# A Cognitive Model of Argumentation

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

**Abstract**

In order to argue effectively one must have a grasp of both the normative strength of the inferences that come into play and the effect that the proposed inferences will have on the audience. In this paper we describe a program, *NAG* (Nice Argument Generator), that attempts to generate arguments that are both persuasive and correct. To do so NAG incorporates two models: a normative model, for judging the normative correctness of an argument, and a user model, for judging the persuasive effect of the same argument upon the user. The user model incorporates some of the common errors humans make when reasoning. In order to limit the scope of its reasoning during argument evaluation and generation NAG explicitly simulates attentional processes in both the user and the normative models.

## Introduction

In order to argue well one must have a grasp of both the normative strength of the inferences that come into play and the effect that the proposed inferences will have on the audience. Our program *NAG* (Nice Argument Generator) is intended to argue well — that is, to present arguments that are persuasive for an intended audience and also are as close to normatively correct as such persuasiveness allows. In order to develop such a system we have had to incorporate two models within NAG: a normative model, for judging the normative correctness of an argument, and a user model, for judging the persuasive effect of the same argument upon the user. The user model should ideally reflect all of the human cognitive heuristics and weaknesses that cognitive psychologists may establish as widespread, such as the failure to use base rate information in inductive reasoning (Tversky and Kahneman, 1982a) and overconfidence (Lichtenstein et al., 1982). The normative model should ideally incorporate as many items of knowledge as we can muster and the best evaluative tools for judging their relationships. Neither the user being modeled by the system nor NAG itself are unlimited cognitive agents, of course. In order to limit the scope of what might be drawn into consideration during argument evaluation and generation we explicitly simulate attentional processes in both the user and the normative models.

In this paper we first sketch the overall architecture of NAG. We then describe those design features specific to implementing the psychological mechanisms mentioned, some possible directions for extending them, and the effects of such psychological modeling on NAG's argumentation.

## An Overview of NAG

NAG is designed to analyze arguments and to compose its own arguments intended to be persuasive for particular interlocutors. Given a user model, a context and a goal proposition, NAG produces an argument supporting the goal which, according to its user model, will be effective in bringing the user to a degree of belief in the goal proposition within a target range. When presented with an argument by the user, NAG will respond either by agreeing or by presenting an effective counterargument. The system is composed of the following modules: Argument Generator, Abduction Engine, Argument Analyzer, and Argument Strategist (Figure 1).[1]

The Argument Strategist governs the argumentation process. In the first instance it either receives a goal proposition or a user argument. Given a goal proposition it invokes the Generator to initiate the construction of an argument. The Generator uses the argumentative context and the goal to construct an *Argument Skeleton*. The Argument Skeleton forms the initial basis for the system's argument, which is represented as a Bayesian network we call an *Argument Graph*. The Strategist passes this initial argument to the Analyzer, which tests the effect of the argument on the goal proposition in both the user and the normative models, using Bayesian network propagation in the submodels corresponding to the Argument Graph (Pearl, 1988; Neapolitan, 1990), while taking into account the psychological mechanisms described below in §*The Psychology of Inference*. In this way the Analyzer may discover that some of the premises employed are insufficiently supported in either the user or the normative model, or that an inference employed in the argument is weak. The Strategist uses the evaluation returned by the Analyzer to determine whether, and how, to strengthen the argument, for example by providing a weak premise in the argument to the Generator as a new goal, so that a supporting subargument may be built. The iterative process of invoking Generator and Analyzer continues until either some Argument Graph is generated which brings the original goal proposition into the target range for strength of belief, the Strategist is unable to fix a problem reported by the Analyzer, some operating constraint is violated which cannot be overcome (e.g., the overall complexity of the argument cannot be reduced to an acceptable level) or time runs out. Finally, the Strategist will report the argument to the user, if a suitable one has been produced.

[1] For a more detailed description of NAG's architecture, see Zukerman, Korb and McConachy (1996) or McConachy, Zukerman and Korb (1996).

