meaning* is what is, if any, that coordination sentence θ among those allowed by the paradigm ("game"), such that *local assignments* of meaning match—namely, $\mathcal{A} \models \theta$ and $\mathcal{B} \models \theta$. Of course, the problem of meaning negotiation arises from the assumption that each agent's meanings, namely guesses, are hidden to all others. (To see this, recall the definition of *SEQ* and of communication function.) We have made such assumption.

(9) EXAMPLE (*Seekers and Providers*) When we view communication as a way of managing a seeker/provider relationship, for instance in an information retrieval scenario, we may be concerned about how to make meaning relative to the seeker's needs—say finding some documents, "usable" by the provider in order to fulfill the seeker's requests, that is, to coordinate by providing him with those and similar documents. How can the agents establish a common language interpretation over a shared context that allows them to communicate meaning? How do such a shared contex eventually emerge by meaning negotiation? An answer to these and related questions, being important *per se*, might also address *knowledge management*, as knowledge management is "the process whereby knowledge seekers are linked with knowledge sources, and knowledge is transferred." (A.J. Murray)[2]

---

[2] Retrieved August 8, 2001 from the World Wide Web: http://www.3-cities.com/~bonewman/what-is.htm.

3 Group Solvability

There is a sense in which some overall evaluation criterion is necessarily implied by the definition of coordination. The most commonly analyzed case of managing dependency in a context occurs when an individual or group decides to pursue a goal from a set of possible goals (in fact, from a context) and then decomposes this goal into activities, or subgoals, which together will achieve the original goal. In this case, we call the process of chosing the goal **goal selection** and the process of chosing the activities **goal decomposition**.

(10) EXAMPLE (*Top-Down Goal Decomposition*) A common kind of dependency among activities is that a group of activities[3] are all "subtasks" for achieving some overall goal. The strategic-planning process in human organizations may be viewed as involving this kind of goal selection and goal decomposition process. In computer systems, we usually think of the goals as being predetermined, but an important problem involves how to break these goals into activities that can be performed separately. In a sense, for example, the essence of distributed information systems and multiagent systems is to decompose a "goal" into elementary versus autonomous activities. Planning in artificial intelligence (see for instance [AHT90, GHT02] and the references cited there) is another example of goal decomposition in multiagent systems. A "plan," intuitively, is what we called an "environment" in this paper.

In what follows, we advance a criterion of coordination where: (a) goal selection from a set π of possibilities corresponds to convergence to a coordination sentence θ (the "selected goal") in π, that is, to π-coordinate; and (b) goal decomposition into sub-goals corresponds to group formation, where for each member of the group that may eventually solve the goal, there is another member that plays a sub-goal with the aim to help towards a complete solution. Goal selection and goal decomposition are instances of the following definition.

(11) DEFINITION Let $\pi \subseteq \mathcal{L}_{sen}$, nonempty class $\mathbf{A}$ of structures, and set Σ of learning agents based on $\mathbf{A}$ be given.

(*a*) Σ $\mathbf{Gr}^{\pi}$**-solves** $\mathbf{A}$ (written: $\mathbf{A} \in \mathbf{Gr}^{\pi}(\Sigma)$) just in case:

  i. for all $\langle \Psi , \mathbf{A} \rangle, \langle \Phi , \mathbf{A} \rangle \in \Sigma$, $\langle \Psi , \mathbf{A} \rangle$ π-coordinates with $\langle \Phi , \mathbf{A} \rangle$;
  ii. for all $\langle \Phi , \mathbf{A} \rangle \in \Sigma$ and for all $\mathcal{A} \in \mathbf{A}$, there are $\langle \Psi , \mathbf{A} \rangle \in \Sigma$ and $k \in N$ such that ${}_k|R(\Psi_0, \Phi)$ is an environment for $\mathcal{A}$.

  In this case, $\mathbf{A}$ is said to be $\mathbf{Gr}^{\pi}$**-solvable**.

(*b*) $\mathbf{Gr}^{\pi} = \{\mathbf{A} \mid \mathbf{A} \text{ is } \mathbf{Gr}^{\pi}\text{-solvable}\}$.

