Minimizing the Length of Non-Mixed Initiative Dialogs

R. Bryce Inouye

Department of Computer Science
Duke University
Durham, NC 27708
rbi@cs.duke.edu

Abstract

Dialog participants in a non-mixed initiative dialogs, in which one participant asks questions exclusively and the other participant responds to those questions exclusively, can select actions that minimize the expected length of the dialog. The choice of question that minimizes the expected number of questions to be asked can be computed in polynomial time in some cases.

The polynomial-time solutions to special cases of the problem suggest a number of strategies for selecting dialog actions in the intractable general case. In a simulation involving 1000 dialog scenarios, an approximate solution using the most probable rule set/least probable question resulted in expected dialog length of 3.60 questions per dialog, as compared to 2.80 for the optimal case, and 5.05 for a randomly chosen strategy.

1 Introduction

Making optimal choices in unconstrained natural language dialogs may be impossible. The difficulty of defining consistent, meaningful criteria for which behavior can be optimized and the infinite number of possible actions that may be taken at any point in an unconstrained dialog present generally insurmountable obstacles to optimization.

Computing the optimal dialog action may be intractable even in a simple, highly constrained model of dialog with narrowly defined measures of success. This paper presents an analysis of the optimal behavior of a participant in non-mixed initiative dialogs, a restricted but important class of dialogs.

2 Non-mixed initiative dialogs

In recent years, dialog researchers have focused much attention on the study of *mixed-initiative* behaviors in natural language dialogs. In general, mixed initiative refers to the idea that control over the content and direction of a dialog may pass from one participant to another.¹ Cohen *et al.* (1998) provides a good overview of the various definitions of dialog initiative that have been proposed. Our work adopts a definition similar to Guinn (1999), who posits that initiative attaches to specific dialog goals.

This paper considers non-mixed-initiative dialogs, which we shall take to mean dialogs with the following characteristics:

1. The dialog has two *participants*, the *leader* and the *follower*, who are working cooperatively to achieve some mutually desired *dialog goal*.
2. The leader may *request* information from the follower, or may *inform* the follower that the dialog has succeeded or failed to achieve the dialog goal.

¹There is no generally accepted consensus as to how initiative should be defined.

3. The follower may only *inform* the leader of a fact in direct response to a *request* for information from the leader, or inform the leader that it cannot fulfill a particular request.

The model assumes the leader knows sets of questions q_{ij}

$$r_1 = \{q_{11}, q_{12}, \dots, q_{1,n_1}\}$$

$$r_2 = \{q_{21}, q_{22}, \dots, q_{2,n_2}\}$$

$\vdots$

$$r_m = \{q_{m1}, q_{m2}, \dots, q_{m,n_m}\}$$

such that if all questions in any one set r_i are answered successfully by the follower, the dialog goal will be satisfied. The sets will be referred to hereafter as **rule sets**. The leader's task is to find a rule set r whose constituent questions can all be successfully answered. The method is to choose a sequence of questions $q_{i_1,j_1}, q_{i_2,j_2}, \dots, q_{i_r,j_r}$ which will lead to its discovery.

For example, in a dialog in a customer service setting in which the leader attempts to locate the follower's account in a database, the leader might request the follower's name and account number, or might request the name and telephone number. The corresponding rule sets for such a dialog would be $\{askName, askPhoneNo\}$ and $\{askName, askAcctNo\}$.

One complicating factor in the leader's task is that a question $q_{i,j}$ in one rule set may occur in several other rule sets so that choosing to ask $q_{i,j}$ can have ramifications for several sets.

We assume that for every question $q_{i,j}$ the leader knows an associated probability $p_{i,j}$ that the follower has the knowledge necessary to answer $q_{i,j}$.² These probabilities enable us to compute an **expected length** for a dialog, measured by the number of questions asked by the leader. Our goal in selecting a sequence of questions will be to minimize the expected length of the dialog.

