ATTENTION DURING ARGUMENT GENERATION AND PRESENTATION*

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

Abstract

We describe the operation of our argumentation system, and discuss its use of attentional focus during both content planning and argument presentation. During content planning, attentional focus guides an abductive process used to build-up arguments. This process is applied to a model of a user's beliefs and a normative model. During argument presentation, attentional focus supports the generation of enthymematic arguments.

1 Introduction

In this paper, we describe the operation of our argument-generation system, *NAG* (Nice Argument Generator). We consider its content planning and argument presentation processes, and discuss the attentional mechanism used in both of these processes.

Given a goal proposition, NAG's objective is to generate a *nice* argument for it, by which we mean one which achieves a balance between what is normatively justifiable and what persuades the interlocutor. To this end, NAG consults a normative model, which contains our best understanding of the domain of discourse, and a model of the user's beliefs. The main modules of the system are shown in Figure 1.

The Strategist drives the argumentation process. During argument generation, it activates a generation-analysis cycle as follows (Section 3). First, it invokes the Attentional Mechanism (Section 4) to activate salient propositions, which are used to construct an initial *Argument Graph* for an argument, or to extend an already existing Argument Graph. (An Argument Graph is a network with nodes that represent propositions, and links that represent the inferences that connect these propositions.) The Strategist then calls the Generator to continue the argument building process (Section 5). The Generator in turn fleshes out the Argument Graph by activating *Reasoning Agents*, which consult several sources of information, and incorporating the inferences and propositions returned by these agents into the Argument Graph. This Argument Graph is returned to the Strategist, which passes it to the Analyzer in order to evaluate its niceness and check for reasoning flaws (Section 6). If the Analyzer indicates that the argument is not nice enough, i.e., there is not sufficient belief in the goal in the user or the normative model, then the Strategist re-activates the Generator in order to find further support for the premises in the argument, and so on. The generation-analysis cycle continues until a sufficiently nice Argument Graph is generated. This graph is then passed to the Argument Presenter, which selects an argumentation strategy and determines propositions to be removed from the argument, aiming to produce a simpler, enthymematic argument. After each removal, the Presenter activates the Analyzer to check whether the argument is still nice and the Attentional Mechanism to determine whether the argument can still be followed by the user.

Thus, the Attentional Mechanism is used in two different stages of the argumentation process. During argument generation, it focuses the argument construction process on concepts which are related to the

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

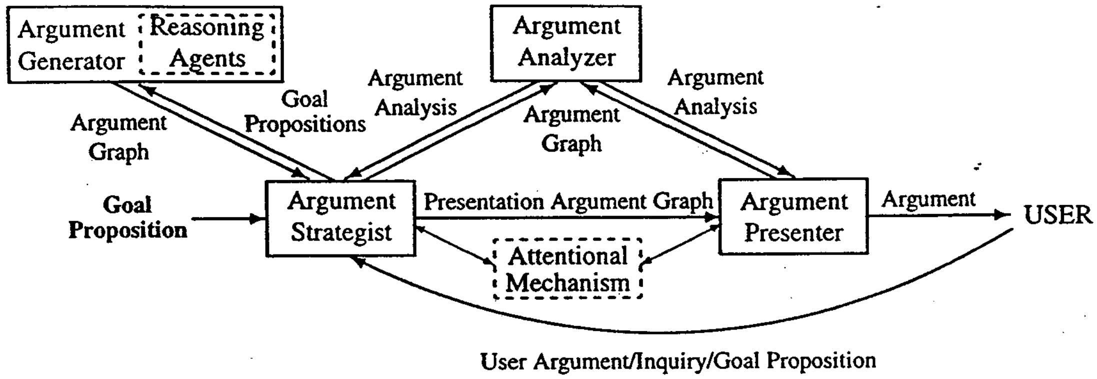

Figure 1: System Architecture

goal, avoiding some distractions. During argument presentation, the Attentional Mechanism supports the generation of enthymematic arguments.

2 Related Research

Charniak and Goldman (1993) describe a Bayesian plan recognition system that uses marker passing as a method for focusing attention on a manageable portion of the space of all possible plans. This is analogous to the way in which NAG uses spreading activation to focus on a small portion of the available data during the content planning process. Walker (1996) points out the effect of attentional focus on discourse comprehension. This effect is taken into consideration by NAG during argument presentation.

