Title
Deciding What Not to Say: An Attentional-Probabilistic Approach to Argument Presentation

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 20(0)

Authors
McConachy, Richard
Korb, Kevin B.
Zukerman, Ingrid

Publication Date
1998

Peer reviewed

Deciding What Not to Say: An Attentional-Probabilistic Approach to Argument Presentation

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

Abstract

Effective arguments must be presented in a cohesive manner: simple collections of believed premises and connecting inferences supporting a goal may not persuade the recipient if they are not well ordered. We use semantic activation and Bayesian propagation in a user model to simulate the effect of presenting an argument generated by our system, NAG, to the user. This simulation is used to select a strategy for presenting the argument to the user. The simulation also identifies superfluous lines of reasoning that may be removed, and enables NAG to determine how multiple subarguments for points should be presented, e.g., as multiple individual supports or collectively. A greedy algorithm is then used to apply probabilistic pruning and semantic suppression to further simplify the argument. Probabilistic pruning removes unnecessary premises from the argument. Semantic suppression is used to select portions of the argument which are within the user's focus of attention, and which are also readily inferred, and hence can be left implicit without damaging the effectiveness of the argument.

Introduction

Effective arguments are rarely simple collections of material supporting a goal. An effective argument must be presented in a cohesive manner, with points following from their predecessors in a sensible order. In addition, obvious points should not be presented, i.e., the argument should be *enthymematic*. NAG (*Nice Argument Generator*), our argument generation-analysis system, attempts to build effective arguments and present them to a user in such a manner. This paper focuses on how NAG achieves its presentation goals.

NAG strikes a balance between arguments that are normatively correct and arguments that are persuasive for a particular user by simulating the effect of its arguments on two models: (1) a normative model, which represents NAG's beliefs, and (2) a user model, which represents a user's presumed beliefs. Each model incorporates a Bayesian network (BN). An argument is represented as an *Argument Graph*, which is a subgraph of relevant elements common to the BN in each model. The nodes in an Argument Graph represent propositions, and their connecting links represent the inferences relating these propositions. The Argument Graph appears in both models so that the effects of presenting the argument can be judged in both its normative and its persuasive aspects. We have chosen BNs as the format in which NAG assembles its arguments, since they support reasoning under uncertainty with multiple (possibly interactive) supporting factors, and can be readily modified to model human cognitive weaknesses (see (Korb et al., 1997)). An Argument Graph starts from *admissible premises* and ends in the goal proposition (the proposition to be argued for). Admissible premises are normatively acceptable propositions (represented in NAG's normative model), which are believed by the user according to NAG's user model or assented to by the user (e.g., drawn from an accepted source). NAG's argument generator builds the Argument Graph, which is considered complete when the anticipated belief in the goal proposition falls within two specified target ranges (or as close to them as possible), one each for the user and for the normative model.[1] These target ranges are given as input to NAG to specify the degree of belief in the goal that is desired for each model after presentation of the argument. Two target ranges are employed, since, for example, it may be sufficient for the user to achieve only a moderate degree of belief in something the normative model shows to be well supported.

NAG's argument generator passes the following information to NAG's presentation module: (1) the goal proposition, (2) the two target ranges of belief to be achieved, and (3) two hierarchical semantic networks (SNs), one built on top of the BN in the normative model and one built on top of the BN in the user model. NAG uses these two hierarchical SNs to consider semantic connections between the items mentioned in the argument. Two semantic networks are necessary, since these connections may differ for NAG and for the user. The presentation module determines a strategy for presenting the argument, organizes the information to be presented in the framework of this strategy, and prunes some superfluous lines of reasoning and what is easily inferred and so may be left implicit.

Related Research

Vreeswijk (1994) describes a system, IACAS, for generating arguments that is designed to be interactive, like NAG. However, IACAS does not attempt to model the user's attentional processes, or tailor the presentation of its arguments to the user. Instead, IACAS shows supporting arguments for the current goal proposition to the user in a sequence until the user is satisfied or chooses a new goal proposition. The chosen sequence of presentation is the order in which IACAS finds its arguments. The argument generation system of Reed and Long (1997), like NAG, takes attention (salience) into consideration when deciding how to present arguments. However, unlike Reed and Long's system, NAG also consid-

[1] Refer to (Zukerman et al., 1998) for a detailed description of NAG's architecture.

ers salience while gathering the information necessary to generate its arguments.

