Formulating Agent Communication Semantics and Pragmatics as Behavioral Expectations*

Matthias Nickles, Michael Rovatsos, Gerhard Weiss

AI/Cognition group, Department of Informatics, Technical University Munich (TUM)
D-85748 Garching bei München, Germany,
{nickles,rovatsos,weissg}@cs.tum.edu

Abstract. Although several approaches to the semantics of agent communication have been proposed, none of them is really suitable for dealing with agent autonomy, which is a decisive property of artificial agents. This paper introduces an observation-based approach to the semantics of agent communication, which combines benefits of the two most influential traditional approaches to agent communication semantics, namely the *mentalistic* (agent-centric) and the *objectivist* (i.e., commitment- or protocol-oriented) approach. Our model makes use of the fact that the most general meaning of agent utterances lays in their expectable *consequences* in terms of agent actions, and that communications result from hidden but nevertheless rational and to some extent reliable agent intentions. In this work, we present a formal framework which enables the empirical derivation of communication meanings from the observation of rational agent utterances, and introduce thereby a probabilistic and utility-oriented perspective of social commitments.
Keywords: *Agent Communication Languages, Open Multiagent Systems, Computational Autonomy, Stochastic Processes, Artificial Sociality*

1 Introduction

Currently, two major approaches to the meaning of agent communication in a broader sense, covering both traditional sentence-level semantics and pragmatics, exist. The *mentalistic* approach (e.g. [5, 6]) specifies the meaning of utterances by means of a description of the mental states of the respective agents (i.e., their beliefs and intentions, and thus indirectly their behavior), while the more recent commitment-based *objectivist* approaches (e.g. [3, 14], also called *social semantics*) try to determine communication from an external point of view, focussing on public rules and inter-agent contracts. The former approach has some well-known shortcomings, which eventually led to the development of the latter: Especially in *open* multiagent systems, agents appear more or less as black boxes, which makes it in general impossible to impose and verify a semantics described in terms of agent cognition. They could only be put into practice making simplifying but unrealistic assumptions to ensure mental homogeneity among the agents, for example that the interacting agents were benevolent and sincere, and it neglects the social context of utterances. Objectivist semantics in contrast is fully verifiable, it achieves a big deal of complexity reduction through limiting itself to a small set

* This article is an extended and revised version of [11]

of normative rules, and has therefore been a significant step ahead. But it oversimplifies social processes, and it doesn't have a concept of meaning indefiniteness, rational attitude (but see [4] for an objectivist approach to modeling the "intuitive" meaning of speech acts) and agent malevolence. In contrast to these approaches, we propose a so-called *empirical semantics* which is based on the assumption that the meaning of agent utterances lies basically in their *consequences* in terms of *expectations*, i.e., *expectable* future agent actions and other events which can be continuously learned and adapted from observed agent actions - a view first introduced in [8, 9]. These consequences are represented as probabilistic *Social Interaction Structures*, which are a variant of *Expectation Networks* [8, 10, 11], and they are learned from ongoing communication processes by a *semantics observer* that can be either an agent participating in the communication himself, or an external agent (e.g., a special middle agent, or a supervision facility [12] of the system designer or application users). This learning task puts two general assumptions about agent communication into practice: i) observed agent interactions within a certain social context are likely to reoccur in similar situations in the future (empirical stationarity assumption), and ii) agents act individually but more or less rationally towards their communicated goals within a *limited sphere of communication* (limiting their commitments' trustability and the predictability of other behavioral characteristics). Therefore, the semantics observer deals with the "intentional stances" [2] of otherwise opaque agents towards their communicated goals and believes (learned empirically from observed utterances) rather than with real "cognitive agents". From these assumptions, we retrieve the following replacements for traditional semantical concepts:

- Verification of semantics according to normative rules as in social semantics → Verification regarding a learned empirical model of observed agent communication processes
- Assumption of a certain mental agent architecture and cognition → revisable, probabilistic expectation of bounded rational behavior (the so called *rational hulls* of communications)
- Social commitments and agent sincerity → revisable, probabilistic expectation of the limited maintenance of communicated goals by the uttering agents

For lack of space, and in order to provide a general, flexible approach, we do not make use of a concrete ACL in this work. Instead, we propose the dynamic semantics of so-called *Elementary Communication Acts* (ECAs) which obtain their concrete meaning not from some pre-defined speech-act typology as usual, but from their usage context. The theoretical assumption behind ECAs is that all kinds of speech acts can be translated into one or more demands to act in pragmatical conformance with a declared course of events (a certain probability distribution of events in the future), in which each ECA can be contextualized with companion social structures resulting from other ECAs to clarify and get accepted the demand (e.g. sanctions). E.g., an assertive act is the request to communicate in conformance with the expressed belief from now on, a command is a request to perform the described actions in order to reach the declared future world state, accompanied with norms and the threatening with sanctions, and the utterance of a performative sentence is also an assertive act which demands to communicate *as if* the proposed (social) consequences this act "makes true" were/would become true in fact. ECAs are represented as pointers to demanded and otherwise rather unlikely world states within some assumably shared world knowledge represented as a so-called *Expectation Network*. Thus, the abstract a priory semantics of ECAs (in contrast to their full meaning which is derived empirically at run-time) can be considered to be the low-level replacement of both traditional a priory sentence semantics (concerned with the propositional content of messages, based on an assumably shared ontology with a semantics given as truth-conditions) and speech act types. In contrast, the "full semantics" of an actually uttered ECA is the probability distribution of expected future events triggered by this utterance.

The remainder of this paper is organized as follows: The next section introduces our novel approach to ACL semantics and pragmatics. Section 3 defines *Expectation Networks* as the data structure used to describe agent communication semantics empirically. Section 4 provides a formal learning and adaptation framework for social (i.e., communication) structures, and finally, section 5 draws some conclusions regarding current limitations of our approach and future work.

2 A novel approach to the modeling of communication

In this section, we provide an informal overview of our approaches called *Empirical Agent Communication Semantics* (or "empirical semantics" for short) [8, 9, 16, 10] and *Empirical-Rational Agent Communication Semantics* [9, 11] in order to motivate the formal framework presented in the following sections.

In its most general sense, the *semantics*[1] of agent communication describes the effect a single communication has in the context of / on its environment. Both the context and the effect can include / affect every changeable aspect of the uttering agents' environment and the agent itself, e.g. agent cognition, other communications, social structures, the "physical" environment, the mental dispositions of the uttering and other agents. Having knowledge about the semantics of agent communication has several obvious advantages, both for the agents (active and passive) and the designer of the agent-based application. Since for truly autonomous black-box agents, every kind of meaningful interaction can be expressed in terms of symbolic, rejectable communications *only* [1, 8] (in contrast to the direct influencing of agents through physical actions or commands), agent communication semantics covers *every* aspect of socially relevant behavior, from social mechanism design (e.g. auctions) and game theory to large artificial societies.