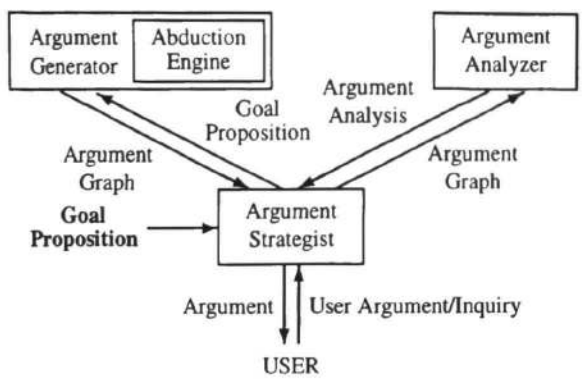


Figure 1: System Architecture

If instead the user has initiated argumentation, the Strategist will invoke the Analyzer in order to uncover flaws in the user's reasoning. Since most (if not all) arguments are enthymematic (suppressing one or more premises), apparent weaknesses are not necessarily treated as errors in the user's argument. Instead, the Strategist will invoke the Generator in an attempt to extend the user argument via abductive processes,[2] until the Analyzer reports an Argument Graph that is both normatively acceptable and acceptable to the user model. If this iterative process fails, the Strategist identifies some proposition in the Argument Graph where the two models differ in their assessments as the basis for the construction of a counterargument. If the process of filling out the argument succeeds, NAG, of course, accepts the user's argument.

NAG, then, employs the model of the user's cognition and the normative model to make a trade-off during argument construction and analysis: the aim is not simply to find arguments which NAG considers to be the best possible normatively, but to find arguments which are sufficiently effective given the user's cognitive apparatus (as represented in the user model) and also sufficiently correct given the normative model. The aim is to find what we call *nice* arguments, in contrast with strictly normatively correct (good) arguments, both for presenting to the user and for understanding what a user has presented to NAG. How NAG should play this trade-off during argument generation, for example whether to rely upon the user's reported cognitive weaknesses largely, moderately or not at all, is determined by a *user profile* associated with the user model and a *system attitude* parameter. The user profile informs NAG to what weaknesses the user is susceptible and the system attitude tells NAG to what extent it is allowed to exploit such weaknesses.

## The Psychology of Inference

### Attention

If we are to predict the effect of our arguments on others, clearly we must have some understanding of how others reason. The first fact about cognition is that, unlike the presumptions of some philosophers, no cognitive agent has access to infinite reserves of time or inferential power.[3] We propose a mechanism that is simple in its basic concept, but difficult in detail, to cope with the combination of a large collection of beliefs and a small slice of time within which to bring the beliefs to bear upon some inferential task: namely, attention.

The exact nature of attention is not yet well understood. Psychologists recognize that attention and cognition are closely related, but the nature of the relationship is far from being established. Nevertheless, there is general agreement that attention serves to apply limited cognitive capacities to problem solving by regulating the flow of information to cognitive processes (e.g., Baars's functions of attention, 1988; cf. also Allport, 1989, and Cowen, 1995). We implement such an attentional process through the following promising candidates for features of cognition that contribute to determining the focus of one's attention: *salience*, the extent to which a cognitive object[4] is prominent within the set of objects being processed by an agent at some time;[5] *recency*, the time elapsed since the cognitive object was last "touched" by a cognitive process; *semantic distance*, the degree of semantic relationship between objects.[6] Additional cognitive features which we expect to employ in the future include: *vividness* and *concreteness*, the association of the object with imagery or with a specific incident or example; *availability*, the ease with which the object can be recalled from memory and used in the current context, potentially mediated by analogical reasoning, pragmatic reasoning schemata (Cheng and Holyoak, 1985), or mental models (Johnson-Laird, 1983).

The above list is hardly exhaustive, however it is already an ambitious list of features to tackle in an initial computational model for argumentation. What we have done is to erect hierarchical semantic networks above the Bayesian networks in our user and normative models (Figure 2). The semantic networks represent the semantic relatedness of items directly, in terms of links between nodes and their association strengths. We take the context in which an argument occurs to provide the salient cognitive objects for our system: for example, if the user presents an argument to NAG, the propositions in the argument, and in the preceding discussion, will be marked as salient. We use activation with decay (Anderson, 1983), spreading from the salient objects (which are clamped), to determine the focus of attention: all items in the Bayesian networks which achieve a threshold activation level while the spreading activation process is iteratively applied will be brought into the span of attention. The spreading activation process passes activation through the pyramidal semantic-Bayesian networks, each node being 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). This itera-

[2] NAG's abductive mechanisms are described more fully in McConachy, Zukerman and Korb (1997).

