Recognizing Intentions from Rejoinders in a Bayesian Interactive Argumentation System

Ingrid Zukerman, Nathalie Jitnah, Richard McConachy and Sarah George

School of Computer Science and Software Engineering
Monash University, Clayton, Victoria 3800, AUSTRALIA
{ingrid,njitnah,ricky,sarahg}@csse.monash.edu.au

Abstract. We describe a mechanism which recognizes a user's intentions from short-form rejoinders to arguments generated from Bayesian networks. The mechanism builds candidate reasoning paths that link the user's rejoinder with a previously presented argument, and considers the following factors to select a path: linguistic clues, the impact of the user's rejoinder on the system's argument along the different paths, the user's attentional focus, and the system's confidence in its representation of the user's beliefs. The results of our preliminary evaluation indicate that the interpretations produced by our mechanism are generally appropriate.

1 Introduction

Ideally, an interactive argumentation system would allow a user to respond to an argument with a counterargument, and it would allow the argumentation process to go on indefinitely, producing a series of arguments and counterarguments. During argumentation, conversational partners often use expressions of doubt, such as "*But the victim was stabbed*", and requests for the consideration of additional facts, such as "*What about the fingerprints on the gun?*". In this paper, we describe a mechanism which interprets such rejoinders to arguments generated from Bayesian networks (BNs) [8]. This mechanism is implemented in a system called BIAS (*Bayesian Interactive Argumentation System*).

Given an argument produced by BIAS followed by a rejoinder posed by a user, our mechanism identifies the user's likely line of reasoning and the proposition(s) in BIAS' argument the user intends to affect. This is done by taking into account the following factors: the linguistic clues in the rejoinder, the impact of the rejoinder on BIAS' argument, the user's attentional focus, and BIAS' confidence regarding its representation of the user's beliefs. Once a line of reasoning has been postulated, BIAS generates a rebuttal to the user's rejoinder [6].

In Section 2, we introduce BIAS' current scenario and show a sample interaction with BIAS. Next, we describe our knowledge representation formalism, followed by the algorithm which identifies the user's line of reasoning. We then discuss results of a preliminary evaluation of BIAS' performance, review related research and present concluding remarks.

2 Scenario

BIAS and the user are partners in solving a crime. They have access to information about the world, e.g., the crime scene, witnesses and forensic reports. At the beginning of the interaction, BIAS and the user receive a preamble that describes the preliminaries of the case (Figure 1). After receiving the preamble, the user can use BIAS' WWW interface

Preamble: *Mr Body was found dead in his bedroom, which is in the second story of his house. Bullet wounds were found in Mr Body's body. The bedroom window was broken, and broken glass was found inside the window. A gun was found at the premises, and some fingerprints were found on the gun. In addition, inspection of the grounds revealed footprints in the garden and circular indentations in the ground outside the bedroom window.*
Initial argument: *Bullets being found in Mr Body's body implies Mr Body was almost certainly shot. This implies Mr Body was murdered.* *Forensics matching the bullets with the found gun implies the gun is almost certainly the murder weapon. Forensics matching the fingerprints with Mr Green implies Mr Green probably fired the gun, which together with the gun almost certainly being the murder weapon implies Mr Green probably fired the murder weapon. This implies he very probably had the means to murder Mr Body.* *The Bayesian Times reporting Mr Body took Mr Green's girlfriend implies Mr Green and Mr Body very probably were enemies, which implies Mr Green probably had a motive to murder Mr Body.* *A witness reporting Mr Green being at the football at 10:30 implies Mr Green almost certainly wasn't in the garden at 11.* *Forensics reporting the time of death being 11 implies the time of death was very probably 11. This together with Mr Green almost certainly not being in the garden at 11 implies he almost certainly wasn't in the garden at the time of death, which implies he almost certainly didn't have the opportunity to murder Mr Body.* *Even though Mr Green very probably had the means to murder Mr Body and he probably had a motive to murder Mr Body, Mr Green almost certainly not having the opportunity to murder Mr Body implies he probably didn't murder Mr Body.*
Rejoinder: *But Mr Green was in the garden at 11*.
Interpretation: *Mr Green being in the garden at 11 implies he was more probably in the garden at the time of death. This implies he more probably had the opportunity to murder Mr Body, which implies he more probably murdered Mr Body.*

Fig. 1. Preamble, initial argument, sample rejoinder and BIAS' interpretation

to obtain additional information about the world, e.g., from witnesses or the crime scene, and to post his/her beliefs about selected propositions (this is done by clicking a belief value for these propositions). BIAS has access to these beliefs and to the obtained information, but it does not necessarily share the user's beliefs. Further, BIAS can investigate the world directly to obtain additional information that will enable it to formulate an argument.

