Dialogue Requirements for Argumentation Systems

Richard McConachy
Ingrid Zukerman
School of Computer Science and Software Engineering
Monash University
Clayton, Victoria 3168, AUSTRALIA
email: {ricky,ingrid}@csse.monash.edu.au

Abstract

We consider the requirements of an ideal automated argumentation system designed to model human interactions and compare them with the capabilities of our argumentation system, NAG. We ground our discussion with a description of NAG's architecture, particularly emphasizing NAG's knowledge representation scheme, which supports the consideration of domain and user model information as well as contextual information when reasoning about arguments.

1 Introduction

Argumentation is an interesting domain for dialogue systems as it highlights aspects of dialogue which may be obscured in systems with simpler and more prescriptive interactions. For example, a system which looks up bus routes is unlikely to require a complex rebuttal-handling capability. The additional requirements brought to light by argumentation systems will in turn help practitioners determine the features required by their particular systems, the level of support for each feature, and the techniques that are appropriate for producing this level of support. To illustrate these requirements, let us consider the fictitious argument in Figure 1 (U stands for User and S for System). According to this dialogue, an argumentation system should do the following.

Understand implicit goals (U1) – The user does not explicitly state the goal of the current argument, but the system determines that the user is interested in speculating about who will win the premiership, rather than discussing how well the Kangaroos have played during the pre-season matches.

Generate enthymematic arguments (S1) – These are arguments which omit easily inferable information [Zukerman et al., 1998]. As shown in our example, the system indirectly refutes the proposition that the Kangaroos will win by proposing Adelaide as a more likely candidate.

U1: "The Northern Kangaroos will probably win the Australian Rules Football League premiership this year, since they have done well in the last few years and were in good form in the pre-season competition."

S1: "In each of the last two years Adelaide has fared poorly in the pre-season competition but has won the premiership."

U2: "But Adelaide has a worse win/loss ratio when playing matches at the MCG [the site where the premiership final is played] than the Kangaroos."

S2: "True, but if a club can win enough matches to make the finals, then their win/loss ratio for finals matches is more important than their win/loss ratio over the whole season."

U3: "Fair enough."

S3: "... and, over the last two years, Adelaide has had a better win/loss ratio than the Kangaroos when playing finals matches at the MCG."

U4: "What if Shaun Rehn [Adelaide's ruckman] needs more knee surgery?"

S4: "Adelaide would be somewhat less likely to win, but I would still favour them."

U5: "OK, what about Wayne Carey [Kangaroos' captain]?"

S5: "He is a good player, but not good enough to carry the Kangaroos."

U6: "That's not what I mean."

S6: "Wayne Carey has a shoulder injury at the moment."

Figure 1: Sample dialogue with an ideal system

Understand enthymematic arguments (U2) – As for S1, the user's rebuttal is enthymematic in that it does not present a complete inference chain, and does not explicitly indicate whether he or she is attacking the system's last statement (S1) or reinforcing his or her own previous argument (U1).

Judiciously admit new information (S2) – The system concedes the user's point about the win/loss ratio (U2). If this information was previously unknown to the system, the system must decide whether it should admit this information, the belief it should associate with this new information, and how this information affects the argument.

Determine an argumentation strategy (S2) – The system uses a hypothetical argument to defeat the user's argument about Adelaide's win/loss ratio (U2). Context may make certain argumentation strategies more appealing than others. For instance, if the user has repeatedly mentioned a particular belief, the system may select an argumentation strategy that takes advantage of this (entrenched) belief.

Handle turn taking and interruptions (U3, S3) – This capability is asymmetrical, since the user should be allowed to interrupt the system, but the system would usually not interrupt the user. In this example, the system notes that the user has accepted its (partial) argument, but decides to present the remainder of its argument anyway.

Handle subdialogues (U4) – Conversational partners normally probe each other's arguments to decide which parts of an argument to accept. This may be done by means of different types of subdialogues, e.g., information-sharing [Chu-Carroll and Carberry, 1995] or clarification [Litman and Allen, 1987]. In our example, the user poses an exploratory *what if* query to speculate on the effect of a particular event on the argument goal.

Track focus (U5) – This utterance may be interpreted as an exploratory query to consider the effect of Wayne Carey on the argument goal or as a shift in topic, where the focus is now on Carey's health. Both interpretations should be considered, and a preferred interpretation selected [Litman and Allen, 1987, Raskutti and Zukerman, 1991]. In this example, the former interpretation is (erroneously) selected.

Recover from misunderstandings (U6, S6) – U5 was previously interpreted as a query about the effect of Wayne Carey on the argument goal. However, when the user indicates that this is not his or her intent, the system adopts an alternative interpretation – that the user wants to shift the discussion to Wayne Carey's health.

