Carnegie Mellon University
Research Showcase @ CMU

Department of Psychology | Dietrich College of Humanities and Social Sciences

November 2008

An ACT-R Model of Adaptive Communication

M. Matessa
Carnegie Mellon University

J. R. Anderson
Carnegie Mellon University

Follow this and additional works at: http://repository.cmu.edu/psychology

This Conference Proceeding is brought to you for free and open access by the Dietrich College of Humanities and Social Sciences at Research Showcase @ CMU. It has been accepted for inclusion in Department of Psychology by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

An ACT-R Model of Adaptive Communication

M. Matessa (matessa@cmu.edu)
J. R. Anderson (ja0s@andrew.cmu.edu)
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA, USA

Abstract

When people communicate they try to establish mutual knowledge. Garrod and Anderson (1987) proposed that a way to minimize effort during this process would be to follow a "output/input coordination" principle, where output to a partner is formulated according to the same principles of interpretation as those needed to interpret input from a partner. A computational model of establishing mutual knowledge efficiently can be given in the ACT-R architecture (Anderson & Lebiere, 1998) where goals that are completed successfully can be retrieved and used later. Applied to communication, goals of presenting and accepting information include semantic and syntactic representations of that information, and these goals can later be retrieved to provide templates for the creation of new utterances. Results from an ACT-R model communicating with human subjects show similar performance to that of human subjects communicating together.

ACT-R Theory

ACT-R (Anderson & Lebiere, 1998) is a theory of human cognition incorporating both declarative knowledge (e.g., addition facts) and procedural knowledge (e.g., the process of solving a multi-column addition problem) into a production system where procedural rules act on declarative chunks. ACT-R has been successful in using mental representations to interact with students in tutoring sessions (Anderson, Corbett, Koedinger, & Pelletier, 1995), in making detailed predictions of errors and latencies in memory retrieval (Anderson & Matessa, 1997; Lebiere & Anderson, 1998), and in accounting for individual differences in working memory capacity (Lovett, Reder, & Lebiere, 1999).

Lebiere and Anderson's model of addition fact learning provides a good example of learning new declarative chunks. The ACT-R theory stipulates that there are only two sources of new chunks: from perception and from completed goals. The goal in addition is to find the sum of two numbers and this can be accomplished by computing the answer (e.g., by counting) or retrieving the answer from memory. ACT-R accounts for the creation of addition fact chunks as follows: initially, the goal of an addition problem is completed by computing the answer and storing the answer in the goal. Once this goal is completed, it is then available as an addition fact. This process of creating new declarative chunks can also be applied in the domain of communication, where the declarative knowledge assumed to be shared by participating individuals is known as common ground.

Common Ground

Common ground can be thought of as the presuppositions an individual involved in conversation makes concerning what information is mutually believed by the participants. Clark and Schaefer (1989) proposed that a speaker cannot believe their contribution is part of common ground until the listener gives evidence that this is the case.

Since new additions to common ground are dependent on the explicit or implicit acknowledgment of the listener, and since ACT-R has a mechanism of creating a declarative chunk upon completion of a goal, it seems natural to model (from the speaker's point of view) the addition of the speaker's information to common ground as the successful completion (marked by listener acknowledgment) of a goal to contribute information. Likewise, from the listener's point of view, a goal of accepting information contributed to common ground by the speaker can be judged as complete if an acknowledgment can be given to the speaker. These two goals would contain the information contributed to common ground and would be able to be retrieved as declarative knowledge.

Dialogue Acts

Clark and Schaefer also claimed that a contribution to common ground is done with an illocutionary act such as making an assertion, asking a question, etc. A set of widely accepted acts comes from the Discourse Resource Initiative (Core & Allen, 1997), developed by an international team of dialogue researchers. These dialogue acts represent ways to introduce new information (forward-looking acts) and ways to respond to previous dialogue acts (backward-looking acts). But how do these dialogue acts relate to the beliefs and intentions of individuals involved in communication? Poesio and Traum (1998) suggest an axiomatisation of the DRI dialogue acts in terms of mental attitudes of individuals where reactions to certain dialogue acts can make certain changes in beliefs and intentions in common ground. Some of these effects can be seen in Table 1.

