GRICE INCORPORATED
Cooperativity in Spoken Dialogue

Laila Dybkjær, Niels Ole Bernsen and Hans Dybkjær
Centre for Cognitive Science, Roskilde University
PO Box 260, DK-4000 Roskilde, Denmark
emails: laila@cog.ruc.dk, nob@cog.ruc.dk, dybkjaer@cog.ruc.dk
phone: +45 46 75 77 11 fax: +45 46 75 45 02

Abstract

The paper presents a consolidated set of principles of cooperative spoken human-machine dialogue which have the potential for being turned into practically applicable design guidelines. The principles have been validated in three ways. They were established from a Wizard of Oz simulation corpus used to develop the dialogue model for a spoken language dialogue system. Developed independently of Gricean theory, some of the principles were refined through comparison with Grice's maxims of cooperativity in conversation. Finally, the principles were tested in the user test of the implemented dialogue system. The paper shows that Grice's maxims constitute a sub-set of the principles. The non-Gricean principles and dialogue aspects they introduce are presented and discussed.

1 Introduction

In the last four years, we have designed and implemented the dialogue component of a spoken language dialogue system (SLDS) prototype in the domain of flight ticket reservation. The aim has been to develop a realistic, application-oriented prototype whose dialogue management allows users to perform their reservation task in spontaneous and natural spoken language. Being well-structured, the ticket reservation task generally lends itself to system-directed dialogue in which the user answers questions posed by the system. The only user initiative our system permits is that users may initiate clarification and repair meta-communication through uttering the keywords 'repeat' and 'change'. In designing such a system, it is crucial to reduce the number of situations in which users are inclined to take other forms of dialogue initiative, such as asking questions when they do not understand the system's dialogue behaviour or providing information which the system did not ask for (Schegloff et al. 1977). This is why the issue of dialogue cooperativity came to play a central role in our design of the dialogue structure. We needed to optimise system dialogue cooperativity in order to prevent situations such as those described above. To this end, we developed a set of general principles to be observed in the design of cooperative, spoken human-machine dialogue. The principles have been validated in three ways. Firstly, they were developed on the basis of a simulated human-machine dialogue corpus collected during dialogue model design. Secondly, we compared the principles with Grice's maxims of cooperative human-human dialogue. Thirdly, the principles were tested against the dialogue corpus from the user test of the implemented system.

This paper analyses the relationship between our principles and Grice's maxims. We first describe how the principles were developed (Section 2). We then justify the comparison between principles and maxims (Section 3). Section 4 compares principles and maxims. Section 5 briefly describes how the principles were tested the on user test dialogue corpus, and Section 6 concludes the paper.

2 Developing and Testing Principles of Cooperative Human-Machine Dialogue

The dialogue model for our flight reservation system was developed by the Wizard of Oz (WOZ) experimental prototyping method in which a person simulates the system to be designed (Fraser and Gilbert 1991). Development was iterated until the dialogue model satisfied the design constraints on, i.a., average user utterance length. The dialogues were recorded, transcribed, analysed and used as a basis for improvements on the dialogue model. We performed seven WOZ iterations yielding a transcribed corpus of 125 task-oriented human-machine dialogues corresponding to approximately seven hours of spoken dialogue. The 94 dialogues that were recorded during the last two WOZ iterations were performed by external subjects whereas only system designers and colleagues had participated in the earlier iterations. A total of 24 different subjects were involved in the seven iterations. Dialogues were based on written descriptions of reservation tasks (scenarios).

A major concern during WOZ was to detect problems of user-system interaction. We eventually used the following two approaches to systematically discover such problems: ( ) prior to each WOZ iteration we matched the scenarios to be used against the current dialogue model. The model was represented as a graph structure with system phrases in the nodes and expected contents of user answers along the edges. If a deviation from the graph occurred during the matching process, this would indicate a potential dialogue design problem which should be removed, if possible. (ii) The recorded dialogues were plotted onto the graph representing the dialogue model. As in (i), graph deviations indicated potential dialogue design problems. Deviations were marked and their causes analysed whereupon the dialogue model was revised, if necessary.

