Adaptiveness in Agent Communication: Application and Adaptation of Conversation Patterns

Felix Fischer[1,3], Michael Rovatsos[2], and Gerhard Weiss[3]

[1] Department of Informatics, University of Munich, 80538 Munich, Germany
fischerf@tcs.ifi.lmu.de
[2] School of Informatics, University of Edinburgh, Edinburgh EH8 9LE, UK
mrovatso@inf.ed.ac.uk
[3] Department of Informatics, Technical University of Munich, 85748 Garching, Germany
weissg@in.tum.de

Abstract. Communication in multi-agent systems (MASs) is usually governed by agent communication languages (ACLs) and communication protocols carrying a clear cut semantics. With an increasing degree of *openness*, however, the need arises for more flexible models of communication that can handle the uncertainty associated with the fact that adherence to a supposedly agreed specification of possible conversations cannot be ensured on the side of other agents.
In this paper, we argue for *adaptiveness* in agent communication. We present a particular approach that combines *conversation patterns* as a generic way of describing the available means of communication in a MAS with a decision-theoretic framework and various different machine learning techniques for *applying* these patterns in and *adapting* them from actual conversations.

1 Introduction

Traditional approaches to agent communication, with their roots in speech act theory [1], do not respect the *autonomy* of individual agents in that they suppose effects of communication on agent's mental states [25, 3] or a normative quality of publicly visible commitments [7, 26]. In environments involving some degree of *openness* like, for example, design heterogeneity or dynamically changing populations, such a "normative" attitude is put into question by the fact that adherence to supposedly agreed modes of communication cannot be ensured on the side of other agents. While this can be seen as a witness of a fundamental conflict between agent autonomy and the need for cooperation (and communication) with other agents toward a joint goal, there is also a more practical side to this problem.

Compared to the long-established areas of interaction protocol and agent communication language (ACL) research (see, e.g., [12, 9]), the development of agent architectures suitable for dealing with provided communication mechanisms in practical terms has received fairly little attention. As yet, there exists no uniform framework for defining the interface between the inter-agent communication layer and intra-agent reasoning, i.e. how specifications of interaction protocols and communication semantics influence agent rationality or, in turn, are influenced themselves by agents' rational decision-making processes. Moreover, there is a growing concern that most specification methods for ACLs and interaction protocols do not provide sufficient guidance as to which part of the semantics of communication should be specified at a supra-agent level and which part of them is only a result of agents' mental processing and cannot be captured without knowledge of their internal design. Clearly, concentrating on one of these two sides may either overly constrain agent autonomy (i.e., agents would merely "execute" centralised communication procedures that modify their internal states) or lead to uncertainty about the consequences of communication (e.g. in terms of adherence to previously created commitments) and loss of social structure altogether. This poses two central questions:

1. If strict adherence to communication languages and protocols cannot be taken for granted, how can meaningful and coherent communication be ensured?
2. Observing the course of conversations that take place in a MAS, how can agents effectively organise this kind of knowledge and relate it to existing specifications, so that they can actually benefit from it?

An obvious answer to these questions would be to devise a *probabilistic* model of agent conversation, and update it in order to maximise communicative success. There are two problems, though. Firstly, generic "purely" probabilistic models are not very well suited to describe intelligent agents (including symbolic agent communication), since their behaviour is not at all "random". Instead, one would rather like to identify patterns and relational properties of communication (like communication protocols containing variables, for example). The resulting view resembles decision-theoretic learning and reasoning, where the classical paradigm of direct control exerted on an uncertain environment is replaced by a more indirect influence via communication between (and hence via the allegedly rational reasoning processes of) intelligent agents. Secondly, agent communication cannot exist on its own, but is only a means to the end of coordinating or cooperating with respect to some "physical" actions (i.e., communication works as a kind of mediator between actions). Hence, success (or optimality) in communication will somehow have to be defined in terms of the actions it entails.

This view is in line with *empirical communication semantics* [21], where the meaning of an utterance (or sequence thereof) is defined solely in terms of its expected consequences as given by past experience (to say it in terms of speech act theory [1], the meaning of illocutions is defined solely in terms of their expected perlocutions). Currently two different "flavours" of empirical communication semantics exist, borrowing from two different sociological schools of thought. Interaction frames [23] view empirical semantics from the perspective of symbolic interactionism (particularly [8]), thus focusing on how an individual deals with the communication mechanisms available in a given social system, while expectation networks [14] take the (more global) point of view of social systems theory (see, e.g., [13]) to develop methods to analyse the evolving semantics of communication across an entire society of agents.

In this paper, we focus on a particular instance of the interaction frame approach, which is formally defined in section 2. In section 3, we introduce a formal framework for decision-theoretic reasoning about communication, using interaction frames to represent different classes of conversation and thus to structure the reasoning process hierarchically. In section 4, we further use methods from the fields of case-based reasoning, inductive logic programming and cluster analysis to devise a formal scheme for the adaptation of interaction frames from the actual conversations conducted in a MAS, enabling agents to autonomously (i.e., independent of users and system designers) create and maintain a concise model of the different classes of conversation on the basis of an initial set of ACL and protocol specifications. To our knowledge, the work described in this paper constitutes the first approach to adaptive communication management for deliberative, knowledge-based agents, which is an important prerequisite for building agents that communicate and act in full appreciation of the autonomy of their respective peer.

2 Conversation Patterns

The greatest common denominator of the multitude of different methods for specifying ACL semantics and interaction protocols (see, e.g., [15, 27] for examples in this volume) is that they describe the *surface structure* of possible dialogues and logical *constraints* for the applicability of these. The former corresponds to a set of admissible message sequences, the latter may include statements about environmental conditions, mental states of the participating agents, the state of commitment stores, etc. In the most simplistic case, these structure/constraint pairs can be represented as a set of *conversation patterns*, i.e. combinations of a conversation trace and a set of conditions. For example,

$$\left\langle \mathtt{request}(\mathtt{a},\mathtt{b},\mathtt{pay}(\$100)) \rightarrow \mathtt{do}(\mathtt{b},\mathtt{pay}(\$100)),\ \{can(\mathtt{b},\mathtt{pay}(\$100))\}\right\rangle$$

expresses that a request of agent $\mathtt{a}$ is followed by an action if the requestee $\mathtt{b}$ is able to execute the action, i.e. pay $\mathtt{a}$ an amount of \$100. The question serving as a point of departure for the research presented in this paper is how we can build agents that are capable of processing a set of such (conditioned) conversation patterns in a goal-oriented and adaptive fashion, given that the reliability of these specification is contingent on others' (and the agent's own) adherence to their prescriptive content.