The probabilities may be estimated by aggregating the results from all interactions, or a more sophisticated individualized model might be maintained for each participant. Some examples of how these probabilities might be estimated can be

²In addition to modeling the follower's knowledge, these probabilities can also model aspects of the dialog system's performance, such as the recognition rate of an automatic speech recognizer.

found in (Conati et al., 2002; Zukerman and Albrecht, 2001).

Our model of dialog derives from rule-based theories of dialog structure, such as (Perrault and Allen, 1980; Grosz and Kraus, 1996; Lochbaum, 1998). In particular, this form of the problem models exactly the "missing axiom theory" of Smith and Hipp (1994; 1995) which proposes that dialog is aimed at proving the top-level goal in a theorem-proving tree and "missing axioms" in the proof provide motivation for interactions with the dialog partner. The rule sets r_i are sets of missing axioms that are sufficient to complete the proof of the top-level goal.

Our format is quite general and can model other dialog systems as well. For example, a dialog system that is organized as a decision tree with a question at the root, with additional questions at successor branches, can be modeled by our format.

As an example, suppose we have top-level goal t_1 and these rules to prove it:

$(t_2 \text{ AND } q_{12}) \text{ implies } t_1$

$(q_{11} \text{ OR } q_{21}) \text{ implies } t_2.$

The corresponding rule sets are

$$r_1 = \{q_{11}, q_{12}\}$$

$$r_2 = \{q_{21}, q_{12}\}.$$

If all of the questions in either r_1 or r_2 are satisfied, t_1 will be proven. If we have values for the probabilities p_{11}, p_{12} , and p_{21} , we can design an optimum ordering of the questions to minimize the expected length of dialogs. Thus if p_{11} is much smaller than p_{21} , we would ask q_{21} before asking q_{11} . The reader might try to decide when q_{12} should be asked before any other questions in order to minimize the expected length of dialogs.

The rest of the paper examines how the leader can select the questions which minimize the overall expected length of the dialog, as measured by the number of questions asked. Each question-response pair is considered to contribute equally to the length. Sections 3, 4, and 5 describe polynomial-time algorithms for finding the optimum order of questions in three special instances of the question ordering optimization problem. Section 6 gives a polynomial-time method to approximate optimum behavior in the general case of n rule sets which may have many common questions.

3 Case: One rule set

Many dialog tasks can be modeled with a single rule set $r = \{q_1, q_2, \dots, q_n\}$. For example, a leader might ask the follower to supply values for each field in a form. Here the optimum strategy is to ask the questions first that have the least probability of being successfully answered.

Theorem 1. *Given a rule set $r = \{q_1, \dots, q_n\}$, asking the questions in the order of their probability of success (least first) results in the minimum expected dialog length; that is, for $i = 1, \dots, n-1, p_i \leq p_{i+1}$ where p_i is the probability that the follower will answer question q_i successfully.*

A formal proof is available in a longer version of this paper. Informally, we have two cases; the first assumes that all questions q_i are answered successfully, leading to a dialog length of n , since n questions will be asked and then answered. The second case assumes that some q_i will not be answered successfully. The expected length increases as the probabilities of success of the questions asked increases. However, the expected length does not depend on the probability of success for the last question asked, since no questions follow it regardless of the outcome. Therefore, the question with the greatest probability of success appears at the end of the optimal ordering. Similarly, we can show that given the last question in the ordering, the expected length does not depend upon the probability of the second to last question in the ordering, and so on until all questions have been placed in the proper position. The optimal ordering is in order of increasing probability of success.

4 Case: Two independent rule sets

We now consider a dialog scenario in which the leader has two rule sets for completing the dialog task.