Huang and Fiedler (1997) use a limited implementation of attentional focus to select which step in a proof should be mentioned next. However, unlike NAG, their system does not generate enthymematic arguments. Two other systems that can turn an existing fully explicit argument into an enthymematic one are described in (Horacek, 1994) and (Mehl, 1994). Neither of these two systems can generate an argument from constituent propositions.

Fehrer and Horacek (1997) take advantage of mathematical properties to structure certain types of mathematical proofs. They model a user's inferential ability by means of specialized substitution rules, but offer no mechanism (such as the semantic activation in NAG) to limit the number of applications of their rules. NAG, while developed to be a general argumentation architecture, could be provided with similar heuristics for the restricted domain of mathematical argumentation.

Selecting a Presentation Strategy

The selection of an argumentation strategy determines a partial order of presentation of the argument's propositions. Refinements such as handling subarguments and omitting easily inferred information are discussed later (§ *Refining the Presentation Strategy*). Two basic argument presentation strategies are premise-to-goal and goal-to-premise. A premise-to-goal argument starts from the premises which support intermediate propositions, and eventually reaches the goal. In contrast, goal-to-premise arguments start from the goal[2] and then show how the argument's intermediate propositions and premises lend support to it. A combination of these strategies may in general be used to present different parts of a larger argument. Presently, however, NAG chooses a single argument presentation strategy for the whole argument.

Hypothetical arguments and *reductio ad absurdum* arguments are two applications of the goal-to-premise strategy. Hypothetical arguments are effective in situations where the goal, *were* it true, explains otherwise unlikely premises (which are sufficiently believed). *Reductio ad absurdum* arguments, instead of starting from the goal per se, start from its negation and some admissible premises and reach a contradiction; to resolve the contradiction while preserving the premises, the goal itself is then asserted.

The selection of a strategy to use for the presentation of the argument is made by performing a coarse-grained analysis of the Argument Graph. If NAG, which is designed to be an interactive system, finds itself short on time, then it will not attempt this analysis and will instead fall back to the strategy of building a simple premise-to-goal argument. This coarse-grained analysis examines separately the impact of each individual line of reasoning contributing to the belief in the goal. Sometimes, the impact of certain lines of reasoning cannot be assessed in isolation, since two or more lines may contribute jointly towards the belief in a proposition in a mutually dependent manner. Often however, at least a portion of the contributing lines of reasoning are independent or nearly so, and the coarse analysis can proceed. NAG uses the process outlined below to select a presentation strategy:

Coarse Analysis – Strategy Selection Algorithm

1. For each (approximately) independent line of reasoning $line_i$ with premise $premise_i$, compute the values $P(premise_i)$, $P(premise_i|goal)$ and $P(goal|premise_i)$.
2. If there exists some $line_i$ where $P(premise_i|\neg goal) \approx 1 - P(premise_i)$ and $premise_i$ is correctly and strongly believed by the user, then set the argumentation strategy to be *reductio ad absurdum*.[3] Select that $line_i$ which maximizes the difference between $P(premise_i|\neg goal)$ and $P(premise_i)$. Exit.[4]
3. If there exists some $line_i$ where $P(premise_i|goal) - P(premise_i|\neg goal) \approx 1$ and $premise_i$ is correctly and strongly believed by the user, then set the argumentation strategy to be hypothetical. Go to step 5.
4. Set the argumentation strategy to be premise-to-goal.
5. If the argumentation strategy is premise-to-goal, then rank the set of separable lines of reasoning using $P(goal|premise_i)$ from lowest to highest. Otherwise rank the set using $P(premise_i|goal)$ from lowest to highest.
6. Set $line_j$ to be the first element of the ordered set of separable lines of reasoning.
7. *Initial Pruning*: Tentatively remove the current line of reasoning, $line_j$, from the Argument Graph and check whether the resulting graph is still satisfactory with respect to both the normative model and the user model. If it is not, reinstate $line_j$.
8. If the previous step removed a superfluous line of reasoning, then go to step 7, otherwise exit.

The *reductio ad absurdum* strategy is the first option checked by NAG since this type of argument, when available, is very effective. This is because, once a contradiction between the negated goal and admissible premises is established, any response other than accepting the goal requires a prior belief of the user (an admissible premise) to be retracted, which is unlikely to be the user's preferred response to the argument. This relationship is guaranteed by the probabilistic requirement of step 2 above, that $P(premise_i)$ be high and so $P(premise_i|\neg goal)$ be low.