2.1 Demands and issues

Traditionally, the comprehensive semantics mentioned above is assumed to have two dimensions that need to be covered by a comprehensive approach to the semantics of agent communication: First, the *sentence level*, which is the aspect of meaning that is traditionally subject of linguistical semantics. This aspect of meaning is contextualized

[1] If not stated otherwise, we use the term "semantics" in the computer scientific sense, not as a linguistic term. Linguists would talk about "meaning" instead, covering both linguistic sentence semantics and pragmatics.

with an environmental description in the form of an *assumably* (not necessarily actually) shared *ontology*. In addition, a calculus to describe objects and events within the environment the respective utterance refers to has to be provided, for example predicate logic and temporal modalities. The second dimension of meaning is its *pragmatics*, i.e., the actual use and effect of utterances in social encounters. Contemporary approaches to agent communication language (ACL) semantics go pretty far in their claimed area of coverage, since they attach either far-reaching mentalistic or social-normative assumptions to single ACL sentences. This leads to a mixture of traditional sentence-level semantics and pragmatics. Even though also sentence-level semantics largely depends from use-dependant contextual information, required e.g. for the resolution of anaphora, most linguists carefully prevent the mixing of (socio-)pragmatic issues and sentence-level semantics. In contrast, most approaches to ACL semantics are in fact "pragsemantics" since they include elements which traditionally belong to pragmatics, mostly borrowed from speech-act theory and socio-normative theories. In principle, there is nothing wrong with such an hybrid approach (at least for the case of formal languages, where things are less complicated than with human languages), and our approach follows this direction too. But there are several problems with the mentioned mentalistic and objectivist approaches to put "pragsemantics" into action, as discussed now.
Following [9], we have identified the following demands and issues for ACL semantics, for which we aim to provide a basic approach:

Expressibility Communications are basically (possibly false) demands directed to other agents to bring about or to act in accordance with a certain projected (respectively asserted) world state (respectively point of view), in which a “world state” is a expected course of events or a proposed view of history. Thus, the means in order to bring these states about, possible reactions from other agents (including bystanders) and other aspects and implications of the initial and resulting states need to be modeled. In our opinion, neither the notation of social commitment nor the specification of mental agent properties are adequate order to do so. The former not, because it essentially reduces interaction meaning to contracting protocol semantics (leaving the term “commitment” itself rather under-specified [15]), the latter not also, because communication meaning external to the agents’ minds can be modeled only indirectly.

Verifiability Whereas most approaches to ACL semantics use this term in order to check if normatively imposed regulations are observed (i.e., if agents think and behave “correctly”), we use the term “verifiability” in a model-theoretical sense to express that a model of agent communication corresponds with observable processes of agent interactions, in which this model is to its largest part learned from observations itself, and only to a small, abstract part imposed normatively.

Flexibility and support for meaning indifference, emergence and change Current approaches to ACL semantics work if the set of speech act types is known a-priori and each locution denotes a fixed and known illocutionary act. Tackling these issues, Empirical-Rational Semantics restricts itself to a very small predefined core part, decomposes all kinds of communication acts into a single type of elementary act, and determines much of the actual semantics empirically at run-time.

Consideration of heterogeneous agent architectures and agent insincerity If the semantics of utterances is given in terms of mental agent states, it can either not be validated (from an agent-external point of view, e.g. of the responding agent, or of the system designer), or requires the restriction of agent autonomy (e.g. demanding sincerity). The latter also affects some "objectivist" approaches, if these require norm fulfilling, or make additional mentalistic assumptions (e.g. "whole-hearted satisfaction" [17]).

Consideration of agent intentionality and rational attitude Communication has an unique property: It constructs a social situation, which is inherently consistent and reasonable, even if it opposes the "real world" outside communication and the cognitive beliefs of the agents: 1) Communicated information is supposed to be consistent with information previously communicated by the same agent, or this agent at least justifies his change of mind, 2) the agent defends and asserts his utterances by means of argumentation or other rational means like rewards and sanctions, and 3) information not expressed explicitly can be deduced from information communicated before and background knowledge. If, for example, in an open auction on the internet some agent a asserts "I will deliver the goods if you win the auction.", an observer does not need to believe him. But the observer believes that the further communication of a complies with this assertion. To make communication work, this belief is to some extent independent from reasoning about the true motives "within the agents mind". Agent a is supposed to act at least for some time in a rational manner in accordance with the social image that he *projects* for himself by means of communication (e.g., a sanctions the denial of his proposal, rewards its acceptance etc). The information about such (bounded-)rational attitude is implicitly associated with each communication of a self-interested agent, and is thus part of communication semantics. Commitment-based approaches largely neglect this kind of intentionality, moving ACL semantics towards contract making instead.

Interaction process generalization and social structures Social structures like norms strongly influence the semantics of communications. If for example an agent appoints another agent to be group leader, an explicit acceptance by other agents is not necessary (in contrast to the joint acceptance of a commitment as in commitment-based semantics) if the appointing agent already has the necessary power granted by existing social structures. Empirical-Rational semantics supports such pre-structuring, and the extrapolation of past interaction experiences (if, e.g., the appointing agent has been successful in the past, it becomes more likely he also will with his new appointment, even before this proposal has been accepted explicitly).

Support for agent generalizations and mass communication Current approaches to ACL semantics are intended primarily for dyadic situations. Some of them allow for message broadcasts, but they lack a concept for unification and weighting of multiple messages or, respectively, responses, to reflect a (possibly inconsistent) common point of view of multiple agents, or to enable collaboration in joint communicative action. It is hardly imaginable, how thousands or even millions of agents shall contribute to, e.g., the Semantic Web, if social agents are unable to generalize upon their communications by means of statistical evaluation. Whereas our current formal framework still focusses on 1:1 communication, and does not yet support generalization, it allows for the stochastic representation of communication processes, providing a basis for the future inclusion of the described features.

2.2 Empirical-Rational Semantics

The three central assumptions underlying our approach are that 1) the meaning of communications lies primarily in their expectable, observable consequences (a view which was for artificial agency first formulated in [8]), that 2) these consequences can be learned from the observation and as extrapolations of past communication processes (without too much reasoning about what is "inside the agents heads", which significantly reduces the complexity of the learning task), and that 3) the meaning of communications might evolve during the interaction processes. Please refer to [1, 8, 9] for theoretical justifications of these assumptions. The basic requirements in order to put these assumptions into practice are the presence of a *semantics observer* which derives communication semantics from observations, and a knowledge medium which represents the assumably shared semantics among the agents as interrelated stochastic expectations (*Expectation Network*).