The definition identifies the key components of a paradigm of **group-learning** (or learning by groups).[4] We call a **group** ("based on $\mathbf{A}$") the set Σ of learning

---

[3] A group of agents, by argument (5).

[4] In fact, it identifies a *family* of paradigms over the parameter π.

agents that satisfies clause (a).ii of the definition. Goal selection is modeled by clause (a).i. Goal decomposition follows from a composition of clauses (a).i and (a).ii. The definition also gives a formalization of certain "matching problems."

(12) EXAMPLE (*Model-Matching Problem*, I) Suppose that a coordination context π is partitioned into two parts, say π_p and π_d. Let countable class $\mathbf{A}$ of structures be such that for each $\mathcal{A}$ in $\mathbf{A}$ there is a π-coordination sentence θ satisfiable in $\mathcal{A}$. Then we can represent the following problem of "model matching" [MBDH02]. Given **models** (or "classifications") $(\mathbf{A}, \pi_p)$ and $(\mathbf{A}, \pi_d)$ with their associated data, namely, the elements in the union of all structures-in-$\mathbf{A}$' domains, *for each* $\vartheta \in \pi_p$ *(i.e., "concept"), find the most similar* $\vartheta' \in \pi_d$*, for a predefined similarity measure.* π-coordination provides a similarity measure in the strict sense of the existance of a sentence θ in π such that $\vartheta = \vartheta' = \theta$. Less strict measures are possible as variants of π-coordination.

Given a problem of model-matching, a natural question is: "Is the problem solvable?" According to the interpretation of goal selection and goal decomposition we have given in the example above in terms of matching, the following theorem helps to address answer.[5]

(13) THEOREM Let $\pi \subseteq \mathcal{L}_{sen}$ and let countable class $\mathbf{A}$ of structures be such that for every $\mathcal{A} \in \mathbf{A}$ there is $\theta \in \pi$ such that $\mathcal{A} \models \theta$. Then $\mathbf{A}$ is $\mathbf{Gr}^{\pi}$-solvable.

The proof of the theorem provides a procedure of coordination between two agents sharing a "convention of leadership." Here is a sketch.

Let $\pi = \{\theta_i \in \mathcal{L}_{sen} \,|\, i \in N\}$ and let $\{\mathcal{A}_j \,|\, j < \omega\}$ be an infinite-repetition enumeration of the countably many members of $\mathbf{A}$. Let $e^{\mathcal{A}_n}$ be an environment for $\mathcal{A}_n \in \mathbf{A}$ and let $\theta_m \in \pi$ be such that $\mathcal{A}_n \models \theta_m$. We describe a limiting game of a learning agent $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ who interacts with a learning agent $\langle \Psi_{<k,j>}, \mathbf{A} \rangle$ in order to π-coordinate. To make π-coordination possible, we assume that the agents stipulate a convention regulating leadership. This is assumed to happen before the beginning of the game. The convention we refer to in this proof is *lexicographic order on the agents indexing.*[6] According to this convention, the leader is who among the agents has the lower index. So, in particular, $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ is the leader iff $<n,m>$ is lower than $<k,j>$ according to lexicographic order. The assessment of the leadership is done by $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ soon after the first two stages of the agents' interaction are played. Now the game begins. By using the communication function, $\langle \Psi_{<n,m>}, \mathbf{A} \rangle$ starts moving $v_n \doteq v_n$ to communicate to $\langle \Psi_{<k,j>}, \mathbf{A} \rangle$ that "if

---

[5] Because of the importance of the model-matching problem, we deserve its full study as an independent problem to future work. As the scope of the present work is only to show how the problem can be formalized within the more general setting of group-learning, we present hints to problem solution as an example (See Example (14)).

