Proceedings of the Eighteenth German Conference on Artificial Intelligence, September 1994.

Cooperating to Be Noncooperative: The Dialog System PRACMA

A. Jameson, B. Kipper, A. Ndiaye, R. Schäfer, J. Simons, T. Weis, and D. Zimmermann

Department of Computer Science, University of Saarbrücken
PO Box 151150, D-66041 Saarbrücken, Federal Republic of Germany
{jameson, kipper, ndiaye, ralph, jsimons, tweis, detlev}@cs.uni-sb.de

Abstract. The modeling of noncooperative dialogs, as opposed to dialogs in which the goals of the participants coincide, presents novel challenges to a pragmatically oriented dialog system. PRACMA models noncooperative sales dialogs. In the role of the potential buyer of a used car, the system tries to arrive at a realistic evaluation of the unknown car in spite of biased information presentation on the part of the seller. In the role of the seller, PRACMA tries to form a usable model of the buyer even while using this model to manipulate the buyer's impressions. To realize this behavior, heterogeneous modules and representation formalisms cooperate within a multi-agent architecture.

1 Introduction

1.1 Issues

Almost all computational modeling of the pragmatic aspects of human dialogs has examined the ideal case of cooperative dialog partners. Typically, the system has taken the role of an information-provider whose sole motivation is to help a human information-seeker achieve the latter's goals. Research on the dialog system PRACMA[1] broadens this perspective by modeling noncooperative information-providing dialogs, in which the participants, though constrained by dialog conventions, pursue conflicting goals. In the specific example domain, the two participants are a person $\mathcal{S}$ who is trying to **s**ell her used car and a potential **b**uyer $\mathcal{B}$.[2]

A long-term practical motivation of this research is the observation that the application of information-providing dialog systems will not always be restricted to purely cooperative situations. But the more immediate aim is to yield general insights into the pragmatics of dialog by investigating how several issues, which are in part familiar from research on cooperative dialogs, can be handled in the context of noncooperative motivation. As it is not possible within this paper to discuss all of the questions being

This research is being supported by the German Science Foundation (DFG) in its Special Collaborative Research Program on Artificial Intelligence and Knowledge-Based Systems (SFB 314), project N1, PRACMA. Substantial programming support has been provided by A. Fiedler, R. Knop, T. Wagner, A. Werner, and C. Wirtz.

[1] **PR**ocessing **A**rguments Among **C**ontroversially-**M**inded **A**ctors.

[2] We arbitrarily use feminine pronouns to refer to $\mathcal{S}$ and masculine pronouns for $\mathcal{B}$.

investigated, we concentrate here on the following central issues. For concreteness, we formulate them in terms of the example domain introduced above.

1. How can an information-seeker elicit and interpret information from a non-cooperative information-provider and cope with the uncertainty inherent in the dialog situation? (Section 2.) Whereas a cooperative information-provider often offers *more* useful information than the information-seeker explicitly requests, the $\mathcal{B}$ in PRACMA's situation requires considerable sophistication to arrive at a well-founded decision in the face of the largely evasive and potentially misleading formulations of $\mathcal{S}$.

2. How can a biased information-provider reconstruct, anticipate, and manipulate the inferences and evaluations of a largely unknown information-seeker? (Section 3.) Even the most sophisticated $\mathcal{S}$ is limited by her ignorance about the interests, general knowledge, and specific prior expectations of the specific $\mathcal{B}$ that she is dealing with—and $\mathcal{B}$ is not in general motivated to help her to solve this problem. So $\mathcal{S}$ must exploit her model of $\mathcal{B}$ with caution and continually try to improve it on the basis of largely unreliable evidence.

3. What are the benefits and costs of simulating two conflicting dialog roles? PRACMA has been designed to be *transmutable* (cf. [13]), i.e. to be able to simulate either $\mathcal{B}$ or $\mathcal{S}$ in the example domain. This allows us to investigate the extent to which components that appear to be primarily relevant to one role are also useful when the system assumes the other role. It also raises the issue of how to represent as much knowledge as possible in such a way that it can be used in both roles. One advantage of transmutability is that the system's behavior in each role can be explained and motivated in terms of an explicit model of how the dialog partner in the other role behaves.

4. What type of architecture best supports the necessary pragmatic processing? (Section 4.) A system that addresses the above issues requires an architecture that supports flexible interaction among modules which differ widely in their internal structure and representational techniques. In PRACMA, we evaluate the suitability of a new multi-agent architecture for these purposes.

Each of these issues will be discussed in one of the later sections of the paper, as indicated above (except for the third issue, which is touched on at various points). But first we present an example dialog and give an overview of PRACMA's agents, with special attention to the central task of dialog planning.