Table 1: Dialogue Act effects on common ground (Poesio & Traum, 1998)

A's Dialogue Act	B's Action	Change to common ground
Any	recognize	B is obligated to produce an Understanding-act
Directive(A,B,x)	recognize	B is obligated to Address the directive
	accept	B is obligated to perform x
Statement (A,B,K)	acknowledge	A is socially committed to B to K being true
Assert(A,B,K)	acknowledge	A is trying to cause B to believe K
	accept	B is socially committed to A to K being true

Since common ground is declarative knowledge, since productions are the only way to change declarative knowledge in the ACT-R theory, and since Poesio and Traum suggest reactions to dialogue acts that change common ground, it seems natural to represent the reactions to dialogue acts as productions in ACT-R. The goals of these productions contain knowledge of dialogue acts and public beliefs, intentions, and social commitments, and when completed these goals become part of declarative memory and can become part of common ground.

Communication Task

Communication is usually motivated by the desire to complete a certain task. To begin our modeling efforts, we were interested in a simple two-participant task where both participants have the same abilities and unique knowledge to be communicated. At first we considered using a letter sequence task (Novick, Hansen, & Lander 1994) where subjects are given different sequences of letters with some missing letters with the goal of creating a whole letter sequence. Any letter that is missing for one subject is known by the other subject, and some letters are known by both subjects. So one interesting aspect of this task is that initially there is some information that is mutually known, some that is only known by one subject, and some that is only known by the other subject. We wanted to look at how subjects talked about shared and unique information, but without the one-way linear constraint of reading a sequence, so we created a two-dimensional task where subjects are given parts of a graph with

the goal of creating a whole graph. The graphs are colored circles connect by lines (similar to those used by Levelt (1982) to study communicative reference) and are designed so that similar colors on the parts can overlap and form a larger graph. So like the letter sequence task there is common information and information unique to each subject, but unlike the letter sequence there is no linear constraint to the information and so subjects must agree on how to communicate information about their graph parts and how the parts of the graph overlap. We are interested in communication using text, so the subjects send messages by way of a chat window from two different computers.

In addition to creating a whole graph from two parts, subjects also have the goal of confirming each of the circles. This is done by each subject selecting one circle at a time -- if the circles are the same, their score is increased, but if the circles are different, the score is decreased. This confirmation goal gives an objective measure of task performance in terms of a score, and it allows for the use of more complicated dialogue acts such as requesting that the other person confirm a circle or committing to confirming a circle.

Communication Interface

In a similar spirit to the COLLAGEN project (Rich & Sidner, 1998), we are not interested in modeling the processing of unrestricted English syntax but in modeling the higher-level communicative acts accomplished with English. So like the COLLAGEN project we eventually intend for our model to interact with people with a restricted set of English phrases. This restricted interface need not drastically hinder the communication process or task performance. In a study comparing a restricted interface to an unrestricted interface for students solving physics problems, Baker and Lund (1997) showed that the restricted communication interface did not interfere with task performance. In fact, it promoted a more task-focused and reflective interaction. Still, for our task we want to compare unrestricted to restricted communication to see if the restricted interface has any effect on task performance.

Our restricted interface allows the composition of a text message by first choosing a topic of discussion and dialogue act to address the topic. The topics of conversation are paired connections (how one circle relates to another), multiple connections (rows or columns of circles), numbers (how many of a specific kind of circle there are), correspondences (what circle in one person's graph corresponds to in the other person's graph), confirmations (talking about mutually confirming a circle), and experiment phases. For paired connection, multiple connection, and number topics, the Assert, Info-request, and Answer dialogue acts can be initiated with the Make Statement, Ask Question, and Answer buttons (respectively). For correspondences, the Assert and Agreement dialogue acts can be initiated with the Propose and Assess buttons, and for confirmations and experiment phases, the Action-directive and Agreement dialogue acts can be initiated with the Request and Assess buttons. These choices bring up sentence templates where words can be chosen for the sentences from pull-down menus.

Communication Model

A preliminary model was created that can participate in the graph completion task. The model reacted to statements with related statements, to questions with answers, to proposals of correspondence between circles with an assessment of the proposal, and to requests for confirmation with an assessment of the request. If a valid correspondence is proposed for the last circle, the model started to request confirmations of circles. Proposals of correspondence were rejected if they created a conflict in circle color. Requests for confirmation of a circle were rejected if there was no correspondence for that circle, and if the request was accepted, the circle was confirmed. The reactions of the model represent obligations (Traum & Allen, 1994) to answer questions and to address requests.

The predictions of this model that are compared to subject data were that the obligation to answer questions will be followed (Info-request dialogue acts will be followed by Answer

dialogue acts) and that requests for confirmation will be grounded before the confirmation is acted on (Action-directive dialogue acts by the speaker will be followed by Agreement dialogue acts by the listener before confirmation actions by the speaker are made). Agreement dialogue acts include Accept, Reject, Maybe, and Hold acts. As these predictions involve only simple interpretations of dialogue acts, no inter-rater reliability testing was performed on the interpretation of subjects' messages as dialogue acts.

