

Conversation-Based Search

(How to Build an Intelligent Corporate Memory)

Roger C. Schank¹

Introduction

Conversation is such a natural part of everyday life that we often fail to reflect on its properties. We just engage in it without explicitly understanding the rules that we use unconsciously every time we have a conversation. Now this doesn't matter much when we are talking to people, because people all have pretty much the same ideas about conversation and follow the same rules for conducting them. But when you try to talk to a computer, everything changes.

¹ CEO Socratic Arts

But why would you try to talk to a computer? Who does that exactly?

Actually everyone does that in a sense, but because the computer doesn't know the rules of conversation and because the computer has no actual desire to engage in a conversation and because, in a deep sense, the computer has nothing to say, human-computer conversation doesn't go so well.

The beginning and end of the problem is one of understanding the nature of human memory. When people engage in a conversation, they are, quite frequently, simply telling a story that have told before, which is being retrieved from their memory as they speak. And when they listen to someone else speak in a conversation, they are trying to match what that person is saying with their own stories to ready themselves to have something to say or simply just because they are excited by the match in stories between the speaker and themselves.

But when it comes to computers all this changes. We don't expect computers to have stories to tell. Instead we have come to believe that if a computer has stories we will have to search for them.

This disconnect between the rules of everyday human conversation and the practice of key word search that we have all gotten used to, explains why corporations and

other large institutions are having such a hard time gathering and preserving what they collectively know. They should be making what they have learned over the years available to their employees and possibly to the outside world. But, they all seem to think what they know must be contained within the form of documents because that is what Google searches through and that is what filing cabinets and manila folders have always contained.

The desktop metaphor and the document-based mentality that went with that metaphor is dead. Now, lets go on to wondering what should replace it. I suggest here that it must be something based on the rules of everyday human conversation.

Talking with a Computer

We don't exactly have conversations with computers. We point and click because early designers of computers using CRT terminals lived in a physical world where everyone had a desk, and those desks had desktops with papers on them. Everyone had file cabinets filled with manila folders to put away things no longer best stored on their desktop. They had trash cans when they were done with stuff. Point and click made sense as a way to communicate one's needs to a computer with respect to all these things in the age of desks.

Before those heady point and click days, people communicated with computers by typing commands. They entered a command in a line editor, and line by line the computer responded to the command.

Before those days, people communicated with computers by the insertion of a whole slew of commands submitted all at once. Often, these were embodied in a deck of punch cards.

For many years, there has been a dream, held by AI researchers and others, of a conversational computer, one where a natural dialogue was the means of communication. This dream was subverted by Google in a sense. Google's success at responding to an input of keywords made users think that they were having a conversation with a computer. Often Google users type in complete English sentences because Google has made them feel as if it is indeed a conversational machine. The fact that the conversation is bizarre from a human point of view doesn't seem to matter. No one asks a friend a question and is happy when he gets back 1000 search results any one of which might or might not answer his question. And certainly no one is happy that for the next question he might have, the computer in no way remembers the context set up by the last one.

So, computers have never really been conversational devices. But what if they were? What do we know about

human conversation that would inform us about how to design a conversational computer?

I will take as an example a typical cocktail party conversation between two people who have just met. What can we say is typically true of such a conversation? Here are some rules that participants in such a conversation typically follow:

1. They take turns, one talks a while and then the other talks a while.
2. A speaker must answer a question posed by the other speaker.
3. If there is no question to respond to, the next speaker must say something that relates in some way to what the last speaker has said.
4. Responses must either make the same point (“Something like that happened to me”) or make an opposing point (“No, its not that way, its this way”) or else they must match if no point has been made at all (“Your wife likes to do this, well my wife likes to do this other thing.”)
5. Responses cannot be too long. If the intended response is estimated to be more than two minutes long, the speaker must warn the listener in some way (“This may take a little explanation...”)
6. Responses are in the form of previously told stories, ones that the speaker has told many time before and one that can be told quickly and efficiently.

7. A story must have a point. If it does not have a point it must be a non-complex response to the previous remark.

Considering the above rules, rules that are known implicitly by all adult speakers, it is clear why computers are not good conversationalists:

1. Computers have no stories to tell.
2. Computers do not understand the stories being told to them.
3. Even if they did understand the stories being told to them, computers would have no way of matching the points of those stories to the points of the stories they might be able to tell.

The first point is a little strange. Computers do have stories to tell in the sense that any Wikipedia article or New York Times article is a story to which the computer has access. But, the computer doesn't know the stories it has. It can find them through key words and statistics but when it retrieves a story it isn't exactly telling it because it has no idea what it has just said.

Now, let's consider how we might change all this.

One issue would be giving the computer a means of retrieving stories by something other than key words, a

method more in accord with the notion of a point of a story.

A second issue would be giving the computer stories to tell.

The third issue would be giving computers the ability to know when a situation warranted the telling of one its stories (and which one.)

People tell stories in a variety of ways. Novels, newspapers, movies, and TV shows are types of stories. But people have always held one kind of story as the real thing, and that is the face to face, spoken, story. Our conversations use this normal everyday notion of story so often during our daily lives that we often fail to notice its significance. If we were to give computers stories to tell, they would naturally have to start with these kinds of stories. One reason why that would be important to do is that most of the wisdom that any individual possesses is contained in the form of these kind of stories, one that he or she likes telling as often as he can. Typically such stories are funny or emotional or interesting in some way and they almost always serve to define some important piece of knowledge or experience of the teller.

Capturing Knowledge

When videotape equipment became easy to own and use, many large companies began taping their experts. They figured that they could preserve their corporate memory in that way. Today, in all the cases that I know about, those tapes lie on a shelf somewhere. No one ever looks at them. What went wrong?

