

Using Student Explanations as Models for Adapting Tutorial Dialogue

Pamela W. Jordan

Learning Research and Development Center
University of Pittsburgh
Pittsburgh PA 15260
pjordan@pitt.edu

Abstract

We describe a dialogue management approach that adapts dialogue strategies to a changing user model under the constraint that the approach be able to scale up rapidly for intelligent tutoring systems applications. Dialogue strategies are expressed as constraints on how to move through a student model that is composed of linked nodes. By orienting the dialogue to a linked model of the student's likely beliefs and expressing the dialogue strategies as constraints on movement within the model, the approach facilitates adaptation of the dialogue strategy as the student model changes.

Introduction

One of the many challenges in building intelligent tutoring systems that interact with students via natural language dialogue is selecting a dialogue management approach that can rapidly scale up while still allowing a number of pedagogically valid tutorial dialogue strategies to be flexibly realized for the same subject matter. Our initial approach to dialogue management in the Why-Atlas tutoring system (VanLehn *et al.* 2002) focused on scalability so that the system would be able to cover enough subject matter for student learning gains to be measurable. The initial approach proved successful in this regard (Jordan, Rosé, & VanLehn 2001; Rosé *et al.* 2001) but not with respect to flexibility in presentation.

The dialogue management approach taken in Why-Atlas can be loosely categorized as a finite state model. Although the Why-Atlas approach does allow states to be skipped if the goal of a state has already been achieved and backtracking and retries when a plan fails (Freedman 2000), it does not facilitate a large scale rearrangement of its hand-scripted dialogue plans as would be needed to flexibly realize a variety of tutorial strategies for the same content (Jordan, Makatchev, & Pappuswamy 2003). Because of this limitation and a goal that the approach be able to scale up rapidly, the content and knowledge of how to present the material for tutorial purposes are intertwined within each dialogue plan. So the same pedagogical strategies and content have to be expressed many times in order to enable the same material to be presented in a variety of ways and likewise

the same set of presentation techniques have to be applied to a wide range of subject matter.

In this paper, we describe an alternative dialogue management approach that avoids intertwining content and tutorial strategy within a single plan (as happens in (VanLehn *et al.* 2002; Evens *et al.* 2001)) in favor of one that separates domain knowledge and dialogue strategies (as in (Zinn, Moore, & Core 2002)), while still retaining the potential to scale up rapidly. The domain knowledge in this approach is represented separately in a student model and a domain model as is typically done in intelligent tutoring systems (Wenger 1987), while the dialogue strategies (these are related to pedagogical models in a traditional intelligent tutoring system) are expressed as constraints on how to move through a student model that is composed of linked nodes. By orienting the dialogue to a linked model of the student's likely beliefs and expressing the dialogue strategies as constraints on movement within the model, the approach facilitates adaptation of the dialogue strategy as the student model changes.

First we will briefly describe the Why-Atlas system and how the student model is constructed. Then we will describe our approach for using the student model to drive a dialogue.

The Why-Atlas System

The Why-Atlas system covers 5 qualitative physics problems on introductory mechanics. When the system presents one of these problems to a student, it asks that she type an answer and explanation and informs her it will analyze her final response and discuss it with her. After the discussion, the system asks that she revise her explanation and the cycle

Question: Suppose you are running in a straight line at constant speed. You throw a pumpkin straight up. Where will it land? Explain.

Explanation: While I am running, both the pumpkin and I are moving with the same speed. Once I throw the pumpkin up, it no longer has anything to thrust it forward in the either the horizontal or vertical direction. Therefore, it will fall to the ground behind me.

Figure 1: The statement of the problem and an actual student explanation.

of explanation revision and follow-up discussion continues until all the flaws in the student's response have been addressed.

One such problem is shown in Figure 1 and the student response is from our corpus of students' problem-solving sessions, some of which are with human tutors and some with the Why-Atlas system¹. The student response in this case is an instance of the often-observed *impetus* misconception: If there is no force on a moving object, it slows down.

The dialogues shown in Figure 2 and Figure 3 are from the same corpus and are follow-up dialogues with two other students who have also exhibited instances of the *impetus* misconception. By comparing the dialogue in Figure 2, which is between a human tutor and a student, with the one in Figure 3, which is between Why-Atlas and a student, one can see that the human-human dialogue is adapted to the student's beliefs as expressed in her explanation as well as in her dialogue responses. In contrast, the Why-Atlas dialogue addresses the *impetus* misconception in the context of a new, simplified problem and does not adapt its high-level strategy relative to the explanation nor make any reference to the student's explanation.