1.2 Example Dialog and Overview

Figure 1 shows an example dialog typical of those that PRACMA can conduct in the role of either $\mathcal{S}$ or $\mathcal{B}$ (cf. Sect. 5). Although the other role is always taken by a human user, we will describe the processing of both $\mathcal{S}$ and $\mathcal{B}$ as they are simulated by the system. Figure 2 shows the agents of PRACMA and the nature of the communication that takes place between them.

The agent NATURAL LANGUAGE ANALYZER translates German-language input strings into semantic representations in the language $\mathcal{NLL}$ ([7]; cf. 4.2), sending these representations to the PRAGMATIC DIALOG MEMORY, in which descriptions of all utterances of the two dialog partners are stored. The main control over dialog processing is exerted by the DIALOG PLANNER (1.3), which may attempt to execute any of various dialog plans, depending on the current context; the context is determined in part by the simple

$\mathcal{B}$: `Hello! How much gas does this car consume?`

$\mathcal{B}$ begins by asking a question about an aspect of the car that is likely to have a major impact on his evaluation.

$\mathcal{S}$: `Hello. It consumes 7 liters per 100 km.`

$\mathcal{S}$ has now answered $\mathcal{B}$'s question; but since she interpreted the question as a sign of interest in environmental-friendliness, she takes the opportunity to volunteer the following additional information, which ought to make a good impression:

$\mathcal{S}$: `It has a catalytic converter.`

To be sure, this fact has somewhat negative implications for the dimension "sportiness"; but these implications are not known to everyone, and besides $\mathcal{S}$ has no evidence that $\mathcal{B}$ is especially interested in sportiness.

$\mathcal{B}$: `That's too bad.`

Explicit reactions like this are usually associated with negative evaluation shifts. So $\mathcal{S}$ now suspects that $\mathcal{B}$ is in fact interested in sportiness and probably also fairly knowledgeable about aspects of cars relevant to that dimension.

$\mathcal{B}$: `What are the chances that the car will pass the next inspection?`

$\mathcal{B}$ switches to another aspect that is important to his evaluation.

$\mathcal{S}$: `It ought to pass the inspection.`

$\mathcal{S}$ actually judges the car's chances to be only about 55%. She figures that saying this explicitly would have an unfavorable impact on $\mathcal{B}$; but on the other hand, to give a definitely optimistic characterization of the car's chances would be too untruthful. As a compromise, $\mathcal{S}$ chooses a vague, positive formulation, with which she expects at least to prevent $\mathcal{B}$'s evaluation from becoming much less favorable.

$\mathcal{B}$: `Good.`

$\mathcal{B}$'s evaluation of the car with respect to this aspect has apparently improved, contrary to $\mathcal{S}$'s expectation. $\mathcal{S}$ explains this mainly by postulating that $\mathcal{B}$'s initial impression of the car's chances of passing the inspection was unusually pessimistic.

Fig. 1. *Initial portion of an example dialog with* PRACMA.

agent EGO, which stores information on which role the system is playing and what its motivational parameters are.

Most of the pragmatic-level inferences are performed by the COMMENT AND QUESTION HANDLER and the EVALUATION HANDLER. The former agent is invoked by the DIALOG PLANNER when the system needs to generate or interpret a comment or a question about the car (cf. Sect. 2 and Sect. 3). When PRACMA takes the role of $\mathcal{S}$, this agent needs specific information about the car; it gets this from the DOMAIN BELIEFS (4.2), where $\mathcal{S}$'s knowledge about the car is represented with the knowledge representation system MOTEL (see, e.g., [2]). When the system takes the role of $\mathcal{B}$ (Sect. 2) the EVALUATION HANDLER is responsible for forming evaluative impressions of the car; when the system takes the role of $\mathcal{S}$, the EVALUATION HANDLER engages in meta-level reasoning about this evaluation process of $\mathcal{B}$. The EVALUATION HANDLER calls upon the EVALUATION EXPRESSION HANDLER to help determine the evaluative implications of natural language formulations (e.g., "That's not bad"). The NATURAL LANGUAGE GENERATOR (currently under development) is responsible for verbalizing the $\mathcal{NLL}$ expressions produced by other agents.

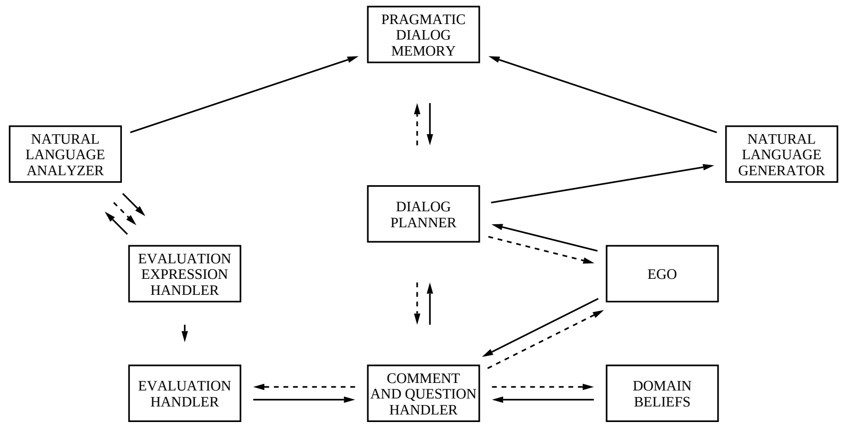

Fig. 2. *Overview of* PRACMA*'s agents. (Dashed arrows correspond to requests for information, solid arrows to the provision of information.)*