Before turning to practical reasoning with and adaptation of conversation patterns, though, we introduce interactions frames as a slightly more complex form of conversation pattern, quoting [4] for a formal definition of a particular instance of the interaction frame data structure. This definition uses a language $\mathcal{M}$ of speech-act [1] like message and action patterns of the form $\mathtt{perf}(A,B,X)$ or $\mathtt{do}(A,Ac)$. In the case of messages (i.e., exchanged textual signals), $\mathtt{perf}$ is a performative symbol (e.g. $\mathtt{request}$, $\mathtt{inform}$), A and B are agent identifiers or agent variables and X is the content of the message taken from a first-order language $\mathcal{L}$. In the case of physical actions (i.e., actions that manipulate the physical environment) with the pseudo-performative $\mathtt{do}$, Ac is the action executed by A (a physical action has no recipient as it is assumed to be observable by any agent in the system). Both X and Ac may contain non-logical substitution variables used for generalisation purposes (as opposed to logical "content" variables used by agents to indicate quantification or to ask for a valid binding). We further use $\mathcal{M}_c \subset \mathcal{M}$ to denote the language of "concrete" messages that agents use in communication (and that do not contain variables other than content variables). Frames are then defined as follows:

Definition 1 (Interaction frame). *An* interaction frame *is a tuple* $F = (T, \Theta, C, h_T, h_\Theta)$, *where*

- $T = \langle p_1, p_2, \ldots, p_n \rangle$ *is a sequence of message and action patterns* $p_i \in \mathcal{M}$*, the* trajectory*,*
- $\Theta = \langle \vartheta_1, \ldots, \vartheta_m \rangle$ *is an ordered list of* variable substitutions*,*
- $C = \langle c_1, \ldots, c_m \rangle$ *is an ordered list of* condition sets*, such that* $c_j \in 2^{\mathcal{L}}$ *is the condition set relevant under substitution* ϑ_j*,*
- $h_T \in \mathbb{N}^{|T|}$ *is a* trajectory occurrence counter *list counting the occurrence of each prefix of the trajectory* T *in previous conversations, and*
- $h_\Theta \in \mathbb{N}^{|\Theta|}$ *is a* substitution occurrence counter *list counting the occurrence of each member of the substitution list* Θ *in previous conversations.*

While the trajectory $T(F)$ models the surface structure of message sequences that are admissible according to frame F, each element of $\Theta(F)$ resembles a past binding of the variables in $T(F)$, and the corresponding element of $C(F)$ lists the conditions required for or precipitated by the execution of F in this particular case. $h_T(F)$ finally indicates how often F has been executed completely or just in part, $h_\Theta(F)$ is used to avoid duplicates in $\Theta(F)$ and $C(F)$. What hence distinguishes interaction frames from the methods commonly used for the specification of ACL and protocol semantics is that they allow for an explicit representation of *experience* regarding their practical use.

The semantics of frames has been defined accordingly as a probability distribution over the possible continuations of an interaction that has started with $w \in \mathcal{M}_c$ and is computed by summing up over a set $\mathcal{F}$ of known frames:

$$P(w'|w) = \sum_{\substack{F \in \mathcal{F} \\ ww' = T(F)\vartheta}} P(\vartheta|F, w) P(F|w) \tag{1}$$

This equation views $\mathcal{F}$ as a compact yet concise representation of the interactions that have taken place so far and projects past regularities into the future. This global view, however, will hardly be computationally feasible in realistic domains, and it also contradicts the way conversation patterns are used in practice. One would rather expect different protocols for different purposes, and not all of them need to be reasoned over at the same time while engaging in a particular kind of interaction.

In the following section, we will instead introduce a framework for conducting decision-theoretic reasoning about frame selection, as well as action selection within a single frame. For this hierarchical approach to be reasonable as well as successful, it is required that the different frames concisely capture the different classes of conversations that can take place. This requirement has to hold as well for frames used by external observers to model, analyse or describe the interactions in a MAS. Particular emphasis will hence have to be put on the acquisition and adaptation of communication patterns from the actual interactions in a MAS, such that the resulting set of patterns corresponds to the different classes of interactions as perceived by the agent or external observer. Methods for the adaptation of interaction frames will be explored in section 4.

3 Reasoning with Conversation Patterns

The distinguishing feature of interaction frames as compared to (the methods commonly used for the specification of) interaction protocols is their ability to capture instance information, i.e. information about how particular conversation patterns have been used in the past according to the agent's experience. This additional information provides agents with a facility to reason about the semantics of communication in an adaptive fashion. In accordance with the empirical semantics view that considers the meaning of communication as a function of its consequences as experienced through the eyes of a subjective observer, agents can adapt existing frame conceptions with new observations of encounters and project past regularities into the future. In *open systems*, in which agents may or may not obey a set of pre-defined conversation patterns, this can be expected to improve agents' communication abilities significantly, particularly with respect to a strategic use of communication.

3.1 Frame Semantics

To gain deeper insight into adaptive agent communication in general and reasoning about communication patterns in particular, we will now take a procedural view on the probabilistic semantics of interaction frames defined by equation 1.