[6] One can identify several different types of conventions, or "agreement conditions," which may be used in different coordination scenarios. It is always assumed, however, that the agents will settle on the convention to be used *before* the actual coordination process proper begins. The selection of a piece of knowledge as shared by the agents is thus an important but *meta-coordination* issue, whose general study falls outside the scope of the present work.

she will be the leader then she completely describe $\mathcal{A}_n$.” This is done by enumerating environment $e^{\mathcal{A}_n}$. At the second interaction stage and, again, by the first component of communication function (since guesses are “hidden,” communicating them to the partner is useless), $\langle \Psi_{<n,m>}\, , \mathbf{A} \rangle$ moves $v_m \doteq v_m$ to communicate to $\langle \Psi_{<k,j>}\, , \mathbf{A} \rangle$ that “if she will be the leader then she guesses (interacting with $\langle \Psi_{<k,j>}\, , \mathbf{A} \rangle$) θ_m forever.” From the third stage of the interaction onwards, both the agents will know who is the leader. If $\langle \Psi_{<n,m>}\, , \mathbf{A} \rangle$ is not the leader, namely, $\langle \Psi_{<n,m>}\, , \mathbf{A} \rangle$ is “the follower” so that $\langle \Psi_{<k,j>}\, , \mathbf{A} \rangle$ results to be the leader, then $\Psi_{<n,m>}$ simply copies the moves of $\Psi_{<k,j>}$, precisely those moves $\Psi_{<k,j>}$ advances by the first component of his communication function. Also, $\Psi_{<n,m>}$ outputs the guess θ_j in π produced by $\Psi_{<k,j>}$. Because guesses are “hidden,” in particular to $\langle \Psi_{<n,m>}\, , \mathbf{A} \rangle$—in fact $\Psi_{<n,m>}$’s input domain is *SEQ*, θ_j is communicated by $\langle \Psi_{<k,j>}\, , \mathbf{A} \rangle$ to $\langle \Psi_{<n,m>}\, , \mathbf{A} \rangle$ via a literal of the form $v_j \doteq v_j$. On the other hand, if $\langle \Psi_{<n,m>}\, , \mathbf{A} \rangle$ is the leader, then she is expected to act consistently to what announced in advance by her first two moves, namely, she starts enumerating an environment for $\mathcal{A}_n$ (specifically: $e^{\mathcal{A}_n}$), and starts guessing sentence θ_m forever after. The game ends in the limit.

Set $\Sigma = \{ \langle \Psi_{<n,m>}\, , \mathbf{A} \rangle \,|\, \mathcal{A}_n \models \theta_m \}$. Our claim that Σ witnesses $\mathbf{A} \in \mathbf{Gr}^{\pi}(\Sigma)$ is then easy to verify by the following argument. Observe that Σ is nonempty. For all $\mathcal{A} \in \mathbf{A}$ and for all $\langle \Psi_{<n,m>}\, , \mathbf{A} \rangle \in \Sigma$, let $k > n$ be such that $\mathcal{A}_k = \mathcal{A}$. Since $\{\mathcal{A}_j \,|\, j < \omega\}$ is an infinite-repetition of the structures in $\mathbf{A}$, such k exists. Then, $\langle \Psi_{<k,m>}\, , \mathbf{A} \rangle \in \Sigma$ is the leader who enumerates an environment for $\mathcal{A}$. Moreover, for all but finitely many $t \in N$, $(R(\Psi_{<k,m>1}, \Psi_{<n,m>}))_t = \theta_m$. On the other hand, it is easy to verify that the convention on leadership stipulated by any two agents in Σ ensures that each pair of learning agents in Σ π-coordinate. Hence, $\mathbf{A} \in \mathbf{Gr}^{\pi}(\Sigma)$. This completes the sketched proof.

The technique for group formation provided by the proof of the theorem suggests new ways of structuring goal selection and, expecially, goal decomposition in multiagent systems. It may also applied to the problem of model matching.

(14) EXAMPLE (*Model-Matching Problem*, II) Think of seekers and providers of Example (9) in terms of leadership. The following convention of leadership is used. A seeker is always the leader of meaning, that is, who first chooses & fixes the coordination sentence, or “concept,” for which a set of domain instances (“documents”) is sought. A provider is always a leader in choosing the domain of documents that satisfy the request, provided it has been understood correctly. This is done by the help of seekers, who consistently use their communication ability to formulate a request, namely, play an environment for a structure which satisfies it. Then, the model-matching problem between $(\mathbf{A}, \pi_p)$ and $(\mathbf{A}, \pi_d)$ is $\mathbf{Gr}^{\pi}$-solvable. The construction of the group of providers that solves the problem for each coordination sentence in π_d by π-coordinating with the seeker, as well as the definition of the seeker, are similar to that we have sketched in the proof of Theorem (13), but both are based on a different convention.