When the user asks for BIAS' opinion about the case, BIAS calls a Bayesian argument generator [11] to produce a preliminary argument in support of a goal proposition. In our scenario, this proposition is either *Mr Green is guilty* or *Mr Green is innocent*, whichever is most likely. If possible, our argument generator produces an argument that is compatible with both BIAS' beliefs about the world and the user's presumed beliefs. Otherwise, BIAS' beliefs take precedence. Figure 1 shows the argument generated by BIAS for Mr Body's innocence in light of the preamble and information gathered from the world.

After receiving BIAS' argument, the user can formulate a rejoinder or continue investigating the world. At present we consider two types of rejoinders (which are formulated by making selections from a dynamic menu in our WWW interface).

– Expressions of doubt ("But **R**", where **R** is a proposition). We focus on one type of expression of doubt where the user asserts or negates a proposition to undermine a proposition stated or implied in the system's argument [2].
– Requests for the consideration of a proposition ("Consider **R**"). Unlike expressions of doubt, which have negative implications, this type of rejoinder just implies that BIAS has omitted a factor that could be relevant.

For example, after receiving the rejoinder in Figure 1, BIAS postulates the line of reasoning shown at the end of Figure 1.[1] This line of reasoning takes into account the user's beliefs in other nodes that were mentioned in BIAS' argument, e.g., *the time of death was 11*, and the user's belief in nodes s/he investigated through the WWW interface. BIAS then generates a rebuttal which addresses the user's rejoinder [6] or produces an updated argument which acknowledges the impact of the rejoinder proposition. The user can now continue inspecting the world or pose another rejoinder, and so on.[2]

3 Knowledge Representation

We have chosen BNs as our main representational formalism owing to their ability to represent normatively correct reasoning under uncertainty. The information about the world and the models of belief consulted by BIAS during the argumentation process are represented as BNs. These models are a normative model and a user model. The normative model contains information that is presumed to be correct according to the world, i.e., observable facts obtained from the world or beliefs inferred by means of Bayesian propagation from the observable facts. The user model stores propositions that are presumed believed by the user, where the probability values of these propositions represent the user's beliefs. These propositions are obtained from a variety of sources, such as BIAS' arguments or the beliefs entered by the user through the interface, and are labelled according to their source(s). For example, a proposition is labelled `accepted` if a belief in it has been entered by the user (either through the interface or by confirming BIAS' interpretation of the user's rejoinder), while it is labelled `seenObservation` or `seenIntuition` if it has been shown to the user, but the user has not indicated a belief in it (`Observations` represent observable events, while `Intuitions` represent inferable propositions). The labels are ranked according to the trustworthiness of the source from which the user's belief was obtained. For instance, `accepted` propositions rank higher than `seenObservations`, which in turn rank higher than `seenIntuitions`. In addition, in order to support a numerical process for rating the candidate reasoning paths (Section 4), these labels are associated with a numerical score according to their ranking.

The interpretation process, which is the focus of this paper, is performed in the context of the user model, since the rejoinder should "make sense" in light of the user's beliefs. In contrast, the processes for generating the initial argument and the rebuttals consult the user model and the normative model in order to produce arguments that rely on beliefs held by both BIAS and the user if possible [11]. These arguments are represented by means of a sub-network of the normative model BN, called an *Argument Graph*, which ideally also contains nodes from the user model BN.

[1] The implications in this line of reasoning may be causal or evidential.
[2] At present, we assume that a user's rejoinder addresses BIAS' initial argument. The dialogue features that determine whether a rebuttal is being addressed are yet to be implemented.

