Conversational Adequacy: Mistakes are the Essence*

Donald Perlis†and Khemdut Purang
Department of Computer Science
University of Maryland
College Park, MD 20742
(301) 405-2685
perlis,kpurang@cs.umd.edu

Abstract

We argue that meta-dialog and meta-reasoning, far from being of only occasional use, are the very essence of conversation and communication between agents. We give four paradigm examples of massive use of meta-dialog where only limited base dialog may be present, and use these to bolster our claim of centrality for meta-dialog. We further illustrate this with related work in active logics. We argue moreover that there may be a core set of meta-dialog principles that is in some sense complete. If we are right, then implementing such a set would be of considerable interest. We give examples of existing computer programs that converse inadequately according to our guidelines.

Introduction

In (McRoy 1993) McRoy urges that the ability to deal with mistakes is central to communication, and moreover that it is best treated as part and parcel of the same reasoning and linguistic abilities as the rest of the communication process, rather than as a separate and optional refinement. We agree. In fact, much of the research in active logics (Miller 1993; Gurney, Perlis, & Purang 1995; Purang, Perlis, & Gurney 1996) is based on a very similar point of view, not only toward communication but toward situated reasoning in general.

This is also related to the grounding problem (Traum 1994) and the misidentification problem (Miller 1993; Maida 1991; 1992)

Here we will explore the thesis that meta-dialog and meta-reasoning are, in some sense, both necessary and sufficient for communication. Just as the commonsense world has proven to be sufficiently complex to require that any intelligent commonsense agent be able to carry on appropriate meta-reasoning with respect to uncertainty, incomplete data, error and inconsistency, conversation too is fraught with complex and frequent surprises requiring the same degree of commonsense reasoning. Mistakes are inevitable in dialog; indeed, there can be no such thing as error-free dialog even in principle, let along in practice: language itself has too many ambiguities for that. In practice, conversation breaks down almost instantly in the absence of facility to recognize and repair errors, ask clarificational questions, give confirmation, and perform disambiguation. Military commands and air traffic control to the contrary: these are highly reduced lingoes. Yet even there we find negotiation over meaning. The closest thing to real error-free speech may be in fact a speech, where the audience listens silently and probably misunderstands a great deal. Yet even there the speaker often resorts to speech repairs.

A good test of this claim can be gotten via a typed conversation in which there is no backspacing, no repair of any sort, no back reference to what was said: one is required to go on and let errors lie in peace. Try it: you'll find that you do not get far, even if you type very accurately, there will be errors of thought, things you decide you did not mean, but you are not allowed to correct or comment on them.

Moreover, sufficient facility with meta-reasoning can overcome large amounts of deficiency in other matters, whether in problem-solving or in conversation. Rieger (Rieger 1974) makes a similar point; here we are concerned to apply this idea to the problem of miscommunication in particular, with attention as well to the Turing Test.

Related work on miscommunication

The related work described below is concerned, in one form or another, with detection and repair of miscommunication. This brief review will motivate the rest of the paper.

*This research has been supported in part by grants from the National Science Foundation and the Army Research Office. We thank Betsy Klipple for helpful discussion.

†Institute for Advanced Computer Studies

T1	Mother:	Do you know who's going to that meeting?
T2	Russ:	Who?
T3	Mother:	I don't know.
T4	Russ:	Oh. Probably Mrs. McOwen and probably Mrs. Cadry and some of the teachers.

Figure 1: Russ notices an inconsistency at T3 and makes a repair at T4

McRoy and Hirst

McRoy and Hirst (McRoy 1993; Hirst & McRoy 1995) consider misunderstanding and repairs in dialog. In their model of conversation, misunderstanding is signalled by an inconsistency between the expectations of a dialog participant and an utterance. The agent must then reason about and explain this inconsistency. This can lead to a change in the interpretation of previous parts of the dialog and trigger a repair utterance. McRoy discusses Figure 1 where T3 is inconsistent with the expectations of Russ and is repaired in T4.

The language of the model includes terms for quoted propositions representing the beliefs of the agent and terms for turn sequences which record the interpretation of the utterances in the discourse. This enables the agent to reason about its beliefs, and about the discourse, its possible interpretations and repairs.

Repairs are treated as part of the core linguistic capabilities of an agent and the processing of repairs is not distinguished from the processing of normal speech.