At the end of the WOZ design phase, we began a more theoretical, forward-looking exercise. All the problems of interaction uncovered during WOZ were analysed and represented as violations of principles of cooperative dialogue. Each problem was considered a case in which the system, in addressing the user, had violated a principle of cooperative dialogue. The principles of cooperative dialogue were made explicit, based on the problems analysis. The WOZ corpus analysis led to the identification of 14 principles of cooperative human-machine dialogue (Section 4) based on analysis of 120 examples of user-system interaction problems. If the principles were observed in the design of the system's dialogue behaviour, we assumed, this would serve to reduce the occurrence of user dialogue behaviour that the system had not been designed to handle.

3 Maxims and Principles of Cooperative Dialogue

We had developed our principles of cooperative human-machine dialogue independently of Gricean cooperativity theory (Bernsen et al., 1996a). Prior to the user test (Section 5), we compared the principles with Grice's Cooperative Principle and maxims. In this process the principles achieved their current form as shown in Table 1. Their original expression is presented in Section 4. Grice's Cooperative Principle (CP) is a general principle which says that, to act cooperatively in conversation, one should make one's "conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which one is engaged" (Grice 1975). Grice proposes that the CP can be explicated in terms of four groups of simple maxims which are not claimed to be jointly exhaustive. The maxims are marked with an asterisk in Table 1.

Grice focuses on dialogues in which the interlocutors want to achieve a shared goal (Grandy 1989, Sarangi and Slembrouck 1992). In such dialogues, he claims, adherence to the maxims is rational because it ensures that the interlocutors pursue the shared goal most efficiently. Task-oriented dialogue, such as that of our SLDS, is a paradigm case of shared-goal dialogue. Grice, however, did not develop the maxims with the purpose of preventing communication failure in shared-goal dialogue. Rather, his interest lies in the inferences which an interlocutor is able to make when the speaker *deliberately* violates one of the maxims. He calls such deliberate speaker's messages 'conversational implicatures'. Grice's maxims, although having been conceived for a different purpose, nevertheless serve the same objective as do our principles, namely that of achieving the dialogue goal as directly and smoothly as possible, e.g. by preventing questions of clarification. It is exactly when a human or, for that matter, an SLDS, *non-deliberately* violates a maxim, that dialogue clarification problems are likely to occur. Thus, the main difference between Grice's work and ours is that the maxims were developed to account for cooperativity in human-human dialogue, whereas our principles were developed to account for cooperativity in human-machine dialogue. Given the commonality of purpose, it becomes of interest to compare principles and maxims. We want to show that the principles include the maxims as a subset and thus provides a corpus-based confirmation of their validity for spoken human-machine dialogue. Moreover, the principles manifest aspects of cooperative task-oriented dialogue which were not addressed by Grice.

4 Comparison between Maxims and Principles

In this section we analyse the relationship between Grice's maxims and our principles of dialogue cooperativity. A first aim is to demonstrate that a sub-set of the principles are roughly equivalent to the maxims. We then argue that the remaining principles express additional aspects of cooperativity. The distinction between *principle* and *aspect* (Table 1) is theoretically important because an aspect represents the property of dialogue addressed by a particular maxim or principle. One result of analysing the relationship between principles and maxims is the distinction, shown in the tables, between *generic* and *specific* principles. Grice's maxims are all generic. A generic principle may subsume one or more specific principles which specialise the generic principle to certain classes of phenomena. Although important to SLDS design, specific principles may be less significant to a general account of dialogue cooperativity.

4.1 Principles which are Reducible to Maxims

Grice's maxims of truth and evidence (GP3, GP4) have no counterparts among our principles but may simply be included among the principles. The reason is that one does not design an SLDS in the domain of air ticket reservation which provides false or unfounded information to customers. In other words, the maxims of truth and evidence are so important to the design of SLDSs that they are unlikely to emerge during dialogue design problem-solving. During system implementation, one constantly worries about truth and evidence. It cannot be allowed, for instance, that the system confirms information which has not been checked with the database and which might be false or impossible. Grice (1975) observed the fundamental nature of the maxims of truth and evidence in general and GP3 in particular (cf. Searle 1992).

Table 1. The generic and specific principles of cooperativity in dialogue. The generic principles are expressed at the same level of generality as are the Gricean maxims (marked with an *). Each specific principle is subsumed by a generic principle. The left-hand column characterises the aspect of dialogue addressed by each principle.