1.3 Dialog Planning

The Dialog Planner. As Fig. 2 shows, the DIALOG PLANNER plays a central role, invoking either directly or indirectly the help of other agents and using their replies in the construction and execution of a dialog strategy. The planner is realised in the planning system VIPER[3], a hierarchical (cf. [8]) incremental planner extended with dependency-directed backtracking and a nonlinear plan expansion strategy. Planning structures for open (sub)goals are built up with generic schemes for problem decomposition, called *plan operators*. For each subgoal stated in a plan operator, a separate subplanner is created which tries to achieve that subgoal. The order in which subgoals are pursued is a priori undetermined; it is controlled by heuristics or by explicit constraints (as illustrated in Fig. 3). This nonlinear plan strategy, augmented with declarative constraints on operator selection and subgoals, supports flexible planning strategies. Plan generation is intertwined with plan execution so that the requirements of a dynamic dialog environment can be met. The execution of (sub)plans obeys the linearization dictated by the left-to-right-bottom-up strategy of the generated global plan structure. Backtracking is adapted to this incremental planning strategy in that it takes into account the execution state of subplans.

Dialog strategies. The main dialog strategies implemented so far for $\mathcal{S}$ and $\mathcal{B}$ result in the sort of sequence illustrated in Fig. 1: $\mathcal{B}$ asks a question about some aspect of the car; $\mathcal{S}$ answers the question and then perhaps volunteers some comments on related aspects of the car; after each of $\mathcal{S}$'s utterances, $\mathcal{B}$ may express an evaluation verbally. With a view to expanding PRACMA's repertoire of dialog strategies, various knowledge elicitation techniques are now being applied with professional car salespersons. Initial results show, for example, that such an $\mathcal{S}$ not only presents facts about the car but also emphasizes the positive *consequences* of these facts with respect to $\mathcal{B}$'s evaluation standards—in some cases, even when it seems likely that $\mathcal{B}$ is already aware of these

[3] **V**erteilter **I**nkrementeller **P**lan**ER** ("distributed incremental planner"), see [14].

```
(define−viper−operator
:NAME ANSWER-QUESTIONS-AND-TRY-TO-VOLUNTEER-COMMENTS
:POSTCONDITIONS     (...)
:APPLICABILITY-CONDITIONS                                      ; For this operator to be usable, there has to be a
 ((PDM (:find−recent (:func is−from−partner−and−class question ; question from the dialog partner in the
          (whole topic−of−question)) (?ask ?topic) 1)))        ; Pragmatic Dialog Memory.
:SUBGOALS
 (...
  (CQH (:answer−question ?ask ?answer) :after :linear)         ; This is a request to the Comment and Question Hand-
                                                               ; ler to answer the question from the dialog partner.
  (NLG−PDM (verbalise ?answer) :after :linear)                 ; Here the answer is sent to the Natural Language
                                                               ; Generator and to the Pragmatic Dialog Memory.
  (ABS (volunteer−comments ?topic ?remarks ?buyer ?seller)     ; This abstract goal triggers operators that will
     :after :linear)                                           ; try to volunteer some extra comments, if the
                                                               ; Comment and Question Handler can supply some.
  (ABS (advertise ?handling−of−topic−next ?buyer ?seller)      ; This abstract goal calls plan operators which try to
     :after :linear)                                           ; answer another question from the dialog partner.
  ... )
:ABSTRACT-TRIGGERS                                             ; This slot contains the trigger which has to match for
 ((advertise ?handling−of−topic ?buyer ?seller))               ; this operator to be taken into consideration.
:SELECTION-HEURISTICS                                          ; This slot states that this operator should be tried
 ((:before STOP−ADVERTISING)))                                 ; before the operator named STOP-ADVERTISING.
```

Fig. 3. *An example of a dialog plan operator.*

consequences. Also, $\mathcal{S}$ often engages in diagnostic question-asking, to determine $\mathcal{B}$'s interests and knowledge and (especially toward the end of the dialog) to find out what reservations $\mathcal{B}$ has about the specific car in question.

To illustrate the main principles realized in PRACMA to date, the system's processing in the roles of $\mathcal{B}$ and $\mathcal{S}$, respectively, will now be discussed in the context of the central dialog strategy mentioned above.