In our communication model (which does not follow speech act theory), a single communication can be seen as a request to act in conformance with a desired state declared by its utterance, in which this state is given as a probability distribution of future events, and the meaning of the utterance is the probability distribution of expected events subsequent to the utterance. If one agent e.g. utters "Close the door" to another agent, the desired world state is the door being closed by the addressed agent, and the meaning of that utterance is if and how the addressed agent works towards this state, possibly together with side-effect as the sanctioning in case of non-compliance. As another example, if an agent performs the act "You are the group leader now", then this act demands that other agents act as if the addressed agent would perform like a group leader from now on. In a strict sense, even this performative act will become successful only a posteriori, but if the nominating agent has been assigned the necessary social power in the past, its success can be derived immediately from past successes empirically.

In contrast to non-communicative events, an utterance has no (significant) direct impact on the physical environment. Instead, its physical consequences are achieved socially and indirectly, and, most important, the addressee is free to deny the communicated proposition. Since an utterance is always explicitly produced by a self-interested agent to influence the addressee which is not already convinced from the necessity of the proposal, communicated content will very likely not "believed" immediately, but needs to be accompanied with communicated reasons given to the addressee to increase the probability of an acceptance of the communicated content. This can be done either explicitly by previous or subsequent communications (especially *reciprocally*: "If you comply, I'll comply too"), or implicitly by means of generalizations from past events (e.g., trust) or given social structures. The whole of the expectations which are triggered by a communication in the context of the preceding communication process we call its *rational hull*. The rational hull specifies the rational social relationships which steer the acceptance or denial of communicated content according the rational attitude the agents exhibit. Typically, a rational hull is initially very indefinite and becomes increasingly definite in the course of interaction, provided that the agents work towards

$$
\begin{aligned}
Expect &\in [0;1] \\
Agent &\rightarrow \mathit{agent\_1} \quad | \quad \ldots \quad | \quad \mathit{agent\_n} \\
PhysicalAction &\rightarrow \mathit{move\_object} \quad | \quad \mathit{touch\_agent} \quad | \quad \ldots \\
Action &\rightarrow ECA(Agent, Projection) \\
& \qquad | \quad do(Agent, PhysicalAction) \\
ActionPattern &\rightarrow Action \quad | \quad ? \\
Projections &\rightarrow (Conditions, GoalStates) \\
Conditions &\rightarrow SimplePath \\
GoalStates &\rightarrow SimplePath \\
SimplePath &\rightarrow Action\ SimplePath \quad | \quad \varepsilon
\end{aligned}
$$

Table 1.

A grammar for event nodes of ENs, generating the language $\mathcal{M}$ (the language of concrete actions, starting with $Action$).

mutual understanding. The utterances themselves are modeled as pointers pointing to the desired/proposed states within the Expectation Network (thus denoting subjective expectation directed to other agents in contrast to the objective expectations maintained by the semantics observer).

3 Expectation Networks

Expectation Networks (ENs) are the graphical data structures we want to use for the stochastic modeling of Social Interaction Structures, which in turn represent the semantics of utterances in the form of EN branches. The formal EN definition we present in this work is an improved yet simplified version of the definition presented in [10], itself based on the definition of Expectation Networks provided in [8].
The central assumption that is made in ENs is that observed events like agent actions (especially symbolic agent messages) may be categorized as expected continuations of other observed event sequences. An edge leading from event m to event m' is thought to reflect the probability of m and m' being correlated from the observer's point of view (the descriptive power of ENs is thus similar to Markov processes, but in contrast edges in ENs relate events, not states).

As for $\mathcal{M}$, this is a formal language that defines the events used for labeling nodes in expectation networks. Its syntax is given by the grammar in table 1. Agent actions observed in the system can be either "physical" actions of the format (a, ac) where a is the executing agent, and ac is an domain-dependent symbol used for a physical action, or symbolic elementary communication acts $ECA(a, c)$ sent from a to another agent with content c. We do not talk about "utterances" or "messages" here, because a single utterance might need to be decomposed into multiple ECAs. The symbols used in the $Agent$ and $PhysicalAction$ rules might be domain-dependent symbols the existence of which we take for granted. For convenience, $agent(eca)$ shall retrieve the acting agent of an ECA *eca*.

In addition to normal node labels, we use the symbol $(\triangleright_{EN})$ to denote the root node of an specific EN. The special symbol ? marks pseudo-nodes which are just graphical abbreviations for the so-called $completeEN$ which models the uniform distribution of all possible combinations and sequences of observable events (see below). A "node" labeled with ? thus stands for a branch with infinite depth. The content c of a non-physical action is given by type $Projections$. The meaning of $Projections$ will be described later.

Syntactically, expectation networks are here represented as lists of edges (m, p, n) where m and n are actions, and p is a transition probability (*expectability*) from m to n. We use functions $in : V \to 2^C$, $out : V \to 2^C$, $source : C \to V$ and $target : C \to V$ which return the ingoing and outgoing edges of a node and the source and target node of an edge, respectively. $children : V \to 2^V$ returns the set of children of a node, with $children(v) = \emptyset$ in case v is a leaf. $\prec : V \times V \to \{true, false\}$ returns $true$ iff there is a path leading from the first argument node to the second and the event associated with the second node is expected to occur after the event of the first node. C is the set of all edges, V the set of all nodes in the EN. Edges denote correlations in observed communication sequences. Each cognitive edge is associated with an expectability (returned by $Expect : C \to [0; 1]$) which reflects the probability of $target(e)$ occurring after $source(e)$ in the same communicative context (i.e. in spatial proximity, between the same agents, etc.).

Sometimes we denote a path p in an EN leading from $v_0 \in V$ to $v_n \in V$ as concatenations of message labels (ECAs) $Label(v_0) \sqcup ... \sqcup Label(v_n)$. The $\sqcup$ are sometimes omitted for shortness. $|p| := n$. $Node : SimplePath_{\mathcal{EN}} \to V$ results in the last node of a certain path given as a string of labels. Nodes or corresponding messages along a path p will be denoted as p_i. $\mathcal{EN}(\mathcal{M})$ is the set of all possible expectation networks over $\mathcal{M}$.

Definition 1. An *Expectation Network* is a structure

$$EN = (V, C, \mathcal{M}, Label, Expect) \in \mathcal{EN}(\mathcal{M})$$

where

- V with $|V| > 1$ is the set of nodes,
- $C \subseteq V \times V$ are the edges of EN. (V, C) is a tree called *expectation tree*. (V, C) shall have a unique root node called $\triangleright_{\mathtt{EN}} \in V$ which corresponds to the first ever observed action[2]. The following condition should hold:

$$\forall v \sum_{e \in out(v)} Expect(e) = 1$$

- $\mathcal{M}$ is the *action language.* As defined in table 1, actions can be symbolic ($ECA(...)$) or physical actions ($do(...)$). While we take the existence and the meaning of the

[2] Of course, there are semantics observers imaginable which maintain multiple ENs to model different social systems, states of knowledge or environmental domains at the same time.

latter in terms of resulting observer expectations as granted and domain-depended, the former will be described in detail later.