Dialogue Aspect	GP no.	SP no.	Generic or Specific Principle
Group 1: Informativeness	GP1		*Make your contribution as informative as is required (for the current purposes of the exchange).
	GP1	SP1	Be fully explicit in communicating to users the commitments they have made.
	GP1	SP2	Provide feedback on each piece of information provided by the user.
	GP2		*Do not make your contribution more informative than is required.
Group 2: Truth and evidence	GP3		*Do not say what you believe to be false.
	GP4		*Do not say that for which you lack adequate evidence.
Group 3: Relevance	GP5		*Be relevant, i.e. Be appropriate to the immediate needs at each stage of the transaction.
Group 4: Manner	GP6		*Avoid obscurity of expression.
	GP7		*Avoid ambiguity.
	GP7	SP3	Provide same formulation of the same question (or address) to users everywhere in the system's dialogue turns.
	GP8		*Be brief (avoid unnecessary prolixity).
	GP9		*Be orderly.
Group 5: Partner asymmetry	GP10		Inform the dialogue partners of important non-normal characteristics which they should take into account in order to behave cooperatively in dialogue.
	GP10	SP4	Provide clear and comprehensible communication of what the system can and cannot do.
	GP10	SP5	Provide clear and sufficient instructions to users on how to interact with the system.
Group 6: Background knowledge	GP11		Take partners' relevant background knowledge into account.
	GP11	SP6	Take into account possible (and possibly erroneous) user inferences by analogy from related task domains.
	GP11	SP7	Separate whenever possible between the needs of novice and expert users (user-adaptive dialogue).
	GP12		Take into account legitimate partner expectations as to your own background knowledge.
	GP12	SP8	Provide sufficient task domain knowledge and inference.
Group 7: Repair and clarification	GP13		Initiate repair or clarification meta-communication in case of communication failure.
	GP13	SP9	Provide ability to initiate repair if system understanding has failed.
	GP13	SP10	Initiate clarification meta-communication in case of inconsistent user input.
	GP13	SP11	Initiate clarification meta-communication in case of ambiguous user input.

The following principles have counterparts among the maxims:

1. Avoid 'semantical noise' in addressing users.

(1) is a generalised version of GP6 (non-obscurity) and GP7 (non-ambiguity). Its infelicitous expression was due to the fact that we wanted to cover observed ambiguity and related phenomena in one principle but failed to find an appropriate technical term for the purpose. (1) may, without any consequence other than improved clarity, be replaced by GP6 and GP7.

2. Avoid superfluous or redundant interactions with users (relative to their contextual needs).

(2) is virtually equivalent to GP2 (do not overdo informativeness) and GP5 (relevance). Grice observed the overlap between GP2 and GP5 (Grice 1975). (2) may, without any consequence other than improved clarity, be replaced by GP2 and GP5.

3. It should be possible for users to fully exploit the system's task domain knowledge when they need it.

(3) can be considered an application of GP1 (informativeness) and GP9 (orderliness), as follows. If the system adheres to GP1 and GP9, there is a maximum likelihood that users obtain the task domain knowledge they need from the system when they need it. The system should say enough and address the task-relevant dialogue topics in an order which is as close as possible to the order expected by users. If the user expects some topic to come up early in the dialogue, that topic's non-occurrence at its expected "place" may cause a clarification sub-dialogue which the system cannot understand. In WOZ Iteration 3, for instance, the system did not ask users about their interest in discount fare. Having expected the topic to come up for some time, users therefore began to inquire about discount when approaching the end of the reservation dialogue. (3) may be replaced by GP1 and GP9 without significant loss.

> 4. Reduce system talk as much as possible during individual dialogue turns.

(4) is near-equivalent to GP8 (brevity).

Summarising, the generic principles (1)-(4) may be replaced by maxims GP1, GP2 and GP5-GP9. These maxims are capable of performing the same task in guiding dialogue design. In fact, as argued, the maxims are able to do the better job because they, i.e. GP6 and GP7, and GP1 and GP9, respectively, spell out the intended contents of two of the principles. This provides corpus-based confirmation of maxims GP1, GP2 and GP5-GP9, i.e. of their stating basic principles of cooperative, task-oriented human-machine dialogue. However, for dialogue design purposes, the maxims must be augmented by *task-specific* or *domain-specific principles,* such as the following.

> 5 (SP3). Provide same formulation of the same question (or address) to users everywhere in the system's dialogue turns.