Traum and the TRAINS project

Traum (Traum 1994) observes that most approaches to natural language understanding assume that utterances are heard and understood correctly by the dialog participants. However, an appreciable portion of dialog is involved in coordinating understanding between the participants. This accounts for half the utterances in the TRAINS domain (Allen & Shubert 1991).[1] These utterances are not about the domain of the conversation, but about the conversation itself, and are used for acknowledgement, acceptance, clarification and repair. An example from the TRAINS system (Heeman & Allen 1994) is:

and pick up um the en- I guess the entire um p- pick up the load of oranges at Corning

where the speaker meant:

and pick up the load of oranges at Corning.

[1] TRAINS is an ongoing project combining dialog and planning.

T1	Agent 1:	Did you hear that John broke his leg?
T2	Agent 2:	No really? What a shame!
T3	Agent 1:	Yes, and his wife now has to do everything for him.
T4	Agent 2:	Wife? John isn't married. Which John are you talking about?
T5	Agent 1:	I am talking about John Jones.
T6	Agent 2:	Oh, I don't know him. I thought you meant John Smith.

Figure 2: The meaning of "John" changes from John Smith to John Jones

Traum discusses a finite state model for grounding (grounding is the process by which participants in a dialog reach mutual understanding of what was intended by the speaker) which uses a set of grounding acts including acknowledgement and repairs. These grounding acts are then incorporated into a theory of action in conversation (see also (Traum & Hinkelman 1992)) which is then used in a theory of dialog management which was implemented in the TRAINS system.

Miller

Miller (Miller 1993) discusses reasoned change in belief and terminological change in the context of active logic[2]. An agent has to be able to change the meaning of a term and later reason about the meanings. Figure 2 illustrates this.

Agent 1 acquires a set of beliefs between turns T1 and T3. At that point, it notices an inconsistency: T3 presupposes that John is married, but agent 1 believes that is not the case. It resolves that by challenging agent 2 at T4. Agent 2 clarifies the misunderstanding at T5 by pointing out that they were talking about different Johns. This causes agent 1 to change its meaning of "John" from John Smith to John Jones which resolves the inconsistency.

Strategic Competence

Canale (?) distinguishes 4 aspects of communicative competence: grammatical competence, sociolinguistic competence, discourse competence and strategic competence. Of these, strategic competence is most closely related to our concerns.

Strategic competence consists of communication strategies that are put to use when communications break down or when the conversational participants lack the other competences (as when one interlocutor speaks in a language he is learning). The strategies invoked in these cases include paraphrase, emphasis, rep-

[2] See a later section below for a brief description of active logic.

In a crowded New York deli a visiting Frenchman has ordered a Swiss cheese sandwich.

Waitress:	What kind of bread do you want for your sandwich, white, whole wheat or rye?
Frenchman:	(Wh)ye.
Waitress:	White?
Frenchman:	*(Wh)ye.*
Waitress:	*White?*
Frenchman:	Whole wheat.

Figure 3: A linguistic deficiency resulting in message abandonment.

etition, avoidance (including message abandonment) and clarification seeking. This is illustrated in Figure 3 from (Savignon 1983) where breakdown of communication leads to repetition, emphasis and finally to the abandonment of the original message.

Tarone (Tarone 1981) notes that communication strategies go beyond those a speaker uses to express himself in the face of a deficient language system. They crucially include negotiation of meaning by both interlocutors to ensure that they reach common understanding in situations where they do not share meaning structures.

Presuppositions and Implicatures

In (Gurney, Perlis, & Purang 1995; Purang, Perlis, & Gurney 1996), the authors consider active logics applied to the processing of presuppositions and of implicatures. The common theme of these papers is the need for tracking over time the set of beliefs an agent has as it participates in a dialog. Figure 4 makes the point for presuppositions.

T2 presupposes that there are roses and that there is a fridge. At T3, though, the presupposition that there are roses is defeated but the presupposition that there is a fridge is unscathed. However, in the standard discourse updating theory (Heim 1983), at the end of T3, the presupposition that there is a fridge would not be present. Using an active logic to process the discourse a step at a time, and to handle contradictions gives the correct results.

Figure 5 illustrates implicatures being made and later withdrawn.

At T2, Kathy infers the implicature that the roses are fresh, since Bill says that they are in the fridge.

T1	Mary:	Are the roses in the fridge?
T2	John:	No the roses are not in the fridge.
T3	John:	Because there are no roses.

Figure 4: A presuppositional example.

However, at the very next moment, she needs to retract this when Bill states that the roses are not fresh after all. The active logic used enables contradictory beliefs to be held (e.g., that the roses are fresh and that they are not fresh) temporarily, which serves as a trigger for subsequent repair (e.g., settling on one of the two).