– $Label : V \rightarrow \mathcal{M}$ is the *action label* function for nodes, with $\forall v \in V : \forall e, f \in children(v) :$
$\neg unify(Label(e), Label(f))$ (where $unify$ shall be $true$ iff its arguments are syntactically unifiable. Cf. [10] for the use of variables in ENs),
– $Expect : C \rightarrow [0; 1]$ returns the edges' expectabilities. For convenience, we define $Expect(label|path) = Expect(in(v))$ if $Node(path \sqcup label) = v$.

Paths starting with $\triangleright$ are called *states* (of communication)[3]

4 Social Interaction Structures

Based on the definition of ENs, we can now define *Social Interaction Structures* as a special kind of communication structures. Social Interaction Structures capture the regularities of externally observed communication processes and other assumably publicly observable events (the latter can be considered as being communicated "by doing", or as projected information). The basic ideas behind this concept are that 1) agent sociality emerges from agent communication, and that 2) communications form a so-called *social system* which is closed in the sense that, to some degree, communication regularities come into being from communications themselves [1], such that the semantics observer does not need to have to "look inside the agents' heads" to derive these structures. Because of that, communication structures can meaningfully be learned from observations. Nevertheless, this learning process needs to be continuously repeated to adapt the EN to new perceptions (since open systems with truly autonomous agents with unknown life spans have no final state), and does always imply the possibility of failure of its prediction task (yet the term "expectation"). The Social Interaction Structures (respectively the probabilistic distribution it represents, as, e.g., an EN branch) following an utterance (the node denoting the ECA which is part of this utterance, to be precise) [4] is called the *semantics* of this utterance.

4.1 Social Interaction Systems

In [10], we've introduced *Communication Systems* as a universal means for the description of social dynamics of multiagent systems. The two main purposes of a Communication System are i) to capture the social expectations (represented as an EN) in the current state of a multiagent system under observation, and ii) to capture changes to these expectations. Whereas the EN models the meaning of communicative action sequences at a certain time (i.e., their expected, generalized continuations in a certain context of previous events), the communication system models the way the EN is build

[3] Actually, two different paths can have the same semantics in terms of their expected continuations, a fact which could be used to reduce the size of the EN by making them directed graphs with more than one path leading to a node instead of trees as in this work.

[4] Usually, this context is build up from previous events, but it would also be possible that utterances become contextualized (e.g., more specific) by *succeeding* utterances.

up, and, if necessary, adapted according to new observations of events. We introduce now *Social Interaction Systems* (SIS) as a concrete kind of general Communication Systems. The difference between general Communication Systems and Social Interaction Systems is that the latter come with a concrete EN learning algorithm, whereas for general Communication Systems we just demand unspecifically that the expectations within learned ENs shall reflect the expectation of the semantics observer regarding the future course of events [10], not specifically taking into account agent rationality and social commitment. The term "interaction system" comes from social systems theory [1], where it denotes the most basic kind of communication (=social) system.
As seen in table 1, we also allow purely physical, non-symbolic events to be contained, like agent actions, but without projections. So the EN of an SIS comprises physical states of the domain too, as far as these are visible for the semantics observer, and of course physical events projected by ECAs.
The SIS maintained by the semantics observer is also the *assumably* shared world knowledge the agents use as the common ground for their uttered ECAs. Social Interaction Systems are thus two dimensional, in the sense that they do not only contain expectations regarding actual agent behavior including utterances (first dimension), but also descriptions of the imaginative behavior which the uttering agents tries to bring about or demand, i.e., which they expect other agents to do.

Definition 2. A *Social Interaction System* at time t is a structure

$$SIS_t = (\mathcal{M}, f, \varpi_t, \rho)$$

where

- $\mathcal{M}$ is the formal language used for agent actions (according to table 1),
- $f : \mathcal{EN}(\mathcal{M}) \times \mathcal{M} \rightarrow \mathcal{EN}(\mathcal{M})$ is the *expectations update function* that transforms any expectation network EN to a new network upon experience of an action $m \in \mathcal{M}$. $f(\perp, m)$ returns the so-called *initial EN*, transformed by the observation of m. This initial EN can be used for the pre-structuring of the social system using given e.g. social norms or other a-priori knowledge which can not be learned using f. Any ENs resulting from an application of f are called *Social Interaction Structures*.
  As a non-incremental variant we define $f : \mathcal{M}^+ \rightarrow \mathcal{EN}(\mathcal{M})$ to be
  $f(m_0 \sqcup m_1 ... \sqcup m_t) = f(...(f(f(\perp, m_0), m_1)...), m_t)$,
- $\varpi_t = m_0 \sqcup m_1 ... \sqcup m_t \in \mathcal{M}^*$ is the list of all actions observed until time t. The subindexes of the m_i impose a linear order on the actions corresponding to the times they have been observed[5],
- $\rho \in \mathbf{N}$ is a time greater of equal the expected life time of the SIS. We require this to calculate the so-called *spheres of communication* (see below). If the life time is unknown, we set $\rho = \infty$.

We refer to events and EN nodes as *past*, *current* or *future* depending on their timely position (or the timely position of their corresponding node, respectively) before, at or

[5] For simplicity, we assume a discrete time scale with $t \in \mathbb{N}$, and that no pair of actions can be performed at the same time, and that the *expected* action time corresponds with the depth of the respective node.

after t. We refer to $EN_t = f(\varpi_t)$ as the *current EN* from the semantics observer's point of view, if the semantics observer has observed exactly the sequence $m_0 m_1 ... m_t$ of events so far.
The intuition behind our definition of SIS_t is that a social interaction system can be characterized by how it would update an existing expectation network upon newly observed actions $m \in \mathcal{M}$. The EN within SIS_t can thus be computed through the sequential application of the structures update function f for each action within ϖ, starting with a given expectation network which models the observers' a-priori knowledge. ϖ_{t-1} is called the *context* (or *precondition*) of the action observed at time t.
To simplify the following formalism, we demand that an EN ought to be implicitly complete, i.e., to contain *all* possible paths, representing all possible event sequences (thus the EN within a social interaction system is always infinite and represents all possible world states, even extremely unlikely ones). If the semantics observer has no a-priori knowledge about a certain branch, we assume this branch to represent uniform distribution and thus a very low probability for every future decision alternative ($\frac{1}{|M|}$), if the action language is not trivially small.
Note that any part of an EN of an SIS does describe exactly one time period, i.e., each node within the respective EN corresponds to exactly one moment on the time scale in the past or the future of observation or prediction, respectively, whereas this is not necessarily true for ENs in general. For simplicity, and to express the definiteness of the past, we will define the update function f such that the a-posteriori expectabilities of past events (i.e., observations) become 1 (admittedly leading to problems if the past is unknown or contested, or we would like to allow contested assertive ECAs *about* the past). There shall be exactly one path pc in the current EN leading from start node $\triangleright_{en_t}$ leading to a node pc_t such that $|pc| = t$ and $\forall i, 0 \leq i \leq t : Label(pc_i) = m_i$. The node pc_i and the ECA m_i are called *corresponding*.
The *semantics* of ϖ_t (i.e. m_t within context ϖ_{t-1}) is defined as the probability distribution Δ_{EN_t,ϖ_t} represented by the branch starting with the node within EN_t that corresponds to ϖ_t:

$$\Delta_{EN_t,\varpi_t}(w') = \frac{\prod_{i,1\leq i\leq |w'|} Expect(w'_i|\varpi_t w'_1 ... w'_{i-1})}{\sum_{m\in M^+} \prod_{i,1\leq i\leq |m|} Expect(m_i|\varpi_t m_1 ... m_{i-1})}$$

for all $w' :\Leftrightarrow \varpi_t \sqcup w' \in M^+$. The w'_i denote single event labels along w'.