(5) represents an additional precaution against the occurrence of ambiguity in machine speech. It can be seen as a special-purpose application of GP7 (non-ambiguity).

> 6 (SP1). Be fully explicit in communicating to users the commitments they have made.

> 7 (SP2). Provide feedback on each piece of information provided by the user.

These principles are closely related. The novel cooperativity aspect they introduce is that they require the cooperative speaker to produce a specific dialogue contribution which explicitly expresses an interpretation of the interlocutor's previous dialogue contribution(s), provided that the interlocutor has made a dialogue contribution of a certain type, such as a commitment to book a flight. We propose that these principles be subsumed by GP1 (informativeness).

4.2 Principles lacking Equivalents among the Maxims

The principles discussed in this section appear irreducible to maxims and thus serve to augment the scope of a theory of cooperativity.

4.2.1 Dialogue Partner Asymmetry

Dialogue partner asymmetry occurs, roughly, when one or more of the dialogue partners is not in a normal condition or situation. For instance, a dialogue partner may have a hearing deficiency or be located in a particularly noisy environment. In such cases, dialogue cooperativity depends on the taking into account of that participant's special characteristics. For obvious reasons, dialogue partner asymmetry is important in SLDS dialogue design. The machine is not a normal dialogue partner and users have to be aware of this if communication failure is to be avoided. The following two principles address dialogue partner asymmetry:

> 8 (SP4). Provide clear and comprehensible communication of what the system can and cannot do.

> 9 (SP5). Provide clear and sufficient instructions to users on how to interact with the system.

Being limited in its task capabilities and intended for walk-up-and-use application, our SLDS needs to protect itself from unmanageable dialogue contributions by providing users with an up-front mental model of what it can and cannot do. If this mental model is too complex, users will not acquire it; and if the model is too simplistic, its remaining details must be provided elsewhere during dialogue. (8) adds an important element to the analysis of dialogue cooperativity by aiming at improving user cooperativity. It shows that, at least in human-machine dialogue, cooperativity is a formally more complex phenomenon than anticipated by Grice. In addition to principles stating how a speaker should behave, principles are needed according to which the speaker should consider transferring part of the responsibility for cooperation to the interlocutor. (9) has a role similar to that of (8).

The principles examined in this section introduce a new aspect of dialogue cooperativity, namely partner asymmetry and speaker's consequent obligation to inform the partner(s) of non-normal speaker characteristics. Due to the latter, the principles cannot be subsumed by any other principle or maxim. We propose that (8) and (9) are both *specific* principles subsumed by a new *generic* principle:

> GP10. Inform the dialogue partners of important non-normal characteristics which they should take into account in order to behave cooperatively in dialogue.

4.2.2 Background Knowledge

> 10 (GP11). Take users' relevant background knowledge into account.

GP11 is expressed at the level of generality of Grice's theory. The principle explicitly introduces two notions: the notion of interlocutors' background knowledge and that of possible differences in background knowledge between different user populations and individual users. GP11 appears to be *presupposed* by maxims GP1, GP2 and GP5-GP9 in the sense that it is not possible to adhere to any of these maxims without adhering to GP11. Moreover, in order to adhere to GP11, it is necessary for the speaker to recognise relevant differences among interlocutors and interlocutor groups in terms of background knowledge. Based on this recognition, a speaker either already has built prior to the dialogue, or adaptively builds during dialogue, a model of the interlocutor which serves to guide speaker cooperativity. Increased user adaptivity in this sense is an important goal in SLDS design (Bernsen et al. 1994).

GP11 cannot be reduced to GP1 (informativeness) because, first, GP1 does not refer to the notions of background knowledge and differences in background knowledge among interlocutors. Second, a speaker may adhere perfectly to 'exchange purpose' (cf. GP1) while ignoring the interlocutor's background knowledge. For instance, in the user test a user wanted to order a one-way ticket at discount price. The system, however, knew that discount is only possible on return tickets. It therefore did not offer the discount option to this user nor did it correct the user's misunderstanding. At the end of the dialogue, the frustrated user asked whether or not discount had been granted. Third, as argued above, GP11 is presupposed by maxims GP1, GP2 and GP5-GP9. Grice, however, does not argue that GP1 is presupposed by those maxims whereas he does argue that GP3 (truth) and GP4 (evidence) are presupposed by them (Grice 1975). For similar reasons, GP5 (relevance) (Sperber and Wilson 1987), cannot replace GP11. Informativeness and relevance, therefore, are not only functions of the purpose(s) of the exchange of information but also of the knowledge of the interlocutor.