Overview of the Processing

In order to find useful dialogue topics, Why-Atlas parses the student answer and explanation into propositional representations. It uses a syntactic grammar and lexical semantics to create a representation for a each sentence (Rosé *et al.* 2002), then resolves temporal and nominal anaphora (Jordan & VanLehn 2002) and constructs proofs using abductive reasoning. *Abduction* is a process of reasoning from an observation to possible explanations for that observation. In this application the observations are the propositions that represent the student's answer and explanation and the proof is the abductive reasoning steps that explain the propositions (Jordan, Makatchev, & VanLehn 2003). The motivation for constructing a proof is that it gives more insight into the line of reasoning the student may be following across multiple sentences since proofs of many propositions should share subproofs. Indeed, one proposition's entire proof may be a subproof of the next proposition. Moreover, subtle misconceptions such as *impetus* are revealed when they must be used to prove a student's explanation. The proof-based approach also opens the possibility of implementing interactive proof generation via a dialogue with the student. This interaction can serve the dual purpose of revealing to the student the conjectured argumentation behind her statement, and disambiguating the student's intended meaning when there are multiple plausible proofs.

Currently, Why-Atlas only uses the first plausible proof tree it finds to select high-level tutorial dialogue goals such as remediation of a misconception or elicitation of particular parts of a prescribed explanation that were not expressed in the student's explanation. The main mechanism for addressing these goals are what we call a knowledge construction

¹Language and domain errors in excerpts from the corpus have not been corrected. We denote deletions that are irrelevant for making a point with ... to save space.

T: Does an object always need something to thrust it to be moving?
S: An object always needs a thrust or push to keep it accelerating but from this problem, we know that it has constant speed which means no acceleration.
T: Yes, the object needs a thrust to accelerate. But does it need a thrust to be moving at constant speed?
S: No, Newton's first law tells us that an object in motion will continue in motion until an outside force acts upon it. ...
T: Good. Now consider the pumpkin just after throwing it up. What forces are acting on it?
S: After the pumpkin is thrown up, gravity is working on the pumpkin.
:
:
T: Now tell me whether this will affect the horizontal motion of the pumpkin
S: This will not affect the horizontal motion of the pumpkin because gravity only acts in the vertical direction.
T: Right. What was the horizontal velocity of the pumpkin just before it was thrown up?
S: ... The pumpkin velocity just before it was thrown up was constant with my velocity.
T: Right again. Since you and the pumpkin have the same horizontal velocity, which is constant, where do you think the pumpkin will land?
:
:
S: Ok, if I were to throw the pumpkin straight up while I was running with a constant speed, the pumpkin would land on me—ouch!

Figure 2: An actual follow-up dialogue between a human tutor and a student who had exhibited the *impetus* misconception

T: OK, let's consider the relationship between force and acceleration in a different context. A hockey puck is resting on the rough ice surface of a skating rink ... You push with a force, F ... Assume that the only force acting horizontally on the puck after the push is friction due to the ice surface ... What will happen to the velocity of the puck after your push?
S: It slows down over time.
T: Yes. ... Now assume a Zamboni smooths the ice so that the friction ... is reduced. You push ... with the same force as before (and for the same amount of time). How does the deceleration of the hockey puck after this push compare ... ?
S: It slows down more slowly.
T: Correct. Now assume ... that the friction caused by the ice surface is zero. ... what ... would happen to the velocity of the puck following the push?
S: It doesn't slow down.
T: Right. In general, according to Newton's first law, if there are no forces acting on a body, what happens to the body's velocity?
S: It is constant.

Figure 3: An actual Why-Atlas system follow-up discussion with a student who had exhibited the *impetus* misconception




Figure 4: Example of Two Alternative Simplified Proofs for “The pumpkin slows down.”

dialogue specification. This specification is a hand-authored push-down network. Nodes in the network are either the system’s questions to students or pushes and pops to other networks. The links exiting a node correspond to anticipated responses to the question. In order to achieve scalability, each question is a canned string, ready for presentation to a student. The last state of the network is saved in the history in order to classify the student dialogue responses and determine which transition link to follow to the next tutor node.