Should a *reductio ad absurdum* strategy not be applicable, NAG next looks at the possibility of presenting a hypothetical argument. Hypothetical arguments are often successful, so we implement them in NAG. If the Argument Graph contains lines of reasoning that allow a hypothesized belief in the goal proposition to explain some otherwise unlikely premises, then such an appeal to the user's (correct) belief in those premises will be tried. In the fictional example of Figure 1, node N_6, which represents documentation from trade

[2]Or, from the negation of the goal. NAG can, of course, argue either for or against a goal proposition. Here we assume a positive bias; if a goal is to be argued against, some expressions below will need to be altered accordingly.

[3]Here, and in the next step, thresholds are used to determine the degree of approximation to equality the criteria must satisfy. The selection of optimal values for these thresholds is a current topic of investigation.

[4]If the resulting *reductio* fails to achieve NAG's argumentative goals, this coarse analysis will be redone, but step 2 will be skipped.

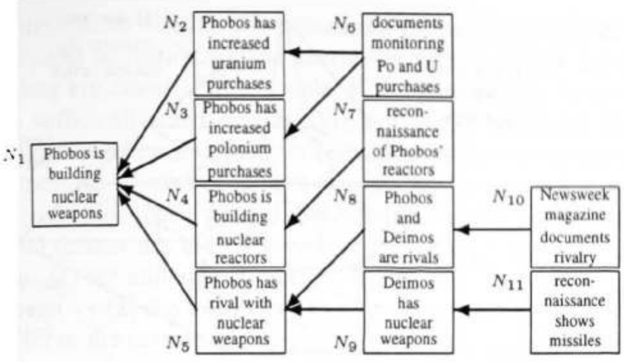

Figure 1: Argument Graph for the Phobos example

monitoring organizations showing Phobos' increased polonium (Po) and uranium (U) purchases, is believed to be true. Polonium is used in only a small number of processes, the production of initiators for nuclear weapons being amongst them. Hence, the likelihood of a nation having documented increases in its intake of polonium, *P(premise)*, is small in general, but the likelihood that a nation building a stockpile of nuclear weapons would do so, $P(premise|goal)$, is large.[5] It follows, then, that $P(premise|\neg goal)$ is small, and so the formal conditions of step 3 for NAG to adopt a hypothetical form of argument are satisfied.

Lines of reasoning that make only a small contribution to the belief in the goal in both the normative model and the user model will be dropped if NAG shows that after their deletion the belief in the conclusion will still be inside the specified normative and user model target ranges.[6] Step 7 of the algorithm above performs an initial round of pruning which removes such minor contributors to the Argument Graph. This initial pruning simplifies the resulting argument and accelerates the process of refining the chosen presentation strategy described in the next section. This simplifying process is currently limited to lines of reasoning that operate largely independently of other lines of reasoning; in the future, it may be extended to remove more complex branches of argument.

Refining the Presentation Strategy

Having determined an overall presentation strategy for the argument, NAG addresses three more problems: (1) dealing with subarguments; (2) maximizing the semantic cohesiveness of the argument by ordering the propositions in the optimal manner consistent with the chosen global presentation strategy; and (3) further pruning of the Argument Graph. We discuss each of these problems in turn.

[5]This assumes that the trade monitoring organizations are performing their job properly.

[6]It is possible that the initial Argument Graph passed to the presentation module is not anticipated to generate a belief in the goal proposition falling within these target ranges of belief. (This is a kind of failure of the argument generation process, which may occur for a variety of reasons, including running out of time; cf. (Zukerman et al., 1998).) For example, the anticipated effect of presenting the argument may be too strong and exceed the upper boundary of one of the target ranges. In this case, lines of reasoning may be deleted if their deletion brings the Argument Graph closer to conformance with the target ranges.

Presenting Subarguments

Most arguments will be complex, i.e., they will be composed of subarguments supporting chains of intermediate conclusions. For example, node N_5 in Figure 1 is an intermediate conclusion supported by nodes N_8 and N_9. How should the support these two nodes offer N_5 be presented? Currently, we consider two policies for presenting subarguments: (1) *collective* and (2) *individual-sequential*, described below.