The semantics of a set $\mathcal{F} = \{F_1, \ldots, F_n\}$ of frames is as follows: Given an *encounter prefix* $w \in \mathcal{M}_c^*$, i.e. a sequence of messages already uttered in the current encounter (possibly the empty sequence) and a *knowledge base* $KB \in 2^{\mathcal{L}}$ of beliefs currently held by the reasoning agent[4], $\mathcal{F}$ defines a set of possible *continuations* $w' \in \mathcal{M}_c^*$, which can be computed as follows:

1. Filter out those frames whose trajectories do not prefix-match w.
2. For each remaining frame F, consider the possible postfixes of $T(F)$ for prefix w, each of them corresponding to a particular variable substitution (where w has already committed certain variables to concrete values).
3. Only consider those substitutions for which at least one of the context conditions in $C(F)$ is satisfied under KB.

For each of these possible continuations, we can then compute a *continuation probability* by virtue of similarity, frequency and relevance considerations. The resulting probability distribution over continuations w' is the *semantics* of w under $\mathcal{F}$.

Definition 2. *Let $\vartheta_f(F,w) = unifier(w, T(F)[1{:}|w|])$ be the most general unifier of w and the corresponding trajectory prefix $T(F)[1{:}|w|]$ of F. Then, the* set of possible substitutions *under frame F, beliefs KB, and conversation prefix w is defined as*

$$\Theta_{poss}(F, KB, w) = \left\{\vartheta \middle| \exists \vartheta'.\vartheta = \vartheta_f(F,w)\vartheta' \wedge \exists i.KB \models C[i]\vartheta\right\}.$$

[4] In equation 1, the agent's knowledge is implicit in the terms $P(\vartheta|F,w)$ and $P(F|w)$. More precisely, we could have written $P(\vartheta|F,w,KB)$ and $P(F|w,KB)$. For notational convenience, we further assume that knowledge bases use the same logical language as is used in the content language of messages.

In this definition, $\mathit{unifier}(v,w)$ denotes the most general unifier of two message pattern sequences v and w, $S\vartheta$ denotes application of substitution ϑ to a (set or list of) logical formula(e) or message(s) S (depending on the context). In other words, Θ_{poss} is the set of substitutions that are extensions of ϑ_f for which at least one condition in $C(F)$ is satisfied. Accordingly, the continuations w' of w that should be expected to occur with non-zero probability (according to F and under KB) are exactly those that result from the application of a substitution in Θ_{poss} to the postfix of $T(F)$.

In order to conduct (quantitative) decision-theoretic reasoning about frames, however, the exact quantities of the probabilities $P(\vartheta|F,w)$ have to be determined. In order to obtain well-defined probabilities even for substitutions ϑ that have never occurred before in actual interactions, we avail ourselves of a method commonly used in the area of *case-based reasoning* [11]. Starting from a *similarity measure* σ defined on message pattern sequences, we compute the similarity of any possible substitution to a frame by taking into account the frequencies of previous cases and the relevance of their corresponding condition sets in a single frame.

Definition 3. *Let $\sigma : \mathcal{M}^* \times \mathcal{M}^* \to [0,1]$ be a similarity measure on message pattern sequences. Let $c_i(F,\vartheta,KB)$ denote the relevance of the ith condition of F under ϑ and KB. Then, the similarity of substitution ϑ to frame F is defined as*

$$\sigma(\vartheta,F) = \sum_{i=1}^{|\Theta(F)|} \Big(\overbrace{\sigma(T(F)\vartheta, T(F)\Theta(F)[i])}^{\textit{similarity}} \cdot \overbrace{h_\Theta(F)[i]}^{\textit{frequency}} \cdot \overbrace{c_i(F,\vartheta,KB)}^{\textit{relevance}}\Big)$$

In other words, $\sigma(\vartheta,F)$ assesses to which extent ϑ is "applicable" to F. Definition 4 in section 4.1 will introduce a distance metric d_* on the set $\mathcal{M}_c^*$ of finite-length message sequences, such that $d_*(v,w)$ is the distance between message sequences v and w. Using this metric, we can define $\sigma(v,w) = 1 - d_*(v,w)$. A possible way to define c_i would be to let $c_i(F,\vartheta,KB) = 1$ if $KB \models C(F)[i]\vartheta$ and 0 otherwise, such that only those substitutions of F contribute to the similarity whose corresponding conditions are satisfied under ϑ and under current belief KB.

The conditional probability $P(\vartheta|F,w)$ in equation 1 can be computed by assigning a probability

$$P(\vartheta|F,w) \propto \sigma(\vartheta,F) \tag{2}$$

to all $\vartheta \in \Theta_{poss}(F,KB,w)$ and a probability of zero to any other substitution. $P(F|w)$ simply corresponds to the number $h_T(F)[|h_T(F)|]$ of successful completions of F normalised over all frames that prefix-match w.

3.2 Decision Making with Frames

In the introductory section, we have argued for the integration of agent communication with decision-theoretic reasoning, by which agents strive for long-term maximisation of expected utility. We hence assume that agents are equipped with a *utility function* $u : \mathcal{M}_c^* \times 2^{\mathcal{L}} \to \mathbb{R}$, such that $u(w,KB)$ denotes the utility associated with executing a message (and action) sequence w in belief state KB. As we have pointed out, substantial positive or negative utility can only be assigned to physical actions in the environment (though messages may be given a small negative utility to express the communication cost incurred by them).