Definition 4.1. *Two rule sets r_1 and r_2 are independent if $r_1 \cap r_2 = \emptyset$. If $r_1 \cap r_2$ is non-empty, then the members of $r_1 \cap r_2$ are said to be **common** to r_1 and r_2 . A question q is **unique** to rule set r if $q \in r$ and for all $r' \neq r, q \notin r'$*

In a dialog scenario in which the leader has

multiple, mutually independent rule sets for accomplishing the dialog goal, the result of asking a question contained in one rule set has no effect on the success or failure of the other rule sets known by the leader. Also, it can be shown that if the leader makes optimal decisions at each turn in the dialog, once the leader begins asking questions belonging to one rule set, it should continue to ask questions from the same rule set until the rule set either succeeds or fails. The problem of selecting the question that minimizes the expected dialog length $E(l)$ becomes the problem of selecting which rule set should be used first by the leader. Once the rule set has been selected, Theorem 1 shows how to select a question from the selected rule set that minimizes $E(l)$.

By expected dialog length, we mean the usual definition of expectation

$$E(l) = \sum_x (\text{prob. of } x)(\text{length of } x).$$

Thus, to calculate the expected length of a dialog, we must be able to enumerate all of the possible outcomes of that dialog, along with the probability of that outcome occurring, and the length associated with that outcome.

Before we show how the leader should decide which rule set it should use first, we introduce some notation.

The **expected length in case of failure** for an ordering $o = q_1, \dots, q_n$ of the questions of a rule set r is the expected length of the dialog that would result if r were the only rule set available to the leader, the leader asked questions in the order given by o , and one of the questions in r failed. The expected length in case of failure is

$$\frac{1}{1 - \prod_{i=1}^n p_i} \sum_{i=1}^n i \left(\prod_{j=1}^{i-1} p_j \right) (1 - p_i).$$

The factor $\frac{1}{1 - \prod_{k=1}^n p(a_k)}$ is a scaling factor that accounts for the fact that we are counting only cases in which the dialog fails. We will let f_i represent the minimum expected length in case of failure for rule set r_i , obtained by ordering the questions of r_i by increasing probability of success, as per Theorem 1.

The **probability of success** s of a rule set $r = \{q_1, q_2, \dots, q_n\}$ is $\prod_{i=1}^n p_i$. The definition

of probability of success of a rule set assumes that the probabilities of success for individual questions are mutually independent.

Theorem 2. *Let $\mathcal{S} = \{r_1, r_2\}$ be the set of mutually independent rule sets available to the leader for accomplishing the dialog goal. For a rule set r_i in $\mathcal{S}$, let s_i be the probability of success of r_i , n_i be the number of questions in r_i , and f_i be the minimum expected length in case of failure. To minimize the expected length of the dialog, the leader should select the question with the least probability of success from the rule set r_i with the least value of $n_i + f_i(\frac{1}{s_i} - 1)$.*

Proof: If the leader uses questions from r_1 first, the expected dialog length $E(l_1)$ is

$$s_1 n_1 + (1 - s_1) s_2 (f_1 + n_2) + (1 - s_1) (1 - s_2) (f_1 + f_2)$$

The first term, $s_1 n_1$, is the probability of success for r_1 times the length of r_1 . The second term, $(1 - s_1) s_2 (f_1 + n_2)$, is the probability that r_1 will fail and r_2 will succeed times the length of that dialog. The third term, $(1 - s_1) (1 - s_2) (f_1 + f_2)$, is the probability that both r_1 and r_2 fail times the associated length. We can multiply out and rearrange terms to get

$$\begin{aligned} E(l_1) &= s_1 n_1 + (1 - s_1)(f_1 + s_2 n_2 + (1 - s_2) f_2) \\ &= s_1 n_1 + f_1 - s_1 f_1 + s_2 n_2 - s_1 s_2 n_2 + f_2 \\ &\quad - s_1 f_2 - s_2 f_2 + s_1 s_2 f_2 \end{aligned}$$

If the leader uses questions from r_2 first, $E(l_2)$ is

$$s_2 n_2 + f_2 - s_2 f_2 + s_1 n_1 - s_2 s_1 n_1 + f_1 - s_2 f_1 - s_1 f_1 + s_2 s_1 f_1$$