Collective Policy. In this policy, all the direct antecedents of a conclusion are mentioned immediately prior to mentioning the conclusion itself. For example, "Given N_8 and N_9, N_5 follows." (The sequencing of these premises is determined by attentional ordering; see § *Attentional Processing* below.) This policy is used when: (1) all the antecedents provide similar levels of support for the consequent; or (2) the antecedents are *not* conditionally independent. By grouping these antecedents together, NAG is reflecting the fact that they are of roughly equal import or that their import can only be judged jointly.

Individual-sequential. In this policy, the effect of one antecedent on the consequent is mentioned separately from the other antecedents. For example, the partial argument, "Given N_5, it is possible that N_1 is true. In addition, N_3 offers strong independent support for N_1." This policy may be used when one antecedent offers support that is conditionally independent from the other premises. The more supportive line of argument is mentioned last, since we expect that what is mentioned last will have greater impact on the audience (this is suggested by the anti-primacy effect in argumentation; see (Bailenson, 1997)). Hence, in this example, the stronger line of reasoning offered by Phobos' increased purchases of polonium is presented after pointing out that Phobos' rival Deimos has nuclear weapons.

These two presentation policies may be combined. For example, when some antecedents are presented using the individual-sequential policy, the remaining antecedents may be presented using the collective policy, e.g., "Given N_2, N_4 and N_5, it is possible that N_1 is true. In addition, N_3 offers strong independent support for N_1." Here N_2 and N_4 are mentioned together since they are conditionally dependent on each other (reactors require fuel such as processed uranium to run). The node N_5 is mentioned in conjunction with N_2 and N_4 since it offers a similar level of support for the goal, N_1. Finally, N_3 is separate and last because it offers conditionally independent and strong support for the goal.

In addition to this refinement of presentation order, the collective versus individual-sequential policies may (in the future) be used to select natural language expressions that signify relations between propositions, such as "together" and "independently" above.

Attentional Processing

Attentional processing is used during presentation planning for two purposes: (1) to maximize the semantic cohesiveness of the presentation order of the argument; and (2) to decide whether a proposition in the Argument Graph can be left implicit when NAG presents the argument.

Where a strategy or policy does not dictate a complete ordering, presentation order is based upon the activation levels of the propositions. Usually propositions within an argument are related semantically as well as probabilistically, so the mention of a premise may bring immediate intermediate conclusions into the user's focus of attention, even before any inferential process comes into play. In addition, the presentation of one line of reasoning for some (possibly intermediate) conclusion may bring the premises forming the beginning of a second line of argument into the focus of attention. For example, for many people the presentation of $N_6 \rightarrow N_2$, about uranium purchases, is semantically connected with the presentation of $N_7 \rightarrow N_4$, about nuclear reactors. That is, presenting information about uranium semantically primes people to think about nuclear devices, such as reactors. Our use of activation to further specify the presentation order relies upon the fact that high activation is self-sustaining, whereas low activation needs reinforcement immediately prior to any relevant (intermediate) conclusion. So, keeping the presentation linked to the user's likely focus of attention allows us to omit some references to previously presented subarguments and to improve the flow of the argument.

NAG uses two hierarchical SNs to capture connections between the items mentioned in the argument. Figure 2 illustrates one such semantic-Bayesian 'pyramid'. The SN (upper levels of the pyramid) and the BN (base of the pyramid) are used by NAG to simulate attentional focus in each model during argument generation as described in (Zukerman et al., 1998). When checking the semantic connectivity of a proposed presentation ordering for an Argument Graph, the propositions composing the argument are activated in the user model in the order in which they will be presented to the user. If the ordering is not completely defined by the presentation strategy and policies adopted to present subarguments, then all the possible orderings consistent with these restrictions are tested. We use activation with decay, similar to that described in (Anderson, 1983), spreading from salient objects, to model the focus of attention. This process passes activation through the Argument Graph and the encompassing pyramidal semantic-Bayesian network. During this process each node in the pyramid is activated to the degree implied by the activation levels of its neighbors, the strength of association to those neighbors, and its immediately prior activation level (vitiated by a time-decay factor). For example, in the situation described in the previous paragraph, the presentation of node N_6 activates N_2, which in turn activates N_7, making the beginning of the second line of reasoning semantically connected to the first. Other considerations being equal, NAG selects the most strongly semantically connected presentation order, i.e., the order maintaining the strongest activation of propositions as they are about to be mentioned.