The approach of "interpretation as abduction" used in [Hobbs *et al.*, 1993] aims to recover the premises and inferential links which lead to an argument's conclusion. This is similar to NAG's argument analysis. The most important difference between NAG and the work by Hobbs *et al.*, in addition to NAG being a system that reasons under uncertainty, is that NAG performs both analysis and generation. A generative system based on the work of Hobbs *et al.* is described in [Thomason *et al.*, 1996]. This system deals with what can be readily inferred, and so deleted, during communication, but the generated discourse does not argue in favour of a proposition. Mehl (1994) describes a system which can turn an existing fully explicit argument into an enthymematic one, but it cannot generate an argument from constituent propositions. The system described in [Horacek, 1997] generates its own arguments and presents them enthymematically. However, neither process models explicitly a user's attentional state.

Like NAG, the systems described in [Reed and Long, 1997, Huang and Fiedler, 1997] consider focus of attention during argument presentation. NAG differs from these systems in that NAG also uses attentional focus to guide the content planning process. In addition, Reed and Long consider attention in order to generate additional information that makes a concept salient, and Huang and Fiedler use a limited implementation of attentional focus to select which step in a proof should be mentioned next. In contrast, NAG uses attentional focus during argument presentation to convert a fully explicit argument into an enthymematic one. Finally, 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 attention in NAG) to limit the number of applications of their rules.

3 The Generation-Analysis Cycle

NAG receives the following inputs: (1) a proposition to be argued for, (2) an initial argument context, and (3) two target ranges of degrees of belief to be achieved (one each for the normative model and the user model). The argument context is composed of salient propositions and concepts appearing in the discussion preceding the argument or in the current Argument Graph. The initial argument context, $context_0$, is composed of salient propositions and concepts mentioned in the preamble to the argument plus the argument's goal. The degrees of belief to be achieved are expressed as ranges of probabilities, e.g., [0.5, 0.6], in order to be able to represent a variety of goals, e.g., inducing indifference or assent. Two target ranges are required since the degree of belief to be reached by the user may differ from that reached by the system.

When constructing an argument, NAG relies on two collections of information: a normative model composed of different types of Knowledge Bases (KBs), 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), rule-based system, or database. The KBs in the normative and user models are consulted by specialist Reasoning Agents which are activated when trying to fill gaps in a partial argument (Section 5). The KBs in the user model are consulted to make an argument persuasive for the target audience, while the normative KBs are consulted so that the generated argument is normatively correct. This distinction between the normative correctness and the persuasiveness of an argument enables the system to control the extent to which it will sacrifice normative correctness in order to be persuasive. During argument generation, relevant material from several KBs may need to be combined into a common representation. We have chosen BNs for this purpose because of their ability to represent normatively correct reasoning under uncertainty, and because simple alterations of the normal Bayesian propagation rules allow us to model various human cognitive phenomena [Korb *et al.*, 1997].

The content planning process produces an Argument Graph which starts from *admissible premises* and ends in the goal proposition. Admissible premises are normatively acceptable propositions that are believed by NAG and are either believed the user (sufficiently for the argument to work) or assented to by the user. The resultant Argument Graph is then passed to the Presenter. The argument generation process is implemented by the following algorithm, which is executed by the Strategist.

Generation-Analysis Algorithm

1. $i \leftarrow 0$.
2. Clamp nodes in the current context, $context_i$, and perform spreading activation. This yields an Argument Graph containing: the clamped nodes, the activated nodes (whose activation exceeds a threshold), and the links connecting these nodes. **(Section 4)**
3. Identify new subgoals in the current Argument Graph. These are nodes which have not been previously passed to the Reasoning Agents, and have a path to the goal in the Argument Graph or a high level of activation (higher than a subgoaling threshold).
4. Pass the argument subgoals identified in Step 3 to the Generator, which adds the new information returned by its Reasoning Agents to the current Argument Graph. **(Section 5)**
5. Pass the Argument Graph generated in Step 4 to the Analyzer for evaluation. **(Section 6)**
6. If the Analyzer reports that the current Argument Graph is sufficiently nice, then pass the Argument Graph to the Presenter. Otherwise, continue. **(Section 7)**
7. $i \leftarrow i + 1$.
8. $context_i \leftarrow context_{i-1}$ + new nodes connected to the goal during cycle i–1 +
   highly activated nodes that are not connected to the goal.
9. Go to Step 2.