Student Model Construction

The proofs that the system builds of the student’s explanation represent the student’s knowledge and beliefs about physics with respect to the problem to which she is responding. Acquiring and reasoning about student beliefs and knowledge is one of the central issues addressed by work in student modelling. A student model is a type of user model and in general a user model provides information the system can use in adapting to the needs of its user (Wahlster & Kobsa 1989). In the case of Why-Atlas, the system uses this representation to select relevant remediation and elicitation dialogue specifications.

To illustrate how abductive proofs are created, we will step through a simplified example for sentence (1). Many plausible proofs are possible for this sentence but we will only examine how the two shown in Figure 4 were constructed. First, we will take it as a given that the air resistance is 0^2 and that it has already been established in another part of the proof (not shown here) that the runner is not applying a horizontal force to the pumpkin after he throws it.

(1) The pumpkin slows down.

Each level of downward arrows from the gloss of a proposition in the two alternative proofs shown in Figure 4 represents a domain rule that can be used to prove that proposition. Or the proposition can be unified with a right hand side of a rule (e.g. the top level inference in both proofs).

²Students often overlook relevant givens, so proofs that ignore a given can be generated as well.

Now we will walk through the two proofs starting at the point at which they diverge; there are two ways of proving that the horizontal component is decreasing. First let’s consider just the case of the left proof in Figure 4. Here the proposition is proved using a buggy physics rule that is one manifestation of the *impetus* misconception; the student thinks that a force is necessary to maintain a constant velocity. The rule gives rise to two new goals that need to be proved. In this proof it is assumed that the student has the bug and no further attempts are made to prove the student believes it although there is the option of asking a diagnostic question at this point. Next it proves the goal that the total force on the pumpkin is zero by proving that the possible addend forces are zero. Since it is a given that air resistance is negligible this part of the proof is complete. Likewise, since we said that it had already been proved elsewhere that the man is applying a horizontal force of 0 to the pumpkin after he throws it, this branch of the proof is complete as well. The proof contains just one assumption, that the student has the *impetus* bug. In the current system, this motivates the initiation of the dialogue shown in Figure 3.

Looking now at the right alternative proof in Figure 4, we see that to prove the horizontal component of the velocity is decreasing it is necessary to prove that the horizontal component of the acceleration is non-zero and to prove this it is necessary to prove that the total horizontal force on the pumpkin is non-zero. One way to prove this is to prove that exactly one of the addend forces is non-zero. Next although it is given that wind resistance is negligible, this is ignored at this point in order to try to prove that there is exactly one non-zero addend force on the pumpkin. Namely it tries to prove that wind resistance is not negligible but since this cannot be proved it must be assumed. This representation means that the student has erroneously ignored the given.

The system now has two plausible proofs involving two different bugs and no means of choosing between them without interacting with the student. However, for this paper we will leave aside the issue of finding the most plausible proof in order to focus on how to use this type of student model to




Figure 5: Relationships between nodes for which high certainty of student belief is attributed (marked with *) and nodes for which there is low certainty: a) ancestor b) sibling c) transitive

drive remediation and elicitation dialogues while still meeting the scalability constraint.

A Model-walking Approach for Realizing and Adapting High-level Dialogue Strategies

The model-walking approach that we will describe here is a type of model-tracing technique used in intelligent tutoring systems (Anderson, Boyle, & Reiser 1985), but instead of only imposing pedagogical constraints such as scaffolding and fading (VanLehn *et al.* 2000) there are also additional constraints particular to natural language dialogue. Natural language generation work on argumentation and explanation (e.g. (Zukerman, McConachy, & Korb 2000; Horacek 1997; Zukerman & McConachy 1993)) will give us a starting point for deriving constraints for both elicitation and remediation dialogue strategies relative to proof-structures and causal graphs. But there are still challenges to address since generation work on argumentation and explanation has focused on generating text as opposed to dialogue and deals mainly with correct representations.

The model-walking approach is independent of the techniques used to create the student model. So a system could build a student model using bayesian inference or classify the student's utterances relative to hand-crafted stereotypical student models. The minimum requirement for the model is that it be constructed of propositional nodes that are informationally linked (e.g. causes, enables), that all the nodes are cognitively plausible (e.g. no nodes embedded about problem solving control) and there is a certainty measure on each node for attributing belief of the proposition to the student.