[3] See Cherniak (1986) for an interesting investigation of the difficulties with such philosophical commitments.

[4] By cognitive object we mean whatever representation is being used in cognition. In the case of NAG, this means a node in an argument graph, representing a proposition.

[5] Salience might be considered a quantitative measure of attention itself, to be sure, rather than a determiner or component of attention. In any case, we do not claim that the various determining features we describe here are either ideal or orthogonal to each other.

[6] By labeling the relationship semantic we do not mean to restrict ourselves to relations that might be identified in a dictionary. *Tweety* may be closely semantically related to *Sylvester* even though *canary* and *cat* are not.

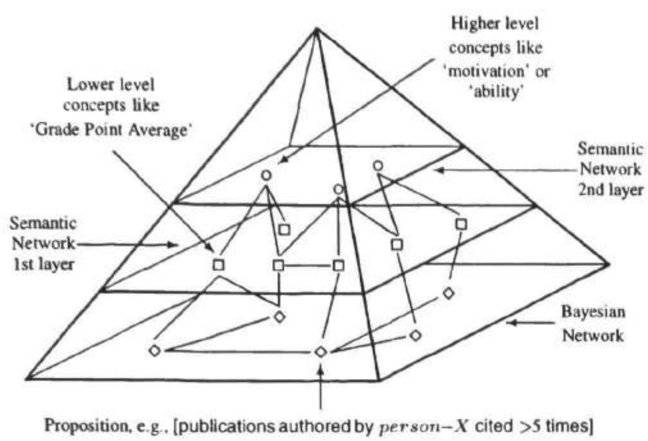


Figure 2: Semantic and Bayesian Networks

tive process ceases when an activation cycle fails to activate any new node. By these means we have a direct, and useful, implementation of attention via salience, semantic relatedness and recency.

The simulation of attention is clearly necessary when modeling users, since, for example, their attention determines what meanings they intend their sentences to have (disambiguation) and the strengths which they attribute to various inferential steps during argumentation. But, as mentioned above, attention is also necessary for modeling normatively correct inference, since any such model necessarily will also cope with finite resources and preferentially will cope with large amounts of information. One consequence of an attention-based model of cognition is that the cognition is inherently incomplete: the system must be prepared to deal with new items of information swimming into attentional focus at any time. This allows us to represent in both the normative and the user models what Bayesian network propagation proper does not acknowledge: that the import of evidence may not be fully realized when that evidence is acquired, that its impact on related propositions may be only partially resolved upon its acquisition and then further accommodated as time and circumstance allow (partial propagation has previously been investigated by Draper, 1995). Hence, our probability propagation scheme is partial, and in consequence a user argument, even by bringing into focus only propositions previously believed by NAG, may cause a change in how NAG normatively assesses conclusions.

## Error

We suggest that the attentional process described above, or some other process fulfilling like functions, is necessary for any complex cognitive agent. Since NAG needs also to model specifically its human interlocutors, we also provide mechanisms specific to NAG's user model to deal with some of the cognitive errors or illusions that psychologists have found in human reasoning. The phenomena of this kind are too diverse and complex for us to attempt modeling many of them in an initial automated arguer (see Evans, 1989, for a useful review of both deductive and inductive cognitive illusions). What we currently model are belief bias, overconfidence and the base rate fallacy.