4 Simulating Attention

NAG uses a hierarchical semantic network built on top of BNs in both the user and normative models to capture connections between the items mentioned in the discourse (Figure 2 illustrates a three-level semantic-Bayesian network). The semantic network portion (upper levels of the pyramid) and the BN portion (base of the pyramid) are used by NAG to simulate attentional focus in each model.

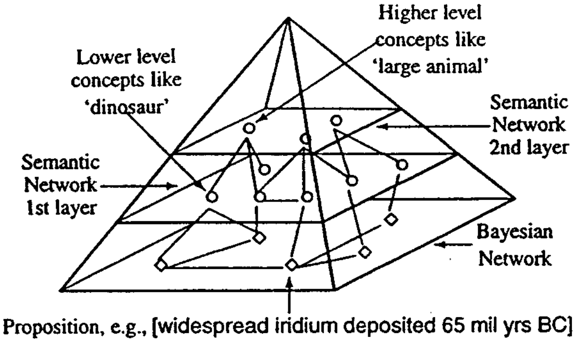

Figure 2: Semantic and Bayesian Networks

The Attentional Mechanism receives as input a context consisting of a set of salient objects. For content planning, the initial context consists of the goal proposition and salient concepts and propositions mentioned in the preamble of the argument; as the content planning process progresses, the context is extended with concepts and propositions included in the argument. For argument presentation, the context initially contains the propositions of the first sub-argument to be presented under a particular presentation strategy, and moves on to different propositions as the presentation of the argument progresses (Section 7).

We use activation with decay [Anderson, 1983], spreading from the current context, to model the focus of attention. All items in the semantic-Bayesian networks which achieve a threshold activation level during the spreading activation process are brought into the current span of attention. This process passes activation through the pyramidal 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). The spreading activation process ceases when an activation cycle fails to activate any new node. By these means we have a direct implementation of attention which we use to identify portions of the pyramidal semantic-Bayesian networks that are strongly related to the argument being built or presented.

During content planning, spreading activation is applied to both the semantic-Bayesian pyramid in the user model and that in the normative model. This supports the retrieval of information that is semantically connected, and hence likely to be useful for building an argument, from each of the models. Further, the items in the initial context and in subsequent, extended contexts are clamped, since they are used for reasoning, and their level of activation should not fade as the system reasons about the argument. In contrast, during argument presentation, spreading activation is performed only in the semantic-Bayesian pyramid in the user model, and salient information items are not clamped. This is because we are trying to anticipate the effect of the information being presented on the addressee, and information presented earlier will fade from the addressee's focus of attention [Walker, 1996].

The Attentional Mechanism offers the following advantages. During content planning, it allows NAG to restrict its search to information semantically or evidentially connected with the propositions already in focus, and it enables NAG to analyze its arguments with respect to just these propositions (Section 6). During argument presentation, it supports the generation of enthymematic arguments by assessing the omission of propositions that are in the addressee's span of attention (as a result of the presentation of other propositions).

5 Argument Extension

The Generator activates the Reasoning Agents to collect information relevant to each of the subgoals in the current Argument Graph (these subgoals were identified in Step 3 of the Generation-Analysis algorithm). The Reasoning Agents determine the relevance of a piece of information to a goal proposition by applying a