In this paper, we discuss two main aspects of our argumentation system NAG (*Nice Argument Generator*): its reasoning and argumentation capabilities, which support the generation and understanding of enthymematic arguments (Section 3.2); and its dialogue capabilities, which support the handling of exploratory queries as a first installment of a more comprehensive dialogue facility (Section 4). Underpinning these capabilities is NAG's knowledge representation scheme, which integrates belief representation (both NAG's and the user's) with attentional focus to enable the consideration of contextual information during the argumentation process (Section 3.1).

2 Related Research

Our work is similar in some ways to IACAS [Vreeswijk, 1994], an interactive system for generating arguments. However, IACAS does not tailor its arguments to a particular user. The system described in [Jönsson, 1995] contains a dialogue manager which directs the interaction between the user and the system by taking advantage of observations of the user's behaviour during information-seeking interactions. However, the reasoning and interactive capabilities required for these interactions are simpler than those required for argumentation.

Aspects of understanding enthymematic discourse have been considered by several researchers. The abductive process applied by NAG to fill reasoning gaps during argument analysis is similar to that described in [Hobbs et al., 1993]. Specific aspects of understanding enthymematic discourse, such as understanding indirect speech acts and expressions of doubt, have been investigated in [Green and Carberry, 1992, Carberry and Lambert, 1999]. The generation of enthymematic descriptions, which convey some of their material implicitly, has been considered in [Zukerman and McConachy, 1993, Horacek, 1997] for example.

The work of Reed and Long (1997) considers attentional focus when generating additional information to make concepts in a presentation salient for the user. Huang and Fiedler (1997) use a limited implementation of attentional focus to select the step in a proof to be mentioned next. NAG uses attentional focus to guide the content planning and presentation processes, updating the attentional focus as the interaction with the user progresses.

As indicated in Section 1, several researchers have considered specific dialogue phenomena in isolation. Different types of subdialogues have been considered in [Chu-Carroll and Carberry, 1995, Litman and Allen, 1987]. Procedures for selecting an interpretation among candidate options are described for example in [Litman and Allen, 1987, Raskutti and Zukerman, 1991]. An abductive mechanism for identifying and recovering from misunderstandings is discussed in [McRoy and Hirst, 1995].

3 Reasoning Facilities

We define a *nice* argument to be both normatively correct and persuasive for the target audience. In addition, a good argument should also be concise. Satisfying the requirements for niceness sometimes involves trade offs, e.g., some normative correctness may have to be sacrificed in order to make an argument more persuasive. The generation of correct arguments requires normative domain knowledge, while the generation of persuasive arguments requires a model of the audience's beliefs and inferential ability. The latter is also required for the production of concise arguments, to support the omission of easily inferable information. In this section, we describe NAG's knowledge representation scheme and discuss its use in argument analysis and generation.

3.1 Knowledge Representation Scheme

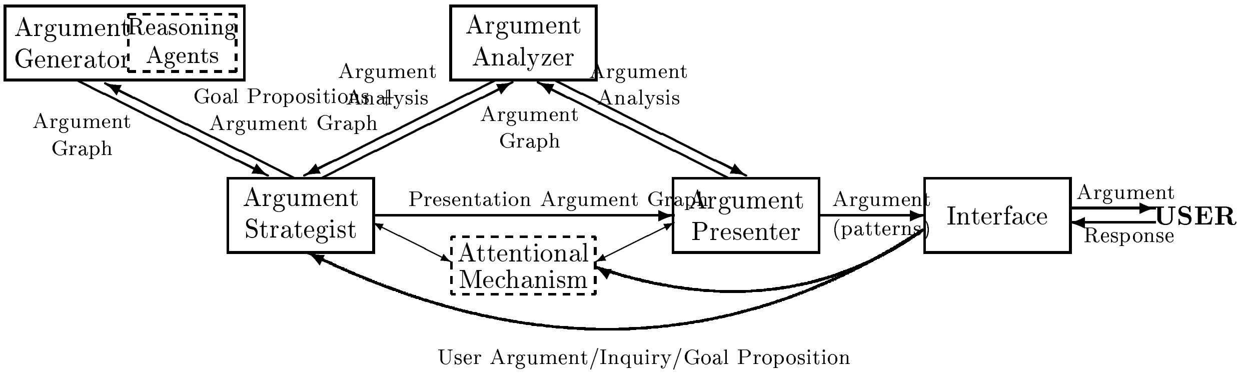

Figure 2: System architecture

3.1.1 Domain Knowledge and the User Model

When constructing or analyzing an argument, NAG relies on two collections of information: a normative model composed of different types of Knowledge Bases (KBs) which represent NAG's best understanding of the domain of the argument, and a user model also composed of different types of KBs which represent the user's presumed beliefs and inferences. A KB represents information in a single format, e.g., semantic network (SN), Bayesian network (BN) or rule-based system. The KBs in the normative and user models are consulted by specialist Reasoning Agents (Figure 2), which are activated to fill gaps in a partial argument (Section 3.2). The KBs in the user model are consulted to make an argument persuasive for the target audience, while the normative KBs are consulted to generate a correct argument. This distinction between normative correctness and persuasiveness allows NAG to control the extent to which it will sacrifice one of these factors for the other.