The main problem was that television confused the makers of these videos in every case I have seen. They were all done as if they were doing an hour long TV show with an interviewer asking questions and the interviewee looking at the interviewer and responding. Whether this makes good hour long TV I cannot say, but it is not how people talk to each other.

But, if we are interested in capturing corporate memory, we need to re-think. Preserving a corporate memory on video would have to be done entirely differently. Why?

1. People are interested in hearing an expert's stories when they have a problem that that story addresses.
2. When someone wants an expert to tell a story they want that story to be short and to the point.
3. They also want that story to be related to the problem they have.
4. They also want the expert to be available for follow up questions that they might have.
5. They also want other experts to be available for follow up questions or even rebuttals. The conversation

needn't be with only one person from the corporate memory.

Therefore it seems obvious, that the right way to build a corporate memory, or any other usable knowledge base, is to record experts telling short (no more than 2 minutes long, preferably one minute long) stories that make a clear point in an interesting way. These should be told with the expert facing the camera directly so that he seems to be in a conversation with the user. And, after the expert's story has been heard, it should be obvious to the user how to get more relevant stories.

Doing this requires a massive effort at interviewing, collecting, and editing the stories of experts must be undertaken. If the finished system has a few hundred stories in it, it will not be able to sustain an interesting conversation about multiple topics nor will it be able to express multiple points of view.. The finished system must be able to surprise the user and must have a great range of wisdom available. This requires thousands of stories.

Indexing Stories

Every story collected must be indexed. Indexing is the trick to making such a system work properly. The indexing is the intelligence in the system.² The indexing cannot be

² I first described this in *Dynamic Memory* (Cambridge Univ Press, 1981)

according to key words, or combinations of key words, nor can it be alphabetical, organized by speaker, or topic, or any other of the standard indexing schemes that have been in use for years for books and articles. The indexing scheme must be the one that has been in use for millennia, namely the one that humans use, unconsciously, to store and find stories in their own minds with no obvious effort. The index must be identical to what a person would have done when they exclaim after hearing someone say something: *that reminds me of a story!*

What indexing scheme do people use? We can learn about human indexing by looking at how reminding works. When people get reminded of a story by something someone has just said, it is because there is a match between how they indexed what they heard and how they had previously indexed the story they now intend to tell.

We can gain insight into what these indices look like by examining how any story on is reminded of matches the incoming story. We have worked on this problem for years and the answer is that people unconsciously attempt to determine certain the answer to certain questions whenever they are listening to each other. These questions are:

1. what is the actor trying to accomplish?
2. what method is he using to accomplish it?

3. what might be preventing his accomplishing it?
4. what do we know about the particular circumstances?
5. What wisdom is he sharing about what happened?
6. Is there are lesson to be learned from what happened?

As listeners, we don't necessarily realize we are asking these questions. As retrievers of our knowledge from own our memory we don't realize that we use the answers previously generated to these questions as away of labeling what we have previously experienced. But this is how it happens.

Our goal is to make computers into conversational devices that can retrieve stories from a properly indexed memory. The computer must be capable of using the index to make full or partial matches with other stories indexing in the same or similar ways.

Any corporate memory that could be used to help people make decisions would have to have indexed the stories from the company's experts in such a way that they matched the decisions being considered at any given time. This means that, in essence the stories from the corporate memory would find you because it knew what you were doing. This is the very opposite of search.

The User's Role

In order for stories to find you, the computer must know who you are and what you are doing. Therefore, a user must be able to make it clear to the computer who he is, what he is doing at the moment, and what his concerns are, in the same language or in some method that uses the same language, that the indexing scheme employed by the computer uses.

When we tell a story to a friend, we know where that friend is, we know what his concerns are, we know what he is doing in life, and we know what stories we have already told him. We choose to tell a story because we understand that this listener might benefit from it, according to the model we have constructed of that listener. To be useful to people, the computer also needs to know to whom it is speaking. When Expedia or any other travel site, suggest flights to me, it acts as if it knows nothing about me and assumes that my biggest concern is the air fare (since it orders possible flights by price.) But it could know a lot about me, it just doesn't, and it could propose flights according to my concerns which are more about seat comfort and food and ease of transfer, then they are about price.

Viewing flight possibilities as yet another kind of story, Expedia could choose to tell a story that is relevant to me. But this would involve a different kind of technology than the search-based one that we have all gotten used to. The next generation of storytelling computers will not only have stories to tell, they will also have to know who they are

talking to and have a idea about what stories that user already knows. In other words, a conversation-based machine would behave very differently than a search-based machine.

EXTRA

We have done all these things in EXTRA, a program that uses cognitive indexing to retrieve stories just in time. These stories can find a user who needs them if it knows what that user is doing. EXTRA can also serve as a system that a user can navigate to find stories of interest.

The computer's job in EXTRA is to easily integrate new stories and self-organize its memory to make those new stories available. It is not a static system. Its intelligence is in the indexing and in its conception of matching within the indexing. Whenever you hear a story in EXTRA, EXTRA is searching to find one or more related stories to next present.

EXTRA gets smarter over time. Each story it receives is entered into memory and automatically connected to similar stories. The role of the human in EXTRA is simply to index the story. EXTRA does the rest.

This allows the creation of a true alternative to

keyword-based search. Stories are searched for by their conversational relevance which is measured by the closeness of fit of their indices.

EXTRA is meant to capture the corporate memory of any large company with experts who rarely meet each other or who may retire or quit. Corporate Memory cannot be captured by a set of dull documents that no one reads or a set of hour long interviews that no one will ever watch.

Corporate knowledge, in the form of short video stories has to find the user when the user needs that knowledge, A dynamic conversation-based indexing system allows that to happen.