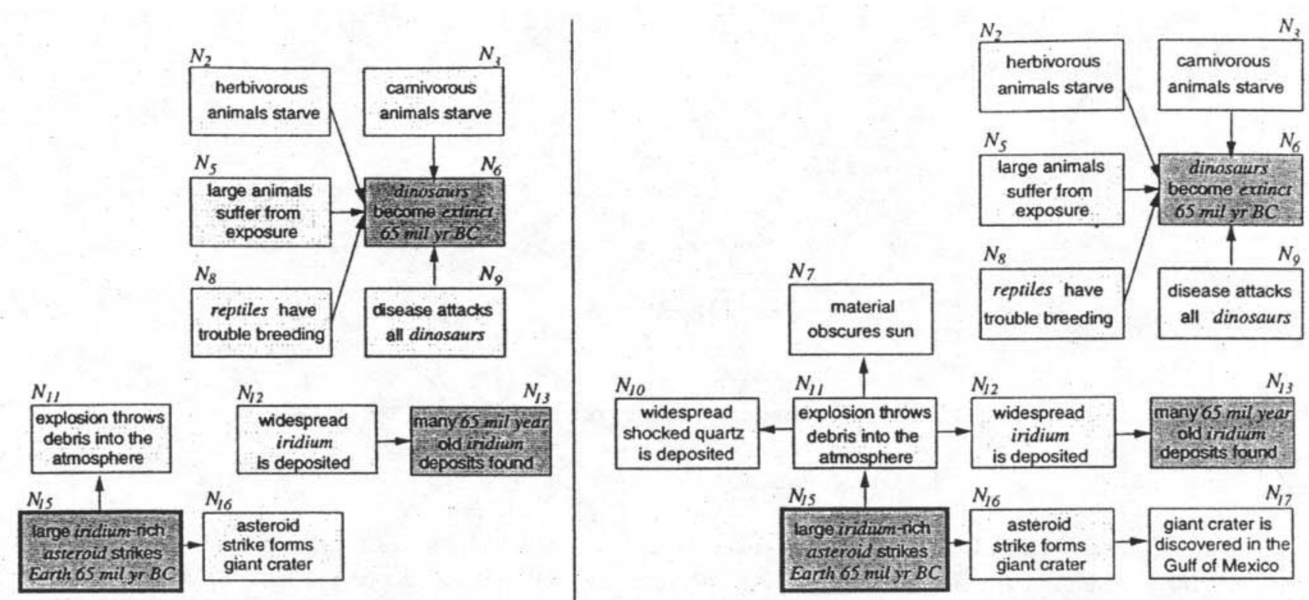

(a) Initial Argument Graph (b) Argument Graph after one cycle

Figure 3: Argument Graphs for the Asteroid Example during Content Planning

procedure that is appropriate for the KB under consideration. For instance, if the goal proposition matches the consequent of a rule in a rule-based system, the Reasoning Agent for rule-based systems selects this rule. The relationships found by the Reasoning Agents represent inferences. The strength of each inference is represented by means of conditional probabilities which reflect the degree of belief in the consequent of an inference given full belief in its antecedents.

Example

Throughout this paper, we consider the generation of an argument for the goal proposition "A large *iridium-rich asteroid* struck *Earth* about *65-million-years BC*," preceded by the preamble "Approximately *65-million-years BC* the *dinosaurs*, large *reptiles* that dominated the *Earth* for many millions of years, became *extinct*". Initially, the goal proposition and the preamble activate any propositions containing two or more of the italicized concepts, i.e., nodes N_6, N_{13} and N_{15} (the goal node) in Figure 3(a) (shown in dark grey boxes).

After clamping the nodes that correspond to this discourse context and performing spreading activation, additional nodes become activated in the semantic and Bayesian networks. All the nodes whose activation level exceeds a threshold are added to the Argument Graph. In this example, this yields the nodes shown in light grey boxes in Figure 3(a). The links between the nodes in Figure 3(a) were obtained from the BN, but the activation of these nodes involved spreading activation through both the BN and the SN.

Since none of the nodes in the current Argument Graph (Figure 3(a)) have been passed to the Reasoning Agents, the following nodes are passed to these agents (through the Generator): those in the subgraph containing the goal node (N_{11}, N_{15} and N_{16}), plus the two clamped (highly active) nodes in the graph fragments not connected to the goal node (N_6 and N_{13}). The information returned by the Reasoning Agents, which is either causally or evidentially connected to the nodes passed to the Generator, is then incorporated into the Argument Graph (Figure 3(b)). Some of the nodes found by these agents have already been activated through spreading activation (shown in light grey in Figures 3(a) and 3(b)), while others are new to the Argument Graph (shown in white in Figure 3(b)). In addition, the Reasoning Agents added a new link between the previously activated nodes $N_{11} \rightarrow N_{12}$.