> 11 (SP8). Provide sufficient task domain knowledge and inference.

(11) may appear trivial as supportive of the design of usable information service systems. However, designers of such systems are continuously confronted with questions about what the system should know and what is just within, or barely outside, the system's intended or expected domain of expertise. The system should behave as a perfect expert vis-à-vis its users within its declared domain of expertise, otherwise it is at fault. In WOZ Iteration 7, for instance, a subject expressed surprise at not having been offered the option of being put on a waiting list in a case in which a flight was already fully booked. We became aware of the problem during the post-experimental interview. However, the subject might just as well have asked a question during the dialogue. Since (11) deals with speaker's knowledge, it cannot be subsumed by GP11. We therefore propose to introduce a new *generic* principle which mirrors GP11:

> GP12. Take into account legitimate partner expectations as to your own background knowledge.

(11), then, is a *specific* principle subsumed by GP12.

> 12 (SP6). Take into account possible (and possibly erroneous) user inferences by analogy from related task domains.

(12) is a *specific* principle subsumed by GP11 (background knowledge). It was developed from examples of user misunderstandings of the system due to reasoning by analogy. For instance, the fact that it is possible to make reservations of stand-by tickets on international flights may lead users to conclude (erroneously) that this is also possible on domestic flights.

> 13 (SP7). Separate whenever possible between the needs of novice and expert users (user-adaptive dialogue).

(13) is another *specific* principle subsumed by GP11. Interlocutors may belong to different populations with correspondingly different needs of information in cooperative dialogue. For instance, a user who has successfully used the dialogue system on several occasions no longer needs to be introduced to the system but is capable of launching on the ticket reservation task right away. A novice user, however, will need to listen to the system's introduction to itself. This distinction between the needs of expert and novice users was introduced in WOZ Iteration 7 when several users had complained that the system talked too much.

4.2.3 Meta-communication

Even if an SLDS is able to conduct a perfectly cooperative dialogue, it will need to initiate *repair and clarification meta-communication* whenever it has failed to understand the user, for instance because of speech recognition or language understanding failure:

> 14 (SP9). Provide ability to initiate repair if system understanding has failed.

(14) states what the cooperative speaker should do in case of failure to understand utterances made by the interlocutor. Our system adheres to (14) in that it communicates its failure to understand what the user just said. (14) cannot be subsumed by GP1 (informativeness) which ignores communication failure. Together with the new specific principles from the user test SP10 and SP11 (Section 2), (14) is a specific principle of human-machine dialogue which may be subsumed by:

> GP13. Initiate repair or clarification meta-communication in case of communication failure.

5. The user test

The set of principles presented in Table 1, apart from SP10 and SP11 (see below), were tested against the user test dialogue corpus (Dybkjær et al., 1996). In the user test the speech recogniser was simulated whereas the other system components were the implemented ones. The system was tested with 12 external users. The 57 recorded and transcribed dialogues were scenario-based and covered the full functionality of the system. We systematically analysed the transcribed dialogues to detect those deviations between expected and actual user utterances that would signal problems of user-system interaction caused by non-cooperative dialogue design. An example is presented below of a piece of dialogue from the user test in which two system cooperativity problems occur (S means system and U means user):

S1: Do you want return tickets?
U1: Yes please.
S2: Are you particularly interested in discount?
U2: Yes.
S3: On which date will the journey start?
U3: Monday.
S4: Monday January 16. At which time?
U4: 7:20.
S5: There is no departure at 7:20. The closest other red departure is at 10:50. Do you want this departure?
U5: No thanks.

Each identified problem in the dialogue corpus was categorised according to which principle had been violated and described in terms of the symptom (**S**), a diagnosis (**D**) and a cure (**C**). Below is shown the description for the two problems in S5.

GP1: Make your contribution as informative as is required (for the current purposes of the exchange).
S: U: interested in discount (red) + out departure time at 7:20. S: no departure at 7:20.
D: The system provides insufficient information. It does not tell that there is a blue departure at 7:20.
C: The system should provide sufficient information, e.g. by telling that there is no red departure but that there is a blue departure at the chosen hour.