2 PRACMA as Buyer: Evaluation Under Uncertainty

2.1 Representation of Beliefs and Evaluations

$\mathcal{B}$ starts the dialog with little definite information about the car. He has mainly uncertain *impressions* based on the information in the advertisement and on general knowledge about cars. Furthermore, since $\mathcal{S}$ is not especially motivated to make $\mathcal{B}$'s impressions as definite as possible (but rather to invoke favorable evaluations), $\mathcal{S}$ often makes vague comments that serve only to replace one indefinite impression with another slightly different one.

Accordingly, when PRACMA takes the role of $\mathcal{B}$ it represents $\mathcal{B}$'s beliefs and evaluations probabilistically in the form of a Bayesian belief network (see, e.g., [10]), which will be referred to as $\mathcal{B}$'s *evaluation network* (cf. Fig. 4).

Beliefs about the car. Each *factual node* in this network includes a probability distribution representing some *aspect* of the car. For example, the node INSPECTION CHANCES in Fig. 4 might store the fact that $\mathcal{B}$ is a priori quite uncertain, and rather pessimistic, about the car's chances of passing the next inspection.[4]

[4] For aspects (like this one) that involve an infinite number of possible states, a finite set of *possibilities* is distinguished (cf. [4] and [9]).

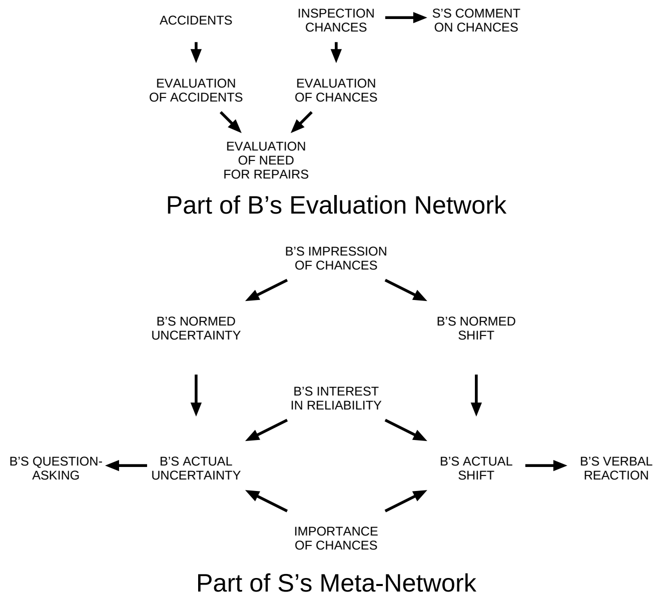

Fig. 4. *Parts of the Bayesian networks used by $\mathcal{B}$ and $\mathcal{S}$.*

Evaluations. The model of how $\mathcal{B}$ evaluates the car is based on the metaphor of an *evaluation form* (cf. [3]). For each aspect, if $\mathcal{B}$ knew exactly which possibility was realized, $\mathcal{B}$ would assign to the car some number of points for that aspect (cf. the first two columns in Fig. 5).[5] He could then add up the points assigned so as to arrive at an overall evaluation, or an evaluation of some more general aspect such as *need for repairs*. Since $\mathcal{B}$ in fact usually has only an impression of most aspects, he likewise has only impressions of the evaluation to be assigned to most aspects. These latter impressions are represented by *evaluation nodes* in his network. An evaluation node corresponding to a specific aspect (e.g., EVALUATION OF CHANCES) is a *child* of the corresponding factual node, i.e., the evaluation is treated as dependent on the factual belief. A node representing a more global evaluation (e.g., EVALUATION OF NEED FOR REPAIRS) is a child of several more specific evaluation nodes.

2.2 Role of the Evaluation Network in the Dialog

Selection of questions. $\mathcal{B}$'s principal motivation is to reduce the uncertainty in his evaluation nodes to the point where he can confidently make a decision as to whether to buy the car. He can only do this by eliciting information from $\mathcal{S}$, mainly by asking questions. To determine which questions are likely to yield the most valuable information, $\mathcal{B}$ compares the evaluation nodes with respect to the amount of *uncertainty*

[5] Many aspects of a car are relevant to two or more *evaluation dimensions*, e.g., for the existence of a catalytic converter, "Environmental Friendliness" and "Sportiness" (cf. Fig. 1 and [4]). To simplify exposition, we consider here only an example where a single evaluation dimension, "Reliability", is involved.