During the generation of arguments and rebuttals and the interpretation of a user's rejoinders, we model the user's attentional focus. The use of attentional focus during argument generation is described in [11]. In this paper, we focus on its impact on the interpretation process. We postulate that the interpretation intended by a user is likely to contain propositions in his/her focus of attention. In order to determine whether a proposition is in the user's focus of attention, we use a model of attention which follows associative links (rather than causal or evidential links) and invoke a process called *spreading activation* [1]. This process passes activation from active propositions, e.g., recently seen propositions, to propositions and concepts that have an associative relation to the active propositions. For instance, after reading the fragments "Bayesian Times" and "Mr Body is dead", the concepts "time" and "death" get activated, in turn activating propositions pertaining to the "time of death". Our model of attention is implemented by incorporating in the user model an associative semantic network which includes the propositions in the user model BN (such a network is also incorporated in the normative model).

4 Path Identification

Algorithm *IdentifyPath* proposes paths in the user model BN that represent possible lines of reasoning from the user's rejoinder. The algorithm receives two inputs: a linguistic clue ("but" or "consider") and the rejoinder proposition (**R**).

Algorithm *IdentifyPath(linguisticClue,***R***)*

1. **Path construction** – Find paths that connect the rejoinder node **R** to the goal proposition in the Argument Graph generated by BIAS.
2. **Path evaluation** – Compute a score for each path based on its effect on the argument according to the user model, its presence in the user's focus of attention and BIAS' confidence in it.
3. **Path selection** – Select one path if possible. Otherwise, present the candidate path(s) to the user for confirmation.

4.1 Path Construction

The user may select a rejoinder node **R** that is not in the user model BN. In this case, **R** is added to the user model (and to the normative model if necessary). When trying to connect **R** to the Argument Graph, BIAS iteratively expands the user model BN from **R** and from the Argument Graph. In each iteration, BIAS adds to the user model nodes and links from the world BN that connect with **R**, and nodes and links from the world BN that connect with the Argument Graph (the corresponding conditional probability tables are also copied from the world BN).[3] If this process is successful, upon its completion **R** will be connected to the Argument Graph by one or more paths. These paths may contain causal links (from the rejoinder node towards the goal) or evidential links (in the opposite direction). This process is performed in a *temporary user model*, since at most one of the paths found by BIAS is intended by the user. This path will be incorporated in the user model after performing path selection (Section 4.3).

[3] During this process, BIAS drops from consideration d-separated paths [8]. A path is d-separated when the presence or absence of some evidence in the user model prevents the rejoinder node from affecting the goal proposition.

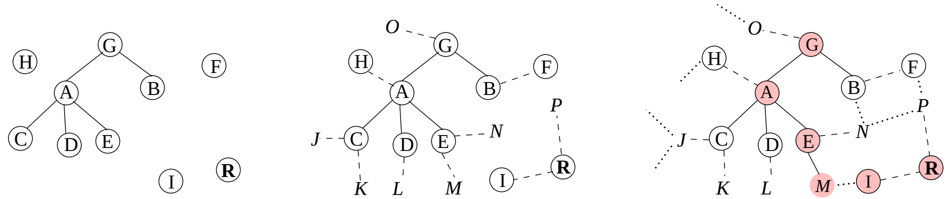

(a) Initial user model (b) After one expansion (c) After two expansions

Fig. 2. Path construction in a sample Argument Graph

To illustrate the path construction process, consider the user model in Figure 2(a), which contains the Argument Graph corresponding to an argument presented by BIAS (the tree whose root is the goal node G), the nodes investigated by the user in the interface (F, H, I), and the user's rejoinder proposition **R**. After one expansion, propositions F, H and I become connected to **R** or to the Argument Graph, and additional propositions, *J, K, L, M, N, O* and *P* (shown in italics in Figure 2(b)), are added to the temporary user model and also linked to **R** or the Argument Graph (these links are drawn with dashed lines). After the second expansion (drawn with dotted lines), **R** becomes connected to the original Argument Graph and the goal through several paths in the temporary user model: **R**-I-*M*-E-A-G (composed of grey nodes in Figure 2(c)), **R**-*P*-*N*-E-A-G, **R**-*P*-F-B-G and **R**-*P*-F-B-*N*-E-A-G.