As indicated above, NAG uses attentional processing for a second purpose, namely, to support the generation of enthymematic arguments. Using the semantic activation process described above, NAG can anticipate which propositions enter the user's span of attention as a result of the presentation of other material. NAG utilizes this information to decide when propositions are sufficiently activated that they may be left implicit in the presentation of the argument, using "semantic suppression" (described below).

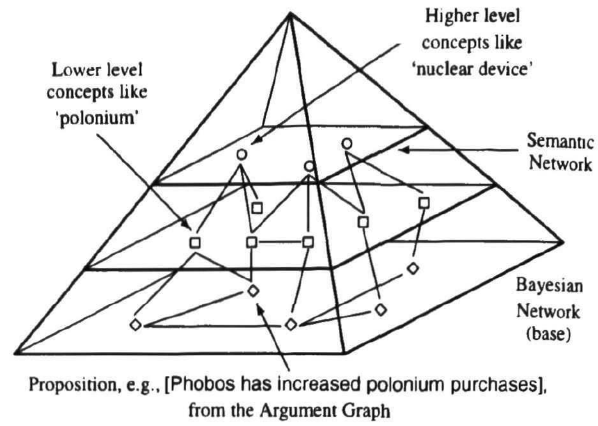

Figure 2: Semantic and Bayesian Networks

Additional Pruning of the Argument Graph

In subsequent passes through the Argument Graph, the presentation module attempts to further simplify the final argument by invoking probabilistic pruning and semantic suppression as appropriate. Both pruning techniques are abandoned after any complete pass fails to effect a change to the Argument Graph (or when time runs out). During this additional pruning phase, NAG continues to check whether the belief in the goal proposition in the reduced Argument Graph remains within the target ranges. The simulation of the user's attentional state described above is performed after each proposed pruning operation to determine whether the propositions in the argument remain in focus when they are needed. The multi-pass greedy algorithm outlined below implements this process (details are provided in the subsections following).

Pruning Algorithm

1. Take the presentation ordering $< N_1, \ldots, N_k >$ of propositions, let i index the current proposition. Set $i \leftarrow 0$.
2. $i \leftarrow i + 1$.
3. *Pruning the Argument Graph.*
   If N_i is a premise:
   (a) Invoke probabilistic pruning.
   (b) If N_i is retained, activate it for attentional processing.
   Otherwise, N_i is an intermediate proposition so:
   (c) Simulate the user's attentional state (§ *Attentional Processing*).
   (d) Invoke semantic suppression to determine whether N_i may be left implicit.
4. If this is the end of a complete pass through the ordering in which no change has been made (or if time has run out), then exit. Otherwise, if this is the end a complete pass, set $i \leftarrow 0$. Go to step 2.

Probabilistic pruning removes two kinds of unnecessary propositions: (1) those that alter belief in the argument goal to only a small degree, so removing them does not endanger the targets for belief; and (2) those forming a line of reasoning that takes the goal outside the target belief ranges.

After each deletion, NAG checks whether the anticipated belief in the goal proposition is still within the target ranges. In addition, affected subgraphs in the Argument Graph are reordered according to the policies mentioned in § *Attentional Processing* and the user's attentional focus is again simulated to determine whether the remaining propositions still have sufficient activation for the user to be able to follow the resulting argument. For example, if the user already believes N_8 sufficiently, and leaving N_{10} out of the argument does not badly affect the semantic activation of the remaining argument (in which node N_8 is now a premise), then N_{10} will be probabilistically pruned. Probabilistic pruning fails when NAG determines after a proposed removal of a proposition that: (1) the anticipated belief in the goal is outside a target range,[7] or (2) the level of activation of a subsequent proposition in the revised argument has fallen below a threshold. The maximum allowable decrease in activation is governed by a predetermined threshold. In either of these cases, the last removed proposition is reinstated, and the pass through the current ordering continues (via step 3 above).

Semantic suppression checks intermediate conclusions in the Argument Graph to see if they can be left implicit, rather than being explicitly presented to the user. To be left implicit a proposition must be *easily inferred*. Propositions are easily inferred when (1) they result from probabilistically strong inferences, and (2) their antecedents are highly activated in the user model during the simulated presentation.