There are three types of primitive model-walking constraints that we define in terms of the structure of the links in the student model between a target node of low certainty, that is to be elicited from the student, and a node of high certainty: 1) a simple ancestor relationship, as in Figure 5a 2) a sibling relationship, as in Figure 5b and 3) a transitive relationship as in Figure 5c. The questions that can be generated given one of these relationships vary depending on the orientation of the target node with respect to the high certainty nodes. For example, in the case of the relationship in Figure 5a, the query is of the form "what follows from A?", but if the certainty levels are instead reversed so that A is now the target node, the query is of the form "how do you know

T: What does a horizontal force on the pumpkin tell you about the pumpkin's acceleration?
S: it is non-zero.
T: What then does the non-zero acceleration tell you about the pumpkin's horizontal velocity?
S: It is not constant
T: What does the problem description say?
S: Oh, so there must not be a horizontal force on the pumpkin.

Figure 6: A hypothetical dialogue that uses reductio ad absurdum as a tutorial strategy

that B holds?". In the case of Figure 5b, queries are of the form "what else do you have to consider in addition to A to know that B holds?", and for Figure 5c, they are of the form "how does knowing that A holds lead you to know that C holds?".

Using these primitives we can define higher level elicitation and remediation strategies by specifying constraints on the distance between the nodes in the three structural relationships, on the orientation between the target node and the high certainty nodes, and on the distance and orientation when choosing the next target node in the model. So if there are no bugs in a student model but most of the model had to be inferred and just a few student propositions are of high certainty, then the student explanation is considered incomplete and a higher level elicitation strategy is needed. One strategy for resolving the incompleteness is to select a proposition of low certainty that is at some distance N on a path in the student model from a proposition of high certainty. For example, if a student said "The pumpkin lands on me because the velocity of the pumpkin is constant." and N is 1, the system responds with "What does the constant velocity tell us about the acceleration of the pumpkin?". Then in selecting the next target node, there are constraints to prefer that a recently elicited node be part of the structural relationship and to prefer that the orientation of the target node and the nodes of high certainty remain the same. As the primitive strategies succeed in eliciting correct nodes, the certainty of the nodes in the student model are increased. In this way, if many nodes are candidates for a particular primitive query then the higher-level dialogue strategy will adapt according to the changes.

Additional constraints beyond those used for elicitation strategies are needed for remediation strategies since correct models of the problem must be consulted as well as the buggy student model. Two possible strategies, drawn from argumentation theory, are *premise to goal* and *reductio ad absurdum* (Zukerman, McConachy, & Korb 2000). In both cases, contradictory nodes in the student model must be identified relative to the correct problem model. In the case of *premise to goal*, the final target node is the contradictory node in the correct model, but the initial and intermediate target nodes are all the nodes of low certainty (relative to the student model) that are between nodes of high certainty and the final target. In moving from the initial target to the final target, the constraints on distance and direction are the same as for the high-level elicitation strategy.

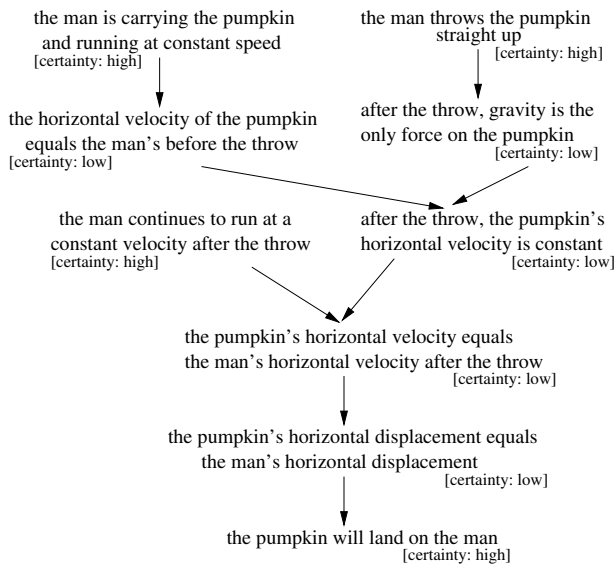


Figure 7: An example of a student model for the proof-walking approach

The *reductio ad absurdum* case is more difficult because it treats the student's actual contradiction as being correct and instead leads the student to a contradiction of one of her strongly believed propositions (that is true), as with the hypothetical dialogue in Figure 6. In this hypothetical dialogue the starting point is "There was a force acting on the pumpkin before the throw" with the contradiction being "the horizontal velocity is not constant".