We assume that the user intends a straightforward interpretation of his/her rejoinder, and that s/he has bounded inferential capacity. BIAS implements these assumptions during path construction by looking for *direct* paths that represent *small inferential leaps*. A direct path represents a simple line of reasoning between the rejoinder node and the goal. A direct path is not necessarily the shortest path to the goal. Rather, it is a path that does not meander unnecessarily, i.e., if there are several routes between two neighbouring nodes, a direct path takes the link between these nodes. An inferential leap occurs when a user intends to affect a proposition in the argument that is different from the rejoinder proposition. For instance, if the user says "*But Mr Green was near the house at 11:15*" after the argument in Figure 1, s/he licenses an implication regarding Mr Green's opportunity to kill Mr Body (through a small inferential leap). However, if the user had said "*But Mr Green owns a blue car*", a system should not be expected to find an interpretation that connects this rejoinder to the goal, since the inferential leap is too large.

BIAS also copes with two types of inaccuracies in its user model: incompleteness (when BIAS is not aware of some of the user's beliefs) and granularity discrepancies (when the user makes a 'complex' inference that connects non-adjacent nodes in the user model BN). During path construction, these inaccuracies show up as *reasoning gaps* in the user's presumed line of reasoning, i.e., the line of reasoning includes propositions that are not in the user model. BIAS considers only paths with *small and easily inferable reasoning gaps*. A gap should be small in order to avoid attributing to a user overly complex inferences (across a large gap); the inference across a gap should be "easy", meaning that is plausible that the user is engaged in this reasoning activity. We adopt the "ease of inference" definition from [7], which requires that the node at the tail of the gap have a strong effect on the node at the head of the gap, and that the nodes at both ends of the gap be in the user's focus of attention. This last requirement approximates the idea that thinking of the antecedent of an inference will make one think of its consequent (across the gap).

BIAS finds reasoning paths which represent small inferential leaps by performing only three iterations to connect **R** to the Argument Graph, and by restricting the length of the paths being built. The small size of the reasoning gaps is ensured by allowing an inferred path to contain at most two consecutive nodes that are not in the user model. For instance, path **R**-I-*M*-E-A-G in Figure 2(c) has a gap of length 1 between I and E. The ease of inference across a reasoning gap is implemented by requiring that (1) the nodes at both ends of the gap, e.g., I and E, have a high activation level; and (2) some value of the node at the 'head' of the gap have a high level of belief given the value of the node at the 'tail' of the gap (inferred from the user's rejoinder).

4.2 Path Evaluation

The path evaluation process produces a score for each path returned by the path construction process. This score, called *pathValue*, incorporates the following factors: (1) the impact of **R** along this path on BIAS' argument, (2) the linguistic clue of the rejoinder, (3) whether the nodes in this path are in the user's focus of attention, and (4) BIAS' confidence regarding its representation of the user's beliefs in the nodes along this path. These factors are assessed in the context of the user model, since BIAS is trying to determine what the user means by the rejoinder. The calculation of *pathValue* is performed by means of heuristics which combine Bayesian principles (Factor 1) with dialogue-related and user modeling aspects (Factors 2, 3 and 4).

Impact of R along a path on BIAS' argument. The impact of a rejoinder **R** along $path_j$ on a proposition $\mathcal{X}$, denoted $Impact_j(\mathbf{R}, \mathcal{X})$, represents the change in belief in proposition $\mathcal{X}$ in light of the value of **R** stated in the rejoinder (denoted *userVal*). This change is relative to the previous belief in $\mathcal{X}$ (in light of a different value of **R** or no information about **R**). The impact of **R** on $\mathcal{X}$ is calculated using the following formula.

$$Impact_j(\mathbf{R}, \mathcal{X}) = \log \frac{\Pr_j(\mathcal{X} = x | \mathbf{R}{=}userVal)}{\Pr(\mathcal{X} = x)} \quad (1)$$

$\Pr_j(\mathcal{X} = x | \mathbf{R}{=}userVal)$, the probability of node $\mathcal{X}$ along $path_j$ given the user's value of **R**, is calculated by propagating the user's value of **R** over a temporary BN consisting of the current user model BN plus the nodes in $path_j$ (the conditional probability tables are marginalized to take into account nodes that are absent from the user model).

