

Conversational Implicatures

Luciana Benotti and Patrick Blackburn

Abstract This paper introduces Paul Grice’s notion of conversational implicature. The basic ideas — the cooperative principle, the maxims of conversation, and the contrast between implicature and presupposition — make it clear that conversational implicature is a highly contextualized form of language use that has a lot in common with non-linguistic behavior. But what exactly is its role? Step by step, we invite the reader to view conversational implicature as a way of *negotiating meaning in conversational contexts*. Along the way, the reader will learn something about the theoretical properties of implicatures, why they are tricky to work with empirically, what can be done with them computationally, and (perhaps) where future research on the topic may lead. But the basic message of the paper is actually quite simple: context and conversational implicature are highly intertwined, and unravelling their interactions is a challenging and worthwhile research goal.

1 Introduction

The notion of **conversational implicature (CI)** is important in both *philosophy of language* [Grice, 1989; Davis, 2010] and *pragmatics* [Horn, 2004; Levinson, 1983], the branch of linguistics which studies how human languages are actually used. The key ideas were first presented in 1967 in Paul Grice’s William James lectures at Harvard, and eventually appeared in the paper *Logic and Conversation* [Grice, 1975, p. 311]. The paper draws our attention to the fact that in typical linguistic exchanges, many things are *meant* without being explicitly *said*, and attempts to explain how this is possible. Let’s start with an example from Grice’s paper:

Luciana Benotti
LIIS Team, FAMAF, Universidad Nacional de Córdoba, Ciudad Universitaria, Córdoba, Argentina.
e-mail: benotti@famaf.unc.edu.ar

Patrick Blackburn
Department of Philosophy, Roskilde University, Denmark. e-mail: patrickb@ruc.dk

- (1) *Man standing by his car: I am out of petrol.*
Passer-by: There is a garage around the corner.

Grice's analysis runs as follows. The utterance made by the passer-by (let's call him B) wouldn't have been relevant (to the conversational exchange) if B knew that the garage was closed or that it had run out of petrol. If B is a local person who knows about local garages, it is thus reasonable to assume that B is directing the man standing by the car (let's call him A) to a garage that is open and currently selling petrol. That is, according to Grice, during the exchange (1), B made the conversational implicature (2):

- (2) *The garage is open and has petrol to sell.*

Well, so far, so good. But isn't this just the sort of thing that goes on all the time? Where's the mystery? If this is all there is to conversational implicature, what exactly requires explanation? What makes the concept so important? And why is it considered difficult?

As we hope to make clear, even this little exchange conceals many problems. Conversational implicature involves highly contextualized inferences that draw on multiple sources of information. For instance, in the garage example, presumably the visual information provided by A standing beside his stationary car plays an important role in initiating the exchange. Moreover, by their very nature — we will soon explain what we mean by this — implicatures tend to be resistant to the usual tools of empirical linguistic investigation. Furthermore, they are ubiquitous: get two people talking, and the conversational implicature flies thick and fast. In short, Grice's garage example is the tip of a large iceberg concerning meaning and inference in context. In the pages that follow, we first show just how large the iceberg is below the water's surface, and then discuss recent attempts to reveal its contours theoretically, empirically and computationally.

We proceed as follows. In Section 2 we present conversational implicature as a form of contextualized inference, and make a first pass at explaining why it tends to resist the usual tools of empirical linguistic analysis. In Section 3 we sketch some of Grice's ideas on the subject, notably his **cooperative principle** and his **conversational maxims**. In Section 4 we note five key theoretical properties of conversational implicature, and in Section 5 we discuss further difficulties with empirical work. This leads us to one of the main points we wish to make: conversational implicature is a form of **negotiation**. In Section 6, we sketch the relevance of **clarification requests** to this conception, and in Section 7 we briefly note some recent computational work. Section 8 concludes with a nod to the future.

Computational implicature is a huge subject, one that has been investigated from many angles. We cannot hope to cover them all, and have opted instead to present a birds-eye-view of relevance to researchers interested in context. But there are many good points of entry to the topic. For a start, Grice's own *Logic and Conversation* is a must: it is clear, accessible and covers many topics we do not have space to consider. Moreover, excellent surveys exist: we particularly recommend Chapter 3 of Levinson's textbook *Pragmatics* [Levinson, 1983]. Furthermore, searching for

‘conversational implicature’ in the (free) online *Stanford Encyclopedia of Philosophy* will lead to several informative and authoritative articles (for example, [Davis, 2010]) and detailed bibliographical information.

2 Implicature as contextual inference

In this section we have two main goals: to convince the reader that conversational implicatures are a highly general form of contextualized inference, and to make a first attempt to explain why conversational implicatures tend to be resistant to standard empirical linguistic techniques. A good way into the discussion is to compare conversational implicatures with what linguists call **presuppositions**.

Presuppositions are another ubiquitous form of inference, one that all known human languages exploit. Consider the following sentence. Imagine it is uttered to you out of the blue, with no prior conversational context:

(3) *Anthony regrets that Brenda is pregnant.*

Now, you know nothing about Anthony or Brenda or who they are or what their life is like. But by reading this sentence, if you trust its source, you have internalized one piece of information: namely, that Brenda is pregnant. But now consider the following sentence:

(4) *Anthony does not regret that Brenda is pregnant.*

Once again, simply by reading this sentence, you have once again internalized the information (or: **accommodated** the information, as a linguist would say) that Brenda is pregnant. And this is surprising because it shows that we are dealing with a very strange form of inference indeed. It is certainly not an ordinary logical inference. In ordinary logical inference, replacing a positive premise with its negation will typically destroy the inference. But here replacing ‘regrets’ with ‘does not regret’ has no effect: in both cases the (immediate and intuitive) inference to Brenda’s pregnancy goes through.