This has illustrated some forms of miscommunication and approaches to studying them. In the next subsection we discuss active logic (alluded to above) which we intend to use as a tool to formalize and explore conversational adequacy.

Active logic

Active logic (Elgot-Drapkin & Perlis 1990; Miller & Perlis 1993) is a family of formalisms developed for the purpose of modeling the reasoning process in a way that respects the passage of time as reasoning proceeds. These formalisms have been applied to a number of domains, from multi–agent interaction to deadline–coupled planning, from fully–decidable default reasoning to reasoning in the presence of contradictions, from correcting misidentification errors to meaning change.

Rather than proceeding from one nonmonotonic theory (with one set of axioms) to another nonmonotonic theory (with an updated set of axioms) there is one discretely evolving theory in active logic. This models a process of thinking that takes a reasoner from one belief state to the next. The sequence of belief states the reasoner experiences is recorded into a history. That history, together with a quotation mechanism enables the logic to reason about its own reasoning in time.

The belief state transitions are effected by a set of rules which map the history up to time t onto a new state at $t+1$. The inheritance rule, for instance, will, by default, cause all the beliefs at time t to be inherited to time $t+1$. One of the situations where this rule is defeated is the presence of direct contradictions at time t—if both p and $not(p)$ are beliefs at t, then neither is inherited to $t+1$. This in effect removes the contradiction from the current belief state and prevents it from infecting future states. Such a contradiction will also trigger the belief $contra(p, not(p))$ at $t+1$ to indicate a contradiction occurred. The logic can then reason about the contradictands and reinstate the one judged correct (if any).

T1	Kathy:	Are the roses fresh?
T2	Bill:	They are in the fridge.
T3	Bill:	But they are not fresh.

Figure 5: The implicature triggered at T2 is absent after T3

Conversational Adequacy

Here we outline a preliminary list of items that we think are part of an essential core of conversational adequacy. First, we distinguish four broad categories of capacities: basic inference, basic learning, basic language, and miscommunication competence. The first three are the "basic" capacities, and the fourth is a "meta" capacity. Basic inference includes deductive, nonmonotonic, and abductive inference; basic learning includes advice-taking, grouping (concept formation), word learning; basic language includes syntax, semantics, and pragmatics, including aspects of conversational analysis (Levinson 1983).

We suppose an agent to have at least a *weak* ability in the three basic capacities; what is critical is that these be sufficient to support the fourth capacity: miscommunication competence.[3] Thus the agent may be rather poor at parsing, for example, but must be able to parse such meta-sentences as "what does X mean?" and "can you repeat that?" and "didn't you say X?". Notice that the use of quoted or reified text, is a special ability not always included in parsers. The agent must furthermore be able to reason appropriately on the basis of such parses. It is helpful to think of the agent as translating parsed sentences into an internal logical language on which its inference engine can run; and among its inferential tasks is the assessment of the parse as plausible or not: if not, then a miscommunication is presumed and a repair effort is initiated.

Our thesis can now be stated as follows:

(i) SUFFICIENCY: as long as there is at least a *weak* ability in the "basic" capacities (inference, learning and language) then effective conversation can proceed if there is a *strong* miscommunication capacity.

(ii) NECESSITY: no matter how strong the basic capacities, effective communication cannot proceed if there is not a strong miscommunication capacity.

We now suggest what miscommunication competence might consist of. It will be general-purpose meta-reasoning, involving at least the following five features:

1. **Time:** In a series of studies of commonsense reasoning, including some aspects of dialog, we have found, again and again, that a high degree of situated temporal reasoning provides tremendous gains in representational and inferential capacity, and in particular in the capacity to deal with error, contradiction, and miscommunication. This appears to result in part from the ability to refer back in time to what has been said (or thought) and thereby examine its meaning and correctness.

   This item is fundamental, perhaps the most fundamental of all, underlying even the use/mention distinction and appearance/reality distinction (see (Miller 1993)). Simply noticing, let alone repairing, an error of miscommunication, requires seeing it as a reified entity rather than as one's current view of the world. The agent must come to see it as part of its previous view of the world; this in turn requires a robust situated temporal-reasoning facility, in which the present moment is constantly being moved into the past as a new moment becomes the present.

2. **History:** Time-representation permits there to be a history of the discourse so far, that is constantly developing during the conversation. Since agents can disagree on the history, it must be possible to revise the history, but in such a way that the fact of the revision as having taken place is also stored, along with the revised and un-revised versions.