Experiment 1

Subjects

Fourteen pairs of Carnegie Mellon University undergraduate and graduate students attempted the graph completion task, with seven pairs using an unrestricted interface and seven pairs using a restricted interface.

Method

Each pair was told that they would each be given part of a graph and their goal was first to create a whole graph as a result of circles overlapping from each part of the graph, and then to confirm each circle in the whole graph. They were told they would be sitting in different rooms and would be using a chat window to talk to each other. They were shown a drawing pad which contained an example graph part consisting of connected colored circles, and were shown how to add and erase circles representing circles from the partner's graph. They were also shown a chat window which could send eighty-character messages and only displayed the partner's last message. In the restricted interface condition, subjects were told that messages were composed in a communication window that allowed the creation of restricted sentences and were led through the creation of each kind of message. After making sure subjects understood the task, they were then given individual practice problems which used the addition, erasing, and confirming functions of the drawing pad. Finally, the subjects were given their graph parts and were told there were no time constraints in solving the problem.

Results

Of fourteen total pairs, one pair in each of the unrestricted and restricted conditions were unable to complete the task in the hour provided. To compare task performance between the unrestricted and restricted interface conditions, the number of turns to complete the task, the time to complete the task, and the final score were measured (Table 2). There was no significant difference in the number of turns (t=0.798), time (t=0.1551), or final score (t=1.185). The low number of students gives these results a low power, but as a pilot result it appears there are no drastic differences between the two interfaces.

Table 2: Performance in Unrestricted and Restricted conditions

	Unrestricted			Restricted			T	df	p
	mean	SD	min,max	mean	SD	min,max			
turns	21.5	(5.8)	[11,28]	24.7	(7.8)	[12,33]	0.798	10	0.44
time	22.3	(6.6)	[15,34]	28.3	(6.8)	[15,33]	1.551	10	0.152
score	90.0	(16.7)	[60,100]	98.3	(4.1)	[90,100]	1.185	10	0.264

To examine model predictions, paired dialogue acts in the two conditions were examined. The model predicts that all Info-request dialogue acts will be followed by Answer dialogue acts. This prediction is supported by the data, where in both restricted and

unrestricted interface conditions, 97% of Info-request dialogue acts were followed by Answer dialogue acts.

The model also predicts that confirmation Action-directive dialogue acts will be followed by Agreement dialogue acts (which include Accept, Reject, Maybe, and Hold acts) before confirmation actions are made. This prediction is supported in the restricted interface condition, where 95% of the confirmation Action-directive dialogue acts were followed directly by explicit Agreement acts (a text message) or implicit Agreement acts (confirmation of the circle by the other subject). In the unrestricted interface condition, only 54% of the Action-directive acts were followed directly by explicit or implicit Agreement acts. Part of the reason for this low number is that 70% of these Action-directive acts occurred after an explicit plan had been made on the sequence of circles to be confirmed. This planning was not supported by the restricted interface or the model. But even with this planning, all (three) of the incorrectly confirmed circle errors in the unrestricted interface occurred as a result of a subject not waiting for an explicit or implicit Agreement act after a confirmation Action-directive. The only error in the restricted interface condition occurred as a result of a "group hallucination" when both partners created and confirmed a circle that neither of them had as part of their original graph. Since each problem had ten circles to confirm and twelve pairs of subjects completed the task, there were 120 chances overall to incorrectly confirm circles. An example of how waiting for an Agreement act prevented an error in the restricted condition can be seen in the following example: Subject A produced an Action-directive act with "Let's confirm our third green circle." Subject B produced an implicit Reject act with "I have two green circles." Subject A then did not confirm the circle but produced an Accept act with "OK."

Experiment 1 Conclusions

Since subjects' performance in the graph completion task (as measured by score, turns to completion, and time to completion) was not unusually different between the restricted and unrestricted interface conditions, the restricted interface seems to be an appropriate tool in studying this task. For subjects using this restricted interface, the ACT-R model was successful in its predictions of the obligation to answer questions and of waiting for an Agreement dialogue act before confirming circles. For subjects using the unrestricted interface, the model was only successful in its prediction of the obligation to answer questions. Most subjects using the unrestricted interface who did not wait for agreement before confirming circles had previously agreed on a sequential plan to confirm circles, and this strategy was not supported in the restricted interface or the model. Subjects using this strategy apparently assume their reference to a particular circle and their decision to confirm that circle will be acceptable to their partner because of their previous plan.