Chances of passing inspection	Evaluation according to $\mathcal{B}$'s actual evaluation form	$\mathcal{S}$'s normed evaluation form
95%	+1.4	+2.5
85%	+0.9	+2.0
75%	+0.6	+1.5
65%	+0.4	+1.0
55%	+0.2	+0.5
45%	−0.2	−0.5
35%	−0.4	−1.0
25%	−0.6	−1.5
15%	−0.9	−2.0
05%	−1.4	−2.5

Fig. 5. *Actual and normed evaluation forms used by $\mathcal{B}$ and $\mathcal{S}$, respectively.*

they exhibit (operationalized as the standard deviation of the probability distribution). For example, the fact that EVALUATION OF CHANCES has a broad distribution is responsible for $\mathcal{B}$'s decision in Fig. 1 to ask about the inspection chances.

Interpretation of comments. When a comment by $\mathcal{S}$ specifies unambiguously which of the possibilities for a given aspect is realized (e.g., "The chances are 55%"), it is straightforward for $\mathcal{B}$ to update the belief for the corresponding factual node, assigning a probability of 1.0 to the realized possibility. This change propagates from parent to child in the network, affecting some of the evaluation nodes. But for interpreting vague comments, $\mathcal{B}$ needs a more general method: For each comment by $\mathcal{S}$ concerning a factual node, a *dummy node* (cf. [10], p. 170) is linked to that factual node in $\mathcal{B}$'s network—as the factual node's child, because the comment can be viewed as having been probabilistically "caused" by the truth about the corresponding aspect of the car. The conditional probabilities linking the comment node to the factual node are determined by the meaning of the comment. The meaning can be represented by a *membership function* in the sense of fuzzy set theory. For German-language adverbs and modal verbs expressing judgments of the chances of a given event, we obtained experimental data on the membership functions of German speakers (see [6]). For example, for the modal verb *dürfte* (*ought to*, as in "The car ought to pass the inspection"), the average membership curve is monotonically increasing and sigmoid, with a steep rise at 50%. These membership values are not probabilities, but within a given pragmatic context they can serve as the basis for probability assessments (cf. [6]). One reasonable assumption is that, all other things being equal, the higher the membership value of a possibility for a given comment, the greater the likelihood that $\mathcal{S}$ will make that comment. $\mathcal{B}$ applies this assumption when interpreting $\mathcal{S}$'s comment in our example dialog.[6] The resulting child-to-parent propagation in $\mathcal{B}$'s network causes the more positive possibilities in INSPECTION CHANCES to be judged more probable. In turn, by parent-to-child propagation,

[6] A more sophisticated approach to comment interpretation, which takes into account the alternative comments that $\mathcal{S}$ might have made but did not, was introduced in the system IMP (see [3]) and is presently being adapted for use within PRACMA. This approach also involves an assessment by $\mathcal{B}$ of aspects of $\mathcal{S}$'s dialog motivation, such as bias.

the evaluative impression in EVALUATION OF CHANCES is shifted upward.

Generation of reactions to comments. This sort of shift in an evaluation impression is often reflected in an overt verbal reaction, as shown by an unpublished empirical study conducted in the PRACMA context (cf. [4]). When simulating $\mathcal{B}$, PRACMA selects such verbal reactions in a fashion consistent with the results of that study.

3 PRACMA as Seller: Forming and Exploiting Meta-Level Impressions

$\mathcal{S}$'s main task is to reconstruct, anticipate and manipulate the processing by $\mathcal{B}$ described in the previous section. When deciding what to say, her goal is not to reduce the uncertainty in $\mathcal{B}$'s evaluation nodes (as a cooperative speaker would do), but rather to shift the evaluation impressions upward. To do this effectively, $\mathcal{S}$ must try to derive a reasonably accurate model of $\mathcal{B}$'s network, making use of any clues in $\mathcal{B}$'s behavior that reflect properties of that network. It might at first be thought that $\mathcal{S}$ could use a model of $\mathcal{B}$'s network that was isomorphic to it. But in fact, to deal adequately with all of $\mathcal{S}$'s uncertainty about various aspects of $\mathcal{B}$'s network, a more complex type of model is required, which we call a *meta-network* (cf. the lower half of Fig. 4).

The first complication arises when $\mathcal{S}$ tries to take into account the various initial impressions that $\mathcal{B}$ might have in each of his factual nodes. In INSPECTION CHANCES, for example, $\mathcal{B}$ might have any of a potentially infinite number of impressions; $\mathcal{S}$ can in general only have an impression of $\mathcal{B}$'s impression. In the meta-network, a node like B'S IMPRESSION OF CHANCES distinguishes 16 classes of possible impressions, which differ in terms of both their expected value and their uncertainty. In our example, the meta-level impression in B'S IMPRESSION OF CHANCES represents the belief that it is most probable that $\mathcal{B}$ will expect the car's inspection chances to be around 75% and that $\mathcal{B}$ will have only a moderate amount of uncertainty about this. The type of initial impression that is in fact initially present in $\mathcal{B}$'s network in our example—centered around 35% with a high degree of uncertainty—is considered by $\mathcal{S}$ to be a priori unlikely.