3. **Meaning:** Words have meanings, and the meanings are different from the words, making it possible to hear a word, and therefore know it has been uttered, without knowing what it means. Moreover, meaning is plastic, at least partially under the control of speakers, who may be asked what meaning they have in mind. A key facility in meta-discourse is the ability to refer to a word, as opposed to its meaning; quotation is one device that can be used, but tone and other devices are also available.

4. **Contradiction:** The examples in the section on related work can be viewed as involving contradictions in understanding between the two speakers. Indeed, three of the examples are explicitly analyzed by their authors in such terms. We consider the discovery of a direct contradiction to be one of the chief ways for detecting miscommunication. Of course, the presence of a contradiction spells problems for a reasoning agent, and we are arguing here that a conversational agent must very much be a reasoner too. See (Miller 1993; Roos 1992; Perlis 1996) for more discussion of the problem of reasoning with inconsistent information.

5. **Deixis:** Without the ability to point, verbally, to *this* utterance and the present time, to ground a discussion in here and now, confusion can hardly be set aright.

[3] For instance, the following would be part of such weak abilities: modus ponens, quotation or reification, and change of belief. These are required in particular to facilitate the five features below that characterize miscommunication competence.

Time allows for the possibility of reasoning about conversational history; and for the possibility of reasoning in the face of contradictions. The latter are important as cues that miscommunication has occurred: recognition of miscommunication is the first step in the meta-reasoning that can lead to repair. The next step is often communication of the fact of miscommunication, to enlist the help of the other agent in the repair;[4] at times however this is unnecessary, if the recognizer can effect the repair alone. In any event, reification of (typically) recent portions of the ongoing conversation is an essential step in the meta-reasoning of recognition and repair. There must be something seen as needing repair, and that something is linguistic rather than "in the world": utterances are taken as things in themselves, not identified with their presumed meanings, as long as they are under inspection with respect to a miscommunication. Thus the recent past is turned into an object. This can often recurse, as the present becomes reified into the past as well; see the McRoy example above, where each agent initiates a repair (Mother at T3 and Russ at T4)—in this case, each agent does the repair alone, but the presumed reasoning involves internal reifications of what has just been said. A similar example could be given in which Mother and Russ talk out loud as they do their repairs:

T3
Mother: You say "Who?"? That doesn't answer my question. You misunderstood me. I want to know who myself, and hope you can tell me.

We now illustrate our thesis with four examples.

Extreme examples of conversational adequacy

Each of the following examples appears to be a case of an individual *not* well understanding a conversation, but nevertheless—by utilizing considerable meta-linguistic and meta-cognitive skills—being able to make some headway and, in a very real sense, being a conversational participant. These are speculative; that is, we are not certain they appropriately support our claim, but we think that there is much to be gained from airing them alongside it.

Apprentice: An apprentice or novice privy to a conversation between experts is often lost in a bewildering array of unfamiliar terminology and inferences, mixed in with segments that are relatively clear. By noting unclear segments, the apprentice can ask for clarification when needed. And when the proferred clarification is itself not clear, this too can be noted; this may or may not lead to eventual understanding of the underlying topic, but in the process of negotiation over meanings, there is real communication.

Child: Much like an apprentice, a child is often lost in adult conversation, yet can—via questions—seek clarification. Key to this is the child's understanding that it does not understand, together with social-cognitive skills such as noting who is saying what. Again, as with the apprentice, new phrases can be learned, as well as new meanings for old phrases. See (Ochs 1991).

Foreigner: A foreigner with *no* grasp of a language is at a tremendous disadvantage. But with only a modicum of familiarity with very elementary features, conversation becomes possible (albeit slow and painful). These features are those we characterize as conversational adequacy, and are the same (we suspect) for the foreigner as for the child and the apprentice. See (Gass & Varonis 1991).

SLI: Persons suffering from SLI (specific language impairment) (Stark & Tallal 1981) have considerable difficulty with ordinary spoken communication, and resort to various strategies to achieve communication, such as repetition, memory, seeking help from others, and heavy use of conscious analysis of utterances. There appears to be no underlying non-linguistic cognitive deficit (although this is controversial), and communication does occur, again slowly and with difficulty. See (Pinker 1994; Fletcher & Hall 1992; Watkins & Rice 1994).

Our point in these examples is not that excellence in parsing and other aspects of language are unimportant frills, but that there is an almost pre-linguistic core of capabilities, essential to any effective communication, even very slow and inefficient communication. In a sense this is definitional; for instance, below we consider current Turing Test agents to be ineffective, because they cannot, no matter how much time they are allowed, come to deal with their interlocutor's concerns.