In order to incorporate this goal planning strategy into the restricted interface and the model, Experiment 2 was created with an improved restricted interface. Also, to look at the development of communication over time, the problems were simplified to have six circles with one marked as a common circle. From previous research (Clark & Wilkes-Gibbs, 1986) it was expected that the message length would decrease over time in the unrestricted interface. To facilitate this decrease in the restricted interface, the manner of composing messages was changed from choosing words from a pull-down menu to typing words that were displayed in a menu. The menu for the word choice could be skipped over with the Tab key, and in this way shorter messages could be produced. This new method permits a closer correspondence to the unrestricted interface (unrestricted typing). As with the older restricted interface, message templates were chosen from a communication menu. In order to provide more redundant information in the problem that could later be left out of messages, resulting in a shorter message length, additional dimensions of size and shape were added to the color dimension of the circles. These dimensions were redundant, so that red objects were always small and thin, green objects were always medium and round, and blue objects were always large and fat.

The ACT-R model used in Experiment 1 was improved to be able to communicate with subjects using the restricted interface in real time. In addition to communicative obligations (such as answering questions), task obligations (such as making sure all objects are confirmed) were added to the model. This allows the model to passively react to requests for information or actions if the subject takes an active role or actively request information or actions if the subject takes a passive role. The concept of goal plan was also added to the model. Experiment 1 showed that some references were not grounded locally, but instead relied on previously mentioned goal plans, such as confirming by rows. Giving this concept to the model allows it to handle even extreme cases where, after establishing a goal plan of confirming by rows, individual objects to be confirmed are not referenced by messages but instead just confirmed in order. The output/input coordination principle of Garrod and Anderson (1987) comes naturally from the model when it uses previous successful goals of presenting and accepting information to produce templates for the creation of new utterances.

Experiment 2

Subjects

Fifty-nine Carnegie Mellon University undergraduates attempted the newer graph completion task. Twenty were paired and used the unrestricted interface, twenty-two were paired and used the restricted interface, and sixteen were told they would be paired with a partner but instead were paired with an ACT-R model.

Method

The procedure was the same as Experiment 1 except that multiple problems were used and subjects were told that there would be a cash incentive for completing as many problems as possible in the time given (one hour and forty minutes).

Results

Figure 1 shows average performance for pairs of subjects using the unrestricted interface, pairs of subjects using the restricted interface, and subjects using the restricted interface with the ACT-R model. Performance for all groups tends to improve with each problem. The model had some problems referencing objects that were not explicitly grounded by subjects and where no previous plan for referencing (e.g. sequentially by rows) had been established.

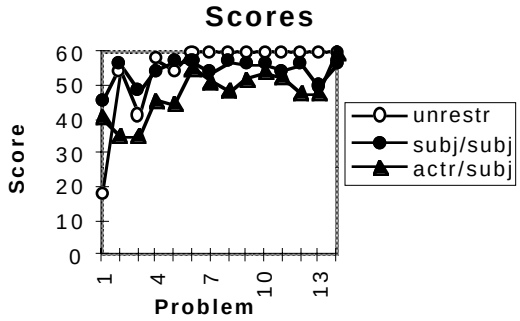

Figure 1: Performance

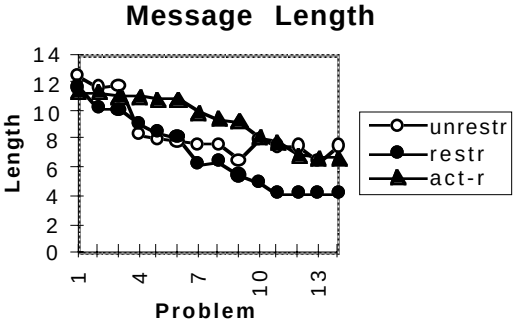

Figure 2: Length of messages

Figure 2 shows the decrease in message length by problem. Messages such as "The small thin red object is above our large fat blue object" in the first problem could be reduced to messages such as "red above blue" by the fourteenth problem. Message length in each group tends to decrease with each problem. In order to see if subjects are following the

output/input coordination principle by matching the message length of their partners, the within-pair difference in message length can be compared to the between-pair difference (Figure 3). The lower within-pair difference shows that subjects are using the same message length as their partner. The ACT-R model used syntactic information contained in previous goals to accept input to create templates for producing output, therefore following the output/input coordination principle. The result of this coordination is shown as decreased message length in Figure 2 and as a lower within-pair difference in message length in Figure 4.