Comparing $E(l_1)$ and $E(l_2)$, and eliminating any common terms, we find that (r_1, r_2) is the correct ordering if

$$\begin{aligned} E(l_1) &< E(l_2) \\ -s_1 s_2 n_2 - s_1 f_2 + s_1 s_2 f_2 &< -s_2 s_1 n_1 - s_2 f_1 + s_2 s_1 f_1 \\ s_1 (-s_2 n_2 - f_2 + s_2 f_2) &< s_2 (-s_1 n_1 - f_1 + s_1 f_1) \\ s_1 (-f_2 + s_2 (f_2 - n_2)) &< s_2 (-f_1 + s_1 (f_1 - n_1)) \\ \frac{-f_2 + s_2 (f_2 - n_2)}{s_2} &< \frac{-f_1 + s_1 (f_1 - n_1)}{s_1} \\ n_2 + f_2 \left(\frac{1}{s_2} - 1 \right) &> n_1 + f_1 \left(\frac{1}{s_1} - 1 \right) \end{aligned}$$

Thus, if the above inequality holds, then $E(l_1) < E(l_2)$, and the leader should ask questions from r_1 first. Otherwise, $E(l_2) \leq E(l_1)$, and the leader should ask questions from r_2 first.

We conjecture that in the general case of m mutually independent rule sets, the proper ordering of

rule sets is obtained by calculating $n_i + f_i(\frac{1}{s_i} - 1)$ for each rule set r_i , and sorting the rule sets by those values. Preliminary experimental evidence supports this conjecture, but no formal proof has been derived yet.

Note that calculating s and f for each rule set takes polynomial time, as does sorting the rule sets into their proper order and sorting the questions within each rule set. Thus the solution can be obtained in polynomial time.

As an example, consider the rule sets $r_1 = \{q_{11}, q_{12}\}$ and $r_2 = \{q_{21}, q_{22}\}$. Suppose that we assign $p_{11} = 0.4, p_{12} = 0.6, p_{21} = 0.3$, and $p_{22} = 0.8$. In this case, $n = 2$ and $s = 0.24$ are the same for both rule sets. However, $f_1 = 1.21$ and $f_2 = 1.08$, so evaluating $n_i + f_i(\frac{1}{s_i} - 1)$ for both rule sets, we discover that asking questions from r_2 first results in the minimum expected dialog length.

5 Case: Two rule sets, one common question

We now examine the simplest case in which the rule sets are not mutually independent: the leader has two rule sets r_1 and r_2 , and $r_1 \cap r_2 = \{q_c\}$. In this section, we will use $E(l_{r_i})$ to denote the minimum expected length of the dialog (computed using Theorem 1) resulting from the leader using only r_i to accomplish the dialog task. The notation $E(l_{r'_i})$ will denote the minimum expected length of the dialog resulting from the leader using only the rule set $r_i - \{q_c\}$ to accomplish the dialog task. For example, a rule set $r_1 = \{q_1, q_2, q_c\}$ with $p_1 = 0.3, p_2 = 0.5$ and $p_c = 0.9$, has $E(l_{r_1}) = 1.45$ and $E(l_{r'_1}) = 1.3$.

Theorem 3. *Given rule sets $r_1 = \{q_c, q_1, \dots, q_n\}$ and r_2 , such that $r_1 \cap r_2 = \{q_c\}$, if the leader asks questions from r_1 until r_1 either succeeds or fails before asking any questions unique to r_2 , then the ordering of questions of r_1 that results in the minimum expected dialog length is given by ordering the questions q_i by increasing v_i , where*

$$v_i = \begin{cases} \frac{p_c + p_c E(l_{r'_2}) - E(l_{r_2})}{1 + p_c E(l_{r'_2}) - E(l_{r_2})}, & q_i = q_c \\ p_i & \text{otherwise.} \end{cases}$$