4.2 Projections

As defined in table 1, ECAs consist of two parts: The uttering agent, and the ECA content (*projections*). Each projection is a set of EN node pairs which are derived from the following two syntactical elements (cf. table 1)[6].

– *Conditions* chooses, using an EN path (without expectabilities), a possibly infinite set of EN states which have to become reality in order to make the uttering agent

[6] Future version of our framework might allow the utterance of whole ENs as projections, in order to freely project new expectabilities or even introduce novel event types not found in the current EN.

start to act towards its uttered goal (e.g. in "*If I deliver the goods*, you must pay me the money"). As shown in table 1, conditions are given as a linear list of node labels. This path must match with paths in the current EN, either absolutely beginning with $\triangleright$, or starting at nodes after the node which corresponds to the ECA. The end nodes of all matches in EN are called the *condition nodes* of the ECA projections. So, if the node list is empty, the only condition node is the node corresponding to the ECA. The path matching is always successful, since in our model, an EN implicitly contains all possible paths, although with a probability near zero for most of them.

– *GoalStates* chooses, using an EN path (without expectabilities), the (possibly infinite) set of states of the expectation network the uttering agent is expected to strive for. The uttered *GoalStates* path must match with a set of paths within the EN such that the last node of each match is a node of an EN branch that has a condition node from *Conditions* as its root. Both in Conditions and GoalStates paths, wildcards "?" for single actions are allowed.

  For the purpose of this paper, we demand that the projections either refer to future interactions or be semantically inactive (i.e., they already failed or have been successful). Theoretically, we could also imagine projections regarding the past. In this case the respective ECA would express that the uttering agent will likely try to change the way other agents communicate about the past, but we do not consider this difficult and rather unusual case here for simplicity.

  Note also that projected goal states possibly describe actions the uttering agent announces to perform *himself*, not just explicit demands directed to other agents. In this case, the rational hull for this goal consists of behavior which likely increases the support from other agents in order to make the goal state come true.

In the context of an EN, every projection implicitly refers to previous or future projections which announce *reasons* or positive or negative *sanctions* the uttering agent would impose on the ECA receiver in case of a positive or negative response to the ECA. So, in our model, sanctions and argumentative reasons are projections also, in order to support the realization of other projections (of course, this can be continued recursively, e.g. projections in order to support sanctions), and learned from previous processes as anticipations of future reasons and sanctions[7]. The projection of accompanying reasons and sanctions is an inevitable part of every elementary communication act, since among self-interested agents it would be unreasonable to make propositions without providing any reciprocative utility to the receiver, with the exception of implicit reasons and sanctions given as pre-existing social structures social structures like social power, laws or other norms (which we do not consider in this work). Such supporting projections can be either unspecified, to be specified later, or already be specified by means of previous events. Of course, like any other kinds of projections, they need not to be "honest", or put into action, or be effective.

Because the projections set can represent arbitrary probability distributions, it is possible for multiple ECAs to express disjunctive statements like "I want you to do either a *or* b", if *a* and *b* are inconsistent events (i.e., events which cannot occur both in the

[7] In order to model explicit argumentation or social reasoning systems as special cases of Social Interaction Systems, we would additionally need to provide an explicit logical interpretation of ENs, which our framework does not yet accomplishes.

same context). Since consistent ECAs uttered by the same agent are interpreted as conjunctively related, and ECAs with redundant projections are allowed (which increases its impact of these projections on the social structures), one can project arbitrary probability distributions using multiple ECAs. The following functions returns the set of projections of a single ECA $ECA(condition, goal) \in M$ with paths $condition \in Conditions$ and $goal \in GoalStates$:

$$
\begin{aligned}
& projections_{\mathcal{EN}(\mathcal{M})} : M \to V \times V \\
& projections_{(V,C,M,Label,E)}(ECA(ce_1...ce_n, ge_1...ge_m)) = \\
& \{(v_n, v_m) : \{(v_i, v_{i+1}) : 1 \leq i \leq n-1\} \subseteq C \\
& \wedge\, unify(Label(v_i), ce_i) \\
& \wedge\, \{(v_i, v_{i+1}) : n+1 \leq i \leq n+m-1\} \subseteq C \\
& \wedge\, unify(Label(v_i), ge_i) \\
& \wedge\, v_n \prec v_{n+m} \wedge unify(Label(v_n), ce_n) \\
& \wedge\, unify(Label(v_{n+m}), ge_m)\}
\end{aligned}
$$

$unify(?, l)$ and $unify(l, ?)$ shall always be true. For convenience, we write $Goal((c, g)) = g$ and $Condition((c, g)) = c$.

4.3 Rational hulls

Per se, a projection has no power to make its goal states become true. In fact, projections don't have to be rational at all. But we consider it to be rational that the uttering agent will act towards the projected events *at least for some significant amount of time* ("allegedly rational", so to say) [8]. This time span and the events within, starting directly after the projecting utterance event, are called *sphere of communication* (cf. figure 1). Theoretically, each ECA could have its own sphere of communication. For simplicity, in this work we assume that the sphere of communication of any ECA eca is simply $\rho - time(eca)$, where the first operand is the expected time of the last observed utterance within the SIS, and the second is the utterance time of the projecting ECA. This setting is assumable realistic for small and simple interaction systems, where the interacting agents likely stick to their opinions and desires for the whole and usually short duration of the SIS (like auctions). For other domains we would have to determine the spheres of communication *a posteriori* from empirical observations.
The actions the uttering agents is expected to perform within the respective sphere of communication in order to make his projections come true is called the *rational hull* of the ECA. Thus, the determination of the rational hulls of observed ECAs constitutes a crucial part of the determination of ACL semantics. The rational hull can be seen as the actual pragmatics and meaning "behind" the more normative and idealistic concept of social commitments.

We assume the manifestation of the following attitudes by means of ECAs *within the respective spheres of communication* and contextualized by means of other ECAs:

– *Information of other agents about desired states of communication.* This information is given as projections as described above.