3.1 Computable Agents

How do information-processing limitations of a learner affect the desirability of different paradigms of solvability? Are some methods of coordination appropriate for coordinating people that would not be appropriate for coordinating computable agents? The method we used to prove Theorem (13) is appropriate for coordinating people, but it is not appropriate for coordinating computable agents, because the undecidable choice of satisfiable goals. We now present an appropriate method for coordinating computable agents by narrowing the interpretation of learning agents to computable objects. We rely on the following somewhat technical definition.

(15) DEFINITION Let $(\mathcal{A}_n : n \in N)$ be a countable class of structures. We say that $(\mathcal{A}_n : n \in N)$ is **uniformly recursive** just in case there is a total computable function E from $N \times N$ to the set of infinite sequences over $\mathcal{L}_{basic}$ such that for every $n \in N$, $\lambda i \,.\, E(n, i)$ is a environment for $\mathcal{A}_n$.

To illustrate, one might think to structure $\mathcal{A}_n$ as a classification structure, for example a tree, a graph, eventually a "concept hierarchy" [BBH98]. $\mathcal{A}_n$ may be used by an agent with background world included in $(\mathcal{A}_n : n \in N)$ to classify n documents, or also a set of documents according to n related concepts. In the latter case, a suitable classification theory $\Gamma_c \subseteq \mathcal{L}_{form}$ should have $\mathcal{A}_n$ as a model.

(16) DEFINITION A learning agent $\langle \Psi \,, \mathbf{A} \rangle$ is **computable** just in case Ψ is computable and $\mathbf{A}$ is uniformly recursive.

We are now ready to set the computable (recursive) version of the paradigm of group-learning $\mathbf{Gr}^{\pi}$.

(17) DEFINITION Let $\pi \subseteq \mathcal{L}_{sen}$ be given. $[\mathbf{Gr}^{\pi}]^{rec} = \{\mathbf{K} \,|\, \mathbf{K} \in \mathbf{Gr}^{\pi}(\Sigma)\}$, where Σ is a set of computable learning agents based on $\mathbf{K}$.

Mimicking the proof of Theorem (13) and taking $\lambda i \,.\, E(n, i)$ to be a computable environment for $\mathcal{A}_n$, we have:

(18) THEOREM Suppose that $(\mathcal{A}_n : n \in N)$ is a uniformly recursive class of structures with infinite repetitions. Let $\pi = \{\theta_i \in \mathcal{L}_{sen} \,|\, i \in N\}$ be such that there is a total computable function $f : N \longrightarrow N$ such that for every $i \in N$, $\mathcal{A}_i \models \theta_{f(i)}$. Then $(\mathcal{A}_n : n \in N)$ is $[\mathbf{Gr}^{\pi}]^{rec}$-solvable.

The proof of the theorem provides a procedure of coalition formation similar to that of Theorem (13), but computable.

4 Discussion

Even though our paradigms omit many important aspects of human coordination and distributed computer systems, they help illuminate a wide range of phenomena. For instance, the models are consistent with a number of previous work about individual learning [OdJMW97, MO98], coordination [Ago00a, MO99] and coalition/group formation [APM00, RS98, Syc89, Ago00b]. These models also help analyze design alternative for distributed systems of "peers" [PSW01] and they suggest ways of analyzing the structural changes associated with introducing the inductive approach to coordination into multiagent systems and organizations' dynamics.

