

Empirical-Rational Semantics of Agent Communication

Matthias Nickles, Michael Rovatsos, Gerhard Weiß
AI/Cognition group, Department of Informatics, Technical University Munich
D-85748 Garching bei München, Germany,
{*nickles, rovatsos, weissg*}@tiki.informatik.tu-muenchen.de

Abstract

The missing of an appropriate semantics of agent communication languages is one of the most challenging issues of contemporary AI. Although several approaches to this problem exist, 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 approach 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, Markov Processes, Artificial Sociality

1. Introduction

Currently, two major approaches to the meaning of agent communication in a broader sense, covering both traditional sentence 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 *objectivist* approaches (e.g. [3], also called *social semantics*) try to determine communication from an external point of view, focussing on public rules. The former approach has some

well-known shortcomings, which eventually led to the development of the latter: At least 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 of normative rules, and has therefore been a significant step ahead. But it oversimplifies social processes favoring traditional sentence-level semantics instead of pragmatics, 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 semantics which is based on the assumption that the meaning of utterances lies basically in their *consequences* in terms of *expectable* future agent actions which can be continuously learned and adapted from observed agent actions [8]. These consequences are represented as probabilistic *Social Interaction Structures*, which are a variant of *Expectation Networks* [9], and they are learned by a *semantics observer*, which can be either an agent participating in the communication himself, or an external agent (e.g., a special middle agent, or a supervision facility 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 model of concrete agent communication processes
- Assumption of mental agent rationality → revisable, probabilistic expectation of limited 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, we do not present the semantical model for a concrete, speech-act based ACL here. Instead, we propose the dynamic semantics of so-called *elementary communication acts* (ECAs) which obtain their meaning not from some given performative 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 future course of behavior (e.g., an informational act would be the request to communicate in conformance with the expressed belief from now on, and a command would be the request to perform the described actions) [9], whereby each ECA can be contextualized with companion social structures resulting from other ECAs to clarify and get accepted the demand (e.g. sanctions).

The reminder of this paper is organized as follows: The next section defines *Expectation Networks* as the data structure used to describe agent interaction. Section 3 extends this definition with a formal learning and adaptation framework for social (i.e., communication) structures, and finally, section 4 draws some conclusions regarding current limitations of our approach and future work.

2. Expectation Networks

Expectation Networks (ENs) are the graphical data structures we will use below for the stochastic modeling of Social Interaction Structures. The formal EN definition we present in this work is an improved yet simplified version of the definition presented in [9].

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).

<i>Expect</i>	∈	[0; 1]
<i>Agent</i>	→	<i>agent_1</i> ... <i>agent_n</i>
<i>PhysicalAction</i>	→	<i>move_object</i> <i>touch_agent</i> ...
<i>Action</i>	→	<i>ECA</i> (<i>Agent</i> , <i>Projection</i>) <i>do</i> (<i>Agent</i> , <i>PhysicalAction</i>)
<i>ActionPattern</i>	→	<i>Action</i> ?
<i>Projections</i>	→	(<i>Conditions</i> , <i>GoalStates</i>)
<i>Conditions</i>	→	<i>SimplePath</i>
<i>GoalStates</i>	→	<i>SimplePath</i>
<i>SimplePath</i>	→	<i>Action SimplePath</i> ϵ

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

Table 1.

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 $in : V \rightarrow 2^C$, $out : V \rightarrow 2^C$, $source : C \rightarrow V$ and $target : C \rightarrow V$ which return the incoming and outgoing edges of a node and the source and target node of an edge, respectively. $children : V \rightarrow 2^V$ returns the set of children of a node, with $children(v) = \emptyset$ in case v is a leaf. $\prec : V \times V \rightarrow \{true, false\}$ 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 communica-

tion sequences. Each cognitive edge is associated with an expectability (returned by $Expect : C \rightarrow [0; 1]$) 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) $Label(v_0) \sqcup \dots \sqcup Label(v_n)$. The \sqcup are sometimes omitted for shortness. $|p| := n$. $Node : SimplePath_{\mathcal{EN}} \rightarrow V$ 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 $\triangleright_{EN} \in V$ which corresponds to the first ever observed action. 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 ($ECA(\dots)$) or physical actions ($do(\dots)$). While we take the existence and the meaning of the latter in terms of resulting observer expectations as granted and domain-dependent, 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. [9] 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. 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.

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) triggered by a certain utterance (or an ECA which is part of this utterance, respectively) in the context of other utterances we call the *semantics* of this utterance.

3.1. Social Interaction Systems

In [9], 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 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 [9], 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.

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

$$f(m_0 \sqcup m_1 \dots \sqcup m_t) = f(\dots(f(f(\perp, m_0), m_1) \dots), m_t),$$

- $\varpi_t = m_0 \sqcup m_1 \dots \sqcup m_t \in \mathcal{M}^*$ 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 observed¹,
- $\rho \in \mathbb{N}$ is a time greater or 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 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 $m_0 m_1 \dots m_t$ 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 action sequences (thus the EN within a social interaction system is always infinite). 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 ($\frac{1}{|\mathcal{M}|}$), 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

¹ 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.

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. To express the definiteness of the past, we will later define the update function f such that the a-posteriori probabilities of past events (i.e., observations) become 1. 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 :

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

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

3.2. 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).

- *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 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 ac-

tions 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. 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 reciprocal utility to the receiver. They can be either unspecified, to be specified later, or already be specified by means of previous events or even predefined social structures like laws or other norms (which we do not consider in this work). 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 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 $goal \in GoalStates$:

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

$$\begin{aligned} & \wedge v_n \prec v_{n+m} \wedge unify(Label(v_n), ce_n) \\ & \wedge unify(Label(v_{n+m}), ge_m)\} \end{aligned}$$

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

3.3. 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*². This time span and the events within, starting directly after the projecting utterance event, are called *sphere of communication*. 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 the 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.
- *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.