In principle, such a utility function could be combined directly with the continuation probabilities of equation 1 to derive utility-maximising decisions in communication. However, it will hardly be feasible to compute the continuation probabilities directly, and this approach would also contradict the role usually played by conversation patterns. As we have said, we will instead use a hierarchical approach based on selecting the appropriate frame for a given situation and then optimising behaviour within this frame. The former activity is referred to as *framing* and will be described in the following section. The latter is standard expected utility maximisation using frames and can be described by the following abstract decision-making procedure:

1. If no encounter is running, consider starting one. If a message m is received, update the encounter prefix: $w \leftarrow wm$.
2. If no frame F has been selected, go to 10.
3. *Validity check:* If $|T(F)| = w$, go to 9; if $\mathit{unifier}(T(F)[1:|w|], w) = \bot$, go to 10.
4. *Adequacy check:* If $\Theta_{poss}(F, w, KB) = \emptyset$, go to 10
5. Compute the expected utility for each *own substitution* ϑ_s:

$$E[u(\vartheta_s, F, w, KB)] = \sum_{\vartheta_p} \Big(u\big(\mathit{postfix}(T(F), w)\vartheta_s\vartheta_p, KB\big) \cdot P\big(\vartheta_p|\vartheta_s, F, w\big) \Big)$$

6. Determine the optimal substitution $\vartheta^* = \arg\max_{\vartheta_s} E[u(\vartheta_s, F, w, KB)]$.
7. *Desirability check:* If $u(\mathit{postfix}(T(F), w)\vartheta^*\vartheta_p, KB) < b$, go to 10.
8. Perform $m^* = T(F)[|w|+1]\vartheta^*$; update the encounter prefix: $w \leftarrow wm^*$
9. If no message arrives until deadline, terminate the encounter; go to 1.
10. *Framing:* Select F, go to 3.

The actual (framing) reasoning cycle is bracketed by steps 1 and 9 which cater for initiating encounters and ending them if no more messages are received (i.e., if the other agent does not reply when expected to, and to make sure we heed additional messages sent by the other party after we considered the encounter completed). We assume encounter initiation on the side of the agent to be spawned by some sub-social reasoning layer, e.g. a BDI [19] engine, which determines whether and about what to converse with whom, depending on the possibility of furthering some private goal through interaction.

Steps 3, 4 and 7 are used to evaluate the usefulness of the currently active frame F. The former two cases are straightforward: If the frame has been completed, if it does not match the encounter prefix w, or if $\Theta_{poss}(F, w, KB)$ is empty, F cannot be used any longer. For the latter case, we have to assess the expected utility $E[u(\vartheta_s, F, w, KB)]$ of any "own" substitution ϑ_s. To this end, we have to conduct an adversarial search over substitutions jointly determined by the agent and her peer, as each of the two agents commits certain variables to concrete values in their turn-taking moves. Using definition 3 and equation 2, the probability for an opponent's substitution ϑ_p in the remaining steps of $T(F)$ can be computed as

$$P(\vartheta_p|\vartheta_s, F, w) = \frac{P(\vartheta_p \wedge \vartheta_s|F, w)}{P(\vartheta_s|F, w)} = \frac{\sigma(\vartheta_f(F, w)\vartheta_s\vartheta_p, F)}{\sum_\vartheta \sigma(\vartheta_f(F, w)\vartheta_s\vartheta, F)},$$

where $\vartheta_p \wedge \vartheta_s$ denotes the event of the peer choosing ϑ_p and the reasoning agent choosing ϑ_p after having committed to the fixed substitution $\vartheta_f(F,w)$, so that the final "joint" substitution will be $\vartheta_f(F,w)\vartheta_s\vartheta_p$.

With this, u can be used to compute the utility of the postfix of $T(F)$ for prefix w (corresponding to application of $\vartheta_f(F,w)$), with ϑ_p and ϑ_s applied to obtain a ground message (and action) sequence still to be executed along $T(F)$. If the utility of the postfix under the optimal substitution ϑ^* is below some threshold b, the frame is discarded. Otherwise, the next step m^* along the trajectory of F is performed.

So far, we have said nothing about the process of updating the frame repository $\mathcal{F}$ upon encounter termination (whether after successful completion or failure of selecting an appropriate frame). This will be done in detail in section 4. What now remains to be specified is a search strategy to decide between different candidate frames in step 10. Effectively, it is this search strategy that determines the degree of complexity reduction achieved by restricting the search space to a single active frame while looking for the optimal next message or action.

3.3 Framing

Given that the frames in $\mathcal{F}$ concisely capture the different classes of conversations that can take place in a MAS, *hierarchical reinforcement learning* (HRL) techniques [2] can be used learn an optimal strategy for frame selection. In HRL, actions available in a "generic" Markov Decision Process (MDP) are combined into macro-operators that can be applied over an extended number of decision steps, the general idea being that compound time-extended policies, which (hopefully) optimally solve sub-problems of the original MDP, help to reduce the overall size of the state space. Using such macro-actions, an agent can use S(emi-)MDP (i.e., state history dependent) variants of learning methods such as Q-learning [28] to optimise its long-term "meta"-strategy over these macro-policies.

An intuitive HRL approach that lends itself to an application to interaction frames particularly well is the options framework [17]. In this framework, an *option* is a triple $o = (I, \pi, \beta)$ consisting of an input set $I \subseteq \mathcal{S}$ of MDP states, a (stationary, stochastic) policy $\pi : \mathcal{S} \times \mathcal{A} \to [0,1]$ over primitive actions $\mathcal{A}$ and states $\mathcal{S}$, and a stochastic termination condition $\beta : \mathcal{S} \to [0,1]$. Option o is admissible in a state s iff $s \in I$. If invoked, o will behave according to π until it terminates stochastically according to β. This definition can be used to re-interpret interaction frames as options, where π is the (deterministic) strategy defined by determining m^*, and I and β are defined by the validity, adequacy and desirability checks performed during the reasoning process of the previous section.

Let $s : \mathcal{M}_c^* \times 2^{\mathcal{L}} \to \mathcal{S}$ be some state abstraction function[5] that returns a state for each pair (w, KB) of perceived encounter prefix w and belief KB. If we regard each frame $F \in \mathcal{F}$ as an option in the above sense, we can apply the SMDP Q-learning update rule

$$Q(s,F) \leftarrow (1-\alpha)Q(s,F) + \alpha\left[\hat{R}(s,F) + \gamma^{\tau} \max_{F' \in \mathcal{F}} Q(s',F')\right],$$

[5] It is unrealistic to assume that $\mathcal{M}_c^* \times 2^{\mathcal{L}}$ itself could be used as state space due to its unmanageable size.

where $s = s(w, KB)$ and $s' = s(ww', KB')$ are the states resulting from the encounter sequences w and ww' and the corresponding knowledge base contents KB and KB' as perceived between two re-framing decisions, α is an appropriately decreasing learning rate and τ is the number of steps during which F was the active frame (i.e., $\tau = |w'|$). Further, $\hat{R}(s, F)$ is the discounted reward accumulated in steps $t+1, \ldots t+(\tau-1)$.

