Cooperative versus argumentative communication

Robert van Rooij[*]

Résumé : En pragmatique, la théorie de l'usage du langage, on suppose habituellement que la communication est une affaire coopérative. Cette conception standard a été récemment attaquée par Ducrot et Merin, et il a été avancé qu'un point de vue argumentatif sur l'usage du langage naturel serait plus approprié. Dans cet article, je discute la question de savoir dans quelle mesure cette attaque est justifiée et si le point de vue alternatif peut fournir une analyse plus adéquate de la « signification pragmatique », à savoir les implicatures.

Abstract: In pragmatics, the theory of language use, it is standard to assume that communication is a cooperative affair. Recently, this standard view has come under attack by Ducrot and Merin, and it has been proposed that an argumentative view on natural language use is more appropriate. In this paper I discuss to what extent this attack is justified and whether the alternative view can provide a more adequate analysis of 'pragmatic meaning', i.e., implicatures.

[*]The research of this work is supported by a fellowship from the Royal Netherlands Academy of Arts and Sciences (KNAW), which is gratefully acknowledged. For writing this paper I profited a lot from discussions with Katrin Schulz on exhaustive interpretation and Merin's analysis of implicatures. Thanks also to Martin Stokhof and an anonymous reviewer for their comments. A longer version of this paper can be found at *http://turing.wins.uva.nl/vanrooy*.

Philosophia Scientiæ, 8 (2), 2004, 195–209.

1 Introduction

It is well established that a speaker in a typical conversational situation communicates more by the use of a sentence than just its conventional truth conditional meaning. Truth conditional meaning is enriched with what is *conversationally implicated* by the use of a sentence. In pragmatics – the study of language use – it is standard to assume that this way of enriching conventional meaning is possible because we assume speakers to conform to Grice's *cooperative principle* [Grice 1967], the principle that assumes speakers to be rational cooperative language users. This view on language use suggests that the paradigmatic discourse situation is one of cooperative information exchange.

[Merin 1999] has recently argued that this view is false, and hypothesized that discourse situations are paradigmatically ones of explicit or tacit debate. He bases this hypothesis on the work of [Ducrot 1973] and [Anscombre & Ducrot 1983] where it is strongly suggested that some phenomena troublesome for Gricean pragmatics can be analyzed more successfully when we assume language users to have an *argumentative* orientation. In section 2 I will describe Merin's analysis of some implicatures which are taken to be troublesome for a cooperative view on language use. Section 3 will be used to analyze these implicature from a cooperative point of view after all.

2 Merin's approach to scalar reasoning

[Anscombre & Ducrot 1983] argue that to account for so-called 'scalar implicatures' an argumentative view is required. Scalar implicatures are normally claimed to be based on Grice's maxim of *quantity*: the requirement to say as much as one can (about a topic of conversation). This gives rise to the principle that everything 'higher' on a scale than what is said is false, where the ordering on the scales is defined in terms of informativity. Standardly, scales are taken to be of the form $\langle P(k), ..., P(m)\rangle$, where P is a simple predicate (e.g. *Mary has x children*) and for each $P(i)$ higher on the scale than $P(j)$, the former must be more informative than the latter. From the assertion that $P(j)$ is true we then conclude by scalar implicature that $P(i)$ is false. For instance, if Mary says that she has two children, we (by default) conclude that she doesn't have three children, because otherwise she could and should have said so (if the number of children is under discussion). Other examples are scales like $\langle a \wedge b, a \vee b \rangle$: from the claim that John *or* Mary will come, we are normally allowed to conclude that they do not come both.

Unfortunately, as observed by [Hirschberg 1985] and others, we see inferences from what is not said to what is false very similar to the ones above, but where what is concluded to be false is *not more informative* than, or does not entail, what is actually said. Such scalar inferences are, according to Anscombre & Ducrot, best accounted for in terms of an argumentative view on language. Ducrot and Anscombre did not formalize their rhetorico-pragmatic theory of argumentation. [Merin 1999] set himself to the task of removing this obstacle and proposes a formalization of their intuitions by making use of the theory of games and decisions. In particular, in [Merin 1999] he proposes to model communication as a *bargaining game* between two agents who have dual preferences with respect to a dichotomic epistemic issue: the question whether a particular proposition h is true or false: If one agent prefers h to be true, the other prefers it to be false. These preferences (together with the beliefs) are used to determine a precise notion of *relevance* of new pieces of information, and he uses this notion to characterize the circumstances in which certain expressions can be used appropriately, and to account for scalar implicatures.