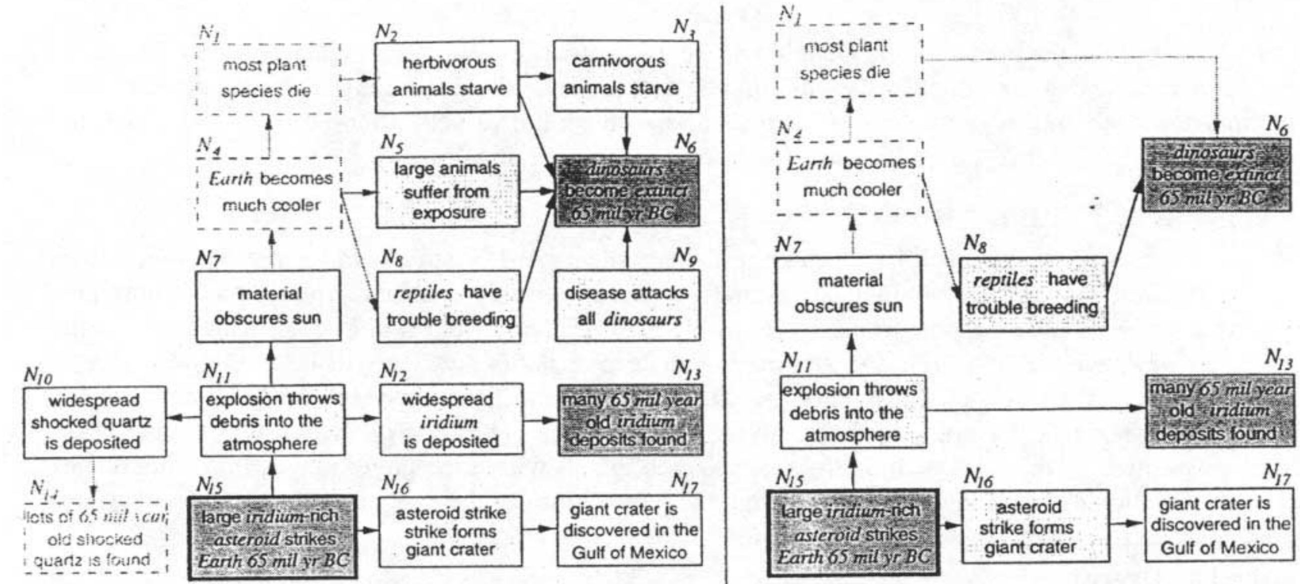

(a) Argument Graph passed to the Presenter

(b) Argument Graph after pruning

Figure 4: Argument Graphs for the Asteroid Example during Presentation

6 Argument Analysis

The process of computing the anticipated belief in a goal proposition as a result of presenting an argument starts with the belief in the premises of the Argument Graph and ends with a new degree of belief in the goal proposition. The Analyzer computes the new belief in a proposition by combining the previous belief in it with the result of applying the inferences which precede this proposition in the Argument Graph. This belief computation process is performed by applying Bayesian propagation procedures to the Bayesian subnetwork corresponding to the current Argument Graph in the user model and separately to the subnetwork corresponding to the current Argument Graph in the normative model.

After propagation, the Analyzer returns the following measures for an argument: its *normative strength*, which is its effect on the belief in the goal proposition in the normative model, and its *effectiveness*, which is its effect on the user's belief in the goal proposition (estimated according to the user model). Of course, an argument's effectiveness may be quite different from its normative strength. When anticipating an argument's effect upon a user, NAG takes into account three cognitive errors that humans frequently succumb to: belief bias, overconfidence and the base rate fallacy [Korb *et al.*, 1997].

If the normative strength or effectiveness of the Argument Graph is insufficient, another cycle of the Generation-Analysis algorithm is executed, gathering further support for propositions which have a path to the goal or have a high activation level (Step 3 of the Generation-Analysis algorithm). In this manner, NAG combines goal-based content planning with the associative inspection of highly active nodes. After integrating the new sub-arguments into the Argument Graph (Step 4), the now enlarged Argument Graph is again sent to the Analyzer (Step 5). Hence, by completing additional focusing-generation-analysis cycles, Argument Graphs that are initially unsatisfactory are often improved.