When reasoning about an argument, relevant material from several KBs may need to be combined into a common representation. NAG uses BNs for this purpose because of their ability to represent normatively correct reasoning under uncertainty. Hence, when reasoning about an argument, information sourced from the user model KBs is added to the user model BN, and information sourced from the normative model KBs is added to the normative model BN. Thus, the nodes in the BNs in the normative and user models and the beliefs in these nodes may change as NAG reasons about an argument (as a result of belief propagation or the incorporation of additional information found in the KBs into the BNs). This is in contrast to NAG's KBs, which are static.

The portions of the normative and user model BNs that are structurally common to both networks and salient in the current context form an *Argument Graph*. For example, consider a situation where propositions P1 and P2 are believed in NAG's normative model. In addition, the normative model has an inference pattern whereby these propositions imply P3. Now, let us assume that the user model shares

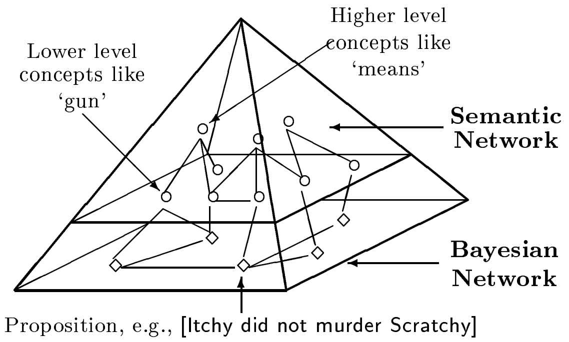

Figure 3: Semantic-Bayesian network

the normative model's belief in P1, is uncertain about P2, and also believes P4. Further, the user model has an inference pattern whereby P1, P2 and P4 affect the belief in P3. In this example, P1 and P2 and their link to P3 are included in the Argument Graph, but P4 is excluded, since it does not appear in the normative model.

3.1.2 Incorporating Context

NAG captures connections between the items mentioned in the discourse by using a *semantic-Bayesian network*, which is composed of a hierarchical semantic network built on top of a BN (Figure 3). Each of the normative model and the user model contains an instance of this structure. The semantic network portion (upper part of the pyramid in Figure 3) and the BN portion (base of the pyramid) are used by NAG to simulate attentional focus in each model. This simulation allows NAG to determine which propositions in its normative and user models are salient in the current context, and hence potentially useful for the argumentation process.

During the argumentation process, NAG's Attentional Mechanism (Figure 2) receives as input a context consisting of a set of salient objects. This context is expanded as the argumentation process progresses. The initial context of an argument consists of the goal proposition and salient concepts and propositions mentioned in the discussion immediately preceding the argument.

In the example used throughout this paper, the argumentation process begins with the presentation of the following preamble, which introduces the fictional Itchy&Scratchy scenario.[1]

> *Scratchy, a notorious spy from Vulcan, was found murdered in his bedroom, which is in the second storey of his house. Indentations were found on the ground right outside Scratchy's window, and they were observed to be circular. In addition, one set of footprints was found outside Scratchy's window, but*

[1] A fictional scenario was chosen to minimize the influence of users' pre-existing beliefs. The information in this preamble is taken to be true by NAG and hopefully also by the user.

the bushes were undisturbed. It is well known that Itchy and Scratchy hated each other, and Itchy's fingerprints were found on the murder weapon (a gun discovered at the scene and registered to Itchy). Itchy, who is the head of the INS, has a ladder with oblong supports, which he was planning to use to paint his house. Poochie, the town's mayor, has a cylindrically-supported ladder which is available only to him.

For this example, the context is composed of the goal proposition, which was initially set to [Itchy did not murder Scratchy], the salient concepts Itchy, Scratchy, ladder and window, and the propositions which include one or more of these concepts. These salient concepts and propositions are activated in the semantic-Bayesian network in each of the user model and the normative model.

We use activation with decay [Anderson, 1983], spreading from the current context, to model the focus of attention. At the beginning of each activation cycle, items already in focus are active. This activation is then passed through the semantic-Bayesian networks, each node being activated to the degree implied by the activation levels of its neighbours, the strength of association to those neighbours, and its immediately prior activation level (vitiated by a time-decay factor). For example, when the node representing guns is activated, the strongly connected node representing Smith & Wesson is passed significant activation, and the weakly connected node representing noisy objects is very weakly activated. The spreading activation process ceases when an activation cycle fails to activate any new node. The items in the semantic-Bayesian networks which achieve a threshold activation level upon completion of the spreading activation process form the current span of attention.