Turing Test Redefined?

In (Hayes & Ford 1995) Hayes and Ford take Turing (Turing 1950) to task for his formulation of what has come to be known as the Turing Test, arguing in particular that it has harmful effects on the long-term development of artificial intelligence. We suggest a redefinition of the Turing Test in light of their critique. We suspect that Turing was led to his formulation as a clever finessing of the tough problem of characterizing

[4] e.g. "That's not what I meant", a deictic reference.

intelligence; as such it does not do much toward solving that problem. However, we think that a slightly more sympathetic reading can suggest a fruitful direction to take, related to what we have been presenting above, namely that of conversational adequacy: the automated agent should be able to engage in discussion including meta-discussion, on any topic whatsoever, as an apprentice if it knows little or nothing about it, learning as it goes. This will include the ability to talk about the conversation itself, any miscommunications that arise, and even about grammatical constructions such as word order, spelling, and meaning.

Turing Test programs

The Loebner Prize is given each year for the program that comes closest to passing the Turing Test, or at least achieves the most plausible conversation. Yet even the winners fall down dramatically in terms of any real command of conversation. We illustrate that with a few of the programs available on the internet. (Note: the transcripts are in `typewriter font` and the inputs to the programs are in *italics*.)

Mystery

Mystery won the second prize in the 1995 competition. It is available via telnet at debra.dgbt.doc.ca on port 3000.

`Yo. I'm here.`
My name is Don.
`I'm Joe Spansen.`
⋮
`I'm a cleaner. I empty trash. Vacuum. Dust. You know.`
So you know what a vacuum is?
`What?`
You just used the word "vacuum" a moment ago.
`I don't know.`
A vacuum is a device you use to vacuum dust.
`Huh?`
What is a vacuum?
`Beats me.`
How do you vacuum?
`What do you mean?`
You told me that you vacuum.
`What?`
Can you vacuum?
`I don't know.`
To vacuum is to clean with a suction device.
`Yeah. I work real hard. No one keeps their building as clean as me. It's not fair that they want to fire me. It's just because my boss, Mr. Johnson, is mad at me.`

The above exchange seems to indicate that the program is unable to reason about its own utterances or about dialog in general. This in turn appears to make it impossible for there to be a "meeting of minds" when miscommunication occurs, or indeed for there to be a recognition of miscommunication. This and most other "Turing Test" programs are very good at "running-on": they keep a superficial semblance of conversation going. The above continues:

If you vacuum with a Hoover, he won't fire you.
`Oh, they'll fire me for sure. I know what they're thinking. They haven't said they're gonna fire me yet because they gotta keep me around to clean until they can find someone else to hire. Then they'll fire me for some other reason that won't make problems with the union. They fire lots of cleaners. They'll just say I'm a bad cleaner. They fired my friend, Fred two years ago because they said he was a bad cleaner.`
Vacuum with a Hoover.
`Huh?`
A Hoover will impress Mr, Johnson so much he will give you a raise.
`Beats me.`
Fred wasn't fired.
`Yeah. He's one of the guys in our group. Just some guys from work. We hang out together. Do stuff. You know.`
You are wrong about Fred. He wasn't fired.
`We usually go down to the club and drink beer and stuff.`

The program is oblivious to the direct contradiction between what it said (Fred was fired) and our rejoinder to the contrary. This is another telltale sign of conversational inadequacy: a meeting of minds requires above all the ability to notice when minds (and utterances) suffer a mismatch.

Hex

Hex was the winner of the 1996 Loebner Prize competition. The source code of the program is available at http://ciips.ee.uwa.edu.au/~hutch. The following dialog again illustrates the same failings as the Mystery program above, but is more tantalizing because it has a much more convincing fluidity, and more careful probing is required to reveal the underlying lack of conversational adequacy. We struggled repeatedly and in various formats and phrasings to get the program to refer to its own past utterances or indeed to our utterances, with no success. Nor was it able to refer to words, as opposed to their referents. Thus no repairs, nor even recognition of miscommunication was possible.

G'day mate.

Hi, I'm Don Smith.

Hello. Now, that's the formalities over!

Not over yet. Did yo forget to tell me your name?

My name is Jason, but you can call me Jas if you like...

Did you forget to tell me your name at first?

I have already told you! My name is JASON.

Yes, I know you told me already. But before you told me your name, you first said "that's the formalities over" --do you remember?