Example – Continued

The argument that can be built at this stage consists of nodes N_7, N_{10}–N_{13} and N_{15}–N_{17}. However, only N_{13} is admissible among the potential premise nodes. Thus, the anticipated belief in the goal node in both the normative and the user model falls short of the desired ranges. This is reported by the Analyzer to the Strategist. Nodes N_7, N_{10}–N_{12}, N_{16} and N_{17} are now added to the context (which initially included N_6,

N_{13} and N_{15}), and the next cycle of the Generation-Analysis algorithm starts. This process continues until the Analyzer reports that the belief in the goal proposition is inside the target ranges in both the user and the normative model. In this example, this happens after two additional generation-analysis cycles, which activate nodes N_1, N_4 and N_{14} among others (in dashed boxes in Figure 4(a)).

7 Argument Presentation

After a successful Argument Graph has been built, the argument must be structured for presentation to the user. This involves selecting an argumentation strategy and pruning unnecessary propositions. These are propositions that lend little support to the belief in the goal or support the goal beyond what is required (removed by *probabilistic pruning*), and also intermediate propositions which will be in the addressee's focus of attention as a result of information presented earlier, and hence may be omitted (removed by *semantic suppression*). After pruning, the Analyzer (Section 6) is invoked to check whether the belief in the goal proposition in the now smaller Argument Graph is still within the target ranges; the Attentional Mechanism (Section 4) is invoked to check whether the propositions in the argument are still in focus when they are needed. The following greedy algorithm implements this process.

Presentation Algorithm

1. Determine an argumentation strategy. **(Section 7.1)**
2. Traverse the Argument Graph according to the strategy selected in Step 1, invoking the Attentional Mechanism to determine the activation level of intermediate propositions. **(Section 7.2)**
3. Alternate between probabilistic pruning and semantic suppression until time runs out or until no pruning has been successful in the last N consecutive iterations. **(Section 7.3)**

At present, the resulting argument is in the form of propositions interleaved with causal or evidential relations. A graphical interface which allows users to build and receive arguments in an annotated network form (similar to that shown in Figure 4(b)) is in preparation. We are also considering the application of microplanning operators for generating paraphrases and aggregations, such as those described in [Huang and Fiedler, 1997], prior to rendering an argument in English.

7.1 Determining an Argumentation Strategy

The argumentation strategy determines the order of presentation of the propositions. Two basic argumentation strategies are premise-to-goal and goal-to-premise. A premise-to-goal argument simply starts at the premises, goes on to support intermediate propositions, and eventually reaches the goal. Two types of goal-to-premise arguments are hypothetical, which assume the goal and go from it to the premises (which are sufficiently believed), and *reductio ad absurdum*, which start from the negation of the goal and reach a contradiction. NAG selects an argument presentation strategy by examining separately the impact of each individual line of reasoning contributing to the belief in the goal in the Argument Graph [McConachy *et al.*, 1998].

7.2 Traversing the Argument Graph

Each argumentation strategy yields a traversal order for the Argument Graph. The premise-to-goal strategy is implemented by means of a post-order traversal modified by policies which select the order of presentation of sub-arguments. We consider two policies for presenting sub-arguments for any conclusion inside the Argument Graph: *collective* and *individual-sequential*. In the collective policy, all the current immediate antecedents of a conclusion, i.e., those that remain after pruning (Section 7.3), are mentioned immediately prior to mentioning this conclusion. For example, this policy yields the following argument for node E in the Argument Graph in Figure 5(a): "A is evidence for B, which together with D strongly supports E."[1]

[1] The phrases that convey the causal and evidential relations in the arguments depend on the numerical values in the conditional probability matrices in the normative BN, the types of the links and the direction of the argument compared to that of the links.