[8] This time span of projection trustability can be very short though - think of *joke questions*.

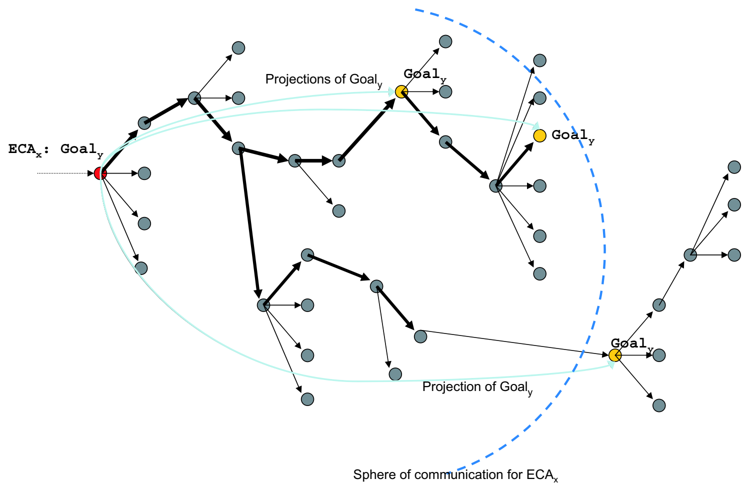

Fig. 1. An EN with projections and a sphere of communication

– *Support of other communicated goals.* The supportive functionality communication has regarding other communications is defined by the rational hulls of the supported elementary communication acts, which will become implicitly more expectable too if supporting rational hulls increase their own expectabilities.
– *Manifestation of understanding.* In case the agents "understand" each other, ECAs cannot express contradiction to the fact that other ECAs pursue the two previous intentions (i.e., Agent 1 does not need to believe Agent 2 is right, but she needs to believe at least that Agent 1 *wants* to be right in a specific case). We do not consider misunderstanding in this work.

Capturing these intentions, and given the set of projections for each ECA eca uttered by an agent a, we calculate the rational hull of a certain ECA using the following two principles.

4.3.1 Bounded rational choice After uttering eca, an agent a is expected to choose an action policy such that, within the respective sphere of communication, his actions maximize the probability of the projected state(-s). Let $p \in projections(eca, EN_t)$ be a projection. Then, considered that p would be a useful state for the uttering agent to be in, the rule of rational choice proposes that for every node v_d with $agent(v_d) = a$ along the path $v_t...p$ leading from the current node v_t to p, $Expect(in(v_d)) = 1$ for the incoming edge of v_d, and that the expectabilities of the reminding outgoing edges of the predecessor of v_d are reduced to 0 appropriately (if no other goals have to be considered). To reduce the complexity of applying this general rule on the possibly infinite projections set, and to observe the bounds of observer rationality, we propose the following constraints:

- Expectabilities will be adapted within the respective sphere of communication of eca only, even if the goal state p is located beyond this sphere.
- Expectabilities will be adapted only for parts of the current EN with a significant evidence regarding actions performed by other agents. Since we represent missing knowledge as uniform distribution, we put this rule into practice by demanding that at decision nodes of other agents (i.e., nodes with children representing actions of agents other than the agent which uttered eca) the *expectabilities entropy* $entropy_{EN} : V \to \mathbb{R}$ should be below some given limit.
  $entropy_{EN}(v) =$
  $\sum_{v' \in children(v)} -Expect(in(v'))log_2 Expect(in(v'))$
- If multiple elements in $projections$ are identical despite their context, and the paths leading to these projections overlap, priority is given to those projections with a higher cumulative expectability. Finding the right paths is a markovian multiple-decision problem from the perspective of the uttering agent a (and thus from the perspective of the semantics observer which models the behavior of a also), which in general cannot simply be solved by pairwise comparison of paths leading from the current node to the competitive projections regarding their maximum expected utilities, if $projections(eca, EN_t) = \{p_1, ..., p_n\}$ contains more than two elements.
- The projections of previously uttered ECAs have to be maintained, so the rule of rational choice needs to do a weighting assessment of previously calculated rational hulls instead of simply outdating them.

We use the following function $u_{\mathcal{EN}(\mathcal{M})} : \mathcal{M} \times V \to [0; 1]$ to calculate the *utility* of an arbitrary node v regarding its supporting function for a specific elementary communication act eca.

$$u_{EN}(eca, v) = \begin{cases} 0 \text{ if } \forall i, 1 \le i \le n : \\ \quad \neg v \prec Goal(p_i) \vee \neg Condition(p_i) \prec v \\ 0 \text{ if } entropy_{en}(v) > \kappa, \text{ or else:} \\ 1 \text{ if } \exists i : v = Goal(p_i) \\ \max\limits_{j, 1 \le j \le c} u_{EN}(eca, vc_j) \\ \quad \text{if } agent(Label(vc_j)) = agent(eca) \\ \max\limits_{j, 1 \le j \le c} Expect(in(vc_j)) u_{EN}(eca, vc_j) \\ \quad \text{otherwise} \end{cases}$$

with $\{p_1, ..., p_n\} = projections(eca)$, $\{vc_1, ..., vc_c\} = children(v)$, and κ being some predefined entropy maximum.
$\max(...)$ could be replaced with $(\sum_{j,1 \le j \le c} ...)/c$ to prefer a high number of paths leading to a goal instead of the highest expectability for one goal node.

Figure 1 shows an EN modeling the future of some communication process. ECA_X is an utterance which encodes $Goal_Y$. This goal itself stands for several (seemingly) desired states of the EN (yellow nodes). Since within the so-called *sphere of communication* of ECA_X (see below) it is expected that the uttering agent rationally strives for these states, certain EN paths leading to these states become more likely (bold edges). Such behavior paths need to be (more or less) rational in terms of their expected utility (e.g. in comparison with competing goal states), and they need to reflect experiences from analogous agent behavior in the past.