If (1) is true — that is, if an intermediate conclusion is potentially greatly strengthened by the inferences connecting it to its immediate antecedents, then, once the values of those antecedents are known, the value of the intermediate conclusion itself is very clear. NAG uses partial propagation to check the strength of the inferences connecting intermediate conclusions to their antecedents. To do this, the Bayesian subgraph comprising the proposition NAG seeks to make implicit together with its immediate antecedents is copied from the Argument Graph for testing. In this copy of the subgraph, the antecedents retain their current probability values and the single consequent proposition is set to a neutral value (i.e., 0.5, since all variables in NAG are propositional). NAG then applies the modified Bayesian belief update rules to this copy of the subgraph of the user model to determine the consequent's posterior probability.[8] Of course, this calculation is performed in the user model only, since it is the user who must follow the implicit reasoning. If the posterior probability of the consequent is sufficiently raised over its prior neutral value ($posterior - 0.5 > threshold$) then NAG accepts that the inference is probabilistically strong enough for semantic suppression.

Given that, the second requirement (that the antecedents of the proposition being considered for semantic suppression be highly activated) is checked using the attentional mechanism described in § *Attentional Processing*. A presentation of the argument represented by the Argument Graph is simulated in the user model, and if the activation levels of the antecedents exceed a threshold value, the proposition in question will be left implicit.

For an example of semantic suppression, the prior probability of the intermediate conclusion N_4, [Phobos is building nuclear reactors], is set to 0.5, and the modified Bayesian belief update rules are applied. The resulting posterior probability in the intermediate conclusion is very high, since here the inference connecting the propositions N_7 and N_4 is just "If you have reconnaissance reports of Phobos having nuclear reactors, then Phobos is building nuclear reactors," which is a strong inference assuming a trustworthy intelligence source, as we did. Given the large posterior value, and the high semantic activation, NAG leaves node N_4 implicit and thus produces a simpler argument that is as persuasive as the original.

Note that only the subgraph leading to the intermediate conclusion is copied; inferences connecting the intermediate conclusion to subsequent conclusions, e.g., $N_4 \rightarrow N_1$, are not examined. The user either will or will not follow the implicit reasoning to the omitted proposition. If the user does follow the reasoning, then, since the omitted proposition is known to the user (according to the user model), s/he will be able to follow the inferences that lead from this intermediate conclusion towards the global argument goal. In this case, the implicit proposition will continue to contribute its probabilistic support to the conclusions built upon it. If the user does not follow the implicit reasoning, then checking any further consequences of the omission is immaterial.

Semantic suppression can be incrementally applied to adjacent propositions in the Argument Graph, possibly creating larger enthymematic gaps in the presented argument. However, since an implicit proposition is unmentioned, it does not achieve the highest level of activation, so the chance of successive suppressions is not large. Semantic suppression fails when omitting a proposition drives the level of semantic activation of any required subsequent proposition below a threshold.[9] In this case, the last removed proposition is reinstated and the pass through the current ordering continues, unless a halt condition is reached, when the Argument Graph is finally presented to the user.

Results and Discussion

The example used throughout this paper can reasonably be presented using two of the strategies implemented in NAG. If the probability $P(N_6|N_1) - P(N_6|\neg N_1)$ exceeds the threshold used in step 3 of the Strategy Selection Algorithm, then a hypothetical argument will be presented to the user, who is known to believe node N_6 and to believe strongly in the links connecting $N_1 \rightarrow N_3 \rightarrow N_6$. The remaining portions of the Argument Graph can be pruned away (N_2, N_4 and N_5 and their supporting lines of reasoning), or suppressed (N_3). The resulting argument may be rendered as: *"If Phobos was building nuclear weapons, that would explain otherwise unlikely trade monitoring reports documenting an increase in Phobos' polonium and uranium purchases."*[10] If the user model does not support a hypothetical argument, as judged in step 3, then NAG will adopt a premise-to-goal strategy

[7] Or, in the event that NAG is working with a defective argument such that the goal starts out outside the target(s), pruning fails if the result is to drive the goal further from the target(s).

[8] The propagation rules used to update beliefs when checking the Argument Graph in the user model are modified to model three cognitive weaknesses commonly observed in human subjects. The three weaknesses modeled are belief bias, overconfidence and the base rate fallacy. See (Korb et al., 1997) for details.

[9] Note that this is distinct from checking for subsequent probabilistic effects, which NAG, as mentioned above, does not do.