Using the long-term utility estimates represented by Q, we can determine the optimal frame to select as

$$F^*(w, KB) = \arg\max_{F \in \mathcal{F}} Q(s(w, KB), F),$$

while applying a "greedy in the limit" infinite exploration strategy to avoid running into local minima. It should be noted that this way of learning a frame selection strategy allows for optimising framing decisions *within* encounters as well as *between* subsequent encounters, at least if there is some utility-relevant connection between them.

4 Adaptation of Conversation Patterns

As we have already said, the need for its acquisition and adaptation from actual interactions is an inherent property of empirical semantics. Using a set of interaction frames for representation, we have further argued that these frames need to model different classes of interactions within a MAS. In particular, this feature is critical with respect to the reasoning framework described in the previous section.

In this section, we will present a method for the adaptation and acquisition of models of empirical semantics using the formalisation of interaction frames given in section 2. For this, we will introduce a metric on the space $\mathcal{M}_c^*$ of finite-length message sequences and then extend it to a metric between frames. This allows us to interpret a frame repository (i.e., a set of known frames) as a (possibly fuzzy) clustering in the "conversation space", and hence to measure the quality of a frame acquisition and adaptation method in terms of the quality of the clustering it produces (referred to as "cluster validity" in [10]). According to this interpretation, adaptation from a new conversation either introduces a new cluster (viz frame) or it adds to an existing one with or without modifying the trajectory of the respective frame. The different alternatives can be judged heuristically in terms of the corresponding cluster validities, which we will use to devise an algorithm for the adaptation of frame repositories. To perform the necessary frame modifications in any of the above cases, we will also present a generic algorithm for merging two frames into one.

Due to lack of space, proofs and examples have largely been omitted from this description. The interested reader is referred to [6] for a more detailed description.

4.1 A Distance Metric on Message Sequences

As a basis of our interpretation of interaction frames as clusters, we will start by introducing a distance metric on the set of possible messages and then extend it to finite-length message sequences. Since messages as defined above are essentially first-order objects, we could simply use a general purpose first-order distance like the one proposed in [24]. In [6], we have instead introduced a family of mappings on messages that are parametrised on two functions d_s and D_s and allow us to add a "semantic" flavour in the form of domain-specific knowledge. The most basic (and domain-independent) instance of this family is in fact a metric on messages (in particular, it satisfies the triangle inequality), and can easily be extended to message sequences.

Definition 4 (Distance between message sequences). *Let $d : \mathcal{M}_c \times \mathcal{M}_c \to [0,1]$ be a mapping on messages with*

$$d(m,n) = \begin{cases} \frac{1}{|m|+1} \sum_{i=1}^{|m|} d(m_i, n_i) & \textit{if } \underline{m} = \underline{n} \\ 1 & \textit{otherwise.} \end{cases}$$

Further, let $|v|$ and v_i denote the length and ith element of sequence v. We define

$$d_*(v,w) = \begin{cases} \frac{1}{|v|} \sum_{i=1}^{|v|} d(v_i, w_i) & \textit{if } |v| = |w| \\ 1 & \textit{otherwise.} \end{cases}$$

As we haven shown in [6], d_* is indeed a metric on the set $\mathcal{M}_c^*$ of finite-length message sequences.

4.2 A Metric on Frames

Having defined a metric d_* on the set of finite-length message sequences, we will now extend this metric (a metric on *points*, so to speak) to a metric on frames by interpreting these as sets of the message sequences they represent (i.e., point *sets*).

[18] proposes a general formalism to define a distance metric between finite sets of points in a metric space. The distance between two sets A and B is computed as the weight of the maximal flow minimal weight flow through a special distance network $N[X,d,M,W,A,B]$ between the elements of the two sets.

Definition 5 (Netflow distance). *Let X be a set with metric d and weighting function W, M a constant. Then for all $A,B \in 2^X$, the* netflow distance *between A and B in X, denoted $d^N_{X,d,M,W}(A,B)$, is defined as the weight of the maximal flow minimal weight flow from s to t in $N[X,d,M,W,A,B]$.*

As further shown in [18], $d^N_{X,d,M,W}(A,B)$ is a metric on 2^X and can be computed in polynomial time (in $size_W(A)$ and $size_W(B)$ and in the time needed to compute the distance between two points) if all weights are integers. Also, this metric is claimed to be much better suited for applications where there is likely a point with a high distance to any other point than, for example, the Hausdorff metric (which only regards the maximum distance of any point in one set to the closest point in the other set).

Additionally, one can assign weights to the elements of A and B in order to alleviate the difference in cardinalities between the two sets. Interpreting (integer) weights as element counts yields a metric on *multisets*, which is ideally suited to measure the distance between interaction frames in which multiple instances of a particular message sequence have been stored (corresponding to a substitution count larger than one). Mapping each frame to the set of messages it represents and weighting each element with the respective substitution count, we directly obtain a metric d_f on frames.

Definition 6 (Distance between frames). *Let*

$$m_f(F) = \{m \in \mathcal{M}_c^* | \exists \vartheta \in \Theta(F).\ m = T(F)\vartheta\}$$

be the set of message sequences stored in frame F. Let

$$W(m_f(F))(m) = h_\Theta(F)[i] \text{ iff } m = T(F)\Theta(F)[i]$$

be a weighting function for elements of $m_f(F)$. Then, the distance between two frames F and G*, denoted $d_f(F,G)$, is defined as the maximal flow minimal weight flow from s to t in the transport network $N[\mathcal{M}_c^*, d_*, 1, W, m_f(F), m_f(G)]$.*

As shown in [6], d_f is a metric on the set of frames, and $d_f(F,G)$ can be computed in polynomial time in $\sum_{i<|\Theta(F)|} h_\Theta(F)[i]$, $\sum_{i<|\Theta(G)|} h_\Theta(G)[i]$ and the time required to compute d_*.