Merin proposes to use Good's notion of 'weight of evidence' [Good 1950] as his notion of relevance. It is defined with respect to a context represented by a probability function P and a goal proposition h as follows $R(h, a) = log \frac{P(a/h)}{P(a/\neg h)}$. Naturally, Merin calls proposition *a positively relevant* to h iff $R(h, a) > 0$. Similarly, a is called *negatively relevant* and *irrelevant* iff $R(h, a) < 0$ and $R(h, a) = 0$, respectively. The fact that informative propositions can be negatively relevant will be important for Merin's analysis of linguistic data.

In terms of $R(\cdot, \cdot)$ Merin can explain scalar reasoning that cannot be accounted for in terms of the standard assumption that scales have to be ordered in terms of informativity. For instance, if Mary answers at her job-interview the question whether she speaks French by saying that her husband does, we conclude that she doesn't speak French herself, although this is not semantically entailed by Mary's answer. If h is the proposition 'Mary gets the job' – which seems only natural –, her actual answer will have a lower '$R(h, \cdot)$'-value than the claim that she speaks French herself. Thus, her actual claim, and the claim that she did not make can be ordered in a scale defined by the $R(h, \cdot)$ values, from which we can derive what is not the case by reasoning with the standard principle that speakers should make their strongest claim possible.

Perhaps somewhat surprisingly, this natural reasoning schema is *not* adopted by [Merin 1999, 2003]. In fact, he doesn't want to account for conversational implicatures in terms of the standard principle that everything is false that the speaker didn't say, but could have said (basically, the principle of exhaustive interpretation). Instead, he proposes to *derive* scalar implicatures from the assumption that *all* conversation is a *bargaining game* in which the preferences of the agents are diametrically opposed. From this view on communication, it follows that assertions and concessions have an 'at least' and 'at most' interpretation. This intuition is formalized in terms of Merin's definition of *relevance cones* defined with respect to contexts represented as $\langle P, h\rangle$ (I minimally changed [Merin 1999] actual definition 8 on page 197.)

Definition 1. The *upward (relevance) cone* ${}^{\geq S}\phi$ of an element ϕ of a subset $S \subseteq F$ of propositions in context $\langle P, h\rangle$ is the union of propositions in S that are at least as relevant to h with respect to P as ϕ is. The *downward (relevance) cone* ${}^{\leq S}\phi$ of ϕ in context $\langle P, h\rangle$ is, dually, the union of S-propositions at most as relevant to h with respect to P as ϕ is.

On the basis of his view of communication as a bargaining game, Merin hypothesizes that while the upward cone of a proposition represents the speaker's claim, the downward cone represents the hearer's default expected compatible counterclaim (i.e., *concession*). Net meaning, then is proposed to be the intersection of speaker's claim and hearer's counterclaim: ${}^{\geq S}\phi \cap {}^{\leq S}\phi$, the intersection of what is asserted with what is conversationally implicated.

In the following I will discuss some phenomena concerning scalar reasoning explicitly discussed by [Merin 1999, 2003] where it is claimed that taking an argumentative view on language use has considerable payoff.

Numerals Merin's analysis of conversational implicatures, and the way it differs from the standard Gricean analysis, can perhaps best be illustrated with numerical expressions. In standard Gricean pragmatics it is assumed that numerals have an 'at least'-meaning. The reason why we conclude from the assertion *John has three children* that John has exactly three children is then due to a conversational implicature: the stronger sentence *John has (at least) four children* is entailed to be false. A well-known problem for this analysis is that it also predicts an 'exactly'-reading when it is explicitly stated that John has *at least* three children, which we don't want. [Merin 2003] proposes that numerical expressions have semantically an 'exactly'-meaning, and takes 'at least' to be a modifier. Such an analysis, of course, has to explain the difference in acceptability between the appropriate *John has three children, in fact four* versus the inappropriate *John has four children, in fact three.* Merin proposes to account for this in terms of the notions of upward and downward cones.

Take ϕ to be the proposition that John has three children, and suppose that, for some reason, the speaker wants John to have as many children as possible. Then, the upper cone of ϕ, i.e. ${}^{\geq S}\phi$, should be thought of as the union of propositions of the form *John has n children*, with $n \geq 3$, and thus claims that John has at least 3 children. The downward cone, ${}^{\leq S}\phi$, is now of course the union of propositions of the form *John has n children*, with $n \leq 3$, meaning that John has at most three children. The appropriateness of *John has three children, in fact four* versus the inappropriateness of *John has three children, in fact two* can now be accounted for in terms of whether what is claimed by the first conjunct, ${}^{\geq S}\phi$, is consistent with the meaning of the second conjunct.