Such presuppositions are ubiquitous, and their ability so survive negations is one of their better known characteristics. Consider the following pair of sentences

(5) *Candy knows that Dave is dead.*

Candy does not know that Dave is dead.

Once again — positive or negative — we infer the same information, no matter the context: Dave is dead.

Now, presupposition is an interesting phenomena. Like conversational implicature it is a ubiquitous fact about language in action. But in at least one crucial respect it is simpler than implicature: *we can point to concrete linguistic triggers that set the presuppositional process in motion, no matter the context*. In the first pair of examples, the trigger is the word ‘regret’. In the second pair of examples, the trigger is the word ‘knows’. Part of the effect of using these words —

part of their *meaning* — is that they induce presuppositional phenomena of this sort, and every competent language user instinctively knows how to make good use of this. Now, it should be stressed that presupposition is not a simple phenomena, nonetheless, we now have detailed theoretical accounts of it [Beaver, 2001; Geurts, 1999] and understand how to incorporate presuppositional phenomena into computational work [Venhuizen *et al.*, 2013; Bos, 2003; Blackburn *et al.*, 2001]. And much of this understanding stems from the basic fact just noted: we can trace presuppositions to specific linguistic locations, namely, certain trigger words (and grammatical constructions).

But conversational implicatures *can't* be linguistically localized in the same way. As we said earlier, conversational implicatures are highly contextualized inferences, capable of exploiting multiple information sources. Consider the garage example. First, and most obviously, it rests on an assumed *common knowledge context*: both A and B need to share the knowledge that “petrol can be bought at open garages which have not run out of petrol”. Secondly, this example draws on the *situational context*, most obviously on B knowing that the garage is around the corner and A can walk there. But a lot also hinges on the fact that we are in a *conversational context*. It is quite obvious that the implicature could not have been triggered without considering the immediate conversational context: if A had said “Where do you come from?” instead of “I am out of petrol”, B’s implicature would have been quite different. But even more basic components of the *interaction context* are crucial: A and B take for granted that the other is a language user, with intentions and goals, who may be prepared to take part in a cooperative exchange in order to overcome undesired states such as “being out of petrol”. In fact, as we shall soon discuss, Grice called one of the central principles driving conversational implicature the cooperative principle. Moreover, other contextual parameters may play a role in how the scenario unfolds, such as information about race, gender and status (is A unshaven, sloppily dressed, and standing by a beat-up old car, or is he wearing a suit and tie and standing by a spiffy new Mercedes?).

So: Conversational implicatures are a highly contextualized form of inference. Moreover (unlike presuppositions) there is no simple linguistic trigger with relatively well defined rules which we can analyze to ‘solve the implicature problem’ in the general case. There is no linguistic trigger in the garage example that sets the implicature process in motion. For example, it would be highly implausible to claim that B’s behavior is induced by the word ‘garage’, on any other word in the exchange for that matter. Words like ‘regrets’ and ‘know’ induce presuppositions, but Grice’s little scenario has more complex origins. The same scenario would work if we used the words ‘service station’ or ‘petrol station’ instead of ‘garage’. The initial trigger may well have been non-linguistic: the sight of A (perhaps in his suit and tie) standing forlornly by his new Mercedes may have been the spark that set the little scene in motion.

Conversational implicatures are, in general, not closely tied to the inner-workings of the lexical system in the way that presuppositions are. Rather, they are a type of behavior exemplified by agents with intentions and goals — but special agents, namely human beings, who have a highly refined form of behavior in their arsenal:

linguistic behavior. We might say: *conversational implicature is a big share of the meaning conveyed by goal seeking, linguistically competent agents when they interact in contexts*. Hopefully this goes some way towards explaining why conversational implicatures are interesting and important — and why they are highly relevant to the study of context. But note the downside: it also tells why they are likely to be resistant to straightforward empirical linguistic investigation. We can't compile a list of trigger words (as we can for presupposition) and explore their effects in a corpus of example sentences. As we have said, in general, it is difficult or impossible to point to critical lexical or grammatical triggers when it comes to conversational implicature, because so much of what is going on in conversational implicature is not specifically linguistic behavior; rather, *it's part of the general human behavior displayed by linguistically adept agents*. It spills over into, draws richly on, and is guided by, the surrounding context in all its crazy variety. So empirical investigation is a tricky business, a point we will return to later.

3 Grice on conversational implicature

We now outline Grice's account of conversational implicature, and in particular his **cooperative principle** and the **conversational maxims**. We draw attention to the possibility of **observing**, **flouting** and **violating** maxims; these are not merely practical distinctions, but are also helpful in understanding Grice's motives. We further note that the conversational maxims are linked to non-linguistic behavior, that they can be used to classify conversational implicatures, and draw the reader's attention to both **relevance implicatures**, the conversational implicatures that embody contextual inference in its most general form, and **scalar implicatures**, those that probably embody it least. Our discussion will establish some standard ideas and terminology and pave the way for our work in subsequent sections. Page references here are to the version of *Logic and Conversation* in [Grice, 1989].