4.3 Validity of Frame Modifications

Based on the metrics defined in the previous sections, we can interpret interaction frames as clusters of points in the space of message sequences, which in particular allows us to define the quality of a set of frames as a model for actual interactions in terms of the quality of the corresponding clustering.

[10] refers to this quality as *cluster validity* and defines the validity of a particular cluster as the ratio between its compactness, i.e. average distance between points within this cluster, and its isolation, i.e. minimum distance to any other cluster. Accordingly, we define the compactness and isolation of a frame using the metrics d_* and d_f on message sequences and frames, respectively.

Definition 7 (Frame compactness and isolation). *Let $\mathcal{F}$ be a set of frames, $F \in \mathcal{F}$ a single frame. The* compactness *of F is then defined as the (normalised) average distance between the individual messages stored in it, weighed by their respective occurrence counts:*

$$c(F) = \Big(\sum_{i=1}^{|\Theta(F)|}\sum_{j=i+1}^{|\Theta(F)|} h_i \cdot h_j\Big)^{-1} \cdot \sum_{i=1}^{|\Theta(F)|}\sum_{j=i+1}^{|\Theta(F)|} h_i \cdot h_j \cdot d_*\big(T(F)\vartheta_i, T(F)\vartheta_j\big)$$

where $\vartheta_i = \Theta(F)[i]$ and $h_i = h_\Theta[i]$ denote the ith substitution of F and the corresponding count. The isolation *of F in $\mathcal{F}$ is defined as the minimal distance to any other frame in $\mathcal{F}$:*

$$i(F,\mathcal{F}) = \min_{G \in \mathcal{F}\setminus\{F\}} d_f(F,G)$$

Since $c(F)$ uses d_* for distances within a single frame F only, there exists a more efficient way of computing it. If we write $w(v,m)$ to denote the *weight* of a variable v in a message pattern m (i.e., the sum of coefficients of $d(v,\cdot)$ in $d_*(m, m\vartheta)$ for some substitution ϑ), then we can precompute $w(v, T(F))$ for any variable v in the trajectory of F, and rewrite $c(F)$ to

$$c(F) \propto \sum_{i=1}^{|\Theta(F)|}\sum_{j=i+1}^{|\Theta(F)|} h_i \cdot h_j \cdot \sum_v w\big(v, T(F)\big) \cdot d_*\big(v\vartheta_i, v\vartheta_j\big)$$

According to definition 7, $c(F)$ is zero for frames with only one distinct substitution, so defining overall validity as the sum or product of individual validities $i(F,\mathcal{F})/c(F)$ is not a good idea. Instead, we define the validity of a frame repository $\mathcal{F}$ as the ratio between average isolation and average compactness for all the frames in $\mathcal{F}$, taking special care of situations where only frames with a single substitution exist.

Definition 8 (Frame validity). *Let $\mathcal{F}$ be a set of frames. The* validity *of $\mathcal{F}$ is defined as*

$$v(\mathcal{F}) = \begin{cases} \frac{\sum_{F\in\mathcal{F}} i(F,\mathcal{F})}{\sum_{F\in\mathcal{F}} c(F)} & \text{if } \exists F \in \mathcal{F}\,.\, |\Theta(F)| > 1 \\ \frac{1}{|\mathcal{F}|}\sum_{F\in\mathcal{F}} i(F,\mathcal{F}) & \text{otherwise} \end{cases}$$

In analogy to cluster analysis we conjecture that the higher the validity $v(\mathcal{F})$ of a frame repository $\mathcal{F}$ built from a particular set of concrete interactions, the better it models the different classes of conversation in a MAS. Facing different alternatives for the incorporation of an interaction that has just been perceived, each of them corresponding to a specific modification of $\mathcal{F}$, we can judge their quality simply by measuring $v(\mathcal{F})$ before and after this modification and hence devise an algorithm that tries to maintain a frame repository with the highest possible validity.

4.4 Frame Abstraction and Merging

Before we can apply the results of the previous section to an algorithm for the acquisition and adaptation of interaction frames from actual interactions, we will first have to make explicit the actual modifications that can be performed on interaction frames and sets thereof in order to adapt them to newly observed interactions. We do so by providing a general algorithm for merging two interaction frames into one. This algorithm can then be used to simply add a new message to an existing frame (by interpreting the message as a "singular" frame with ground trajectory and only the empty substitution) or to reorganise a whole repository. In order to distinguish these two activities, and according to the point in time they are performed relative to the actual interactions, we might refer to them as online and offline merging.

Starting with frame trajectories and following Occam's Razor, the trajectory of the frame obtained by merging F and G should be the least abstract message pattern sequence that can be unified with both trajectories $T(F)$ and $T(G)$ using standard first-order unification, i.e. the *least general generalisation* (lgg) [16] of the two, denoted $lgg(T(F),T(G))$. The following inductive definition of least general generalisation for message sequences can be turned into a simple algorithm for its computation.

Definition 9 (Least general generalisation). *The least general generalisation (lgg) of two terms is given by*

$$lgg(f(s_1,\ldots,s_k),g(t_1,\ldots,t_l)) = \begin{cases} f(lgg(s_1,t_1),\ldots,lgg(s_k,t_k)) & \text{if } f = g \text{ and } k = l \\ x & \text{otherwise,} \end{cases}$$

where x is a new variable (i.e., one that does not occur in any s_i or t_i) such that $lgg(s,t)$ is unique for any subterms s and t throughout the lgg (i.e., equal terms are *replaced with the same variable). The lgg of two messages with the same performative is given by* $lgg(p(a,b,x), p(c,d,y)) = p(lgg(a,c), lgg(b,d), lgg(x,y))$. *It is undefined if the performatives differ. The lgg of two message sequences of equal length is given by* $lgg((m_1,\ldots,m_k),(n_1,\ldots,n_k)) = (lgg(m_1,n_1),\ldots,lgg(m_k,n_k))$. *As before, it has to be ensured that* $lgg(s,t)$ *is unique throughout the lgg for any two subterms* s *and* t.