Once the contradiction in the student model is found, a domain model in which the proposition is used correctly must be found. Now a new contradictory target must be found in the new domain model relative to the student model. Once this is done, the constraints are the same as for a *premise to goal* strategy.

Current Status

We have implemented a proof of concept prototype for the three types of lower-level model-walking constraints and a higher level elicitation strategy that makes use of these lower-level constraints. An excerpt of a dialogue that was produced by this prototype is shown in Figure 8. In this case the initial stereotypical student model that was used is shown in Figure 7 and represents the answer "the pumpkin lands on the man". Nodes in the model that are marked with high certainty for likely student belief are the givens of the problem and any nodes that have been explicitly expressed by the student in their answer and explanation for the problem. In this case there are three givens in the upper part of the tree and one node at the bottom which the student explicitly expressed in her response. Also, in this model we assume that all of the links are causal. This stereotypical student model differs from a proof produced by the Why-Atlas system in that the nodes are higher-level; they are clusters of multiple nodes in a proof-tree.

There are several choices about how to proceed relative

T: How do you know that the pumpkin will land on the man?
 U: I don't know
 T: What is the pumpkin's displacement relative to the man's?
 U: it is equal to the man's
 T: Right. The pumpkin's horizontal displacement will be equal to the man's displacement at all times. So what leads you to know this, given you already know that the man is carrying the pumpkin and is running at a constant speed?
 U: I don't have any idea
 T: What can you say about the velocity of the pumpkin before the runner throws it?
 :
 :
 T: Right. After the release, the pumpkin's horizontal velocity is constant. So what leads you to know this, given you already know that the man throws the pumpkin straight up and given that the horizontal velocity of the pumpkin is the same as the horizontal velocity of the runner at the time of the release?
 U: After the release the only force on the pumpkin is the downward force of gravity
 T: Right. . . .

Figure 8: An actual dialogue produced by a prototype of the proof-walking approach based on the input model shown in Figure 7

to the student model in Figure 7, either the model-walking can start with the givens and proceed downward or start with something the student said and move upward. The later was preferred in the prototype because it was expressed by the student and is in focus.

In this prototype we use simplistic approaches for interpretation of the student's dialogue responses and generation of queries. Queries are generated via templates. The templates we specified for the three general structural relationships require natural language clauses for slot fillers. We then embellished each node of the stereotypical student model with a natural language phrase, that expresses its propositional content, to be used as a filler in the templates. Recognition is done by matching the proposition generated by sentence level recognition against the propositional representation for the node. To deal with cases where the template query fails, we also embellished the nodes with canned questions that are more focused.

The dialogue in Figure 8, which was produced by the prototype, illustrates some of the improvements and future work that are necessary for a better elicitation strategy. We need to experiment with better constraints that prefer retaining the previous direction of movement through the model since the alternation of directions seems unnatural and is potentially confusing. Also constraints on the distance between the nodes that are being elicited and those being brought back into focus need to be considered. When the distance is unconstrained, the questions generated can become overly verbose and difficult to understand, as with the next to last tutor turn.

Our goal now is to expand upon the prototype by implementing remediation strategies, experimenting with a variety of constraint specifications for dialogue strategies and integrating the prototype into the Why-Atlas system in or-

der to verify the scalability of the approach. We are currently reviewing the human-human dialogues in our corpus in order to refine the constraint specifications for dialogue strategies. So far we have confirmed that the argumentation strategies, in which the student is lead to a contradiction, are used frequently in human-human dialogues. We have identified 18 dialogues from our corpus in which students made 17 contradictory statements in their discussions with 3 tutors. 7 of these cases were related to contradictions of the student's initial answer and 10 were related to contradictions of physics concepts held by the student. The 3 tutors represented in the set of dialogues successfully helped the students resolve wrong assumptions, misconceptions and errors by a strategy of leading them to a contradiction.