At the heart of Grice's discussion lies the following principle (see page 26):

The cooperative principle: Make your contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.

In addition, he presents the following maxims (pages 26–27):

Maxims of Quantity:

- (i) Make your contribution as informative as is required (for the current purposes of the exchange).
- (ii) Do not make your contribution more informative than is required.

Maxims of Quality: Try to make your contribution one that is true.

- (i) Do not say what you believe to be false.
- (ii) Do not say that for which you lack adequate evidence.

The Maxim of Relevance: Be relevant.

Maxims of Manner: Be perspicuous.

- (i) Avoid obscurity.
- (ii) Avoid ambiguity.
- (iii) Be brief (avoid unnecessary prolixity).
- (iv) Be orderly.

Grice calls the cooperative principle a “rough general principle” and it is clear from his discussion on pages 26–28 that the maxims are a tentative attempt to understand how human beings interact in conversation. But at this point, the reader may have doubts. Aren’t the cooperative principle and the maxims far too neat and tidy? They seem to describe an ideal world of effective, rational, maximally cooperative conversational interaction. And, all too obviously, real life just isn’t like that at all.


But it is a mistake to dismiss Grice and conversational implicature on these grounds. Grice is *not* suggesting that all human conversational interactions live up to these principles, or even that it would be better if they did. Rather, *Grice is trying to indicate the existence of deep-seated norms of conversational interaction*. Humans are social animals. They interact. Moreover they are social animals graced (or cursed) with the power of speech. Grice is tentatively suggesting (see page 29) that assumptions somewhat like the cooperative principle or the maxims must guide this process. And of course — as Grice himself points out — *what is expected of others swiftly becomes a resource to be exploited*. We may communicate by *observing* the maxims — but sometimes we can communicate more effectively by deliberately *flouting* them; we will give examples of both strategies shortly. Indeed, sometimes we may choose to distance ourselves still further from communicative norms and deliberately *violate* the maxims: liars, tricksters, ironists and cheats thrive in the ambience of the socially expected.

Grice views his conversational maxims as direct analogs of norms governing the way we cooperate in *non-linguistic* settings. He is quite explicit on this point. As he says on page 28:

[O]ne of my aims is to see talking as a special case of purposive, indeed, rational behavior...

and then lists analogies between the conversational maxims and physical actions. For example, with regard to the Maxim of Quality which says we should seek to be truthful, he points out that if we are cooperating to make a cake, and I need sugar, then I expect you to pass me sugar, and not (say) salt, and that if I need a spoon for stirring the cake mixture, I expect you to hand me a real spoon not (say) a trick spoon made of rubber. The strong analogy that Grice draws between linguistic and non-linguistic behavior is important, and we shall return to it when we discuss implicature as negotiation.

The maxims are also useful in that they give us a (somewhat rough-and-ready) way of classifying conversational implicatures. Consider, for example, the maxims of quantity. This has given rise to an extensive literature on what are now called **quantity** or **scalar implicatures** (see, for example, [Hirschberg, 1985; Geurts, 2011]). Let’s briefly consider an example which we will return to in more detail when we discuss empirical work. Consider the following image:



Suppose I then say to you:

(6) *Some of the B's are in the box on the left.*

then I would often be taken as having implicated that

(7) *Not all of the B's are in the box on the left.*

Why? The point is this. If I am observing the maxims, and in particular the Maxim of Quantity that tells us to *make your contribution as informative as is required* then I must be making the strongest claim possible. Logically, the claim that *some* of the B's are in the box on the left is compatible with the (logically stronger) claim that they *all* are. So why did I not make the stronger claim? Well, assuming that I am observing to the maxims, this must be because I was not in a position to (truthfully) do so (and looking at the image we see that I was not). Hence I implicated that *not all* the items are in the box. In short, assuming that conversational agents are *observing* to the maxims gives us explanatory power: it enables us to appeal to and reason about communicative goals and intentions.

But so does *flouting*. Stephen Levinson has a nice example involving Maxims of Manner (see page 104 of [Levinson, 1983]):

(8) *A: Let's get the kids something.*

B: Okay, but I veto I-C-E-C-R-E-A-M-S

Why on earth would B spell out the word ice creams? This is not perspicuous presentation; indeed it's (deliberately) obscure and prolix! And of course, every parent knows why: the maxim has been flouted because B's message is not simply "Yes, I agree we should get the kids something", but the more desperate "For God's sake let's not get them whining for ice creams!", though truly experienced parents have also learnt that the flouting-via-spelling strategy doesn't work for very long. . .

The logic of flouting is interesting. Once again, it involves appealing to and reasoning about communicative goals and intentions, but in a more subtle way: a flouting seems to function as an *invitation* to look beyond the surface level of maxim failure and to search for a deeper vein of cooperativity. Given such possibilities, it is hardly surprising that many authors have found game theory (and related disciplines which focus of strategic thinking) useful tools for exploring conversational implicature; see [van Rooij, 2011] for a useful overview.