SP10: Initiate clarification meta-communication in case of inconsistent user input.
S: U: interested in discount (red) + out departure time at 7:20; S: no departure at 7:20. However, 7:20 does exist but without discount.
D: S gives priority to discount over time without proper reason.
C: S should ask U about priority: 7:20 is not a discount departure. Red discount can be obtained on the departures at x, y and z. Which departure do you want. [If U provides a new departure time: S: do you still want discount. If U: no; S: non-discount departures].

It turned out that almost all of the 86 system dialogue problems identified could be ascribed to violations of the cooperative principles (Bernsen et al., 1996b). We only had to add two specific principles of meta-communication (SP10 and SP11 in Table 1). Since meta-communication had not been simulated during the WOZ experiments, this came as no surprise. The following GPs and SPs were found violated at least once: GPs 1, 3, 5, 6, 7, 10, 11, 12, 13 and SPs 2, 4, 5, 6, 8, 10, 11.

The user test confirmed the broad coverage of the principles with respect to cooperative, spoken user-system dialogue. Less flattering, of course, the test thereby revealed several deficiencies in our cooperative dialogue design.

6 Conclusion

Comparison between our principles and Grice's maxims has shown that there are more generic principles of cooperativity in human-machine dialogue than those identified by Grice. Three groups of principles reveal aspects of cooperative dialogue left unaddressed by the maxims. This produces a total of seven dialogue aspects, each of which is addressed by one or more generic principles (Table 1). Some generic principles subsume specific principles. It may be asked why Grice was not aware of the three generic aspects of dialogue partner asymmetry, background knowledge and meta-communication. It seems obvious that it cannot be because these aspects are absent from human-human spoken dialogue. More plausibly, dialogue partner asymmetry is absent from prototypical cases of human-human dialogue; background knowledge is so pervasive as to be easily ignored; and Grice explicitly was not concerned with dialogue failure pure and simple.

The results from the comparison with Grice's maxims and from the user test suggest that the principles of cooperative spoken human-machine dialogue may represent a step towards a more or less complete and practically applicable set of design guidelines for cooperative SLDS dialogue.

References

Bernsen, N.O., Dybkjær, H. and Dybkjær, L. (1996a). "Cooperativity in Human-Machine and Human-Human Spoken Dialogue," *Discourse Processes*, 21, 2, 213-236.

Bernsen, N.O., Dybkjær, H. and Dybkjær, L. (1996b). "Principles for the Design of Cooperative Spoken Human-Machine Dialogue." To appear in *Proceedings of ICSLP '96*, Philadelphia, October.

Bernsen, N.O., Dybkjær, L. and Dybkjær, H. (1994). "A dedicated task-oriented dialogue theory in support of spoken language dialogue systems design." *Proceedings of ICSLP '94,* Yokohama, September, 875-878.

Dybkjær, L., Bernsen, N.O. and Dybkjær, H. (1996). "Evaluation of Spoken Dialogues. User Test with a Simulated Speech Recogniser." *Report 9b from the Danish Project in Spoken Language Dialogue Systems*. Roskilde University, February.

Fraser, N.M. and Gilbert, G.N. (1991). "Simulating speech systems." *Computer Speech and Language* 5, 81-99.

Grandy, R.E. (1989). "On Grice on language." *The Journal of Philosophy* 86, 10, 514-25.

Grice, P. (1975). "Logic and conversation." In P. Cole & J.L. Morgan (Eds.), *Syntax and semantics Vol. 3: Speech acts* (41-58). New York: Academic Press. Reprinted in Grice, P.: *Studies in the way of words*. Cambridge, MA: Harvard University Press 1989.

Sarangi, A.K. and Slembrouck, S. (1992). "Non-cooperation in communication: A reassessment of Gricean pragmatics." *Journal of Pragmatics* 17, 117-154.

Schegloff, E.A., Jefferson, G. and Sacks, H. (1977). "The preference for self-correction in the organization of repair in conversation." *Language* 53, 361-82.

Searle, J.R. (1992). "Conversation." In Searle, J.R. et al. (Eds.), *(On) Searle on conversation*. Amsterdam: John Benjamin's Publishing Company.

Sperber, D. and Wilson, D. (1987). "Précis of relevance, communication and cognition with open peer commentary." *Behavioral and Brain Sciences* 10, 4, 697-754.