OK. That's good... My name is Jason, but you can call me Jas if you like...

I am curious about your sense of time. Who stated his name first as our conversation began, you or me?

If you are curious about my sense of time, it's your business, mate! I have already told you! My name is JASON.

You told me your name and I told you mine. But who told his name FIRST?

Why do you feel that you have to tell me that I told you my name and you told I yours? Well, I could tell you who told his name first, but you'd never believe me ;^o Would you like to hear a joke? Knock, knock...

No, wait, please, I am confused, let's not do the joke yet. I need to get clear about how we are remembering time. For instance, the last thing I said to you before you started the joke, was "who told his name first?" Do you remember that?

Confusing! I'll show you confusing! Ytu smf eptl piy ejsy O s, dsuomh! OK, here's a joke for you. Knock, knock.

I'm afraid I can't take a joke right now. Please do not go on with a joke. Tell me instead why you did not tell me your name at first.

Alright! Cool. Knock, knock. I refuse to tell you about anything! I dislike your company. I don't know what your motives are, but you certainly aren't treating me with respect!

I am trying to find out if you and I have similar ways of reasoning about conversations.

Wow - you are trying to find out if I and you have similar ways of reasoning about conversations? Me too! Which part of the country are you from?

No, I think we are not understanding one another. Do you know what a converstation is?

I wonder whether you think sometimes! I am feeling like the judge here!

Do you know the word "conversation"?

Sorry, but I don't wish to reveal whether I know the word "conversation". It is kind of personal.

You do know the word "personal" then. You just used it. Yes?

Perhaps I do know the word "personal" then, but it's hardly any of your business... I find it rather insulting that you think I just used it! I'm glad to know that that's the case.

I am confused again. Didn't you say "It is kind of personal" a minute ago?

I didn't want to know that you are confused again, thanks... Maybe I did not say "it is kind of personal" a minute ago. I'll never tell. Hehehe ;^)

What will you never tell?

Well, I know what I will never tell, but I refuse to tell you! So, have you ever visited Australia?

This is a change of topic. We began talking about names, and then I forget, what else did we talk about? Was it the weather?

Perhaps a change of topic is, perhaps not. It's a mystery ;^) We are bloody well talking, aren't we? We have such nice weather all year round back home. I'm not really used to really cold temperatures.

I remember: we talked about names and then about words such as "personal"and "conversation". Do yu remember that?

I always suspected that you remember: we talked about names and then about words such as "personal" and "conversation". Now I know for sure ;^) Look - it doesn't matter if yu remember that. Change the topic! So, what do you do for a living?

You misspelled the word "for" when you typed "dor" just now.

Hey! I don't criticize you, do I? What does one measly mistake mean, anyway?

Can you tell me what "do dor a lliving" means?

I have absolutely no idea what "do dor a lliving" means.

can you tell me what "do for a living" means?

I'm sure someone out there knows what "do for a living" means...

What does "someone out there" mean?

Do *you* know what "someone out there" mean does? I'm unsure myself.

What does "unsure" mean?

I really can't say what "unsure" mean does. Can you give me more info? What do you think of this Loebner contest thing?

can you tell me what "contest" means?

I don't know... sorry...

Can you tell me what "sorry" means?

That's really outside the range of my expertise. Have you seen any good films lately?

can you tell me what a film is?

I do like to go to the movies, but I can't stand most of the Hollywood trash. I actually prefer international cinema. What have you seen recently?

I am a foreigner, I have trouble understanding, please tell me what a film is.

I saw "The City of Lost Children" recently. Have you heard of it?

What does "City" mean?

Hmmm... I really couldn't say. On another note - have you seen any good movies recently?

What does "movie" mean?

If I knew I'd tell you, but I don't so I won't ;^)

But you just said "movie"!

Hehehe. Trying to trick me, eh?

No, I just want to find out if you and I deal with time similarly.Did you say "have you seen any good movies recently" a moment ago?

Oh well, that's a pity! Aw shucks!

Do you know a difference between a movie and the word "movie"?

Hey, I don't know!!! Maybe I know a difference between a movie and the word "movie", but who cares???

Did you type "wordd"?

Most people believe I did type "wordd" - what do you say to that?

Did you type "peole"?

Perhaps I did type "peole", don't ask me! How should I know? What kind of music are you into?

We have similar examples using various other recent Loebner Prize competitors, and we are confident that conversational inadequacy as we have described it earlier in this paper is not present even in rudimentary form in any of these programs.

Toward programs that converse!

