University of Waterloo

Faculty of Mathematics

Computer-Generated Dialogs

WiseUncle Inc.
System Architect
Toronto, Ontario

prepared by

Jordan Naftolin
98096908
2B Computer Science
May 5, 2001

Table of Contents

Executive Summary

This report will examine computer-generated dialogs to determine whether they are a practical alternative to human customer representatives. The examination will analyze computer-generated dialogs from a technical perspective as well as from a user's perspective.

To ascertain whether computer-generated dialogs are currently technically feasible, this report investigates the process of determining the next dialog question to ask as well as several methods for receiving user input. This report then analyzes the overall benefits and drawbacks of computer-generated dialogs from a user point of view.

It is concluded that computer-generated dialogs are currently technically feasible and that their benefits make them a practical replacement for human customer representatives. As a result, it is recommended that companies look into replacing their customer representatives with computer-generated dialogs.

Introduction

"All of our representatives are currently busy assisting other customers at the moment. Please stay on the line for the next available representative." Most of us cringe at the sound of these words and prefer never having to hear them again. Computer-generated dialogs may soon offer us a solution by becoming a viable alternative to human customer representatives. Though the thought of conversing with a computer may seem unnatural, computer-generated dialogs offer many benefits that challenge the shortcomings that we currently experience with human customer representatives. This report will examine the current technical feasibility of computer-generated dialogs by looking at the various ways that questions are generated and the different methods in which user responses are accepted. The report will then analyze the overall benefits and drawbacks that computer-generated dialogs bring to the user. The report will conclude with a recommendation as to whether or not companies should adopt computer-generated dialogs as a realistic replacement to human customer representatives.

Analysis

A practical dialog-generating system can be considered to have two main requirements:

1) The system must have an intelligent mechanism to determine the next question to be asked.

2) The system must have a suitable method for receiving user input.

Both requirements are extremely important in order to ensure that the user accepts the dialog as a sufficient replacement to the human conversations that he or she is accustomed to.

Determining the Next Question

The first requirement implies that the system should decide the next question such that the conversation flows coherently and naturally. There are several methods available for determining the next question in a computer-generated dialog. Two such methods are the 'static' approach and the 'rule-based' approach. These two methods ultimately accomplish the same goal, however their respective ability to manage the coherence of the conversation differs significantly.

The static approach is equivalent to following a dialog through a top-to-bottom flowchart. Each element in the flowchart represents a question. Each question's answers represent different paths down the flowchart. Although this system works and is relatively easy to implement, it can quickly become very difficult to manage. The process of adding, moving, or removing questions is arduous and complicated, especially when trying to maintain the coherence of the conversation. For example, if a question is to be added to a particular path in the flowchart, then each question in that path must be analyzed to ensure

that the new question will remain coherent in the path's context. This becomes difficult when the individual paths become very large. Removing a question is also difficult since the question that is being removed could potentially have many sub-paths that must also be removed, or become integrated with the remaining questions. This involves looking through every question of the affected paths to again ensure coherence. Modifying when questions are asked is the most complicated, as questions must first be removed from the flowchart and then re-added.

The rule-based approach has several advantages over the static approach, however it is a more complex system to debug. In the rule-based approach, the system contains a set of facts and a set of rules. Each rule is comprised of conditions and actions. Each condition indicates the required state of a particular fact in order for the condition to be satisfied. When all of the conditions within a rule are satisfied, the rule's actions get triggered. A triggered action can either change the state of another fact, or schedule a particular question to be asked. When a user responds to a question, the user's answer causes the states of certain facts to change. This causes rules to be triggered, which changes the states of other facts, which in turn triggers more rules, thus starting a chain reaction. When the chain reaction stops, there are three possibilities:

1) One question will have been scheduled by a rule's action during the chain reaction. In this situation the question that has been scheduled will be presented to the user.

2) More than one question will have been scheduled during the chain reaction. In this case, conflict resolution rules are used to determine the most appropriate question to present to the user amongst the questions that were scheduled.

3) No questions will have been scheduled during the chain reaction. This implies that there are no more appropriate questions to be asked.