In addition to the processes described above for managing dependencies within a context, two other processes deserve specific attention: *communication* and *belief revision.*

4.0.1 Communication

One obvious way of generating new coordination procedures is by considering alternative forms of communication for all the places in a coordination process where information, or also "knowledge," needs to be transferred.[7] We restricted attention to communication functions where actions are "basic" (i.e., formulas in $\mathcal{L}_{basic}$) and where the discovery machinery runs over sentences (i.e., formulas in $\mathcal{L}_{sen}$). Moreover, actions available to the agents can be limited to just atomic formulas (no negations), or enriched to include universal or other kinds of formulas. Many alternative definitions of an "environment" are possible. It is worth noting that "real" multiagent environments may suffer omissions, erroneous intrusions, or both omissions and intrusions.

4.0.2 Belief revision

At the borderline of our approach is the problem of belief changes—how an agent should revise her beliefs to coordinate upon learning new information. The arrival of data σ, however, modifies an agent's choice $\mathcal{A}$ of a structure in the agent's background word according to some fixed, but implicit scheme of belief revision. The resulting new beliefs are denoted $\mathcal{A} \dotplus \sigma$ and signifies the impact of σ on $\mathcal{A}$. Belief revision is so far mostly defined for agents whose background world is a set of sentences, possibly a "theory" (see for instance [Lev80, Gär88, MO98]). On the other hand, we have seen that our paradigms consider an agent's beliefs to be represented by a class of structures. New paradigms of coordination among 'revision-based' rather than 'knowledge-based' agents may then be advanced from our work by answering the following question, among others: Is there a natural (justifiable on intuitive grounds), semantic generalization of belief revision in the context of inductive inference, in which revision applies directly to classes of structures? If the answer is "yes", how to extend the mathematics of inductive coordination as presented in this paper so as to be able to frame appropriate conditions on the revision operator so that its use has a "rational" quality to it? We leave these important questions to future work

---

[7] "*Knowledge* is information transformed into capabilities for effective action. In effect, knowledge is action."(P. C. Murray) Retrieved August 8, 2001 from the World Wide Web: http://www.ktic.com/topic6/13-term0.htm.

5 Conclusions

The key point of the foregoing discussion, and indeed of much of this paper, is that the concepts of a coordination theory as that we have presented can help identify similarities among concepts in different disciplines, among other we have mentioned planning, knowledge management, and of course coordination, cooperation and (meaning) negotiation in a multiagent setting. These similarities, in turn, suggest how ideas can be transported back and forth across disciplinary boundaries and where opportunities exist to develop even deeper analyzes.

To summarise, in this paper we have provided a framework to study the important topic of group-learning, where pairwise coordination is modeled over sets of belief-based agents, eventually able to communicate the complete (see "environments") and local (that is, of each agent) meaning of their beliefs by using the literals of a common first-order language. We have provided some examples of coalition formation, and have showed how *coalitions by coordination* may be used to solve some classes of structures, in a precise sense we have made actual within the model-theoretic tradition of formal learning theory. Group formation as depiced in this paper suggests an innovative way to intend group selection and goal decomposition. As an important consequence for real applications to multiagent systems, we have showed how certain classes of structures may be restricted to a computable setting, thus providing a suitable paradigm to analyze coordination and coordination mechanisms, for example "by leadership," among groups of agents with an internal "model," such as a classification, beliefs, a concept hierarchy or any other type of structured or semi-structured background knowledge. Examples of meaning negotiation and matching were given.

References

[Ago00a] A. Agostini. Notes on formalizing coordination. In E. Lamma and P. Mello, editors, *AI*IA 99: Advanced in Artificial Intelligence*, pages 285–296. Springer-Verlag LNAI 1792, 2000.

[Ago00b] A. Agostini. Team-solvability: A model-theoretic perspective. In W. Horn, editor, *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI-00)*, pages 333–337, Amsterdam, 2000. IOS Press.

[Ago01] A. Agostini. *Paradigms of Coordination and Solvability.* PhD thesis, Università di Siena, 2001.

[Ago02] A. Agostini. Inductive meaning negotiation and games. In Luigia Carlucci Aiello, editor, *Research Abstracts*, pages 73–74, Trento, Italy, 2002. ITC-IRST.

[AHT90] J. F. Allen, J. Hendler, and A. Tate, editors. *Readings in Planning*, San Mateo, CA, 1990. Morgan Kaufmann Publishers.