We assume that the user's rejoinder was generated to affect at least one node in the Argument Graph and possibly the goal proposition, e.g., in path **R**-I-*M*-E-A-G in Figure 2(c), the propositions of interest are E, A and G. The effect of the rejoinder depends on its linguistic clue. A request for consideration implies that the rejoinder affects some propositions in the argument (without indicating whether it supports or contradicts these propositions), while an expression of doubt implies that the rejoinder contradicts a proposition in the argument. These considerations are combined with $Impact_j(\mathbf{R}, \mathcal{X})$ to yield the *effective impact* of **R** along $path_j$ on a node $\mathcal{X}_{AG}$ in the Argument Graph.

$$EffImpact_j(\mathbf{R}, \mathcal{X}_{AG}) = \begin{cases} |Impact_j(\mathbf{R}, \mathcal{X}_{AG})| & \text{if request for consideration} \\ Impact_j(\mathbf{R}, \mathcal{X}_{AG}) & \text{if } \Pr(\mathcal{X}_{AG} = argBIAS) < 0.5 \\ -Impact_j(\mathbf{R}, \mathcal{X}_{AG}) & \text{otherwise} \end{cases} \quad (2)$$

where *argBias* is the value of node $\mathcal{X}_{AG}$ resulting from BIAS' argument. According to this formula, if BIAS argued for a high/low degree of belief in node $\mathcal{X}_{AG}$, then a felicitous expression of doubt should reduce/increase the belief in this node. For example, if BIAS argued for $\mathcal{X}_{AG}$=True, yielding $\Pr(\mathcal{X}_{AG} = \text{True}) = 0.8$ in the user model, and the user's expression of doubt reduced this belief to $\Pr_j(\mathcal{X}_{AG} = \text{True}|\mathbf{R}) = 0.6$ when propagated along $path_j$, then the effective impact of this rejoinder on node $\mathcal{X}_{AG}$ would be positive.

BIAS uses the effective impact of a rejoinder along a path to suggest a node in the Argument Graph that the user intended to affect. We submit that this is the *highest-impact* node, i.e., the node with the highest value for *EffImpact* along this path (the goal may also be the highest-impact node). Thus, when considering the effective impact of a rejoinder to assess a proposed path, we focus on two target propositions: the goal (G) and the *highest-impact* proposition (*HI*). For example, assume that the user said "But **R**", and consider path **R**-I-*M*-E-A-G in Figure 2(c). Further, assume that $\Pr(\text{E=True}) = 0.3$, $\Pr(\text{A=True}) = 0.6$ and $\Pr(\text{G=True}) = 0.7$ after BIAS' argument, and that $\Pr(\text{E=True}) = 0.55$, $\Pr(\text{A=True}) = 0.2$ and $\Pr(\text{G=True}) = 0.5$ due to the user's rejoinder. Thus, all three propositions have been contradicted, but A is the highest-impact proposition.

The presence of the nodes along a path in the user's focus of attention. When presenting a rejoinder proposition, the user has in mind a line of reasoning that links this proposition to the system's argument. We postulate that it is more likely that the user will perform this inferential leap if the nodes along his/her intended line of reasoning are in his/her focus of attention. That is, they have a high level of activation.

BIAS' confidence regarding its representation of the user's beliefs in the nodes along a path. The more trustworthy are the sources from which BIAS obtained the user's beliefs in the nodes along a path, the more confident BIAS should be regarding the plausibility of this path. This confidence is a function of the numerical score associated with the label of each node (Section 3).

Determining the overall value of a path. The factors discussed above have the following contribution to the value of a path.

- **Linguistic clue and impact along a path –** Paths with a high combined effective impact on the goal and the highest-impact node are preferred, since it is more likely that the user intended to have a large effect on BIAS' argument.
- **Focus of attention –** Paths with highly activated nodes are preferred, since these nodes are likely to be in the user's mind. We compute the average level of activation over a path (rather than total activation) so shorter paths are not disadvantaged.
- **BIAS' confidence –** Paths whose nodes have labels with high scores are preferred, since the system is more certain about these nodes. As above, we compute the average label score over a path.