In an algorithm, uniqueness of the lgg is usually achieved by means of a table that holds the lggs computed so far for any pair of arguments.

Along with the lgg, definition 9 also yields two substitutions, namely the most general unifier (mgu) of the lgg with each of its arguments, and we use the abbreviation $\vartheta_m(m,n) = mgu(m, lgg(m,n))$. To obtain the substitutions and conditions of the merged frame, the ϑ_m have to be applied to the substitutions and conditions of the respective frame. For this, let F be one of the frames to merge, let t denote the trajectory of the resulting frame and c_j and ϑ_j the condition and substitution of the resulting frame that correspond to $C(F)[j]$ and $\Theta(F)[j]$. If the new frame is to hold all the conversations of F, then $t\vartheta_i = T(F)\Theta(F)[i]$ has to hold for $1 \le i \le |\Theta(F)|$. The definition of ϑ_m implies that $T(F) = t\vartheta_m(T(F),\cdot)$ and thus $t\vartheta_m(T(F),\cdot)\Theta(F)[i] = t\vartheta_i$.

If accordingly ϑ_i is computed as $\vartheta_i = \vartheta_m(T(F),\cdot)\Theta(F)[i]$, however, information might be lost about correlations between multiple conversations originating from the same frame. To retain this kind of information, substitutions should be concatenated rather than applied unless the right side of $\vartheta_m(T(F),\cdot)$ is a variable (which is quite common, as it results from the introduction of a new variable for a variable in the course of computing the lgg). The following definition formalises this concept of selective application of a substitution.

Definition 10. *Let* $\vartheta = [v_1/t_1,\ldots,v_n/t_n]$ *be a single variable substitution and* $\Theta = \langle s_1,\ldots,s_m\rangle$ *a list of substitutions. Then,* $\vartheta \rtimes \Theta$ *denotes the list of substitutions that results from* selectively prepending ϑ *to each element of* Θ *and is given by* $\vartheta \rtimes \Theta = \langle r_1,\ldots,r_m\rangle$ *where* $r_i = [v_1/r_{i1},\ldots,v_n/r_{in}] \cdot s_i$ *and*

$$r_{ij} = \begin{cases} t_j s_i & \text{if } t_j \text{ is a variable} \\ t_j & \text{otherwise} \end{cases}$$

As for the conditions of the merged frame, $c_i\vartheta_i = C(F)\Theta(F)[i]$ has to hold analogously. Replacing ϑ_i with the above result yields $c_i\vartheta_m\Theta(F)[i] = C(F)\Theta(F)[i]$ and thus $c_i\vartheta_m = C(F)$. Writing ϑ^{-1} for the "inverse" of a substitution ϑ (replacing terms by variables), c_i can hence be defined as $c_i = C(F)\vartheta_m^{-1}$. This finally leads us to the following definition of a merging operation on frames:

Definition 11 (frame merging). *Let* F *and* G *be two interaction frames with* $|T(F)| = |T(G)|$. *Then, the result of* merging F *and* G*, denoted by* $M(F,G)$*, is given by*

$$\begin{aligned} M(F,G) = \langle & lgg\big(T(F),T(G)\big), \\ & C(F)\vartheta_m\big(T(F),T(G)\big)^{-1} \cdot C(G)\vartheta_m\big(T(G),T(F)\big)^{-1}, \\ & \vartheta_m\big(T(F),T(G)\big) \rtimes \Theta(F) \cdot \vartheta_m\big(T(G),T(F)\big) \rtimes \Theta(G), \\ & hmax(F,G), \\ & h_\Theta(F) \cdot h_\Theta(G)\rangle, \end{aligned}$$

where $hmax(F,G) = \langle h_1, h_2, \dots \rangle$ *with*

$$h_i = \begin{cases} max\{h_T(F)[i], h_T(G)[i], \sum_k h_\Theta(M(F,G))[k]\} & \text{if } i = |T(F)| \\ max\{h_T(F)[i], h_T(G)[i], h_{i+1}\} & \text{if } i < |T(F)|. \end{cases}$$

The rather complex definition of the step counter values for the merged frame stems from the fact that it is impossible to determine the value $h_T(merge(F,G))$ would have taken if $merge(F,G)$ had been in the repository during all the conversations stored in F and G just from the information provided by F and G. On the other hand, it is also impossible to determine which additional conversations would have been stored in $merge(F,G)$ if this had been the case, so it seems fair to approximate h_T based on the following observations: Obviously, $\max(h_T(F), h_T(G))$ is a lower bound for $h_T(merge(F,G))$. In addition to that, the sum of the values of h_Θ is a lower bound for the value of $h_T[|T|]$, since it resembles the exact number of past conversations stored in the frame. Finally, for each i, $h_T[i]$ is a lower bound for $h_T[j]$ with $j < i$. Hence, as we cannot infer any upper bounds from the counter values alone, we simply choose the values of $h_T(merge(F,G))$ such that the bounds are tight. If only online merging is used, this approximation always yields accurate values for h_T.

4.5 An Algorithm for Learning Frames

Based on the formal notion of validity of a set of frames presented in section 4.3, which extends cluster validity to the space of multi-agent conversations, and on the frame merging procedure given in section 4.4, the following simple algorithm computes the best way to incorporate a newly observed message sequence m into a frame repository $\mathcal{F}$:

function $flea(\mathcal{F}, m)$ **returns** a frame repository
inputs: frame repository $\mathcal{F}$, message sequence m
  /* compute the singular frame F for m ***/**
  $F := (m, C_m, \{\}, \langle 1, \dots, 1\rangle, \langle 1\rangle)$
  /* compute the set $\mathbb{F}$ of alternatives for inclusion of m ***/**
  $\mathbb{F} := \{\mathcal{F} \cup \{F\}\} \cup \bigcup_{F' \in \mathcal{F}} \{\mathcal{F} \setminus \{F'\} \cup M(F', F)\}$
  /* return the most valid frame repository ***/**
  return $\arg\max_{\mathcal{F}' \in \mathbb{F}} v(\mathcal{F}')$

While the surface structure of a particular message sequence equals the message sequence itself, identification of a set C_m of logical conditions that held during a conversation (according to the observer's world model) and that were *relevant* or *crucial* is clearly a nontrivial task. If frames exist, however, the execution of which was hindered due to reasons of context (especially if pre-specified "protocol" frames are used), these can be used to identify conditions other than those (physically) required for the execution of the individual messages.