The proof is in two parts. First we show that the questions unique to r_1 should be ordered by

$$\begin{aligned}
E(l) &= 1 + \sum_{i=1}^a \prod_{j=1}^i p_j + p_c \sum_{i=a}^{n-1} \prod_{j=1}^i p_j + \\
&\quad E(l_{r_2}) \left(1 - \prod_{i=1}^a p_i\right) + p_c E(l_{r'_2}) \left(\prod_{i=1}^a p_i - \prod_{i=1}^n p_i\right)
\end{aligned}$$

Figure 1: A general expression for the expected dialog length for the dialog scenario described in section 5. The questions of r_1 are asked in the arbitrary order $q_1, \dots, q_a, q_c, q_{a+1}, q_n$, where q_c is the question common to r_1 and r_2 . $E(l_2)$ and $E(l'_2)$ are defined in Section 5.

increasing probability of success given that the position of q_c is fixed. Then we show that given the correct ordering of unique questions of r_1 , q_c should appear in that ordering at the position where $\frac{p_c + p_c E(l_{r'_2}) - E(l_{r_2})}{1 + p_c E(l_{r'_2}) - E(l_{r_2})}$ falls in the corresponding sequence of questions probabilities of success. Space considerations preclude a complete listing of the proof, but an outline follows.

Figure 1 shows an expression for the expected dialog length for a dialog in which the leader asks questions from r_1 until r_1 either succeeds or fails before asking any questions unique to r_2 . The expression assumes an arbitrary ordering $o = q_1, \dots, q_a, q_c, q_{a+1}, \dots, q_n$. Note that if a question occurring before q_c fails, the rest of the dialog has a minimum expected length $E(l_{r_2})$. If q_c fails, the dialog terminates. If a question occurring after q_c fails, the rest of the dialog has minimum expected length $E(l_{r'_2})$.

If we fix the position of q_c , we can show that the questions unique to r_1 must be ordered by increasing probability of success in the optimal ordering. The proof proceeds by showing that switching the positions of any two unique questions q_a and q_b in an arbitrary ordering of the questions of r_1 , where q_a occurs before q_b in the original ordering, the expected length for the new ordering is less than the expected length for the original ordering if and only if $p_b < p_a$.

After showing that the unique questions of r_1 must be ordered by increasing probability of success in the optimal ordering, we must then show how to find the position of q_c in the optimal ordering. We say that q_c occurs at position a in or-

dering o if q_c immediately follows q_a in the ordering. $E(l_a)$ is the expected length for the ordering with q_c at position a . We can show that if $E(l_a) < E(l_{a+1})$ then $\frac{p_c - E(l_{r_2}) + p_c E(l_{r'_2})}{1 - E(l_{r_2}) + p_c E(l_{r'_2})} < p_{a+1}$ by a process similar to that used in the proof of Theorem 2. Since the unique questions in r_1 are ordered by increasing probability of success, finding the optimal position of the common question q in the ordering of the questions of r_1 corresponds to the problem of finding where the value of $\frac{p_c - E(l_{r_2}) + p_c E(l_{r'_2})}{1 - E(l_{r_2}) + p_c E(l_{r'_2})}$ falls in the sorted list of probabilities of success of the unique questions of r_1 . If the value immediately precedes the value of p_i in the list, then the common question should immediately precede q_i in the optimal ordering of questions of r_1 .

Theorem 3 provides a method for obtaining the optimal ordering of questions in r_1 , given that r_1 is selected first by the leader. The leader can use the same method to determine the optimal ordering of the questions of r_2 if r_2 is selected first. The two optimal orderings give rise to two different expected dialog lengths; the leader should select the rule set and ordering that leads to the minimal expected dialog length. The calculation can be done in polynomial time.