Acknowledgments

This research was supported by MURI grant N00014-00-1-0600 from ONR Cognitive Science and by NSF grants 9720359 and 0325054. I offer my thanks to the entire natural language tutoring team at LRDC for their many contributions in creating and building the Why-Atlas system. I especially thank Maxim Makatchev and Uma Pappuswamy for their work on abductive reasoning within Why-Atlas and Kurt VanLehn for many helpful discussions and suggestions on the model-walking approach.

References

- Anderson, J.; Boyle, C.; and Reiser, B. 1985. Intelligent tutoring systems. *Science* 456–462.
- Evens, M.; Brandle, S.; Chang, R.; Freedman, R.; Glass, M.; Lee, Y.; Shim, L.; Woo, C.; Zhang, Y.; Zhou, Y.; Michael, J.; and Rovick, A. 2001. Cirsim-tutor: An intelligent tutoring system using natural language dialogue. In *Proceedings of 12th Midwest AI and Cognitive Science Conference*, 16–23.
- Freedman, R. 2000. Plan-based dialogue management in a physics tutor. In *Proceedings of the 6th Applied Natural Language Processing Conference*.
- Horacek, H. 1997. A model for adapting explanations to user's likely inferences. *User Modeling and User-Adapted Interaction* 7(1):1–55.
- Jordan, P., and VanLehn, K. 2002. Discourse processing for explanatory essays in tutorial applications. In *Proceedings of the 3rd SIGdial Workshop on Discourse and Dialogue*.
- Jordan, P.; Makatchev, M.; and Pappuswamy, U. 2003. Extended explanations as student models for guiding tutorial dialogue. In *Proceedings of AAAI Spring Symposium on Natural Language Generation in Spoken and Written Dialogue*.
- Jordan, P.; Makatchev, M.; and VanLehn, K. 2003. Abductive theorem proving for analyzing student explanations. In *Proceedings of Artificial Intelligence in Education*.
- Jordan, P.; Rosé, C.; and VanLehn, K. 2001. Tools for authoring tutorial dialogue knowledge. In *Proceedings of AI in Education 2001 Conference*.
- Rosé, C.; Jordan, P.; Ringenberg, M.; Siler, S.; VanLehn, K.; and Weinstein, A. 2001. Interactive conceptual tutoring in atlas-andes. In *Proceedings of AI in Education 2001 Conference*.
- Rosé, C.; Roque, A.; Bhembe, D.; and VanLehn, K. 2002. An efficient incremental architecture for robust interpretation. In *Proceedings of Human Language Technology Conference*.
- VanLehn, K.; Freedman, R.; Jordan, P.; Murray, C.; Oran, C.; Ringenberg, M.; Rosé, C.; Schultze, K.; Shelby, R.; Treacy, D.; Weinstein, A.; and Wintersgill, M. 2000. Fading and deepening: The next steps for ANDES and other model-tracing tutors. In *Proceedings of the Intelligent Tutoring Systems Conference*.
- VanLehn, K.; Jordan, P. W.; Rosé, C.; Bhembe, D.; Böttner, M.; Gaydos, A.; Makatchev, M.; Pappuswamy, U.; Ringenberg, M.; Roque, A.; Siler, S.; Srivastava, R.; and Wilson, R. 2002. The architecture of Why2-Atlas: A coach for qualitative physics essay writing. In *Proceedings of the Intelligent Tutoring Systems Conference*.
- Wahlster, W., and Kobsa, A. 1989. User models in dialogue systems. In Kobsa, A., and Wahlster, W., eds., *User Models in Dialogue Systems*. Springer Verlag, Berlin. 4–34.
- Wenger, E. 1987. *Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge*. Los Altos, CA: Morgan Kaufmann.
- Zinn, C.; Moore, J. D.; and Core, M. G. 2002. A 3-tier planning architecture for managing tutorial dialogue. In *Proceedings of Intelligent Tutoring Systems Conference (ITS 2002)*, 574–584.
- Zukerman, I., and McConachy, R. 1993. Generating concise discourse that addresses a user's inferences. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 1202–1207. Morgan Kaufmann Publishers, Inc.
- Zukerman, I.; McConachy, R.; and Korb, K. B. 2000. Using argumentation strategies in automated argument generation. In *Proceedings of the 1st International Natural Language Generation Conference*, 55–62.