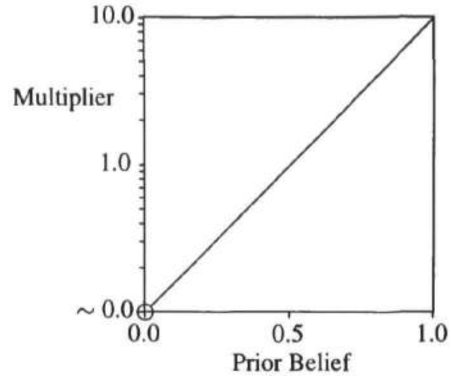


Figure 3: Logarithmic Belief Bias Curve

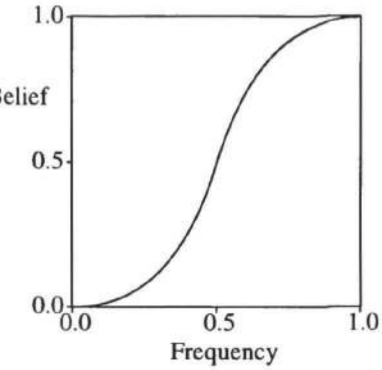


Figure 4: Overconfidence Curve

**Belief bias** is the assessment of an inference as stronger (weaker) than it is normatively because it supports (undermines) an existing belief (Evans et al., 1983). To model this effect we employ a function which converts the user's prior degree of belief in a proposition into a multiplicative factor used in any update to that belief.[7] Figure 3 shows the multiplicative factor in the case of positive evidence, i.e., evidence supporting the proposition in question, plotted against a logarithmic vertical axis. If the user's prior belief is 0.5, then the multiplicative factor is 1, so that the belief is updated in the normal way. An extreme disbelief on the other hand suppresses the impact of evidence via a very low multiplier, whereas a strong belief enhances that impact via a factor greater than 1. As a result, NAG tends to assume that users will require more and better arguments to be persuaded to change their strongly held beliefs than NAG holds to be normatively necessary.

**Overconfidence** is the attitude that people tend to adopt towards very frequent, or very infrequent, phenomena: people tend to exaggerate the probability of likely events and the improbability of unlikely events (Lichtenstein et al., 1982). NAG uses this bias to select prior probabilities for propositions in the user model when it has not been given explicit information about the user's probabilities, but only frequency information. In particular, in a direct implementation of overconfidence NAG applies an S-curve (Figure 4) to convert frequencies into user probabilities. NAG does not use this S-curve directly when the base rate fallacy applies, since in that case both errors are dealt with at once.

[7]The factor is multiplied into the likelihood ratio during belief updates.

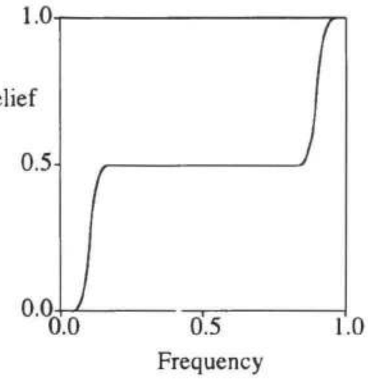

Figure 5: Base Rate Fallacy Curve (2 outcomes)

**The base rate fallacy** is the tendency humans have to ignore objective prior probability information when it is given (Tversky and Kahneman, 1982a). A striking example of this is the cab problem (Tversky and Kahneman, 1982b):

> A cab was involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. You are given the following data:
>
> - 85% of the cabs in the city are Green and 15% are Blue.
> - A witness identified the cab as Blue. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colors 80% of the time and failed 20% of the time.

By Bayes's theorem it is straightforward that the probability that the cab was Blue based on the evidence given is 0.41, because the prior probability being only 0.15 dominates the computation.[8] Nevertheless, most people presented with the story respond that the probability that the offending cab was Blue is 0.8. Not coincidentally, that response would be correct if the prior probabilities were uniformly distributed among the available options: this is characteristic of most people's response to many situations involving uncertainty — they flatten all prior probabilities. This response (as with many other cognitive illusions) is suppressed in areas of reasoning where people are highly experienced or expert. As mentioned earlier, whether the base rate fallacy, or the other cognitive illusions, are modeled and exploited for a given user is determined by that user's profile and the system attitude parameter.[9]

In order to accommodate the tendency to flatten priors, when information about the user's priors is not directly given but made available via frequency data (data that is indicated to be known to the user), NAG computes the user's prior probability using the function in Figure 5. This function incorporates overconfidence by flattening the two ends corresponding to extreme frequencies and incorporates the base rate fallacy by flattening most of the rest in the middle, the result in a binary situation being two S-curves squashed together.

[8]The computation being $0.41 = (.8 \times .15)/(.8 \times .15 + .2 \times .85)$.

[9]In the future we may tie such profile information to submodels, so that arguments in different domains will be modeled differently.