[APM00] L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI-00)*, pages 338–342, Amsterdam, 2000. IOS Press.

[BBH98] A. Büchner, D. Bell, and J. Hughes. A contextualised object data model based on semantic values. In *Proceedings of the Eleventh International Conference on Parallel and Distributed Computing Systems (PDCS-98)*, pages 171–176, 1998.

[BW02] P. Bouquet and M. Warglien, editors. *Meaning Negotiation - Papers from the AAAI Workshop*, volume Technical Report WS-02-09, Menlo Park, CA, 2002. AAAI Press.

[CJ02] C. Castelfranchi and W. Lewis Johnson, editors. *Proceedings the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, New York, NY, 2002. ACM Press.

[Dur88] E. H. Durfee. *Coordination of Distributed Problem Solvers.* Kluwer Academic Publishers, Boston, MA, 1988.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* The MIT Press, Cambridge, MA, 1988.

[GHT02] M. Ghallab, J. Hertzberg, and P. Traverso, editors. *Proceedings of the Sixth International Conference on AI Planning and Scheduling (AIPS-02)*, Menlo Park, CA, 2002. AAAI Press.

[JFL$^+$01] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated Negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.

[JORS99] S. Jain, D. Osherson, J. Royer, and A. Sharma. *Systems That Learn - An Introduction to Learning Theory, 2nd edition.* The MIT Series in Learning, Development, and Conceptual Change, v. 22. The MIT Press, Cambridge, MA, 1999.

[KG02] M. Klusch and A. Gerber. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems*, 17(3):42–47, 2002.

[Lev80] I. Levi. *The Enterprise of Knowledge.* The MIT Press, Cambridge, MA, 1980.

[MBDH02] J. Madhavan, P. A. Bernstein, P. Domingos, and A. Y. Halevy. Representing and reasoning about mappings between domain models. In R. Dechter, M. Kearns, and R. Sutton, editors, *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI-02) and of the 14th Innovative Applications of Artificial Intelligence Conference*

(IAAI-02), pages 80–86, Menlo Park, CA, 2002. AAAI Press/The MIT Press.

[MC94] T. W. Malone and K. Crowston. The Interdisciplinary Study of Coordination. *ACM Computing Surveys*, 26(1):87–119, 1994.

[MO98] E. Martin and D. Osherson. *Elements of Scientific Inquiry.* The MIT Press, Cambridge, MA, 1998.

[MO99] F. Montagna and D. Osherson. Learning to coordinate: A recursion theoretic perspective. *Synthese*, 118(3):363–382, 1999.

[OdJMW97] D. Osherson, D. de Jongh, E. Martin, and S. Weinstein. Formal Learning Theory. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 737–775. Elsevier Science B.V., 1997.

[PSW01] M. Parameswaran, A. Susarla, and A.B. Whinston. P2P networking: An information-sharing alternative. *IEEE Computer*, 34(7):31–38, 2001.

[RS98] S. K. Rustogi and M. P. Singh. The bases of effective coordination in decentralized multi-agent systems. In J. P. Müller, M. P. Singh, and A. S. Rao, editors, *Intelligent Agents V - Agent Theories, Architectures, and Languages*, pages 149–161. Springer-Verlag LNAI 1555, 1998.

[RZ94] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers.* The MIT Press, Cambridge, MA, 1994.

[SK95] O. Shehory and S. Kraus. Task allocation via coalition formation among autonomous agents. In C. S. Mellish, editor, *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 655–661, San Mateo, CA, 1995. Morgan Kaufmann.

[SLA$^+$99] T. W. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1-2):209–238, 1999.

[Syc89] K. P. Sycara. Argumentation: Planning other agents' plans. In N. S. Sridharan, editor, *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, pages 517–523, San Mateo, CA, 1989. Morgan Kaufmann.

[ZS97] D. Zeng and K. Sycara. How can an agent learn to negotiate? In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III - Agent Theories, Architectures, and Languages*, pages 233–244. Springer-Verlag LNAI 1193, 1997.