[10] The English arguments are hand-generated from NAG's final, ordered Argument Graph, since NAG does not (yet) generate natural language.

for this argument. This also suggests that additional support for the conclusion may be needed. In such a case the final presentation ordering of the revised Argument Graph may be: $\{N_8, N_9\} \rightarrow N_1$, $N_6 \rightarrow N_1$. Nodes N_8 and N_9 offer collective support for the goal node, N_1 (with N_5 being left implicit after semantic suppression). In addition, the stronger line of reasoning from N_6 to N_1 is then presented individually (with N_3 semantically suppressed). The lines of reasoning $N_6 \rightarrow N_2 \rightarrow N_1$ and $N_7 \rightarrow N_4 \rightarrow N_1$ were removed during coarse analysis, while nodes N_{10} and N_{11} were removed later by probabilistic pruning. The output in such a case might be: *"Phobos and Deimos are rivals and Deimos has nuclear weapons, suggesting that it is possible that Phobos is building nuclear weapons. Furthermore, trade monitoring reports documenting Phobos' increased purchases of polonium and uranium strongly suggest that Phobos is building nuclear weapons."*

Evaluation

Thus far, our evaluation of NAG has been informal and incomplete. Informally, we can report that NAG has been applied to some dozen argument generation problems from different domains — successfully by the light of human intuition. In this regard, we should point out that the argument generation *methods* we have described in this paper (and elsewhere) have worked across all of the test domains. Application to new problems does involve a fair amount of human intervention, specifically in building the semantic and Bayesian networks characterizing the domain and the target user. However, once those are built, NAG's mechanisms do not in general require human intervention. The automated construction of semantic and Bayesian networks awaits further successes from the machine learning community.

We have conducted a *preliminary* Web-based formal evaluation of NAG's argument generation and the pruning of lines of reasoning deemed superfluous. Pre-test and post-test questionnaires showed a clear tendency among the respondents to shift belief towards the argumentative goal as a result of reading the arguments presented. This suggests that NAG's pruning of some lines of reasoning does not greatly damage the effectiveness of its arguments. A much more rigorous evaluation of the argument generation and presentation techniques presented in this paper, including the selection of a presentation strategy, the different policies for presenting sub-arguments and semantic suppression, will be performed in the near future. This evaluation will also incorporate comparative testing with human-generated arguments, in a kind of Turing test of the adequacy of NAG's argumentation.

Conclusion

NAG uses semantic activation in a user model to simulate the user's attentional processing during argument presentation. A restricted and modified form of Bayesian propagation is used to check the probabilistic strength of user inferences. The strength of the different lines of reasoning supporting an argument are used to select a presentation strategy and to quickly remove superfluous items. A greedy algorithm applies probabilistic pruning and semantic suppression in order to simplify the argument, while retaining its effectiveness. Probabilistic pruning removes premises which are not necessary for the user to achieve the desired degree of belief. Semantic suppression identifies intermediate conclusions which are easily inferred, and hence can be left implicit. Semantic activation is also used to order the remaining propositions in the argument so that the argument proceeds smoothly, minimizing disruptive jumps in the user's simulated attentional state.

Acknowledgements

This work was supported in part by Australian Research Council grant A49531227.

References

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

Bailenson, J. (1997). Claim strength and burden of proof in interactive arguments. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, pages 13–18, Stanford, California.

Fehrer, D. and Horacek, H. (1997). Exploiting the addressee's inferential capabilities in presenting mathematical proofs. In *IJCAI97 – Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pages 959–964, Nagoya, Japan.

Horacek, H. (1994). How to avoid explaining obvious things (without omitting central information). In *ECAI94 – Proceedings of the Eleventh European Conference on Artificial Intelligence*, pages 520–524, Amsterdam, The Netherlands.

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.

Korb, K., McConachy, R., and Zukerman, I. (1997). A cognitive model of argumentation. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, pages 400–405, Stanford, California.

Mehl, S. (1994). Forward inferences in text generation. In *ECAI94 – Proceedings of the Eleventh European Conference on Artificial Intelligence*, pages 525–529, Amsterdam, The Netherlands.

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, G. (1994). IACAS: an interactive argumentation system. User manual version 1.0. Technical Report CS 94-03, Dept. of Computer Science, University of Limburg.

Zukerman, I., McConachy, R., and Korb, K. B. (1998). Bayesian reasoning in an abductive mechanism for argument generation and analysis. In *AAAI-98 – Proceedings of the Fifteenth National Conference on Artificial Intelligence*, Madison, Wisconsin.