## Remarks on Error and Attention

We are not prepared to argue that our methods of modeling either the cognitive errors described immediately above or attention are fully satisfactory; rather, they are admittedly simple and probably naive. One feature of cognitive errors that has repeatedly surfaced in the experimental literature is that such errors are situation-specific: it is often easy to enhance or suppress such errors by relatively subtle manipulations of the experimental context (e.g., Tversky and Kahneman, 1986). Yet there is nothing very subtle about our implementation. Jonathan St. B. T. Evans (1989) argues at length that most of the cognitive illusions uncovered thus far are to be explained as *selection biases*: that is, the errors can often be explained by a failure to pay attention to one or more relevant features of the problem context. Thus, for example, Nisbett, Borgida, Crandall and Reed (1976) argue that the base rate fallacy is explained by the fact that pallid, imageless data tend to be ignored, and the base rates in these cases are given as pallid quantities. When, by contrast, data are presented via vivid examples they are better utilized in problem solving. Again, when the cab story is changed to state that Green cabs *cause* 85% of accidents involving cabs (rather than merely that they make up 85% of the cabs in the city), then people do *not* ignore the information (Evans, 1989). Presumably, where the data fit naturally into some causal schema or mental model, they are utilized.

In consequence, we suppose that the more complex reasoning structures and processes that we have chosen to ignore in our initial work — such as reasoning schemata, mental models and analogies — may be crucial to a full representation of human cognitive failures (and successes, of course). Furthermore, Evans' suggestion that selection is a dominant factor in such errors implies that attention and biases are intimately bound up with one another, as opposed to being separate processes, as modeled by NAG. Since we wish to keep the normative model as free as possible from merely *human* failings, this suggests that the attentional mechanisms for the user and normative models in future versions of NAG will diverge. For the time being, out of sheer complexity of the task at hand, we push these considerations ahead of us: whereas all of these considerations are relevant, our initial simplified model nevertheless shows promise in having produced a variety of plausible arguments and analyses of arguments, a few of which we now illustrate.

## Some Arguments Advanced by NAG

Studies have shown that an individual's future research performance (measured by the admittedly crude technique of counting refereed publications) can be estimated by a number of factors, including but not limited to: the ranking of the institution where she or he completed a Ph.D.; the number of publications produced prior to the Ph.D.; and the researcher's ability (as measured by standardized tests; see Rodgers and Maranto, 1989). Given a model describing the relationships between propositions describing applicants, NAG can generate arguments about whether they are likely to be highly productive in the future. Below we present some sample arguments of this type and show how modeling the cognitive errors described in §*The Psychology of Inference* leads NAG to construct different arguments in order to take their influence

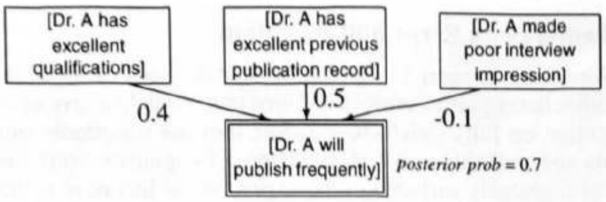


Figure 6: Normative Argument Graph for Dr. A

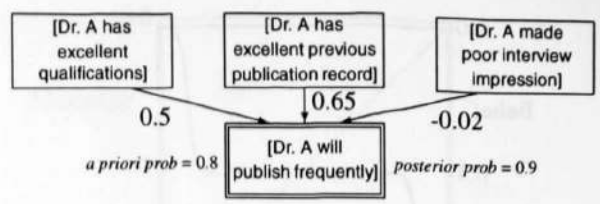


Figure 7: Argument Graph for Dr. A (with belief bias)

into account.

Suppose, for example, a research institution is looking to hire new research staff. Professor H, who is on the hiring committee, knows a recent graduate, Dr. A, and suggests that she should apply for the position. Dr. A has decent qualifications and has published previously. Prof. H has a positive opinion of Dr. A, having known her for several years. Dr. A is brought in for an interview, but unfortunately the interview goes poorly. Given this information a normative argument for why Dr. A will publish frequently in the future would have to be concessive:[10] "While Dr. A interviewed poorly, she has decent qualifications and has already published. Hence, she is reasonably likely to publish often in the future." The normative Argument Graph produced by NAG for this argument is shown in Figure 6.[11]