Finally, we remark that we can now see that conversational implicatures come in all shapes and sizes. For example, the sort of scalar implicature involved in *some of the B's are in the box on the left* example is relatively specific, and closely tied to the meaning of the word *some* — though as we shall learn later, even this seemingly context independent example is not as innocent as it looks. Others, such as the garage example, with which we started, are more general and commonplace. This is because the garage example is a **relevance implicature**, governed simply by the splendidly general: *Be relevant!* Some authors (notably Deirdre Wilson and Dan

Sperber [Wilson and Sperber, 2004]) have insisted that the notion of relevance is the real gold in Grice's work. Indeed, Grice himself seems to be partly of this opinion. As he remarks on page 27:

Though the maxim itself is terse, its formulation conceals a number of problems that exercise me a good deal: questions about what different kinds and focuses of relevance there may be, how these shift in the course of a talk exchange, how to allow for the fact that the subjects of conversation are legitimately changed, and so on. I find the treatment of such questions exceedingly difficult, and I hope to revert to them in later work.

The problems involving relevance implicature remain exceedingly difficult to this day; some of the hardest problems of contextual inference live there.

4 A little theory

So far, our discussion has been relatively informal; in this section we make it more precise. We first give a definition of conversational implicature; this won't play a role in our subsequent work, but it will give the reader a taste of the theorizing typical of formal pragmatics and it constitutes a formal definition in which the role of the context is made explicit. Following this, we will briefly discuss five basic properties of conversational implicatures; three of them will play an important role in the emergence of the concept of negotiability in the following section.

The following definition is adapted from [Hirschberg, 1985] by making explicit the role of context and the role of the conversation participants as agents that can modify context:

Definition 1. Proposition q is a conversational implicature of utterance U by agent B in context C if and only if:

- (i) B believes that it is mutual, public knowledge in C of all the discourse participants that B is obeying the cooperative principle.
- (ii) B believes that, to maintain (i) given U , the hearer will assume that B believes q holds in C or can achieve q by modifying C .
- (iii) B believes that it is mutual, public knowledge of all the discourse participants that, to preserve (i), one must be assume that B believes q holds in C or can achieve q by modifying C .

What does this mean? Let's return to the garage example, where B implicated *The garage is open and has petrol to sell* (that is, q) by uttering *There is a garage around the corner* (that is, U). If B believes it is common knowledge to all the participants (namely A and B) that B is obeying the cooperative principle, and B further believes that A will assume, on the basis of B 's utterance, that the garage is open and has petrol to sell and that A can get to the garage, then this implicature is successful. It's a tricky definition, and we won't pause to discuss it further, but do note the following: the *form* of the definition, with its explicit appeals to an agent's beliefs and what is mutual public knowledge make it clear that implicature is being modeled as a form of epistemic reasoning that draws on contextual knowledge.

Now let's examine the *properties* of conversational implicatures, for these *will* play an important role in our discussion. Hirschberg argues that we need to insist that conversational implicatures have the following five: 1) deniability, 2) reinforceability, 3) non-lexicality, 4) non-detachability, and 5) calculability. In the light of our previous discussion, many of these properties will probably appear familiar, or at least natural.

First, conversational implicatures are **deniable** without contradiction. Let's stick with Grice's garage example. B can append material that is *inconsistent* with the implicature — for example, B can add *but I don't know whether it's open* — and the resulting exchange will *not* be contradictory. Indeed, the resulting exchange would be extremely natural: B would be implicating potentially useful information about the garage, but then expressing some reservation. This is something that happens all the time.

Second, note that B can also add material to the exchange that *explicitly asserts* the implicature — *and I know it's open* — without any sense that he is repeating himself. That is, B can **reinforce** the implicature without redundancy. Indeed, once again, this is very natural language use: we implicate the extra information and then (if it seems important) ram the message home to make sure our conversational partner gets the point.

Third, implicatures are **non-lexical**: they do *not* trace back to particular lexical items. We have already mentioned this. For example, in Grice's garage example, the implicature is not triggered by any particular word in the exchange (such as 'garage') but is a result of the overall semantic content.

Fourth — again something we have discussed — since an implicature is attached to the semantic content of what is said and not to the particular lexical items involved, a conversational implicature **cannot be detached** from the utterance simply by changing the words of the utterance by synonyms. B can replace each word in his utterance with a word with the same meaning — he can say *petrol station* or *service station* instead of *garage* — and the implicature will still go through. Note that non-detachability and non-lexicality are not really two independent properties: non-lexicality can only be tested by evaluating non-detachability. Basically, these two properties are another way of getting at the basic point that conversational implicatures are *not* part of the conventional meaning of the words uttered, but depend on features of the conversational context.

Fifth and last, conversational implicatures are traditionally considered to be **calculable**. Calculability means that the addressee should be able to infer the implicatures of an utterance. For example, in the garage example, A should be able to infer that B conversationally implicates that the garage is open and has petrol to sell.

Three of these properties will shortly return in the guise of *negotiability*. But to properly motivate this, we first need to take a closer look at some intriguing empirical issues.


5 Towards Negotiability

The empirical literature on conversational implicature is based almost entirely on evidence obtained using the **inference method** [Chemla, 2009], a pragmatic-level analog of the introspective method, which has been traditionally used in linguistics and philosophy (in essence, the traditional approach is to obtain native-speaker judgments on linguistic examples). However, Geurts and Pouscoulous [2009] have shown that the inference method is a biased tool when it comes to gathering data on conversational implicatures. Let's briefly consider Geurts and Pouscoulous's (henceforth G&P) argument.