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

- *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 rules:

3.3.1. 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 $v_t \dots p$ 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* $entropy_{EN} : V \rightarrow \mathbb{R}$ 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 $projections(eca, EN_t) = \{p_1, \dots, p_n\}$ 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 outdated them.

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

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

with $\{p_1, \dots, p_n\} = projections(eca)$, $\{vc_1, \dots, vc_c\} = children(v)$, and κ being some predefined entropy maximum.

$\max(\dots)$ could be replaced with $(\sum_{j, 1 \leq j \leq c} \dots) / c$ to prefer a high number of paths leading to a goal instead of the highest expectability for one goal node.

3.3.2. Empirical stationarity assumption If we would use the previous rule 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 uniform distribution. Therefore, we have to provide an initial probability distribution the rule can be applied on³. 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

3 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.

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 stochastic 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 modified 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),
- $\gamma : V \rightarrow \{(d_1, \dots, d_{|M|}) : d_i \in \mathbb{R}\}$ 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* $\nu, |\nu| \leq L$, 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 use the set $samples_{SIS_t} = \{m_0 m_1 \dots, m_t\} \cup \{r_1^1 \dots r_1^{l_1}, \dots, r_n^1 \dots r_n^{l_n}\}$, where $m_0 m_1 \dots, m_t$ 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 $m_0 m_1 \dots, m_t$ 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 $m = m_1 \dots m_n \in (M)^n$ 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 $m_i m_{i-1} \dots m_1$, starting at the root node.

From the probability distribution obtained from P_{PST} , we derive the corresponding EN using the function $\delta : M^+ \rightarrow \mathcal{EN}(\mathcal{M})$:

$$\delta(m_0 m_1 \dots, m_t) = (V, C, M, Label, Expect)$$

with

$$V = \{\triangleright\} \cup \{v_p : p \in paths\},$$

$$Label = \{v_{p_1 \dots p_n} \mapsto p_n : p_1 \sqcup \dots \sqcup p_n \in paths\},$$

$$C = \{(\triangleright, v_p) : |p| = 1, v_p \in V\}$$

$$\cup \{(v_{p_1 \dots p_{n-1}}, v_{p_1 \dots p_n}) : v_{p_1 \dots p_{n-1}} \in V \wedge v_{p_1 \dots p_n} \in V\},$$

$$Expect = \{in(v_{p_1 \dots p_n}) \mapsto \frac{P_{PST}(p_1 \dots p_n)}{P_{PST}(p_1 \dots p_{n-1})}, v_{p_1 \dots p_n} \in V\}, \text{ and}$$

$paths = \{p : p \in M^+ \wedge P_{PST}(p) > P_{min}\}$, where P_{min} is a predefined lower bound for significant expectabilities.

3.3.3. 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.

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

with $EN_{stat} = (V_{EN_{stat}}, C_{EN_{stat}}, \mathcal{M}, Label_{EN_{stat}}, Expect_{EN_{stat}})$ such that $V_{EN_{stat}} = \{v_{m_0}, \dots, v_{m_t}\} \cup V_\delta$, $C_{EN_{stat}} = C_\delta \cup \{(\triangleright_{EN_{stat}} = v_{m_0}, v_{m_1}), \dots, (v_{m_{t-1}}, v_{m_t}), (v_{m_t}, \triangleright_\delta)\}$ 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 $(V_\delta, C_\delta, \mathcal{M}, Label_\delta, Expect_\delta) = \delta(m_0 m_1 \dots m_t)$.

$Expect(in(v_{m_i})) = 1$ 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 3.3.1).

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

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

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

$$\forall v_j \in V : Expect_0(in(v_j)) = \begin{cases} \frac{Expect(in(v_j)) + \Delta U(v)[v_j]}{2} & \text{if } Time(v_j) < \rho \wedge agent(Label(v_j)) = agent(Label(v)) \\ Expect(in(v_j)) & \text{otherwise} \end{cases}$$

and

$$\forall n, 1 \leq n \leq |children(v)| : Expect_n := \varrho((V, C, M, Label, Expect_n) = \varrho((V, C, M, Label, Expect_{n-1}), path \sqcup Label(children(v)_n)).$$

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 .

4. Conclusions

We have introduced an approach to the semantics of agent communication which combines features from traditional mentalistic and objectivist approaches. Being a novel proposal, several important things remain to do. Most important, ECAs have no explicit content level (ontological level) in a traditional sense, which is rather implicitly encoded within the EN the ECAs refer to. More specifically, ECAs and ENs currently do not have the power to *explicitly* model logical statements. We propose the combination of situation calculus with ENs for this purpose, similar to the annotation of EN nodes with states of a knowledge base as introduced in [9]. To be of practical use with common ACLs, ECAs also need to be obtainable from conventional speech acts, which requires a translation of performatives into ECA patterns. Another issue is that the EN learning algorithm does not yet make use of generalizable behavior patterns that multiple agents have in common (like agent roles). And finally, we are currently evaluating an implementation of the presented framework regarding its applicability for open multiagent systems.

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] M. Nickles, G. Weiss. Empirical Semantics of Agent Communication in Open Systems. Procs. of the 2nd International Workshop on Challenges in Open Agent Environments. 2003 (position paper)
- [9] M. Nickles, M. Rovatsos, W. Brauer, G. Weiss. Communication Systems: A Unified Model of Socially Intelligent Systems. In K. Fischer, M. Florian: Socionics: Its Contributions to the Scalability of Complex Social Systems. Springer LNAI, 2004. To appear.