When Prof. H considers whether it is probable that Dr. A will publish frequently in the future, it is likely that the professor will update his beliefs in a different way from that suggested by Figure 6. In an example of *belief bias*, Prof. H will over- or under-weigh the evidence, so that his final degree of belief that Dr. A will publish often in the future is stronger than in the normative case. To show how NAG's output changes depending on the modeling of belief bias, consider what happens when NAG tries to persuade Prof. H that Dr. A is highly likely to publish frequently in the future (i.e., increasing an already positive belief). NAG can take two different approaches. If the system attitude parameter allows it, NAG can use the existing positive belief about Dr. A and the consequential belief bias (by applying the multiplicative factors in Figure 3 to the link coefficients), in which case NAG builds the simple Argument Graph shown in Figure 7. If the system parameter does not allow such a departure from the normative, NAG will be forced to build a larger Argument Graph (not pictured due to space constraints), because the smaller graph is not sufficient to generate a very strong final belief in Dr. A's future productivity (note the strength differences in the posterior probabilities of the double boxed goal nodes in Figures 6 and 7). This situation is similar to that which arises when NAG is called upon to argue on behalf of an applicant about whom Prof. H has no prior impressions.

[10] At present NAG's output takes the form of an Argument Graph. The text presented is an illustration of what NAG's output means in this case.

[11] In the Argument Graphs in the figures the number next to each inference link is a linear coefficient between the two nodes at each end of that link. An Argument Graph is actually a Bayesian network, and as such has a conditional probablility matrix associated with each node, which has more expressive power than a vector of linear coefficients; however, the studies performed on research productivity have restricted themselves to linear models. The probabilities used in NAG's normative model are loosely based upon these studies.

Prof. J[12] is also on the selection committee. He is at present sorting through the applications. He opens an envelope from an applicant, Dr. C, whom he has never heard of, and finds that Dr. C completed his Ph.D. dissertation, and has since worked in a post-doctoral position, at a university with a weak reputation. Prof. J, perhaps in a hurry to get through the pile of applications, throws this applicant's package into the rejection bin, neglecting important facts further down in the application such as that Dr. C's dissertation has led to a book contract, he has published a large number of papers recently, he scored highly on measures of ability and has excellent letters of recommendation. Plausibly, Prof. J has fallen into the *overconfidence* trap: he has seen the initial two pieces of information and, knowing (let us suppose) that applicants from rather ordinary academic institutions are unlikely to perform at a high level, dismisses the remaining information. When given the task to argue in support of Dr. C with Prof. J, bringing the latter's belief in the future high productivity of Dr. C to 0.7, NAG produces the moderately elaborate argument of Figure 8, which can be understood as: "Although the low ranking of the institutions where Dr. C did his Ph.D. and post-doctoral work suggest that he will not be highly productive, Dr. C's excellent ability, previous publications, book contract and excellent letters of recommendation strongly support the opposite. The overall assessment is in favor of Dr. C."

When Prof. J is not modeled as suffering from overconfidence (but starting from the same limited initial information), NAG requires only a subargument from Figure 8, namely: "Although the low ranking of the institutions where Dr. C did his Ph.D. and post-doctoral work suggest that he will not be highly productive, Dr. C's record of previous publications and excellent letters of recommendation strongly support the opposite. The overall assessment is in favor of Dr. C."

The same example may be used to illustrate what happens to NAG when its focus of attention is too restricted. When attempting to cope with the overconfident Prof. J., if NAG is allowed a reasonably broad attentional focus, then it creates the Argument Graph of Figure 8, as described. Given a very tight span of attention, NAG will fail to discover a satisfactory argument, since it will consider only very small Argument Graphs.

Space considerations preclude us from including an Argument Graph from NAG that demonstrates how the system models the *base rate* fallacy. This cognitive error is modeled analogously to the two other errors depicted above, with NAG applying a suitable multiplicative factor (such as from Figure 5) to the user's beliefs.

[12] All those identified by single-character surnames are wholly fictitious and any resemblance to actual people, alive or dead, is entirely coincidental :-).