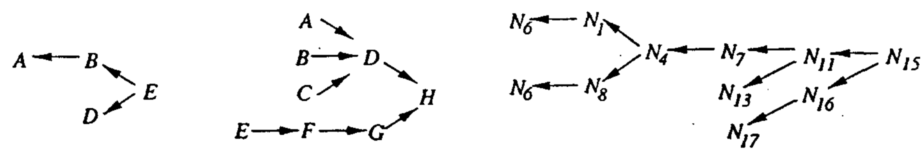

(a) (b) (c)

Figure 5: Sample Argument Graphs

This policy is used when all the antecedents provide similar levels of support for the consequent or when the antecedents are *not* conditionally independent. Since the collective policy requires all antecedents to be mentioned prior to stating the consequent, intermediate conclusions drawn earlier in the argument are mentioned again prior to their actual usage. For example, in the following argument, which corresponds to the Argument Graph in Figure 5(b), proposition D is repeated prior to its use in arguing H: "A, B and C cause D. E causes F which in turn causes G. G and D cause H." In the individual-sequential policy, the effect of one antecedent on the consequent is mentioned first followed by the effect of the other antecedents, e.g., "A is evidence for B, which strongly supports E. D also supports E." This policy is used when one antecedent is conditionally independent from the others, and provides a level of support that is quite different from that provided by the others. If an antecedent is selected according to the individual-sequential policy, the policy selection considerations are again applied to the remaining antecedents.

When the collective policy is applied, the sub-arguments are ordered according to the largest-subgraph-first heuristic, which is designed to reduce the number of intermediate conclusions that fade from the focus of attention (due to decay in their activation). This heuristic states that among several sub-arguments for the same consequent, the longer sub-arguments, i.e., those containing the largest number of propositions, should be presented before the shorter ones.

7.2.1 Activation Levels in the Argument Graph

After determining a traversal order for the Argument Graph, the activation level of each proposition is calculated based on the activation spread from the propositions which NAG has planned to mention so far (taking into account a time decay factor). As each proposition is "mentioned", activation spreads from it in a pattern which depends on the level of activation of its neighbours and the strength of its association with its neighbours (Section 4).[2] For instance, when a highly connected node is activated, it in turn activates a large number of nodes, but the activation spread to each node is relatively small, leading to a quick decay in the activation process. This process forms a small but densely populated "bubble of attention" around the activated node [Taylor, 1998]. In contrast, when a node with a few connections is activated, the activation spread to each neighbouring node is stronger, hence the process takes longer to decay, yielding a larger but more sparsely populated bubble of attention. Such larger bubbles normally encompass several reasoning steps in the Argument Graph, thereby enabling NAG to remove intermediate conclusions (Section 7.3).

7.3 Pruning the Argument Graph

The Presenter interleaves probabilistic pruning with semantic suppression. It starts with probabilistic pruning, which iteratively removes from the Argument Graph premises that have a relatively small contribution to the belief in their consequents, and also removes entire lines of reasoning which provide more support for the goal than what is strictly necessary in order to achieve a belief within the target ranges. After each removal, the Analyzer is called to check whether the 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 Section 7.2 (reordering is necessary if their relative sizes have changed as a result of the removal of

[2]Nothing is actually mentioned at this stage. However, in order to anticipate the effect of an argument, NAG pretends that propositions have actually been mentioned.

propositions), and the Attentional Mechanism is called to determine whether the remaining propositions still have sufficient activation for the addressee to be able to follow the resulting argument. Probabilistic pruning fails when the Analyzer reports that the anticipated belief in the goal is outside a target range or when the level of activation of a proposition falls below a certain threshold. In this case, the last removed proposition is reinstated, and semantic suppression is activated. This process iteratively omits from the Argument Graph intermediate conclusions whose level of activation was already high before they were "mentioned", so long as the probabilities of the links from their antecedent to their consequent are sufficiently strong. The rationale for this omission is that propositions which are readily inferred, because they have a high level of activation and the user attributes them a very high probability on the basis of a simple inference from just mentioned premises, will continue to contribute much of their support to the conclusions that follow them. For instance, if the bubble of attention around node E in the Argument Graph in Figure 5(b) includes nodes F and G, then node F may be omitted, yielding an argument such as "A, B and C cause D. E leads to G, which together with D causes H". Semantic pruning fails when the level of activation of any required subsequent proposition falls below a threshold. In this case, the last removed proposition is reinstated, probabilistic pruning is reactivated, and the interleaved pruning continues. These pruning processes continue until time runs out or until both pruning methods have failed several consecutive times.[3]