Since the above algorithm only considers a single frame at a time for inclusion into the repository, it is unable to detect structures in the space of interactions that develop over time. This corresponds to a more general problem of *order dependence* in incremental unsupervised learning and might in practice result in several frames actually modelling the same class of interactions. This problem can be handled, though, by supplementing the above online merging algorithm with one that periodically checks if two frames in the repository can be merged to increase its overall validity.

5 Conclusion

In this paper, we have presented a novel approach to *adaptive agent communication*. Agents in open environments that communicate according to high-level pre-specified conversational patterns can use the approach to augment these patterns with empirical observation of actual conversations, and conduct decision-theoretic reasoning about them in the framework of empirical semantics. Interaction frames have been used as the central data structure, allowing for the integration with our previous work on interaction frames [4, 5, 20]. The basic principles of the approach, however, could also be applied to other, possibly more complex, forms of representation.

Our current work focuses on an experimental exploration of the benefits and limitations of our approach in real-world "communication learning" tasks. An experimental evaluation in the context of proposal-based and argumentation-based negotiation can be found in [22]. Further applications include performance measurement of a MAS or of individual agents with respect to communication or the design of new interaction protocols. An open issue that will have to be dealt with in future work to allow for the acquisition of conversation patterns from scratch is the discovery of conditions that were relevant or crucial for a particular class of conversation. While inductive logic programming techniques may again be the appropriate means to attack this problem, the transition to relative least general generalisation (which might be required to handle background knowledge already available for a particular class of conversation) would make this one disproportionately harder to solve.

References

1. J. L. Austin. *How to do things with Words.* Clarendon Press, 1962.
2. A. G. Barto and S. Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13(4):41–77, 2003.
3. P. R. Cohen and H. J. Levesque. Communicative actions for artificial agents. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)*, pages 65–72, 1995.
4. F. Fischer and M. Rovatsos. An empirical semantics approach to reasoning about communication. *Engineering Applications of Artificial Intelligence, Special Section on Best Papers of CIA 2004*, 18(7):809–823, 2005.
5. F. Fischer, M. Rovatsos, and G. Weiss. Hierarchical reinforcement learning in communication-mediated multiagent coordination. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1334–1335. ACM Press, 2004.
6. F. Fischer, M. Rovatsos, and G. Weiss. Acquiring and adapting probabilistic models of agent conversation. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 106–113. ACM Press, 2005.

7. N. Fornara and M. Colombetti. Operational specification of a commitment-based agent communication language. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 536–542. ACM Press, 2002.
8. E. Goffman. *Frame Analysis: An Essay on the Organisation of Experience*. Harper and Row, New York, NY, 1974. Reprinted 1990 by Northeastern University Press.
9. M.-P. Huget, editor. *Communication in Multiagent Systems*, volume 2650 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, Germany, 2003.
10. A. K. Jain and R. C. Dubes. *Algorithms for clustering data*. Prentice Hall, Upper Saddle River, NJ, 1988.
11. J. L. Kolodner. *Case-Based Reasoning*. Morgan Kaufmann, San Francisco, CA, 1993.
12. M. T. Kone, A. Shimazu, and T. Nakajima. The state of the art in agent communication languages. *Knowledge and Information Systems*, 2:259–284, 2000.
13. N. Luhmann. *Social Systems*. Stanford University Press, Palo Alto, CA, 1995.
14. M. Nickles, M. Rovatsos, and G. Weiss. Empirical-rational semantics of agent communication. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2004.
15. S. Paurobally, J. Cunningham, and N. Jennings. A formal framework for agent interaction semantics. In this volume.
16. G. Plotkin. A note on inductive generalization. *Machine Intelligence*, 5:153–163, 1971.
17. D. Precup. *Temporal Abstraction in Reinforcement Learning*. PhD thesis, Department of Computer Science, University of Massachusetts, Amherst, 2000.
18. J. Ramon and M. Bruynooghe. A polynomial time computable metric between point sets. *Acta Informatica*, 37(10):765–780, 2001.
19. A. S. Rao and M. P. Georgeff. An abstract architecture for rational agents. In *Proceedings of Knowledge Representation and Reasoning (KR&R)*, pages 439–449, 1992.
20. M. Rovatsos, F. Fischer, and G. Weiss. Hierarchical reinforcement learning for communicating agents. In *Proceedings of the 2nd European Workshop on Multiagent Systems (EUMAS)*, pages 593–604, 2004.
21. M. Rovatsos, M. Nickles, and G. Weiss. Interaction is meaning: A new model for communication in open systems. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2003.
22. M. Rovatsos, I. Rahwan, F. Fischer, and G. Weiss. Adaptive strategies for practical argument-based negotiation. In *Proceedings of the 2nd International Workshop on Argumentation in Multi-Agent Systems (ArgMAS)*, 2005.
23. M. Rovatsos, G. Weiss, and M. Wolf. An Approach to the Analysis and Design of Multiagent Systems based on Interaction Frames. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. ACM Press, 2002.
24. M. Sebag. Distance induction in first order logic. In *Proceedings of the 7th International Workshop on Inductive Logic Programming*, 1997.
25. M. P. Singh. A semantics for speech acts. *Annals of Mathematics and Artificial Intelligence*, 8(1–2):47–71, 1993.
26. M. P. Singh. A social semantics for agent communication languages. In *Proceedings of the IJCAI Workshop on Agent Communication Languages*, 2000.
27. M. Verdicchio and M. Colombetti. A commitment-based communicative act library. In this volume.
28. C. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8:279–292, 1992.