$\mathcal{S}$'s second complication arises when she tries to predict the content of $\mathcal{B}$'s evaluation nodes: $\mathcal{S}$ of course doesn't know exactly what the content of $\mathcal{B}$'s evaluation form is, mainly because $\mathcal{S}$ doesn't know how much interest $\mathcal{B}$ has in the evaluation dimension "Reliability". A solution is to (a) use a *normed evaluation form* that is presumably typical of a $\mathcal{B}$ with a *moderate* level of interest in this dimension (cf. the right-hand column in Fig. 5), and (b) take into account separately (as described below) the possibility that $\mathcal{B}$'s interest might be quite different, or that the numbers in the form might be systematically inaccurate even for $\mathcal{B}$'s who do have a moderate level of interest.

For each of the factual impressions that $\mathcal{B}$ might have in INSPECTION CHANCES, $\mathcal{S}$ can use the normed evaluation form to determine the corresponding evaluation impression, and in particular to assess the uncertainty in $\mathcal{B}$'s evaluation. In this way, the normed evaluation form constitutes the basis for the conditional probabilities that link B'S IMPRESSION OF CHANCES to its child node B'S NORMED UNCERTAINTY, which represents the uncertainty in $\mathcal{B}$'s node EVALUATION OF CHANCES under the assumption that $\mathcal{B}$ has a moderate level of interest in "Reliability".

In a similar way, whenever $\mathcal{S}$ considers making a particular comment, she can predict what sort of evaluation shift it might evoke in $\mathcal{B}$: For each of the possible impressions that $\mathcal{B}$ might have, $\mathcal{S}$ simulates the process of comment interpretation that was described above for $\mathcal{B}$, using the normed evaluation form, and calculates the extent to which $\mathcal{B}$'s relevant evaluation node would shift upward or downward. The results are summarized in $\mathcal{S}$'s node B'S NORMED SHIFT.

The remaining nodes in $\mathcal{S}$'s meta-network will be explained in terms of how $\mathcal{S}$ uses the meta-network to interpret $\mathcal{B}$'s behavior and to select comments.

Interpretation of $\mathcal{B}$'s question-asking. B'S ACTUAL UNCERTAINTY can be predicted as a multiplicative function of B'S NORMED UNCERTAINTY, B'S INTEREST IN RELIABILITY and IMPORTANCE OF CHANCES. In turn, B'S ACTUAL UNCERTAINTY influences the assessment of B'S QUESTION-ASKING, a two-valued variable representing whether $\mathcal{B}$ asks or does not ask a question about the car's chances of passing the inspection. The utility of defining these nodes can be seen in the case where $\mathcal{B}$ does ask such a question, as in Fig. 1: By child-to-parent propagation, $\mathcal{S}$ gets the impression that $\mathcal{B}$'s has more evaluative uncertainty about the inspection chances than $\mathcal{S}$ originally thought. Further child-to-parent propagation represents an attempt by $\mathcal{S}$ to "explain" this unexpected shift in terms of one or more underlying variables: Perhaps $\mathcal{B}$ is especially interested in the dimension "Reliability" (B'S INTEREST IN RELIABILITY), perhaps the aspect "chances of passing the inspection" is especially important (IMPORTANCE OF CHANCES), or perhaps $\mathcal{B}$'s a priori impression of this car's chances is different from what $\mathcal{S}$ thought—in particular, $\mathcal{B}$ may have a much less definite impression than $\mathcal{S}$ would have expected (B'S IMPRESSION OF CHANCES). These adjustments in $\mathcal{S}$'s meta-network will influence $\mathcal{S}$'s further inferences and decisions with respect to this particular $\mathcal{B}$ (and in the case of IMPORTANCE OF CHANCES, with respect to future potential buyers as well).

Judging the desirability of a comment. Prediction of the evaluative shift that a given comment would evoke in $\mathcal{B}$ is important when $\mathcal{S}$ is considering whether to make that comment, either as an answer to a direct question or as a piece of unsolicited information. Analogously to the case with B'S ACTUAL UNCERTAINTY, for example, B'S ACTUAL SHIFT can be predicted as a multiplicative function of B'S NORMED SHIFT, B'S INTEREST IN RELIABILITY, and IMPORTANCE OF CHANCES. Although producing positive evaluation shifts is $\mathcal{S}$'s primary goal, her judgment of the desirability of a given comment also takes into account general dialog conventions. In particular, there is some danger in using a comment whose truthfulness is questionable, e.g. "The car will definitely pass the inspection", for which the actually realized possibility of 55% chances would have only approximately .3 as a membership value. $\mathcal{S}$ therefore assigns to each comment a *penalty* that is a positive function of the extent to which the relevant membership value is lower than 1.0. The result of this policy is that $\mathcal{S}$ avoids outrightly lying, but she is willing to apply a vague expression in a marginally acceptable way (e.g., *ought to* for chances of 55%) if she can thereby evoke a more favorable impression.