During argument generation, the context contains the concepts and propositions included in the current partial argument, which gets expanded as the content planning process progresses. For example, as NAG attempts to create an argument about Itchy's innocence, it tries to demonstrate that Itchy lacked the means, motive or opportunity to commit the crime. Thus, the context is extended with these subgoals. Spreading activation proceeds from this extended context to activate additional concepts and propositions. During argument understanding, the context is extended with the concepts and propositions in the user's argument. The user's interactions with the system also affect the context, e.g., when the user requests that NAG include a proposition in the current argument, that proposition is brought into the current context (Section 4).

The context extension and spreading activation processes provide NAG with a direct implementation of attentional focus, which is used to identify portions of the semantic-Bayesian networks that are relevant to the argument being analyzed or generated [Zukerman et al., 1998].

3.2 The Argumentation Process

NAG's main components and the data flow between them are illustrated in Figure 2. The argumentation process consists of a series of focusing-generation-analysis cycles which are driven by the Argument Strategist. This series of cycles is used both during argument understanding and generation. During understanding, the Analyzer is invoked to assess the correctness of the user's argument, and the Generator is invoked to fill gaps in the user's (enthymematic) argument. During generation, the Analyzer is invoked to assess the correctness and persuasiveness of NAG's argument as it is being constructed, and the Generator is invoked to build up the argument.

3.2.1 Generating Enthymematic Arguments

The argument generation process starts with the selection of an argument goal, which is then passed to the Strategist (at present, NAG is given the initial goal and the user can select subsequent goals). The Strategist first invokes the Attentional Mechanism to focus attention on parts of the BNs in the user model and the normative model that are likely to be useful in the argument. This process generates an initial Argument Graph, and in later cycles extends the Argument Graph. The Strategist then calls the Generator to continue the argument building process by activating the Reasoning Agents to find additional information to incorporate in the Argument Graph. The extended Argument Graph is returned to the Strategist, which invokes the Analyzer to determine the beliefs in the nodes in the Argument Graph under a variety of conditions. The Strategist uses the Analyzer's assessment to select an argumentation strategy. The strategies being considered at present are: premise to goal, inference to best explanation, *reductio ad absurdum* and argument by cases. If several strategies yield a nice argument, i.e., one where the belief in the goal is as intended for both of the models, a strategy which is likely to yield a concise argument is selected. If no strategy yields a nice argument, another focusing-generation-analysis cycle is performed: the Strategist reactivates the Attentional Mechanism, followed by the reactivation of the Generator and then the Analyzer. This process iterates until a successful Argument Graph is built, or NAG is unable to continue. In this manner, NAG uses contextual information to fill in the missing pieces of an argument, saving computation time by favouring lines of reasoning which are related to the propositions currently in focus [Zukerman et al., 1998].

Once an Argument Graph corresponding to a nice argument has been generated, it is passed to the Presenter, which removes from the Argument Graph premises that have a relatively small effect on the belief in their consequents and intermediate conclusions that can be easily inferred. After each removal, the Analyzer checks that the argument still achieves its goal, and the Attentional Mechanism checks the cohesiveness of the argument. After the Presenter determines that no more propositions can be removed from the argument, it invokes the Interface, which presents the argument in hypertext form [Zukerman et al., 1999].

When understanding a user's argument, the focusing-generation-analysis cycles stop when the user's argument has no more gaps. The Analyzer then reports on the belief in the goal proposition as a result of the user's argument. The Strategist should then determine an appropriate course of action, e.g., attacking a flaw in the user's argument or reinforcing NAG's argument. This decision procedure has not been implemented yet.

3.2.2 Analyzing Enthymematic Arguments

The Analyzer performs constrained Bayesian propagation on the portion of the Argument Graph that is connected to the goal. This is done once in each of the normative model and the user model to assess the correctness and persuasiveness of the argument respectively.

During argument generation, propagation is performed under a variety of conditions which match the different argumentation strategies supported by NAG. For example, the Analyzer negates the goal and propagates the resulting belief through the Argument Graph in order to determine whether *reductio ad absurdum* is an applicable strategy. If a proposition which is firmly believed independently of the goal (in both the normative model and the user model) is substantially undermined as a result of this propagation, then the Strategist will consider *reductio ad absurdum* a viable candidate.

During argument understanding, NAG propagates the beliefs in the propositions that are connected to the goal. However, there may be reasoning gaps within a user's argument or between the user's argument and the goal proposition. NAG takes into account the plausibility and availability of the propositions (or inferences) required to fill these gaps to determine whether a user's argument is enthymematic or poor.