4.3.2 Empirical stationarity assumption The following describes the "purely" empirical aspects of our approach. If we would use the previous rule (ostensible rationality) as the only EN updating mechanism, we would face at least three problems: 1) Predicting agent actions according to the rule of rational choice requires some given evidence about subsequent actions of other agents. In case this previous evidence is missing, the rule of rational choice would just "convert" uniform distribution into unform distribution. Therefore, we have to provide an initial probability distribution the rule can be applied on[9]. 2) the set of projections for a single ECA might be infinite. Most of the expectabilities along the paths leading from the current node to these EN branches sum up to very low probabilities for the respective projection. Thus, a pre-selection of likely paths will be necessary. And most important 3), the rule of rational choice does not consider individual behavioral characteristics like (initially opaque) goal preferences of the agents, but treats all projections uniformly. Goal hierarchies need thus to be obtained from past agent practice as well as individual strategies towards these projections. For these reasons, we combine the application of the rule of rational choice with the assumption of some stationarity of past event trajectories, i.e., the assumption that previously observed action sequences repeat themselves in the future in a similar context. We use this assumption to retrieve a probability distribution the rule of rational choice can be applied on and weighted with subsequently.
In order to learn EN stationarity from previous observations, we follow the so-called *variable-memory approach* to higher-order Markov chains using *Probabilistic Suffix Automata* (PSA) introduced for *L-predictable* observation sequences [7]. This approach efficiently models Markov chains of order L (i.e., with a model memory size of L), allowing for rich stochastical models of observed sequences. The applicability of this approach to our scenario is based on the heuristical assumption that many Social Interaction Systems are *short-memory systems*, which allow the empirical prediction of social behavior from a relatively short preceding event sequence (assumedly pre-structuring using social norms , constraints from rational choice etc is done properly). The main characteristic of the PSA-based approach is its straightforward learning method, with expressiveness and prediction capabilities comparable with the more common *Hidden Markov Models* [7].
For the calculation of the PSA from a set of sample agent action sequences, we use an algorithm introduced in [7], originally coming from *PAC-learning*, in a slightly mod-

[9] This probability distribution must also cover projected events and assign them a (however low) probability even if these events are beyond the spheres of communication, because otherwise it would be impossible to calculate the rational hull.

ified version. It constructs a so-called *Prediction Suffix Tree* (PST) (sometimes called *Probabilistic Suffix Tree*) from the samples, which is roughly equivalent to the target PSA, but easier to build up. Its only disadvantage in comparison to the corresponding full PSA is that the time complexity for the predicting task is higher approximately by the factor L.

Definition 3. A *Prediction Suffix Tree* with memory size L over the language of concrete agent actions M is a structure $PST_L(M) = (V, C, Label, \gamma)$ where

- (V, C) defines a tree graph consisting of a set of nodes $V, |V| > 0$ and a set of edges $C \subseteq V \times V$,
- $Label : V \rightarrow M^+$ returns for a node its label (with maximum length L),
- $\gamma : V \rightarrow \{(d_1, ..., d_{|M|}) : d_i \in \mathbb{R}\}$ returns for each node a vector which defines the probability distribution associated with this node. Each element $\gamma_\sigma(v)$ of the resulting vector corresponds to the conditional probability of the particular message σ in M.
  $\sum_{\sigma \in M} \gamma_\sigma(v) = 1$ should hold - nevertheless, vector elements with a very low probability are omitted.

A PST is able to predict the probability of sequences using a tree traversal up to the root, as γ returns for a specific message its conditional occurrence probability given that the largest *suffix* ν, $|\nu| \leq L$, of the message sequence observed before matches with the label of this node. L should depend from the available memory resources, the length of the samples and the expected spheres of communication.
In order to build up the PST from the empirical observations, we need to define the conditional empirical probability within a set of sample action sequences (where actions are either ECA utterances or physical actions). As input we us the set $samples_{SIS_t} = \{m_0m_1..., m_t\} \cup \{r_1^1...r_1^{l_1}, ..., r_n^1...r_n^{l_n}\}$, where $m_0m_1..., m_t$ is the sequence of events observed so far for SIS_t until time t, and the reminder of this set consists of additional samples to improve prediction accuracy. The $r_i^1 r_i^{l_i}$ are optional; we can omit these additional samples and learn the PSA from the single sequence $m_0m_1..., m_t$ only. But as a rule of thumb, the lengths of the sample sequences should be at least polynomial in L[7]. If an a-priori EN is given for pre-structuring, the r_i could be obtained from a frequency sampling of sequences from this EN, which is straightforward and thus omitted here. For lack of space, we also omit the detailed PST-learning algorithm, which can be found in [7].
The probability for the PST-generation of an event sequence $m = m_1...m_n \in (M)^n$ is

$$P_{PST}(m) = \prod_{i=1}^{n} \gamma_{m_i}(v^{i-1})$$

where v^0 is the (unlabeled) root node of the PST and for $1 \leq i \leq n-1$ v^i is the deepest node reachable by a tree traversal corresponding to a prefix of $m_i m_{i-1}...m_1$, starting at the root node.
From the probability distribution obtained from P_{PST}, we derive the corresponding EN using the function $\delta : M^+ \rightarrow \mathcal{EN}(\mathcal{M})$:

$$\delta(m_0m_1..., m_t) = (V, C, M, Label, Expect)$$

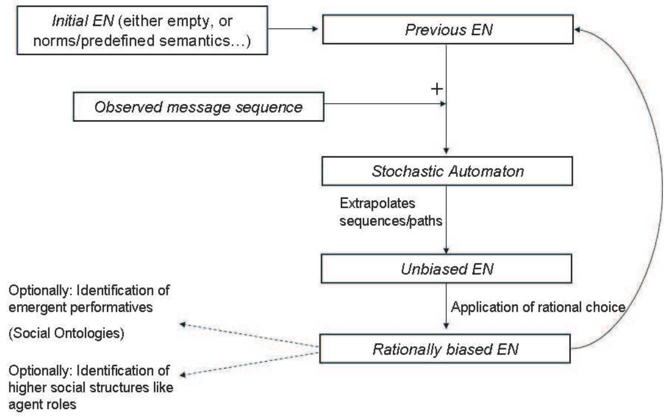

Fig. 2. Iterative version of the algorithm (outline)

with
$V = \{\triangleright\} \cup \{v_p : p \in paths\}$,
$Label = \{v_{p_1 \ldots p_n} \mapsto p_n : p_1 \sqcup \ldots \sqcup p_n \in paths\}$,
$C = \{(\triangleright, v_p) : |p| = 1, v_p \in V\}$
$\cup \{(v_{p_1 \ldots p_{n-1}}, v_{p_1 \ldots p_n}) : v_{p_1 \ldots p_{n-1}} \in V \wedge v_{p_1 \ldots p_n} \in V\}$,
$Expect =$
$\{in(v_{p_1 \ldots p_n}) \mapsto \dfrac{P_{PST}(p_1 \ldots p_n)}{P_{PST}(p_1 \ldots p_{n-1})}, v_{p_1 \ldots p_n} \in V\}$, and
$paths = \{p : p \in M^+ \wedge P_{PST}(p) > P_{min}\}$, where P_{min} is a predefined lower bound for significant expectabilities.