These considerations are incorporated into the following formula, which represents the overall value of a path.

$$\begin{aligned} \mathit{pathValue}_j(\mathbf{R}) = {} & N_G \times \mathit{EffImpact}_j(\mathbf{R}, \mathrm{G}) + N_{HI} \times \mathit{EffImpact}_j(\mathbf{R}, \mathit{HI}) + \\ & N_A \times \log\left(1 + \frac{\sum_{node_i \in path_j} \mathit{Activation}(node_i)}{\mathit{Length}(path_j)}\right) + \\ & N_L \times \log \frac{\sum_{node_i \in path_j} \mathit{LabelScore}(node_i)}{\mathit{Length}(path_j)} \end{aligned} \quad (3)$$

where the weights N_G, N_{HI}, N_A and N_L determine the contribution of the above factors to the value of a path: effective impact of the rejoinder on the goal (N_G) and on the highest-impact proposition (N_{HI}), level of activation of the nodes in the path (N_A), and scores of the labels of these nodes (N_L).[4] In Section 5, we consider the effect of these weights on the performance of the system.

4.3 Path Selection

As indicated in Section 4.2, the value of a path reflects the likelihood that this is the path intended by the user. Thus, when there is a single path with a high *pathValue* or the *pathValue* of a path is significantly higher than that of the other candidate paths, this path is selected and passed to the rebuttal generation procedure [6]. However, when several paths with similar *pathValues* are generated, BIAS cannot discriminate between them, and lets the user choose; a single path with a low *pathValue* is also presented to the user for confirmation.

The recognition process fails if BIAS could not find a path between the rejoinder proposition and the Argument Graph during path construction, or too many paths with similar *pathValues* were found (so they could not be presented for confirmation), or the user does not select any of the presented paths. In the future, our interface will ask the user to further specify his/her rejoinder in these situations.

5 Preliminary Evaluation

Our preliminary evaluation assesses the overall performance of our system and the influence of the factors considered in Equation 3 on this performance.

The overall performance of the system was evaluated by means of the following experiment. Twelve subjects were shown the preamble and argument from Figure 1, and were shown the following rejoinders, each accompanied by the candidate interpretation(s) proposed by BIAS: (1) *But Mr Green and Mr Body had an argument*, (2) *But the forensic analysis of the found fingerprints is reliable*, (3) *Consider that the found gun was not registered to Mr Green*, (4) *But Mr Green's ladder was at Mr Body's window*, and (5) *But Mr Green was in the garden at 11*. Rejoinders 1, 2, 3 and 5 had a single interpretation, while Rejoinder 4 had three interpretations (the interpretations for Rejoinders 1 and 2 are shown in Figure 3, and that for Rejoinder 5 appears in Figure 1). These interpretations were generated in light of the user model which results from the presentation of the preamble and BIAS' argument (the values for the weights in Equation 3 were $N_G = 2, N_{HI} = 3$ and $N_A, N_L = 1$, i.e., the effective impact of the rejoinder on the highest-impact proposition is the most important factor). The subjects were then asked to give each interpretation a score between 1 (very UNreasonable) and 5 (very reasonable), and to propose their own interpretations if they found BIAS' inappropriate. The number of people who gave each score is shown in Table 1 together with the average score for each interpretation. These results indicate that the interpretations proposed by BIAS were considered generally appropriate. In addition, the single interpretations generally had more support than the multiple ones, indicating that users split their support when several plausible interpretations are available. For Rejoinder 4, all the interpretations had

[4] We compute the log of the last two factors so that their contribution to the value of a path is compatible with that of *EffImpact*.