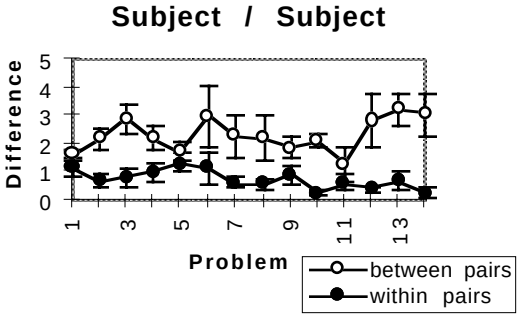

Figure 3: Difference in message length for subjects interacting with subjects

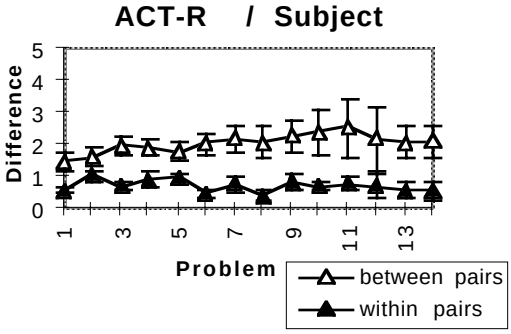

Figure 4: Difference in message length for ACT-R interacting with subjects

Conclusions

Results from these two experiments show that the ACT-R architecture is useful in examining the processes involved in communication. An incorrect prediction of a model in Experiment 1 showed where simple theories of reference grounding may need to take into account goal plans that have been previously grounded. A part of the ACT-R theory involving the retrieval of past goals provided mechanisms for the creation of common ground and the principle of output/input coordination.

One future direction could be an experiment manipulating whether or not a model interacting with a human subject followed the output/input coordination principle. The hypothesis would be that a non-coordinating model (one that uses a different lexicon for referring to objects and directions) would not allow the creation of a stable common ground, and this would result in longer comprehension and production times, fewer solved problems, and longer messages than when a coordinating model was used.

Acknowledgments

This research has been supported by grants N00014-96-I-0491 from the Office of Naval Research and SBR-94-21332 from the National Science Foundation.

References

Anderson, J. R., Corbett, A. T., Koedinger, K., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. Journal of the Learning Sciences, 4, 167-207.

Anderson, J. R., & Lebiere, C. (1998). The Atomic Components of Thought. Hillsdale, Mawhaw, NJ: Erlbaum.

Anderson, J. R., & Matessa, M. (1997). A production system theory of serial memory. Psychological Review, 104, 728-748.

Baker, M., & Lund, K. (1997). Promoting reflective interactions in a CSCL environment. Journal of Computer Assisted Learning**.** 13(3), 175-193.

Clark, H., & Schaefer, E. (1989). Contributing to Discourse. Cognitive Science, 13, 259-294.

Clark, H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. Cognition, 22, 1-39.

Core, M. G., & Allen, J. F. (1997). Coding dialogs with the DAMSL scheme. In Working Notes of the AAAI Fall Symposium on Communicative Action in Humans and Machines, Boston, MA. AAAI.

Garrod, S. & Anderson, A. (1987). Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. Cognition, 27, 181-218.

Lebiere, C., & Anderson, J. R. (1998). Cognitive Arithmetic. In J. R. Anderson & C. Lebiere (Eds.), The Atomic Components of Thought, (pp. 297-342). Hillsdale, Mawhaw, NJ: Erlbaum.

Levelt, W. J. M. (1982). Linearization in describing spatial networks. In S. Peters & E. Saarinen (Eds.), Processes, beliefs, and questions, (pp. 199-220). Dordrecht: D. Reidel.

Lovett, M. C., Reder, L. M, & Lebiere, C. (in press). Modeling working memory in a unified architecture. In A. Miyake & P. Shah (Eds.), Models of working memory: Mechanisms of active maintenance and executive control. New York: Cambridge University Press.

Novick, D., Hansen, B., & Lander, T. (1994). Letter-sequence dialogues. Technical Report CSE 94-007, Department of Computer Science and Engineering, Oregon Graduate Institute of Science & Technology.

Poesio, M., & Traum, D. (1998). Towards an Axiomatization of Dialogue Acts, In J. Hulstijn and A. Nijholt (Eds.), Proceedings of the Twente Workshop on the Formal Semantics and Pragmatics of Dialogues (13th Twente Workshop on Language Technology), Enschede, pp. 207--222.

Rich, C., & Sidner, C. L. (1998). COLLAGEN: A Collaboration Manager for Software Interface Agents. User Modeling and User-Adapted Interaction, Vol. 8, No. 3/4, 1998, pp. 315-350.

Traum, D., & Allen, J. (1994). Discourse Obligations in Dialogue Processing. In ACL94, Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics, pp. 1-8.