Consider the scalar implicature example we gave in Section 3. Experimenters using the traditional inference paradigm might ask the experimental subjects whether they think that sentence (9a) implies the sentence (9b):


- (9) a. *Some of the B's are in the box on the left.*
 b. *Not all of the B's are in the box on the left.*

Now, G&P argue that to ask oneself — or an experimental subject — whether or not (9a) implies (9b) is already to suggest that it *might* be implied. That is, the experimenter's question raises the issue of whether or not all of the B's are in the box on the left, and makes it relevant to establish whether this is the case. Indeed, consider the image we used back in Section 3:




In order to explain the example, we carefully placed a B in the right-hand box. But this makes the issue of whether *all* the Bs are in the box on the left very much at stake. G&P claim that such inference experiments do not necessarily tell us anything about how (9a) is interpreted in situations where (9b) is *not* at stake.

In order to support their claim, G&P compare the inference method with what they call the **verification method**. In the verification version of the previous experiment, subjects have to decide whether (9a) correctly describes the situation shown below. Someone who interprets (9a) as implicating (9b) should *deny* that (9a) gives a correct description of the picture (for note: no Bs in the right-hand box):



Intriguingly, participants derived conversational implicatures almost twice as frequently given the inference condition (62%) as in the verification condition (34%)! The inference task *does* increase the rate of detection of conversational implicatures, and the effect is substantial. Moreover, G&P show that the effect is even more evident in complex sentences such as (10a) describing the following situation:



Someone who interprets (10a) as implicating (10b) should *deny* that (10a) gives a correct description of the picture (for note: one of the squares is connected to all of the circles).

- (10) a. All the squares are connected with some of the circles.
 b. All the squares are connected with some but not all of the circles.

Studies carried out using these more complex sentences result in participants deriving the conversational implicature (10b) from (10a) in 46% of the cases with the inference method, and in 0% of the cases with the verification method!. As we have said several times, the empirical study of conversational implicatures is not easy. G&P’s work provides a useful warning of another source of error.

But what is the lesson we should draw? Is it simply that the verification method is better than the inference method? We think not: if inference methods fail to tell us anything about how utterances should be interpreted when the alleged implicature is *not* at stake, the verification method used by G&P fails in the opposite direction; it does not tell us anything when the issue *is* at stake. In our view, what G&P’s experiments really show is that *whether the issue is at stake or not* is crucial. Moreover their experiment is powerful because it involves scalar implicatures.

Scalar implicatures are usually considered ‘better behaved’ than (say) relevance implicatures. Clearly the above inferences have something to do with how we interpret the logical words ‘some’ and ‘all’, and such inferences are often felt to be ‘less contextually sensitive’ than full-blown relevance implicatures such as the Grice’s garage examples. And this makes G&P’s results even more fascinating: *even allegedly simply and well-understood implicatures such as scalar implicatures involving ‘some’ and ‘all’ can be dramatically manipulated by playing with the context.* G&P’s experiments simply change one pragmatic factor — whether something is at stake or not — and the effects are plain to see. Conversational implicature really *does* involve context — and this is true even of scalar implicatures.

And this brings us, at last, to the heart of the paper. In conversation, whether an issue is at stake or not is naturally determined by what Grice calls “the accepted purpose or direction of the talk exchange”. And the word that sums up best how this is done is the word **negotiation**. Conversational implicature is not a museum of dead examples that needs explaining — it is an practice, an activity, a way of

interacting. Instead of asking (say) “How do we explain Grice’s garage example?” we need to ask: how do we change it? How do we extend it? How do we see it as a living, ongoing, conversational interaction?

Recall the implicature from the garage example — *the garage is open and has petrol to sell*. Also recall the properties of *deniability* and *reinforceability* we isolated in the previous section. Actually, these are not so much two independent properties as two sides of the same coin: *conversational implicatures are not explicitly stated: they are not actual contributions to an exchange, but only potential contributions that the speaker can either deny or reinforce*. Moreover, if we also bring *calculability* into the picture then we see that not only the speaker but also the hearer can deny or reinforce the conversational implicatures. For instance, A can naturally continue Grice’s garage example with — *I went there, it’s closed* — thereby denying the conversational implicature. But A can also continue the exchange by reinforcing the implicature — *oh, and it must be open because it’s only 3pm!*

That is, as promised, deniability, reinforceability and calculability can be summarized by saying that computational implicatures are *negotiable*. Speaker and hearer negotiate — they deny, reinforce, calculate, and switch roles as the conversation precedes. In this fashion they explore the issue at stake until (all being well) they are satisfied. Conversational implicature is truly *conversational*. In the following section we will make these ideas concrete with the help of real conversational examples.

6 Conversational examples

In the previous section we argued that deniability, reinforceability and calculability can be summed up by saying that conversational implicatures are negotiable. The hearer can infer the implicatures of an utterance but cannot be 100% sure that the speaker meant them (and the speaker knows this), so both the speaker and the hearer can talk — negotiate — without redundancy or contradiction.

This account is theoretically satisfying, but it is of additional interest because it suggests a novel *empirical* approach to conversational implicature. The idea, first suggested in [Benotti, 2009], is to track the negotiation process by noting what linguists call **clarification subdialogues**. Take, for instance, the clarification request which can naturally follow Grice’s garage example:

(11) A: *and you think it’s open?*

B will have to answer and support the implicature (for instance with “yes, it’s open till midnight”) if he wants to get it added to the common ground; otherwise, if he didn’t mean it, he can well reject it without contradiction with “well, you have a point there, they might have closed”.