Interpretation for Rejoinder$_1$ (*But Mr Green and Mr Body had an argument*): *Mr Green arguing with Mr Body implies Mr Green and Mr Body more probably were enemies. This implies Mr Green more probably had a motive to murder Mr Body, which implies he more probably murdered Mr Body.*
Interpretation for Rejoinder$_2$ (*But the forensic analysis of the found fingerprints is reliable*): *The forensic analysis of the found fingerprints being reliable and forensics matching the fingerprints with Mr Green imply the found fingerprints more probably belong to Mr Green. This implies Mr Green more probably fired the found gun, which implies he more probably fired the murder weapon. This implies he more probably had the means to murder Mr Body, which implies he more probably murdered Mr Body.*

Fig. 3. Sample rejoinders with BIAS' interpretations

Interpretation for	Number of people who gave a score of					Average score
	1	2	3	4	5	
Rejoinder$_1$	0	2	0	8	2	3.83
Rejoinder$_2$	0	0	0	8	4	4.33
Rejoinder$_3$	0	3	3	3	3	3.50
Rejoinder$_5$	0	2	0	4	6	4.17

Interpretation for Rejoinder$_4$	Number of people who gave a score of					Average score
	1	2	3	4	5	
Interpretation$_1$	0	3	2	4	3	3.58
Interpretation$_2$	1	3	3	4	1	3.08
Interpretation$_3$	2	3	4	2	1	2.75

Table 1. Scores given by subjects to BIAS' interpretations for five rejoinders

support from some subjects (scores of 4 and 5), indicating that all the generated interpretations were worth presenting. It is also worth noting that the ranking produced by the scores of the interpretations of Rejoinder 4 (3.58>3.08>2.75) matches BIAS' ranking for these interpretations according to their *pathValues*.

The influence of the factors considered in Equation 3 on BIAS' performance was assessed by means of three experiments. In the first experiment, we activated BIAS with two settings: (1) `EqualWeights`, where $N_G = N_{HI} = N_A = N_L = 1$; and (2) `ImpactOnly`, where $N_G = N_{HI} = 1$ and $N_A = N_L = 0$. The objective of this experiment was to determine whether attentional focus and BIAS' confidence in a path affect the relative ranking of the paths. Each setting was tested on five rejoinder nodes for which BIAS generated multiple paths. In three of the five cases, the `ImpactOnly` setting yielded different rankings to those produced by the `EqualWeights` setting. This indicates that the weights assigned to attentional focus and the system's confidence in a path affect BIAS' results.

The second experiment was performed with the same five rejoinder nodes as the first experiment. Here we activated BIAS with the `EqualWeights` setting, and simulated user clicks of nodes in lower-ranked paths, thereby increasing their activation and changing their labels. This caused the lower-ranked paths to move up in rank in all the runs. The results of this experiment indicate that by taking into account a user's attentional focus and the system's confidence in a path, BIAS can react appropriately to the user's input.

The third experiment was conducted on four rejoinders, such as "*But a blue car was here*", for which BIAS returned no paths, because the paths found during path construction had reasoning gaps that were too large. We then simulated user clicks to relevant nodes, which resulted in the addition of these nodes to the user model BN (and also affected their activation and labels). This in turn enabled BIAS to propose interpretations for the problematic rejoinders. The results of this experiment show that additional contextual information enables BIAS to propose interpretations which would otherwise be considered far fetched.

6 Related Research

In this section, we focus on related research that specifically pertains to the topic of this paper, viz computational mechanisms for intention recognition during argumentation and applications of BNs to argumentation and plan recognition.

Our research builds on the system described in [11], which generates arguments from BNs, and the system described in [12], which allows a user to explore an argument by performing certain modifications, such as excluding a proposition. However, these systems did not interpret a user's utterances.

Several researchers have dealt with different aspects of intention recognition during argumentation, e.g., [2, 4, 9, 10]. Flowers *et al.* [4] focused on recognizing episodic justifications to historical events, Quilici [9] considered plan-related arguments, and Carberry and Lambert [2] and Restificar *et al.* [10] modeled expert-consultation dialogues.

Our system is closest to Carberry and Lambert's [2] in its focus on rejoinders and its combination of linguistic, contextual and world knowledge to recognize a user's intentions. However, there are significant differences between our models. Carberry and Lambert covered a wider range of linguistic phenomena than those considered in this paper. However, they considered short exchanges where each participant utters a few propositions in each conversational turn, and they used a plan-based inference mechanism to recognize a user's intention. Restificar *et al.* [10] also modeled short exchanges, using simple argument schemata combined with inference rules to detect whether an utterance attacks or supports an argument. The interpretation of rejoinders to complex, probabilistic arguments calls for techniques such as Bayesian propagation and spreading activation to determine the impact of a proposition and to model attentional focus respectively.