It is much easier to add or remove questions in the rule-based approach than in the static approach. To add a question, a rule is created whose action schedules the new question, and whose conditions indicate when the new question should be asked. Removing a question simply involves removing the rules whose actions schedule the particular question. It is also easy to change when a question is asked by simply modifying the conditions. However, the rule-based system can be difficult to debug, as it is not as straight forward as the static approach. When using the rule-based approach, it becomes critical to have a data management tool capable of querying how and when questions are scheduled in order to assist in debugging the dialog.

Both solutions are capable of properly determining the next question to ask and are therefore both sufficient to use in computer-generated dialogs. The main differences lie in the debugging complexity and in question management. To ensure the coherence of the dialog, it is advised that the rule-based system be used when the number of questions being managed is large and when the questions are frequently changing. However, when the number of questions being asked is small and the question structure is going to remain unchanged, then the static approach is most fitting due to its simpler design.

Receiving User Input

The second requirement of a practical dialog-generating system implies that the method used to receive user input is congenial to the user. People are currently comfortable conversing with humans using natural language speech. In order to maintain familiarity for the user, it is ideal for computer-generated dialogs to receive user input in a similar fashion. Companies such as IBM and Nuance have made considerable progress in developing voice recognition technology that is capable of listening to human speech and identifying which words were spoken. However, although we are reasonably close to identifying natural language speech, we are far from being able to accurately interpret it.

The latter requires that the system has complete knowledge of the specified domain. Consider, for example, a system that is intended to determine a user's favourite colours. It is practical to imagine a system capable of recognizing and interpreting input such as "red and white", but suppose a user says "the colours of the Canadian flag". Whereas a human representative would (hopefully) be able to identify that the user is referring to red and white, a computer system would be required to know the colours of the Canadian flag, as well as all other colour-identifiable objects. This is clearly not an easy task and is one of the primary reasons that natural language recognition is far from being perfect. As a result of being unable to fully interpret user input, natural language recognition is currently not a suitable method for receiving user input.

An alternative to natural-language recognition is to provide the user with a set of possible answers to select from. Although this limits what the user is permitted to input, the system is capable of accurately interpreting the user's input. The limited number of possible answers can be seen as a disadvantage, however users will likely not feel limited as long as the answers span all of the options that the user could have selected using natural language. Although "the colours of the Canadian flag" may not be an option, "red and white" might be, thus allowing the user to select their favourite colours while assuring that their input is properly interpreted. As long as the options provided to the user for each question are appropriate, this alternative approach is a suitable method for receiving user input.

Overall User Benefits and Drawbacks

An important aspect when measuring the effectiveness of a computer-generated dialog is the overall user experience. If the user does not feel comfortable conversing with the computer, or if he or she lacks control of the conversation, then the dialog will not be well accepted. In order to meet these requirements, a computer-generated dialog should not

appear as though it was computer-generated. Though this is clearly a difficult task to accomplish, computers are becoming more and more human-like as technology continues to evolve. However, at the present time, computer-generated dialogs still convey an artificial feeling that some users may find discomforting, and this can be seen as a big drawback.

Despite not having a completely authentic feeling, computer-generated dialogs do offer the user many advantages. In current human customer representative settings, users must often wait on hold before being able to speak with a customer representative. As well, not all customer representatives have equal knowledge and experience, and users can be disappointed with less competent customer representatives. However, because multiple instances of a computer-generated dialog system can be running simultaneously, users never have to wait and are always guaranteed equal 'competence' when using a computer-generated dialog. Computer systems are also capable of accessing huge amounts of up-to-date information very quickly, thus offering benefits that humans cannot.

Though users may at first be uncomfortable interacting with computer-generated dialogs, these dialogs will most likely become ubiquitous over time due to the benefits that they have to offer. There have already been many examples in which computers have acted as an alternative to humans. The Automatic Teller Machine, for example, has already become extremely popular due to its indisputable convenience.

Conclusions

There are several viable methods for both determining the next question to ask and receiving user input. As a result, it is currently technically feasible to implement computer-generated dialogs such that they meet the requirements of a practical dialog system. Though computer-generated dialogs still have a sense of being artificial, the benefits and convenience that they offer should convince most users to accept them as a replacement for human customer representatives. It is certain that most people would much rather converse with a competent computer agent than wait on hold for the next available human representative.

Recommendations

It is recommended that companies look into using computer-generated dialogs as an alternative to human customer representatives. The justification is that computer-generated dialogs are less expensive to the company in the long run, and customers will not be inconvenienced by delays and/or possibly sub-par customer representatives.