During both argument generation and understanding, NAG uses salient concepts and propositions as focal points while searching for information to fill gaps in the argument. When the Analyzer reports that an argument is not nice enough, i.e., the belief in the goal is not high enough in the user model or the normative model, the Strategist activates the Attentional Mechanism to expand the reasoning context, and then invokes the Generator to produce supporting information for the salient propositions at the boundaries of the reasoning gaps. If the Generator can produce sub-arguments which readily repair the gaps (e.g., with single, "obvious" inferential steps), then the argument was enthymematic, and the newly filled in argument will be re-examined. However, if there is no easy repair for the gaps in the argument, then the source of the argument determines the system's behaviour. If the argument was generated by the user, it will be flagged as poor; if the argument was produced by the Generator, the iterative process of building support for the goal proposition will continue.

4 Dialogue Facilities

An ideal argumentation system interface should be simple to use, support multimedia interaction (accept and produce speech or text accompanied with diagrams as necessary), and allow the user to interrupt the system at any time. As an evolutionary step towards such ideals, NAG features a hypertext interface where arguments are presented in English and the user may respond by choosing among several options. Figure 4 shows NAG's argument in favour of Itchy's innocence, which was presented to a particular user after showing the preamble presented in Section 3.1. After reading the argument, the user may respond as follows: (1) by clicking on one of the options presented at the bottom of the screen; or (2) by clicking on a portion of the argument to focus upon a particular proposition, which in turn leads to the display of additional response options. After each user response, NAG tries to generate a (revised) argument which incorporates the user's request. The user can either retain this argument or ignore it. The latter choice does not necessarily result in the reinstatement of the previous argument, since the context has now changed as a result of the interaction.

We now describe the operations available to the user in NAG's interface. These operations allow the user to examine different aspects of an argument in piece-meal fashion.

Select a new goal. Since NAG can argue only for the propositions it knows, we need a means to restrict the user to these propositions when selecting a goal. At present, this is done by allowing the user to choose a proposition from a pull-down menu. This solution is not appropriate in the long run, since it requires the interface to know in advance all the propositions accessible to the argumentation process (this is not a reasonable requirement for a system that consults knowledge sources containing many propositions). Alternative goal selection options are currently being investigated.

The newly selected goal proposition is added to the reasoning context, and the Strategist activates the generation-analysis process (Section 3.2).

Ask NAG to argue for or against a proposition in the argument. The objective of this option is similar to that of the information-sharing subdialogues described in [Chu-Carroll and Carberry, 1995]. The proposition selected from the argument is added to the reasoning context, and the generation-analysis process is activated to produce a sub-argument for or against this proposition. NAG presumes that a new sub-argument for a selected proposition must be stronger than the current argument for that proposition. If NAG cannot generate an argument for the selected proposition, it will attempt to generate an argument for its