Let’s see some examples of clarification, rejection and reinforcement of conversational implicatures in real dialogues between two players of a multiplayer online game [Stoia *et al.*, 2008]. In this game, the player DF (Direction Follower) is collaborating with the player DG (Direction Giver) in order to reach the goal of the game.

Since the DF knows how well the game world and how to reach the goal, most of her contributions come in the form of instructions on how to perform actions in the game world. Situated instructions are a kind of language that maintains a tight link between physical actions and conversational maxims and, offer a concrete scenario where to get strong intuitions about the mechanisms governing conversational implicatures and how they interact with different kinds of contextual information.

Clarifying conversational implicatures.

In the dialogue below, the players are trying to move a picture from one wall to another. The utterance being interpreted is the one uttered by the DG in (1). Using the *common knowledge context* that “in order to put something somewhere you need to have it in your hands first” and the *situational context* of “DF is not holding the picture”, the DF infers a conversational implicature that involves *picking up the picture*. This implicature, foreshadowed by (2) and (3), is finally made explicit as a clarification request in (4):

DG(1): well, put it on the opposite wall

DF(2): ok, control picks the .

DF(3): control's supposed to pick things up and .

DF(4): am I supposed to pick this thing?

Rejecting conversational implicatures.

Here, the players are trying to find a gun that is supposed to be hidden. As in the previous example, the inference of the conversational implicature in (2) is allowed by the *common knowledge context* that “in order to take something from a container is has to be open”. However, the *situational context* of “the cabinet is closed” was unknown by the DG and is an obstacle for carrying out instruction (1). As a result, DF rejects the implicature “the cabinet is open” by uttering (2):

DG(1): take the gun from the cabinet .

DF(2): the cabinet is closed

Reinforcing conversational implicatures.

In this fragment, DG asks DF to go back where he started. This instruction relies on its previous *conversational context*, in fact, it requires that the DF remembers the beginning of the game that was also the beginning of the conversation. And the DF remembers, and thus is able to infer a sequence of conversational implicatures that he incrementally grounds making them explicit in (2), (4), and (6) while he is executing it. Uttering the implicatures of an utterance is a frequently used method for performing acceptance acts. As before, the conversational implicatures of (1) are strongly dependent of the *situational context*. In particular, note that if the DF was already near the steps at the moment of uttering (1), he would only need to infer (6), not (2) and (4) as well.

DG(1): let's go back where you started . so
DF(2): ok . so I have to go back through here .
DG(3): yeah
DF(4): and around the corner .
DG(5): right
DF(6): and then do I have to go back up the steps
DG(7): yeah
DF(8): alright, this is where we started

Conversational implicatures are negotiable, and dialogue provides an intrinsic mechanism for carrying out negotiations of meaning, namely clarifications. As these examples show, conversational implicatures are a rich source of clarification subdialogues. When talking, we do not make explicit all those things from the world that serve “the accepted purpose or direction of the talk exchange in which we are engaged”, but only those that are necessary for the addressees to fill in the details in the way we intend them to. If the addressee is not sure that he has filled in the details in the way we intended them to, if he is not sure he had interpreted our intended conversational implicatures, a clarification subdialogue emerges.

A pleasant aspect of this approach is that it brings us closer to standard linguistic techniques: essentially the idea is to find conversational implicatures in corpora by identifying clarification subdialogues. The characteristics and functions of clarification subdialogues have been deeply studied by dialogue system researchers [Gabsdil, 2003; Purver, 2004; Rodríguez and Schlangen, 2004; Rieser and Moore, 2005; Skantze, 2007] and sociolinguists. In discourse analysis, this has been an even more favored theme for almost three decades now; see [Schegloff, 1987] as a representative example. This opens a myriad of possibilities for studying conversational implicatures under a new light.

The approach has the additional merit of anchoring the study of conversational implicatures in their natural habitat, that is, in naturally occurring dialogue. As we have seen, it is often controversial whether something is a conversational implicature or not (people have different intuitions, which is not surprising given that people often have different contexts in mind). But in naturally occurring dialogue, the follow up conversational turns often provides good evidence of the implicatures that have been made, simply because they make implicatures explicit.

7 Recent Computational work

Recall the five properties of conversational implicature we isolated in Section 4. According to [Levinson, 1983], non-detachability (which is closely related to non-lexicality) serves to distinguish conversational implicatures from presuppositions. As we pointed out in Section 2, words like ‘knows’ and ‘regret’ carry, as a part of their conventional meaning, the ability to trigger presuppositions. However conversational implicatures are not part of the conventional meaning of the words involved

but depend on features of the conversational context — indeed (as the experiments of G&P show) this is true even of seemingly-simple scalar implicatures.

But non-detachability (and non-lexicality) have a computational downside. While progress has been made on computing presuppositions, computing implicatures is a heavier task, for work that models the inference of conversational implicatures has to model the (linguistic and extra-linguistic) conversational context. However, in spite of the difficulties involved, there has been some recent work on the topic, we shall briefly note here. These approaches differ not only on the conversational context elements that they model but also on the inference techniques used.