Interpreting $\mathcal{B}$'s verbal reactions to comments. As noted in Sect. 2, $\mathcal{B}$'s verbal reaction to a comment by $\mathcal{S}$ can be viewed as a probabilistic function of the shift in $\mathcal{B}$'s relevant evaluation node. This probabilistic relationship (derived from the empirical study mentioned above) links B'S ACTUAL SHIFT to the node B'S VERBAL REACTION. So any observed reaction by $\mathcal{B}$ constitutes an observation whose effects are subject to child-to-parent propagation and ultimately affect the same nodes that can be affected by the asking of a question. In our present example, the explicitly positive reaction by $\mathcal{B}$ is unexpected: All $\mathcal{S}$ really expected to achieve with her vague comment was to prevent $\mathcal{B}$ from being disappointed. The fact that $\mathcal{B}$'s evaluation apparently does change for the better is hard to explain in terms of a high level of B'S INTEREST IN RELIABILITY or IMPORTANCE OF CHANCES; more likely is that $\mathcal{B}$'s initial impression was more pessimistic than $\mathcal{S}$ originally expected (B'S IMPRESSION OF CHANCES).

4 A Cooperative Architecture

4.1 Overall Multi-Agent Architecture

Processing such as that described in the previous sections requires flexible communication among the system modules, because it is largely unpredictable when particular types of information will be required or become available. To this end, the architecture CHANNELS[7] was developed, integrating techniques from distributed artificial intelligence and the concurrent object-oriented paradigm.

The modules within the system are realized as (semi-)autonomous specialized problem solvers, called *agents*, which interact cooperatively. In CHANNELS, communication and interaction among the agents are achieved through a *communication-act-based protocol* which governs the exchange of messages. Each message is characterized by attributes including: the sender, the recipient(s), the type of communication act, the mode of communication (*synchronous* or *asynchronous*), the actual content of the message, and optionally the history of the message and the agents to whom the answer to the message's query should be forwarded. The communication acts supported at present are *ask*, *reply*, and *inform*. An *ask* message requests the recipient(s) to send back information in a *reply* to the originator of the message, while an *inform* message sends unsolicited information.

With the three basic communication acts and the two modes of communication, we attain some of the flexibility of concurrent object-oriented languages, e.g., ABCL (see [15]). Approaches similar to that of CHANNELS can be found in CARAMEL (see [11]) and TALISMAN (see [12]). They differ with respect to the nature of the communication supported and the mechanism for controlling the interaction among the agents (cf. [9]).

4.2 Interlocking of Heterogeneous Architectures

CHANNELS supports agents with complex internal structures which may be quite different from that of the overall multi-agent architecture. In this way, one can exploit the advantages of alternative languages and architectures for specific purposes, or make use of modules that were developed previously in another context. This interlocking approach is illustrated by the following two examples.

A blackboard for natural language analysis. The agent NATURAL LANGUAGE ANALYZER itself comprises six heterogeneous and relatively independent modules. Direct communication among these modules according to the principles embodied in CHANNELS would be problematic in that a module in this set does not always know which

[7] **C**ooperating **H**eterogeneous **A**gents for a **N**atural-**L**anguage **S**ystem, see [9].

other modules can give it the information it needs or which modules can make use of the results it generates (cf. [5]). In accordance with the criteria proposed, e.g., in [1], a blackboard architecture was chosen for the internal structure of the agent NATURAL LANGUAGE ANALYZER. The following modules have been realized as knowledge sources which use the blackboard (cf. [5]): An HPSG-based chart parser; a module that constructs semantic representations in $\mathcal{NLL}$ (cf. 4.2) on the basis of the parser's output; a rudimentary speech act recognizer; an anaphora resolution module that handles personal pronouns; a component for the disambiguation of modal verbs; and an external communicator that writes input strings onto the blackboard and sends results to other PRACMA agents.

A modal logic representation for domain beliefs. A further example of heterogeneous representation concerns the system's beliefs about (the dialog partner's beliefs about) cars in general and the specific car being discussed. Although these beliefs are in part implicitly reflected in the Bayesian (meta-)networks described in Sect. 2 and Sect. 3, there is also a need for an explicit, declarative representation—for example, to enable PRACMA, when modeling $\mathcal{S}$, to determine the truth value of possible comments. These beliefs are therefore represented within the agent DOMAIN BELIEFS in MOTEL, a multi-agent logic-based knowledge representation system (see, e.g., [2]). MOTEL's belief contexts enable the system to reason using its own specific and general domain beliefs, the corresponding beliefs which it ascribes (tentatively) to the dialog partner, and generally shared conceptual knowledge. Messages to DOMAIN BELIEFS are formulated using the logical semantic representation language $\mathcal{NLL}$ (see, e.g., [7]) which serves as the output language of the NATURAL LANGUAGE ANALYZER and the input language of the NATURAL LANGUAGE GENERATOR. To process *ask* and *inform* messages coded using $\mathcal{NLL}$, the agent DOMAIN BELIEFS has an $\mathcal{NLL}$-query-answerer and a corresponding updating procedure.