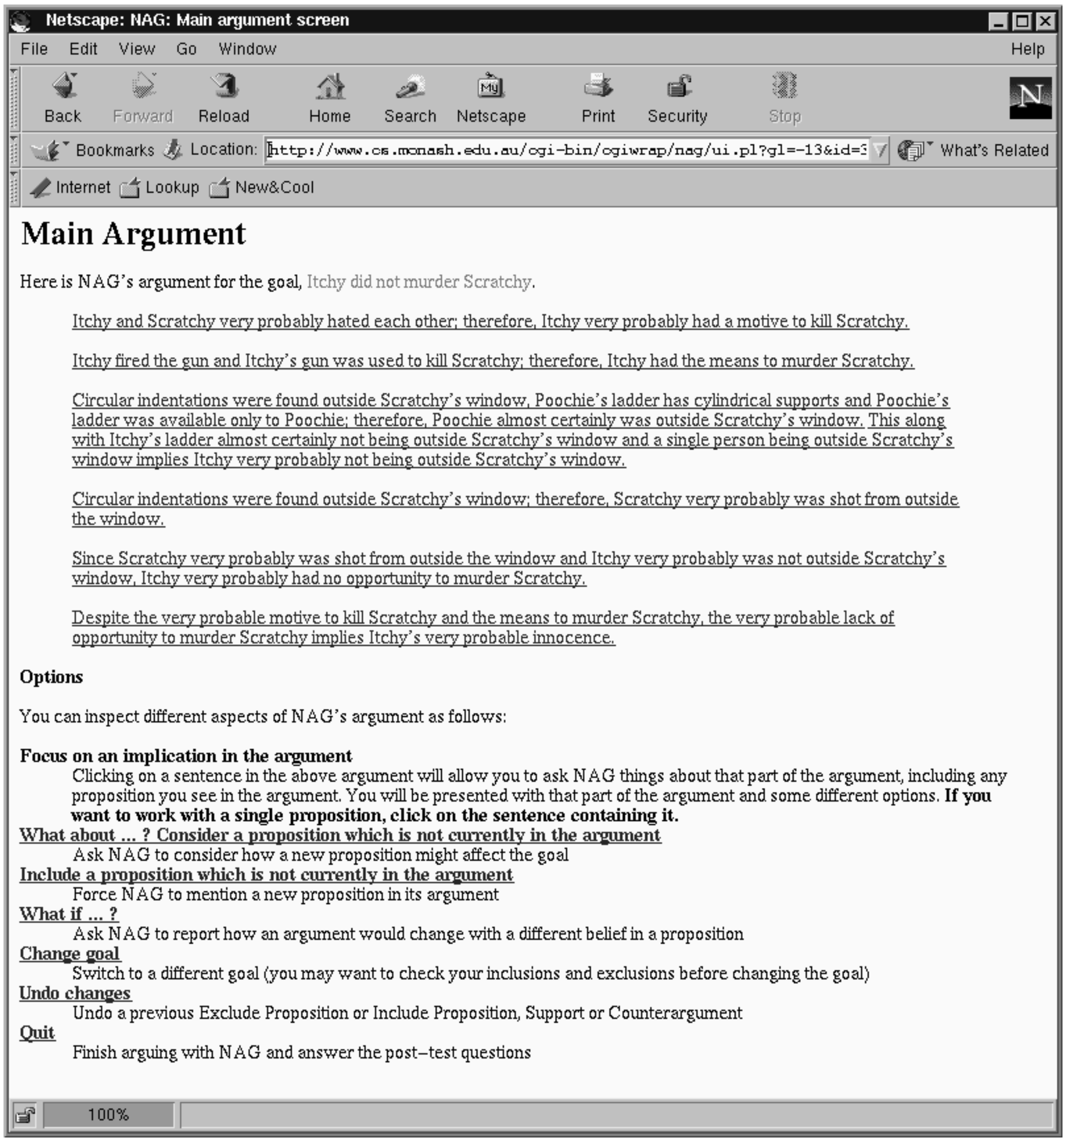

Figure 4: NAG's interface and a sample argument

negation, and inform the user of this fact. If both attempts to generate an argument fail, then NAG reports this failure.

Include/exclude a proposition. As when selecting a new goal, a proposition to be included in the argument is selected from a pull-down menu. The proposition is added to the reasoning context, and the generation-analysis process is activated to produce a revised argument which includes this proposition. The resultant argument may differ substantially from the last argument, in particular if the included proposition has a detrimental effect on the belief in the goal and requires the introduction of additional sub-arguments to maintain the desired level of belief in the goal. If, despite the inclusion of the selected proposition in the context, NAG generates an argument without it, then two additional generation-analysis cycles are activated to try to link the proposition to the argument. If this connection is still not achieved, NAG reports its failure.

A proposition to be excluded is selected from the current argument. NAG tries to construct an argument without it by removing this proposition from the BNs (and those of its ancestors that have no other connections to the BNs). Any children of an excluded proposition that appear in the resulting argument are used as premises. Although any probabilistic impact of excluded propositions is avoided, the exclusion has the opposite effect on attentional processing: asking someone not to think of unicorns has the opposite effect. Thus, excluded propositions are *added* to the reasoning context. This may induce NAG to incorporate related propositions in the argument.

Consider the effect of a proposition (*what about*). This operation is similar to the **include proposition** operation. However, here NAG returns only the reasoning path which connects the selected proposition to the goal (rather than the entire argument), and reports on the effect of this proposition on the goal. This allows the user to investigate the effect of individual factors on the goal. To perform this operation NAG adds the selected proposition to the reasoning context and activates the generation-analysis process. The subgraph that connects the selected proposition to the goal is then returned for presentation. The user can choose to revert to the previous (possibly modified) argument, or to incorporate the examined factor into the argument.

Consider a hypothetical belief (*what if*). In this operation the strength of belief in the selected proposition is set to the value requested by the user in both the normative model and the user model. The proposition is then converted into a premise and added to the reasoning context. These changes in belief are *temporary*, since hypothetical reasoning involves the introduction of beliefs that are not necessarily correct, and hence should not be perpetuated. After producing a revised argument in light of the hypothetical belief, NAG reinstates the original beliefs in both the normative model and the user model. It then returns to the previous (possibly modified) argument.

Undo changes. The user may undo the inclusion or exclusion of propositions and also the inclusion of supporting sub-arguments and counterarguments. Each undo operation brings a proposition into focus. After each undo, the generation-analysis process is reactivated, and NAG presents the resulting argument.

Argue for a proposition (rebuttal). We are currently implementing a very limited capability whereby the user can present a one-step argument for a proposition by selecting antecedents and a consequent from a menu as done for the inclusion of a proposition. NAG then applies the analysis-generation process to fill any gap between the user's argument and NAG's argument, as well as gaps within the user's argument (these gaps happen when the antecedents and consequent in the user's argument are not linked by a single inference in both the normative model BN and the user model BN).