Example – Continued

After traversing the Argument Graph according to the premise-to-goal strategy, pruning is performed. Probabilistic pruning removes node N_9 in Figure 4(a) as its probability is very low, and hence so is its effect on its consequent N_6. The branch containing nodes N_{10} and N_{14} is removed despite the high probability of these nodes, because after their omission the probability of the goal proposition is still within the target ranges. The branch containing N_5 and N_6 is removed for the same reason (without removing node N_6, since it is still connected to other nodes in the Argument Graph). Probabilistic pruning fails after this step, so semantic suppression is attempted, which removes nodes $N_2 \rightarrow N_3$ and N_{12}. Both types of pruning fail after this step, yielding the Argument Graph in Figure 4(b). The expanded tree which represents the final ordering of this Argument Graph appears in Figure 5(c). The post-order traversal of this tree yields an argument which may be roughly paraphrased as follows.

Dinosaurs became extinct about 65 million years ago. This could have been caused in part by most plant species dying. This could also have been caused in part by reptiles having trouble breeding. This and most plant species dying was likely caused by the Earth becoming much cooler, which could have been caused by material obscuring the Sun.

Many 65 million-year-old iridium deposits (which were found) and material obscuring the Sun could have been caused by an explosion throwing debris into the atmosphere.

A giant crater which was discovered in the Gulf of Mexico was likely caused by an asteroid strike forming a giant crater. This and the explosion debris in the atmosphere were likely caused by a large iridium-rich asteroid striking Earth 65 million years BC.

8 Conclusion

NAG uses items in the user's focus of attention to guide a series of generation-analysis cycles which result in the generation of an Argument Graph that is both normatively acceptable and persuasive. Attentional focus also supports the generation of enthymematic arguments.

NAG was written in Common Lisp. It was tested on five sample scenarios with KBs containing up to 200 propositions. The simulation of attention via spreading activation generally led to a significant speed-up in content planning times, with little effect on the generated arguments [Zukerman *et al.*, 1998]. Content planning times slowed down when extremely slow decay factors and low activation thresholds were used

[3] An alternative approach to combining probabilistic pruning and semantic suppression is described in [McConachy *et al.*, 1998].

and when extremely fast decay factors and high activation thresholds were used. The former incorporated into the Argument Graph nodes that were only marginally related to the goal, while the latter incorporated too few nodes, resulting in essentially a goal-based search for an argument.

We conducted a preliminary Web-based evaluation of NAG's content planning component. This evaluation consisted of giving respondents pre-test questions regarding certain propositions in an asteroid argument similar to that discussed in this paper, and then presenting them with hand-generated renditions of an asteroid argument which took into account the responses to these questions. A post-test was used to determine the effect of the argument. This test showed a clear tendency among the respondents to shift belief towards the targets as a result of NAG's argument. A more rigorous evaluation of the content planning component and an evaluation of the argument presentation component will be performed in the next months.

References

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

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

[Fehrer and Horacek, 1997] 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.

[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.

[Korb *et al.*, 1997] Korb, K. B., 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.

[McConachy *et al.*, 1998] McConachy, R., Korb, K. B., and Zukerman, I. (1998). Deciding what not to say: An attentional-probabilistic approach to argument presentation. In *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*, Madison, Wisconsin.

[Mehl, 1994] 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 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.

[Taylor, 1998] Taylor, J. (1998). Constructing the relational mind. *Psyche*, 4. http://psyche.cs.monash.edu.au/v4.

[Thomason *et al.*, 1996] Thomason, R. H., Hobbs, J. R., and Moore, J. D. (1996). Communicative goals. In *Proceedings of ECAI-96 Workshop – Gaps and Bridges: New Directions in Planning and NLG*, pages 7–12, Budapest, Hungary.

[Walker, 1996] Walker, M. A. (1996). The effect of resource limits and task complexity on collaborative planning in dialogue. *Artificial Intelligence*, 85:181–243.

[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*, Madison, Wisconsin.