There is such a thing as talking past one another, where I talk for a bit, then you talk for a bit, with only very slight attention to the details of what the other is saying. This talking-past, or running-on, is typical of Turing Test efforts as seen above; programs are assessed with respect to how well they can chatter on in rough correspondence to a topic, with appropriate defusers such as "I don't know" or "huh?" when the syntax gets too tough. But there is little or no communication, in the sense of success at dealing with the specific points being made by the other speaker. In particular, clarification questions such as "Are we talking about the same person?" are generally not dealt with at all. In general, meaning is not dealt with at all, and in particular not situated meaning (e.g., reference to the dialog itself). Some programs are quite cleverly designed to appear to be engaging in "real" dialog, i.e., engaged in discussion of a topic in concert with their interlocutor; but no amount of effort to get them to deal with meaning succeeds, and this reveals that they are in effect babbling words meaningless to them.

Here we give an example of the sort of "real" discussion that reveals attention to meaning (and hence to errors: error after all is a mismatch between what is meant or intended and what was produced, i.e., the meaning produced is not the meaning desired).

We imagine a "bare" human: stripped of all personal and episodic memory and even perhaps of much vocabulary, but with rationality and error-dealing capacity intact, as well as some bare-bones language.

you: How old are you?

AL: I don't know what "old" means.

you: "Old" refers to the number of years you have been alive.

AL: I don't know what "years" and "alive" mean.

you: OK. Let's start with time. Do you know about time?

AL: "About time" means "almost too late".

you: You misunderstand me. I mean, do you know what "time" means?

AL: Yes: "time" is the measure that keeps increasing second by second.

you: Good. Now, how much time has passed since you were born?

AL: What is "born"?

you: Let's back up. A minute is 60 seconds; an hour is 60 minutes; a day is 24 hours; and a year is 365 days.

AL: OK.

you: So, how many years have you been here, or anywhere?

AL: I can't calculate that; I don't know where to count from.

you: Do you remember ever not being here?

AL: No.

you: How long have we been talking?

AL: 22 seconds; now 23; now 24...

you: Stop! That's enough.

AL: What is enough?

you: What do you mean?

AL: You said "that's enough." What does "that" refer to?

you: Oh, "that's enough" is just an expression, meaning "stop".

AL: OK. So you repeated yourself: you said "stop" and then said "that's enough" as well.

you: Hmm, you are right. I guess there is more meaning to "that's enough" than simply "stop" but I can't explain it easily.

AL: OK.

you: What is your name?

AL: I don't know.

you: Let's call you "AL".

AL: OK, my name will be "AL".

you: So, AL, what is the earliest memory you have?

AL : Hmm, I think it is when we began talking, about a minute ago.

...

Although AL and the interlocutor have not gotten very far, and have not answered the opening question, nevertheless each utterance in the sequence does by and large bear on what has gone before, especially when either party calls the others' attention to that. This dialog illustrates at least four of our five miscommunication competence features. In addition the two agents have made some significant progress in reaching some common ground and in particular in referring to the course of the conversation itself.

The above conversation is not a human/computer dialog. However, it should be! There seems to be a no-man's land between academic NLP and Loebner Prize NLP. Each extreme seems far removed from "genuine" conversation of the sort we are urging here. We think that substantial effort devoted to algorithmic treatments of conversational adequacy may help bridge the gap between theoretically informed single-sentence research and totally unbridled (as well as unprincipled) Turing Test programs.

Conclusions

We do not believe that our preliminary list of principles for conversational adequacy is, in itself, adequate! We are sure that further essential principles can be added, without which effective communication cannot occur. But we think that a complete list may be possible, and that, like a universal Turing machine, a certain degree of capability opens the door to the rest (modulo efficiency). Since we are not working within the stricture of a precise definition of effective communication, proofs are not possible; but for the same reason proofs of the Church-Turing Thesis are not possible. Instead we can hope for suggestive evidence, such as in the form of computer programs with more or less conversational ability, and in the form of human examples such as we have presented.

Our suggested conversational adequacy principles are largely ones of cognitive ability, not specific to conversation. We think this is appropriate: much of the reasoning that underlies effective communication, and especially effective dealing with miscommunication, is general-purpose reasoning. For this reason, Turing was appealing to a good hunch in focusing on conversation as a measure of thought. While thought of course can occur without conversation, it is less plausible that (effective) conversation can occur without thought; this is most salient in the conversational patterns associated with (repair of) miscommunication.

References

Allen, J., and Shubert, L. 1991. The trains project. Technical Report TRAINS Technical Note 91-1, University of Rochester.