To illustrate NAG's exploratory responses we present part of an actual interaction with a user. After presenting the argument preamble shown in Section 3.1.1 to the user, NAG generated the argument shown in Figure 4 for the goal proposition [Itchy did not murder Scratchy] (the goal proposition, grey in Figure 4, is displayed in pink by the interface). The underlined text in the argument corresponds to hypertext links (blue in the interface) which the user may click to inspect particular propositions more closely. The clickable dialogue options (also blue in the interface) are boldfaced and underlined.

After reading this argument, the user asked NAG to **argue for** the proposition [Itchy was not outside Scratchy's window]. NAG responded with the following sub-argument for this subgoal.

> *Circular indentations found outside Scratchy's window and oblong supports in Itchy's ladder imply Itchy's ladder almost certainly not being outside Scratchy's window. A single set of footprints outside Scratchy's window implies a single person being outside Scratchy's window. The circular indentations found outside Scratchy's window and cylindrical supports to Poochie's ladder imply Poochie's ladder almost certainly being outside Scratchy's window, which together with the availability of Poochie's ladder only to Poochie implies Poochie almost certainly was outside Scratchy's window. This together with Itchy's ladder almost certainly not being outside Scratchy's window and the single person being outside Scratchy's window implies Itchy very probably not being outside Scratchy's window.*

While exploring the list of available propositions, the user noticed the proposition [murder weapon was registered to Itchy]. Asking NAG to **consider** how this proposition affected the argument goal led to the following argument.

> *The fact that a gun was used to kill Scratchy and the gun was registered to Itchy imply Itchy's gun being used to kill Scratchy, which together with Itchy's firing of the gun implies a means to murder Scratchy. Despite this and a very probable motive to kill Scratchy, a very probable lack of opportunity to murder Scratchy implies Itchy's very probable innocence.*

5 Discussion

The system discussed in this paper uses contextual information to bridge gaps in both the system's and the user's arguments during argument generation and understanding. NAG updates the context during the argument generation process and as a result of the user's requests and rebuttals, which in turn affects subsequent arguments. Contextual information is taken into account by the Attentional Mechanism, which enables NAG to favour salient concepts and propositions when reasoning about an argument. For instance, frequently mentioned items enter the focus of attention, leading NAG to reason about them. However, NAG does not deliberately select or avoid particular information items in its arguments because the user exhibited a strong preference for or aversion to them.

We now focus our discussion on the capabilities of an ideal argumentation system mentioned in Section 1, and consider the extent to which NAG supports these capabilities.

Understand implicit goals and enthymematic arguments – NAG's use of the Attentional Mechanism to fill gaps in enthymematic arguments (both its own and the user's) is similar to the use of spreading activation for plan recognition in [Charniak and Goldman, 1993]. At the interface level, NAG makes these problems tractable by restricting the way in which the user can present requests and rebuttals. The user can explore specific aspects of NAG's arguments by selecting actions and propositions from available options; rebuttals are restricted to simple (possibly enthymematic) arguments composed of one set of antecedents and a single consequent.

Generate enthymematic arguments – NAG generates enthymematic arguments by omitting information that can be easily inferred from the argument presented so far.

Admit new information – NAG is currently implemented as a closed system. It does not admit new information presented by the user unless it can find corroborating information in its KBs.

Determine an argumentation strategy – NAG can select from several argumentation strategies, e.g., premise to goal and argument by cases, when presenting its arguments. The strategy chosen for a particular argument depends on how well the information in the Argument Graph fits the requirements of the candidate strategies. For example, if NAG cannot show that some crucial proposition is true or false, but can show that the argument goal follows regardless of the truth value of this proposition, then an argument by cases will be selected.

Handle turn taking, interruptions and subdialogues – NAG presents its arguments as a single turn in a conversation. The user cannot interrupt NAG during the presentation of an argument, thereby avoiding the problem of handling interruptions. Handling conversational turn taking is simplified since each interaction step is composed of a request formulated by the user followed by a response from NAG. Our exploratory queries represent a limited form of subdialogue handling, which supports a single operation for each conversational turn. Some users commented that they would have liked to combine several operations, e.g., reasoning with hypothetical beliefs assigned to two different propositions. We are currently considering the extension of NAG's interface to support this capability.

Track focus – Our use of attentional focus implies that the attentional state can go only forward, i.e., even if the user chooses to ignore the outcome of a request, the change in context has already occurred, hence the interaction cannot return to a previous attentional state. In addition, the consideration of the user's request may have caused changes in beliefs in one or both of the models (e.g., due to the propagation of the beliefs in newly added propositions). Thus, if the user chooses to ignore the outcome of a request, the argument that is 'reinstated' may be quite different from the original one. It is unclear whether users require a focus-tracking capability that supports the return to earlier belief states.