BNs have been used in a variety of plan recognition tasks. For example, Heckerman and Horvitz [5] used a BN applied to features extracted from users' queries to infer their software assistance requirements. Charniak and Goldman [3] used BNs and marker passing (a form of spreading activation) for plan recognition during story understanding. They automatically built and incrementally extended a BN from propositions read in a story, so that the BN represented hypotheses that became plausible as the story unfolded. During this process, they used marker passing to restrict the nodes included in the BN. In contrast, we use BNs as a formalism for argument representation, apply Bayesian propagation to recognize a user's intention when posing a rejoinder to an argument, and use spreading activation to model the user's attentional focus. In the future, we intend to investigate the use of spreading activation during path construction (as well as path evaluation).

7 Conclusion

We have described a mechanism for interpreting expressions of doubt and requests for the consideration of information in the context of arguments generated by a Bayesian argumentation system. The interpretation process, which is performed on a BN that represents a model of a user's beliefs, takes into consideration linguistic clues, the impact of the user's rejoinder on the system's argument, the user's attentional focus, and the system's confidence in the candidate interpretations. Our evaluation suggests that the interpretations generated by BIAS are generally appropriate, and shows how the last two factors can improve the relative rankings of multiple interpretations.

In the future, we intend to perform a full system evaluation where users read arguments, pose rejoinders and receive rebuttals to these rejoinders. In addition, we propose to build upon our results to implement more complex types of rejoinders, such as *what about* questions (e.g., "*What about the murder weapon?*") and explicit inferences (e.g., "*Mr Green being in the garden implies he had opportunity*"), as further stepping stones towards a full argumentation capability.

8 Acknowledgments

This work was supported in part by Australian Research Council grant A49927212. The authors also thank the three anonymous reviewers for the insightful comments.

References

1. J. R. Anderson. *The Architecture of Cognition*. Harvard University Press, Cambridge, Massachusetts, 1983.
2. S. Carberry and L. Lambert. A process model for recognizing communicative acts and modeling negotiation subdialogues. *Computational Linguistics*, 25(1):1–53, 1999.
3. E. Charniak and R.P. Goldman. A Bayesian model of plan recognition. *Artificial Intelligence*, 64(1):50–56, 1993.
4. M. Flowers, R. McGuire, and L. Birnbaum. Adversary arguments and the logic of personal attack. In *Strategies for Natural Language Processing*, pages 275–294. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1982.
5. D. Heckerman and E. Horvitz. Inferring informational goals from free-text queries: A Bayesian approach. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, pages 230–237, Madison, Wisconsin, 1998.
6. N. Jitnah, I. Zukerman, R. McConachy, and S. George. Towards the generation of rebuttals in a Bayesian argumentation system. In *Proceedings of the International Natural Language Generation Conference*, Mitzpe Ramon, Israel, 2000.
7. R. McConachy, K.B. Korb, and I. Zukerman. Deciding what not to say: An attentional-probabilistic approach to argument presentation. In *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*, pages 669–674, Madison, Wisconsin, 1998.
8. J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann Publishers, San Mateo, California, 1988.
9. A. Quilici. Arguing about planning alternatives. In *COLING-92 – Proceedings of the Fourteenth International Conference on Computational Linguistics*, pages 906–910, Nantes, France, 1992.
10. A. Restificar, A. Syed, and S. McRoy. Arguer: Using argument schemas for argument detection and rebuttal in dialogs. In *UM99 – Proceedings of the Seventh International Conference on User Modeling*, pages 315–317, Banff, Canada, 1999.
11. I. Zukerman, R. McConachy, and K.B. Korb. 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, 1998.
12. I. Zukerman, R. McConachy, K.B. Korb, and D.A. Pickett. Exploratory interaction with a Bayesian argumentation system. In *IJCAI99 – Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 1294–1299, Stockholm, Sweden, 1999.