First, [Vogel *et al.*, 2013b] show that a team of interacting agents collaborating to maximize a global reward using only local information reach implicature-rich interpretations simply as a by-product of the way they reason about each other beliefs. The technique used to model the interaction is multi-agent Decentralized POMDP which only use local information to maximize joint utility. A similar approach is taken by [DeVault and Stone, 2009] who instead of POMDPs use Maximum Entropy models over abductive interpretations to model the maximization of interpretation success. These two approaches provide a rich representation of the conversational context and the goals shared by the agents. Using such techniques the cooperative principle and the associated maxims of relevance, quality, and quantity have been shown to emerge from agent interaction because they maximize utility. For instance, agents do not lie to each other and do not give more information than necessary to make collaboration more effective. The main problem of these two approaches is their intractability, although [Vogel *et al.*, 2013a] use cognitively-inspired heuristics to simplify the otherwise intractable task of reasoning jointly about actions, the environment, and the nested beliefs of other actors.

An alternative approach is not to explicitly reason about the nested beliefs but to model and reason on the common ground using cheap causal reasoning techniques such as classical planning [Benotti, 2010; Benotti and Blackburn, 2011]. Such have been used to computationally explore the idea of conversational implicature as negotiation, and have also proven to be useful in large scale practical applications [Smith, 2012]. The main problem with this approach is the limited expressive power offered by the classical planing paradigm which is able to model only some kinds of conversational implicatures. An interesting way around this problem would be to use Dynamic Epistemic Planning [Bolander and Andersen, 2011], though, at least at present, the additional expressivity this offers is likely to come with a hefty price in terms of tractability.

8 Towards the future

In this paper we introduced conversational implicature, contrasted this notion with presupposition, explained the cooperative principle and the maxims of conversation on which it rests, and invited the reader to view conversational implicature as a way of *negotiating meaning in conversational contexts*. We discussed recent theoretic,

empirical and computational work, and hope it is now clear to the reader that context and conversational implicature are deeply intertwined, and that unraveling their interactions is a worthwhile research goal.

But what of the future? Where is research in this field likely to go? There has been interesting recent work in pragmatics that points towards new research directions. In this paper we simply contrasted conversational implicature with presupposition in an attempt to make clear just how important context is to conversational implicature. But presupposition and conversational implicature are only two of many phenomena, and two strands of recent work are clarifying our understanding here. First, in [Simons *et al.*, 2011; Tonhauser *et al.*, 2013] and related publications, a determined attempt has been to more accurately map the contours of what are called projectable inferences (presuppositions are an important example here, but there are many others). On the other hand, in [Beaver and Zeevat, 2007] there are signs that linguists are beginning to better understand the process of accommodating new information that is inferred in context. These lines of work both hint at deeper theoretical analyses of some of the key concepts in pragmatics, so the intellectual landscape may well look very different from now.

Nonetheless, despite these new insights, it seems clear that Paul Grice got several things right. We do use language in conversational contexts, this does involve reasoning about the intentions and goals of those we interact with, and this may well be a norm-governed process open to exploitation. And — above all — when we negotiate with our conversational partners and adversaries, it seems that we draw upon and reason about the information made available by the context in deep and intriguing ways. Paul Grice: still relevant after all these years!

[Zeevat, 2014]

References

- [Beaver and Zeevat, 2007] David Beaver and Henk Zeevat. Accommodation. In *The Oxford Handbook of Linguistic Interfaces*, pages 503–539. Oxford University Press, 2007.
- [Beaver, 2001] David Beaver. Pragmatics, and that’s an order. In *Logic, Language and Visual Information*. CSLI Publications, 2001.
- [Benotti and Blackburn, 2011] Luciana Benotti and Patrick Blackburn. Classical planning and causal implicatures. In Michael Beigl, Henning Christiansen, Thomas R. Roth-Berghofer, Anders Kofod-Petersen, Kenny R. Coventry, and Hedda R. Schmidtke, editors, *Modeling and Using Context*, volume 6967 of *Lecture Notes in Computer Science*, pages 26–39. Springer Berlin Heidelberg, 2011.
- [Benotti, 2009] Luciana Benotti. Clarification potential of instructions. In *Proceedings of the 2009 SIGDIAL Conference on Discourse and Dialogue*, pages 196–205, 2009.
- [Benotti, 2010] Luciana Benotti. *Implicature as an Interactive Process*. PhD thesis, Université Henri Poincaré, INRIA Nancy Grand Est, France, supervised by Patrick Blackburn, 2010.
- [Blackburn *et al.*, 2001] Patrick Blackburn, Johan Bos, Michael Kohlhase, and Hans de Nivelle. Inference and computational semantics. In Harry Bunt, Reinhard Muskens, and Elias Thijsse, editors, *Computing Meaning*, volume 2, pages 11–28. 2001.
- [Bolander and Andersen, 2011] Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21:9–34, 2011.