5 Further Research

A PRACMA prototype now exists which can simulate a dialog similar to the one in Fig. 1, taking the role of either $\mathcal{S}$ or $\mathcal{B}$. Some of the more recently developed parts of the system have been tested and demonstrated in separate modules and are at the time of this writing being integrated into the main prototype.

Some of the further issues being investigated have been mentioned in passing above, e.g., the expansion of $\mathcal{S}$'s repertoire of dialog strategies and the problem of how $\mathcal{B}$ can assess and take into account $\mathcal{S}$'s dialog motivation. The work done so far on these extensions suggests that the basic mechanisms described in this paper have sufficient generality to accommodate them. For example, when modeling $\mathcal{B}$, the system can form impressions of the motivational parameters of $\mathcal{S}$ using the same probabilistic techniques currently used by $\mathcal{S}$ to assess $\mathcal{B}$'s interests and knowledgeability.

More generally, the ideas introduced here should be of some use, on an abstract or concrete level, for modeling or analyzing information-providing dialogs where at least one of the participants forms impressions of general dispositions and internal states of the other participant, and/or where the goals of the participants partly conflict.

References

1. Cawsey, A., Galliers, J. R., Reece, S., & Sparck Jones, K. (1992). A comparison of architectures for autonomous multi-agent communication. *Proceedings of the Tenth ECAI*, Vienna, 249–251.
2. Hustadt, U., & Nonnengart, A. (1993). Modalities in knowledge representation. *Proceedings of the Sixth Australian Joint Conference on Artificial Intelligence*, Sydney, 249–254.
3. Jameson, A. (1989). But what will the listener think? Belief ascription and image maintenance in dialog. In A. Kobsa & W. Wahlster (Eds.), *User models in dialog systems* (pp. 255–312). Berlin: Springer.
4. Jameson, A., & Schäfer, R. (1993). Probabilistische Einschätzung von Wissen und Interessen. Die Anwendung der intuitiven Psychometrik im Dialogsystem PRACMA. *Arbeitspapiere des Workshops Adaptivität und Benutzermodellierung in interaktiven Softwaresystemen*, 100–117. Report 30/93, Project BGP-MS, Dept. of Information Science, University of Konstanz, Germany.
5. Kipper, B. (1994). A blackboard architecture for natural language analysis. *Proceedings of the Seventh Florida AI Research Symposium*, Pensacola Beach, FL.
6. Kipper, B., & Jameson, A. (1994). Semantics and pragmatics of vague probability expressions. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, Atlanta.
7. Laubsch, J., & Nerbonne, J. (1991). *An overview of* $\mathcal{NLL}$. Technical report, Hewlett Packard Laboratories, Palo Alto, CA.
8. Moore, J. D., & Paris, C. L. (1989). Planning text for advisory dialogues. *Proceedings of the Twenty-Seventh Annual Meeting of the ACL*, Vancouver, 203–211.
9. Ndiaye, A., & Jameson, A. (1994). Supporting flexibility and transmutability: Multi-agent processing and role-switching in a pragmatically oriented dialog system. *Proceedings of the Sixth International Conference on Artificial Intelligence: Methodology, Systems, Applications*, Sofia, Bulgaria.
10. Pearl, J. (1991). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo, CA: Morgan Kaufmann. (Revised second printing).
11. Sabah, G., & Briffault, X. (1993). CARAMEL: A step towards reflection in natural language understanding systems. *Proceedings of the Fifth IEEE International Conference on Tools with AI*, Boston.
12. Stefanini, M.-H., Berrendonner, A., Lallich, G., & Oquendo, F. (1992). TALISMAN: Un système multi-agents gouverné par des lois linguistiques pour le traitement de la langue naturelle. *Proceedings of the Fourteenth COLING*, Nantes.
13. Wahlster, W., & Kobsa, A. (1989). User models in dialog systems. In A. Kobsa & W. Wahlster (Eds.), *User models in dialog systems* (pp. 4–34). Berlin: Springer.
14. Weis, T. (1994). *VIPER: Ein verteilter inkrementeller Dialogplaner für eine Multi-Agenten-Umgebung*. Master's thesis, Dept. of Computer Science, University of Saarbrücken.
15. A. Yonezawa (Ed.) (1990). *ABCL: An object-oriented concurrent system*. Cambridge, MA: MIT Press.