6 Approximate solutions in the general case

Specific instances of the optimization problem can be solved in polynomial time, but the general case has worst-case complexity that is exponential in the number of questions. To approximate the optimal solution, we can use some of the insights gained from the analysis of the special cases to generate methods for selecting a rule set, and selecting a question from the chosen rule set. Theorem 1 says that if there is only one rule set available, then the least probable question should be asked first. We can also observe that if the dialog succeeds, then in general, we would like to minimize the number of rule sets that must be tried before succeeding. Combining these two observations produces a policy of selecting the question with the minimal probability of success from the rule set with the maximal probability of success.

Method	Avg. length
Optimal	2.80
Most prob. rule set/least prob. question	3.60
Most prob. rule set/random question	4.26
Random rule set/most prob. question	4.26
Random rule set/random question	5.05

Table 1: Average expected dialog length (measured in number of leader questions) for the optimal case and several simple approximation methods over 1000 dialog scenarios. Each scenario consisted of 6 rule sets of 2 to 5 questions each, created from a pool of 9 different questions.

We tested this policy by generating 1000 dialog scenarios. First, a pool of nine questions with randomly assigned probabilities of success was generated. Six rule sets were created using these nine questions, each containing between two and five questions. The number of questions in each rule set was selected randomly, with each value being equally probable. We then calculated the expected length of the dialog that would result if the leader were to select questions according to the following five schemes:

1. Optimal
2. Most probable rule set, least probable question
3. Random rule set, least probable question
4. Most probable rule set, random question
5. Random rule set, random question.

The results are summarized in Table 1.

7 Further Research

We intend to discover other special cases for which polynomial time solutions exist, and investigate other methods for approximating the optimal solution. With a larger library of studied special cases, even if polynomial time solutions do not exist for such cases, heuristics designed for use in special cases may provide better performance.

Another extension to this research is to extend the information model maintained by the leader to allow the probabilities returned by the model to be non-independent.

8 Conclusions

Optimizing the behavior of dialog participants can be a complex task even in restricted and specialized environments. For the case of non-mixed ini-

tiative dialogs, selecting dialog actions that minimize the overall expected dialog length is a non-trivial problem, but one which has some solutions in certain instances. A study of the characteristics of the problem can yield insights that lead to the development of methods that allow a dialog participant to perform in a principled way in the face of intractable complexity.

Acknowledgments

This work was supported by a grant from SAIC, and from the US Defense Advanced Research Projects Agency.

References

- Robin Cohen, Coralee Allaby, Christian Cumbaa, Mark Fitzgerald, Kinson Ho, Bowen Hui, Celine Latulipe, Fletcher Lu, Nancy Moussa, David Pooley, Alex Qian, and Saheem Siddiqi. 1998. What is initiative? *User Modeling and User-Adapted Interaction*, 8(3-4):171–214.
- C. Conati, A. Gerntner, and K. Vanlehn. 2002. Using bayesian networks to manage uncertainty in user modeling. *User Modeling and User-Adapted Interaction*, 12(4):371–417.
- Barbara Grosz and Sarit Kraus. 1996. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357.
- Curry I. Guinn. 1999. An analysis of initiative selection in collaborative task-oriented discourse. *User Modeling and User-adapted Interaction*, 8(3-4):255–314.
- K. Lochbaum. 1998. A collaborative planning model of intentional structure. *Computational Linguistics*, 24(4):525–572.
- C. R. Perrault and J. F. Allen. 1980. A plan-based analysis of indirect speech acts. *Computational Linguistics*, 6(3-4):167–182.
- Ronnie W. Smith and D. Richard Hipp. 1994. *Spoken Natural Language Dialog Systems: A Practical Approach*. Oxford UP, New York.
- Ronnie W. Smith and D. Richard Hipp. 1995. An architecture for voice dialog systems based on prolog-style theorem proving. *Computational Linguistics*, 21(3):281–320.
- I. Zukerman and D. Albrecht. 2001. Predictive statistical models for user modeling. *User Modeling and User-Adapted Interaction*, 11(1-2):5–18.