Elgot-Drapkin, J., and Perlis, D. 1990. Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence* 2(1):75–98.

Fletcher, P., and Hall, D., eds. 1992. *Specific Speech and Language Disorders in Children: Correlates, Characteristics and Outcomes.* Singular Publishing Group.

Gass, S. M., and Varonis, E. M. 1991. Miscommunication in nonnative speaker discourse. In *Miscommunication and Problematic Talk.*

Gurney, J.; Perlis, D.; and Purang, K. 1995. Active logic and Heim's rule for updating discourse context. In *IJCAI 95 Workshop on Context in Natural Language.*

Hayes, P., and Ford, K. 1995. Turing test considered harmful. In *14th International Joint Conference on Artificial Intelligence.*

Heeman, P., and Allen, J. 1994. Detecting and correcting speech repairs. In *32nd Annual Meeting of the Association for Computational Linguistics.*

Heim, I. 1983. Pragmatics. In Davis, S., ed., *On the Projection Problem for Presuppositions.*

Hirst, G., and McRoy, S. 1995. The repair of speech act misunderstandings by abductive inference. *Computational Linguistics* 17(4).

Levinson, S. 1983. *Pragmatics.* Cambridge University Press.

Maida, A. 1991. Maintaining mental models of agents who have existential misconceptions. *Artificial Intelligence* 50(3):331–383.

Maida, A. 1992. Propositionally representing incomplete knowledge about existence. Stanford University: AAAI92 Spring Symposium Series on Propositional Knowledge Representation.

McRoy, S. 1993. *Abductive Interpretation and Reinterpretation of Natural Language Utterances.* Ph.D. Dissertation, University of Toronto.

Miller, M., and Perlis, D. 1993. Presentations and this and that: logic in action. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society.*

Miller, M. 1993. *A View of One's Past and Other Aspects of Reasoned Change in Belief.* Ph.D. Dissertation, Department of Computer Science, University of Maryland, College Park, Maryland.

Ochs, E. 1991. Misunderstanding children. In *Miscommunication and Problematic Talk.*

Perlis, D. 1996. Sources of, and exploiting, inconsistency: preliminary report. In *1996 Workshop on Commonsense Reasoning.*

Pinker, S. 1994. *The Language Instinct.* Harper-Perennial.

Purang, K.; Perlis, D.; and Gurney, J. 1996. Active logic applied to cancellation of gricean implicature. In *AAAI 96 Spring Symposium on Computational Implicature.*

Rieger, C. 1974. *Conceptual Memory: A Theory and Computer Program for Processing the Meaning Content of Natural-Language Utterances.* Ph.D. Dissertation, Department of Computer Science, Stanford University, Palo Alto, California.

Roos, N. 1992. A logic for reasoning with inconsistent knowledge. *Artificial Intelligence* 57:69–103.

Savignon, S. 1983. *Communicative Competence: Theory and Classroom Practice.* Addison-Wesley.

Stark, R., and Tallal, P. 1981. Selection of children with specific language deficits. *Journal of speech and hearing disorders* 46:114–122.

Tarone, E. 1981. Some thoughts on the notion of communication strategy. *TESOL Quarterly* 15(3):285–295.

Traum, D., and Hinkelman, E. 1992. Conversation acts in task-oriented spoken dialogue. *Computational Intelligence* 8(3):575–599.

Traum, D. 1994. *A Computational Theory of Grounding in Natural Language Conversation.* Ph.D. Dissertation, University of Rochester.

Turing, A. 1950. Computing machinery and intelligence. *Mind* 59:433–460.

Watkins, R., and Rice, M., eds. 1994. *Specific Language Impairments in Children.* Paul H. Brookes.

Addendum

After we had contributed the published version of this paper, we noticed that the transcript of the dialog with program Hex had no spelling errors even though the actual dialog had a number of errors which in fact influenced the course of the dialog. These errors are a feature of the program, to make its typing seem more human. We responded to some of them, calling attention to them, and the program's responses pretty much ignored our efforts. In fact nearly all our attempts to get either program to discuss the course of the conversation itself, words, time just past, spellings, meanings, topics, failed. This is precisely the point we make in the paper. But in the particular cases of mis-spellings, program Hex kept an "expurgated transcript" that in fact did not have the spelling errors at all! That is the transcript in the published paper, and as such lends an even odder aspect to the dialog, and lending further weight to our contention that the program has no robust record of the ongoing course of the dialog that it can use in formulating its utterances.