4.3.3 Rationality-biased empirics Putting together the rule of rational choice and the assumption of empirical stationarity, we gain the following (non-iterative) definition for the Social Interaction Structures update function f of an SIS. Figure 2 outlines the iterative counterpart not described here.

$$f(m_0 m_1 \ldots m_t) = \varrho(EN_{stat}, \triangleright_{EN_{stat}})$$

with $EN_{stat} =$
$(V_{EN_{stat}}, C_{EN_{stat}}, \mathcal{M}, Label_{EN_{stat}}, Expect_{EN_{stat}})$ such that
$V_{EN_{stat}} = \{v_{m_0}, \ldots, v_{m_t}\} \cup V_\delta$,
$C_{EN_{stat}} =$
$C_\delta \cup \{(\triangleright_{EN_{stat}} = v_{m_0}, v_{m_1}), \ldots, (v_{m_{t-1}}, v_{m_t}), (v_{m_t}, \triangleright_\delta)\}$
and $\forall i, 1 \leq i \leq t :$
$Expect(in(v_{m_i})) = 1, \forall i, 0 \leq i \leq t : Label(v_{m_i}) = m_i$, with
$(V_\delta, C_\delta, \mathcal{M}, Label_\delta, Expect_\delta) = \delta(m_0 m_1 \ldots m_t)$.

$Expect(in(v_{m_i})) = 1$ reflects the definiteness of already observed events.

Above, $\varrho : \mathcal{EN}(\mathcal{M}) \times SimplePath \rightarrow \mathcal{EN}(\mathcal{M})$ applies the results of the calculation of rational hulls to the entire EN resulting from the PST by means of a recursive top-down tree traversal which is limited by the maximum search depth maxdepth (alternatively, we could apply a entropy-based search limitation criterion similar to the criterion used in 4.3.1).

$$\varrho((V, C, M, Label, Expect), path) =$$

$$\begin{cases} (V, C, M, Label, Expect) \text{ if } |path| > \text{maxdepth} \\ (V, C, M, Label, Expect_{|children(v)|}) \text{ otherwise} \end{cases}$$

using $v = Node(path)$, $\Delta U(v) = \{(v_j, u(Label(v), v_j)) :$
$v_j \in V, agent(Label(v_j)) = agent(Label(v))\}$,

$$\forall v_j \in V : Expect_0(in(v_j)) =$$
$$\begin{cases} \dfrac{Expect(in(v_j)) + \Delta U(v)[v_j]}{2} \\ \text{if } Time(v_j) < \rho \wedge agent(Label(v_j)) = agent(Label(v)) \\ Expect(in(v_j)) \text{ otherwise} \end{cases} \quad \text{and}$$
$$\forall n, 1 \leq n \leq |children(v)| :$$
$$Expect_n :\Leftrightarrow (V, C, M, Label, Expect_n) =$$
$$\varrho((V, C, M, Label, Expect_{n-1}),$$
$$path \sqcup Label(children(v)_n)).$$

Here, $\Delta U(v)$ assigns every node v_j its utility regarding the ECA $Label(v)$, if the acting agent is the same for v and v_j. $Expect_0(in(v_j))$ assigns the node its new expectability (equally weighted with its previous expectability, which might be already be utility biased from another ECA), and $Time(v_j) < \rho$ limits the application to nodes within the sphere of communication. $\Delta U(v)[v_j]$ denotes the utility for reaching v assigned to v_j.

5 Conclusions

We have introduced an approach to the semantics of agent communication which combines features from traditional mentalistic and objectivist approaches. Being a novel and very basic proposal, several important things remain to do:

- ECAs and ENs do not explicitly model logical propositions and their relationships (e.g., in an EN, a low probability for uttering “no” does not automatically increase the probability for uttering “yes”, as it should, and one can express logical statements only indirectly by stating their pragmatical consequences in terms of events).
- To be of practical use with common ACLs, ECAs also need to be obtainable from conventional speech acts, which requires a translation of speech act types into ECA patterns within the EN (and vice versa, in order to learn new speech act types from emergent ENs).
- Related to the previous issue, the explicit emergence of communication symbols as

"shortcuts" for combinations of ECA patterns is not yet supported.
- Meta-communication (communication about communication) is not yet supported.
- The EN learning algorithm does not yet make use of generalizable behavior patterns that *different* agents have in common (like agent roles).

This work has been supported by DFG (German National Science Foundation) under contracts no. BR609/11-2 and MA759/4-2.

References

1. N. Luhmann. Social Systems. Stanford University Press, Palo Alto, CA, 1995.
2. D. C. Dennett. The Intentional Stance. MIT Press, 1989.
3. M. P. Singh. A Social Semantics for Agent Communication Languages. In Procs. of the IJCAI Workshop on Agent Communication Languages, 2000.
4. F. Guerin, J. Pitt. Denotational Semantics for Agent Communication Languages. In Procs. of the 5th International Conference on Autonomous Agents (Agents'01), 2001.
5. P. R. Cohen, H. J. Levesque. Performatives in a Rationally Based Speech Act Theory. In Procs. 28th Annual Meeting of the ACL, 1990.
6. P. R. Cohen, H. J. Levesque. Communicative Actions for Artificial Agents. In Procs. of the First International Conference on Multiagent Systems (ICMAS-95), 1995.
7. D. Ron, Y. Singer, N. Tishby. The Power of Amnesia - Learning Probabilistic Automata with Variable Memory Length. In Machine Learning Vol. 25, p. 117–149, 1996
8. K. F. Lorentzen, M. Nickles. Ordnung aus Chaos – Prolegomena zu einer Luhmann'schen Modellierung deentropisierender Strukturbildung in Multiagentensystemen. In T. Kron (Ed.), Luhmann modelliert. Ansätze zur Simulation von Kommunikationssystemen, Leske & Budrich, 2002.
9. M. Nickles, G. Weiss. Empirical Semantics of Agent Communication in Open Systems. In Procs. of the 2nd International Workshop on Challenges in Open Agent Environments. 2003
10. M. Nickles, M. Rovatsos, W. Brauer, G. Weiss. Communication Systems: A Unified Model of Socially Intelligent Systems. In K. Fischer, M. Florian (Eds.): Socionics: Its Contributions to the Scalability of Complex Social Systems. Springer LNAI, 2004. To appear.
11. M. Nickles, M. Rovatsos, G. Weiss. Empirical-Rational Semantics of Agent Communication. In Procs. of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-04), 2004.
12. W. Brauer, M. Nickles, M. Rovatsos, G. Weiss, K. F. Lorentzen. Expectation-Oriented Analysis and Design. In Procs. of The Second International Workshop on Agent-Oriented Software Engineering (AOSE-01), 2001.
13. R.-J. Beun, R. M. van Eijk, H. Prust. Ontological Feedback in Multiagent Systems. In Procs. of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-04), 2004.
14. R. A. Flores, P. Pasquier, B. Chaib-draa. Conversational Semantics with Social Commitments. In Procs. AAMAS-04 Workshop on Agent Communication (AC-2004), 2004.
15. M. Colombetti, M. Verdicchio. An analysis of agent speech acts as institutional actions. In Procs. of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-02), 2002.
16. M. Rovatsos, M. Nickles, G. Weiß. Interaction is Meaning: A New Model for Communication in Open Systems. In Procs. of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-03), 2003.
17. M. P. Singh. Multiagent Systems - A Theoretical Framework for Intentions, Know-How, and Communications. Springer LNCS 799, 1994.