- [Bos, 2003] Johan Bos. Implementing the binding and accommodation theory for anaphora resolution and presupposition projection. *Computational Linguistics*, 29(2):179–210, 2003.
- [Chemla, 2009] Emmanuel Chemla. Universal implicatures and free choice effects: Experimental data. *Semantics and Pragmatics*, 2(2):1–33, May 2009.
- [Davis, 2010] Wayne Davis. Implicature. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Stanford, CA: CSLI winter 2010 edn., 2010.
- [DeVault and Stone, 2009] David DeVault and Matthew Stone. Learning to interpret utterances using dialogue history. In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 184–192, Athens, Greece, 2009. Association for Computational Linguistics.
- [Gabsdil, 2003] Malte Gabsdil. Clarification in spoken dialogue systems. In *Proceedings of the AAAI Spring Symposium on Natural Language Generation in Spoken and Written Dialogue*, pages 28–35, Palo Alto, California, 2003.
- [Geurts and Pouscoulous, 2009] Bart Geurts and Nausicaa Pouscoulous. Embedded implicatures?!? *Semantics and Pragmatics*, 2(4):1–34, July 2009.
- [Geurts, 1999] Bart Geurts. *Presuppositions and Pronouns*, volume 3 of *Current Research in the Semantics/pragmatics Interfaces*. Elsevier, Amsterdam, 1999.
- [Geurts, 2011] Bart Geurts. *Quantity implicatures*. Cambridge University Press, 2011.
- [Grice, 1975] Paul Grice. Logic and conversation. In P. Cole and J. Morgan, editors, *Syntax and Semantics*, volume 3, pages 41–58. Academic Press, New York, 1975.
- [Grice, 1989] Paul Grice. *Studies in the Way of Words*. Harvard University Press, 1989.
- [Hirschberg, 1985] Julia Linn Bell Hirschberg. *A theory of scalar implicature*. University of Pennsylvania, 1985.
- [Horn, 2004] Larry Horn. Implicature. In *Handbook of Pragmatics*, pages 3–28. Blackwell, Oxford, 2004.
- [Levinson, 1983] Stephen Levinson. *Pragmatics*. Cambridge textbooks in linguistics, 1983.
- [Purver, 2004] Matthew Purver. *The Theory and Use of Clarification Requests in Dialogue*. PhD thesis, King’s College, University of London, 2004. Supervised by Jonathan Ginzburg.
- [Rieser and Moore, 2005] Verena Rieser and Johanna D. Moore. Implications for generating clarification requests in task-oriented dialogues. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics, ACL ’05*, pages 239–246, Ann Arbor, Michigan, 2005. Association for Computational Linguistics.
- [Rodríguez and Schlangen, 2004] Kepa Rodríguez and David Schlangen. Form, intonation and function of clarification requests in german task oriented spoken dialogues. In *Proceedings of 8th Workshop on the Semantics and Pragmatics of Dialogue (SEMDIAL)*, pages 101–108, 2004.
- [Schegloff, 1987] Emanuel Schegloff. Some sources of misunderstanding in talk-in-interaction. *Linguistics*, 8:201–218, 1987.
- [Simons et al., 2011] Mandy Simons, Judith Tonhauser, David Beaver, and Craig Roberts. What projects and why. In *Proceedings of 20th Semantics and Linguistic Theory Conference (SALT 2010)*, volume 20, pages 309–327, 2011.
- [Skantze, 2007] Gabriel Skantze. *Error Handling in Spoken Dialogue Systems*. PhD thesis, KTH - Royal Institute of Technology, Sweden, 2007. Supervised by Rolf Carlson.
- [Smith, 2012] Dustin Smith. Managing implicit assumptions in natural language interfaces. In *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces, IUI ’12*, pages 383–388, New York, NY, USA, 2012. ACM.
- [Stoia et al., 2008] Laura Stoia, Darla Magdalene Shockley, Donna K. Byron, and Eric Fosler-Lussier. Scare: A situated corpus with annotated referring expressions. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC 2008)*, Marrakesh, Morocco, May 2008.
- [Tonhauser et al., 2013] Judith Tonhauser, David Beaver, Craig Roberts, and Mandy Simons. Toward a taxonomy of projective content. *Language*, 89(1):66–109, 2013.
- [van Rooij, 2011] Robert van Rooij. Optimality-Theoretic and Game-Theoretic Approaches to Implicature. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Stanford, CA: CSLI Dec 1, 2006, substantive revision Feb 17, 2011 edn, 2011.

- [Venhuizen *et al.*, 2013] Noortje Venhuizen, Johan Bos, and Harm Brouwer. Parsimonious semantic representations with projection pointers. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013) – Long Papers*, pages 252–263, Potsdam, Germany, 2013.
- [Vogel *et al.*, 2013a] Adam Vogel, Max Bodoia, Christopher Potts, and Dan Jurafsky. Emergence of Gricean maxims from multi-agent decision theory. In *Human Language Technologies: The 2013 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1072–1081, Stroudsburg, PA, June 2013. Association for Computational Linguistics.
- [Vogel *et al.*, 2013b] Adam Vogel, Christopher Potts, and Dan Jurafsky. Implicatures and nested beliefs in approximate Decentralized-POMDPs. In *Proceedings of the 2013 Annual Conference of the Association for Computational Linguistics*, pages 74–80, Stroudsburg, PA, August 2013. Association for Computational Linguistics.
- [Wilson and Sperber, 2004] Deirdre Wilson and Dan Sperber. Relevance theory. In *The Handbook of Pragmatics*, pages 607–632. Blackwell, Oxford, 2004.
- [Zeevat, 2014] Henk Zeevat. *Language Production and Interpretation: Linguistics meets Cognition*, volume 30 of *Current Research in the Semantics/Pragmatics Interface*. Brill Academic Pub, 2014.