Recover from misunderstandings – This problem is largely avoided by NAG's interface, since the available operations leave little room for misunderstandings. The user is responsible for recovering from a misunderstanding, in the sense that if she or he finds a particular argument unsatisfactory, she or he can simply try a different operation. A special type of misunderstanding takes place when the user selects incompatible operations. For example, this happens when the user requests that NAG exclude a particular proposition from all arguments and later asks NAG to generate an argument for a different goal which requires the excluded proposition. A more extreme version of this problem occurs when the user attempts to select the excluded proposition itself as a new goal proposition. We are currently considering actions which will avoid or correct these problems.

Clearly, our interface, while practical, is not as powerful or flexible as an ideal interface. At this stage, it is mainly an exploratory interface in that it allows the user to probe the system's argument, while giving the user only limited capabilities to present his or her own views. One of the main advantages of our interface is in the explicit presentation of the options available to the user. We speculate that such a presentation, which would be absent from a more natural interface, prompts the user to examine the system's argument in more detail than he or she would otherwise. To validate this idea, a comparative study of both modes of interaction would be required. In the mean time, a preliminary evaluation of the system suggests that our exploratory interactions have a beneficial impact on users' beliefs in propositions mentioned in NAG's arguments [Zukerman et al., 1999].

6 Acknowledgments

This work was supported in part by Australian Research Council grant A49531227. The authors are indebted to Kevin Korb for his work on NAG, and to Deborah Pickett for her implementation of NAG's interface.

References

[Anderson, 1983] Anderson, J. R. (1983). *The Architecture of Cognition.* Harvard University Press, Cambridge, Massachusetts.

[Carberry and Lambert, 1999] Carberry, S. and Lambert, L. (1999). A process model for recognizing communicative acts and modeling negotiation subdialogues. *Computational Linguistics*, 25(1):1–53.

[Charniak and Goldman, 1993] Charniak, E. and Goldman, R. P. (1993). A Bayesian model of plan recognition. *Artificial Intelligence*, 64(1):50–56.

[Chu-Carroll and Carberry, 1995] Chu-Carroll, J. and Carberry, S. (1995). Generating information-sharing subdialogues in expert-user consultation. In *IJCAI95 – Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1243–1250, Montreal, Canada.

[Green and Carberry, 1992] Green, N. and Carberry, S. (1992). Conversational implicatures in indirect replies. In *Proceedings of the Thirtieth Annual Meeting of the Association for Computational Linguistics*, pages 64–71, Newark, Delaware.

[Hobbs et al., 1993] Hobbs, J. R., Stickel, M. E., Appelt, D. E., and Martin, P. (1993). Interpretation as abduction. *Artificial Intelligence*, 63(1-2):69–142.

[Horacek, 1997] Horacek, H. (1997). A model for adapting explanations to the user's likely inferences. *User Modeling and User-Adapted Interaction*, 7(1):1–55.

[Huang and Fiedler, 1997] Huang, X. and Fiedler, A. (1997). Proof verbalization as an application of NLG. In *IJCAI97 – Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pages 965–970, Nagoya, Japan.

[Jönsson, 1995] Jönsson, A. (1995). Dialogue actions for natural language interfaces. In *IJCAI95 – Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1405–1411, Montreal, Canada.

[Litman and Allen, 1987] Litman, D. and Allen, J. F. (1987). A plan recognition model for subdialogues in conversation. *Cognitive Science*, 11:163–200.

[McRoy and Hirst, 1995] McRoy, S. W. and Hirst, G. (1995). The repair of speech act misunderstandings by abductive inference. *Computational Linguistics*, 21(4):435–478.

[Raskutti and Zukerman, 1991] Raskutti, B. and Zukerman, I. (1991). Generation and selection of likely interpretations during plan recognition. *User Modeling and User Adapted Interaction*, 1(4):323–353.

[Reed and Long, 1997] Reed, C. and Long, D. (1997). Content ordering in the generation of persuasive discourse. In *IJCAI97 – Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pages 1022–1027, Nagoya, Japan.

[Vreeswijk, 1994] Vreeswijk, G. (1994). IACAS: An interactive argumentation system. Technical Report CS 94-03, Department of Computer Science, University of Limburg.

[Zukerman and McConachy, 1993] Zukerman, I. and McConachy, R. (1993). Generating concise discourse that addresses a user's inferences. In *IJCAI93 – Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 1202–1207, Chambery, France.

[Zukerman et al., 1998] Zukerman, I., McConachy, R., and Korb, K. B. (1998). Bayesian reasoning in an abductive mechanism for argument generation and analysis. In *AAAI98 – Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 833–838, Madison, Wisconsin.

[Zukerman et al., 1999] Zukerman, I., McConachy, R., and Korb, K. B. (1999). Exploratory interaction with a Bayesian argumentation system. In *IJCAI99 – Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, Stockholm, Sweden.