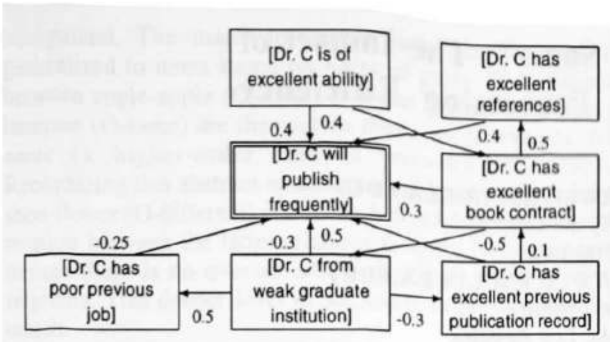


Figure 8: NAG's Argument Graph for Dr. C

## Conclusion

To argue well one must produce arguments that are supportable both factually and within the interlocutor's mind. Using both a normative model and a user model, with the latter reflecting the user's actual beliefs and inferential biases, in principle allows an argument system to generate arguments that are effective while also being justified normatively. Modeling the user's attentional focus enhances the accurate prediction of the effect of a new argument and enhances the system's understanding of a user's argument. Employing an attentional mechanism in the normative model, on the other hand, allows the system to avoid the examination of large volumes of material unlikely to aid argument analysis and construction in a given context.

The evidence we have that the above is a true account of argumentation is, at present, indicative only. We have built and run NAG in the ways reported herein, which demonstrate that there is the possibility that NAG's future development along the lines adumbrated earlier will lead to an effective argument system. Despite our original combination of Bayesian networks, semantic networks with activation, cognitive biases and attentional mechanisms and their application to argumentation, undoubtedly the greater part of the work involved in producing a complete and effective argument system lies ahead of us. In addition, in order to evaluate our system we will require comparative results assessing the system's arguments against good and bad arguments produced otherwise.

## Acknowledgements

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

## References

Allport, A. (1989). Visual attention. In M.I. Posner (Ed.), *Foundations of cognitive science* (pp. 631-682). Cambridge: MIT Press.

Anderson, J.R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Baars, B. (1988). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.

Cheng, P.W. and Holyoak, K.J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology, 17*, 391-416.

Cherniak, C. (1986). *Minimal rationality*. Cambridge, MA: MIT.

Cowen, N. (1995). *Attention and memory: an integrated framework*. Oxford: Clarendon Press.

Draper, D.L. (1995). *Localized partial evaluation of belief networks*. Doctoral dissertation, University of Washington, Seattle.

Evans, J.St.B.T. (1989). *Bias in human reasoning: Causes and consequences*. Hillsdale, NJ: LEA.

Evans, J.St.B.T., Barston, J.L., and Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition, 11*, 295-306.

Johnson-Laird, P.N. (1983) *Mental models*. Cambridge: Cambridge University Press.

Kahneman, D., Slovic, P., and Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.

Lichtenstein, S., Fischhoff, B., and Phillips, L.D. (1982) Calibrations of probabilities: The state of the art to 1980. In Kahneman et al. (1982) (pp. 306-334).

McConachy, R.S., Zukerman, I., and Korb, K.B. (1996). Consulting a user model while generating arguments. In *Proceedings of the Fifth International Conference on User Modeling, Kona, Hawaii* (pp. 153-160).

McConachy, R.S., Zukerman, I., and Korb, K.B. (1997). *Abductive argument generation and understanding*. Technical report 97/302. Department of Computer Science, Monash University.

Neapolitan, R.E. (1990). *Probabilistic reasoning in expert systems: Theory and algorithms*. New York: John Wiley & Sons.

Nisbett, R., Borgida, E., Crandall, R., and Reed, H. (1976). Popular induction: Information is not necessarily informative. In J.S. Carroll and J.W. Payne (Eds.) *Cognition and social behavior* (pp. 113-133). Hillsdale, NJ: LEA.

Nisbett, R., and Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgement*. Prentice-Hall.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers.

Rodgers R.C. and Maranto C.L. (1989). Causal models of publishing productivity in psychology. *Journal of Applied Psychology*, 74(4), 636–649.

Tversky, A. and Kahneman, D. (1982a). Judgments of and by representativeness. In Kahneman et al. (1982) (pp. 84-98).

Tversky, A. and Kahneman, D. (1982b). Evidential impact of base rates. In Kahneman et al. (1982) (pp. 153-160).

Tversky, A. and Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business, 59*, S251-S278.

Zukerman, I., Korb, K.B., and McConachy, R.S. (1996). Perambulations on the way to an architecture for a nice argument generator. In *Proceedings of the European Conference on Artificial Intelligence Workshop, Gaps and Bridges: New Directions in Planning and NLG, Budapest, Hungary* (pp. 31-36).