The net meaning of a sentence was defined as the intersection of its upward and downward cones: ${}^{\geq S}\phi \ \cap \ {}^{\leq S}\phi$. Observe that when ϕ is an element of S, ϕ itself will be a subset of ${}^{\geq S}\phi \ \cap \ {}^{\leq S}\phi$ as well. Thus, if $\phi \in S$, the net meaning of ϕ cannot be any stronger than the semantic meaning of ϕ, though it could be weaker![1] Indeed, with respect to normal numerical expressions, Merin's analysis of implicatures is somewhat unexciting. The analysis is still interesting, because for 'at least' sentences it gives better predictions then the standard Gricean analysis does: the downward cone of 'At least three men came' is the set of *all* states (of the world). As a consequence, the latter sentence doesn't give rise to the unwanted 'exactly'-implicature.

Temperature expressions Temperature scales are problematic for the standard Gricean picture. Intuitively, we conclude from *It is warm* that it is not hot, and from *It is cold* that it is not freezing. But what should the meanings be of these coarse-grained temperature expressions to predict these intuitions? The former inference is easy: we say that it is warm if it is *at least*, say, 15 ^{o}C, and hot if it is, say, *at least* 25 ^{o}C. Thus, the proposition claiming that it is hot entails the proposition that it is warm, and we can by standard Gricean reasoning account for the first intuition. Moreover, by assuming that these expressions have at an 'at

[1] I owe this simple but still rather critical observation to Katrin Schulz.

least'-meaning, we can immediately account for the phenomenon of *scale reversal*: the fact that from the scale ⟨boiling, hot, warm, ...⟩ we can derive ⟨..., not warm, not hot, not boiling⟩. This follows immediately by contraposition: if a entails b, by contraposition we immediately predict that $\neg b$ entails $\neg a$. Thus, from the assertion *It is not hot* we can derive by scalar implication that the stronger *It is not warm* is false, and thus that it is warm but not hot, i.e., somewhere between 15 and 25 ^oC.

Unfortunately, to account for the intuition that from *It is cold* we conclude that it is not freezing, an 'at least' reading of coarse-grained temperature expressions gives exactly the wrong reading. [Horn 1972] observed the problem and proposes to take the following two scales to be basic: ⟨boiling, hot, warm, lukewarm⟩ and ⟨freezing, cold, cool, tepid⟩. Although it is natural to assume that we do have these two scales, we would like to see how these scales are related to the meanings of the coarse-grained temperature expressions and on the basis of what principle the scales should be defined. Ducrot proposes an ordering in terms of argumentative value, and [Merin 1999, 2003] proposes a formal implementation. The reason why we have two scales is due to the fact that some (basic propositions based on) temperature expressions have a positive value with respect to some suitable chosen h (desiring for a high temperature), while others have a negative value (or positive w.r.t. the goal $\neg h$ for a low temperature). Moreover (though this is clearer in Ducrot's work than it is in Merin's), the use of a particular temperature expression, 'warm' versus 'cold', indicates what the speaker's preferences are, h versus $\neg h$. But how can we derive the two scales in terms of the meanings of the expressions? In analogy with his above described analysis of numerals, Merin crucially assumes that the meanings of the set of coarse-grained temperature expressions *partitions* the state space. Thus, 'warm', for instance, doesn't mean that it is *at least* 15 ^oC, but rather that the temperature is *between* 15 and 25 ^oC, leaving temperatures above 25 ^oC to be 'hot'. Although 'it is warm' is now taken to mean that the temperature is *between* 15 and 25 ^oC, what is asserted is still taken to mean that it is above 15 ^oC, because this is the upper cone of its semantic meaning. Still, if we also take the expected concession into account, we receive the desired reading: $^{\geq}[\text{It is warm}] \cap {}^{\leq}[\text{It is warm}] = [\text{It is warm}] = \{w \in T|$ the temperature in w is between 15 and 25 $^oC\}$.

If scales are not defined in terms of entailment, the standard explanation for scale reversal can no longer be used. How then to account for it? At first, this seems rather straightforward, for according to the relevance function used by Merin, it holds that a is *more* relevant than b with respect to h if and only if a is *less* relevant than b with respect to

$\neg h$. Unfortunately, given the 'exactly'-readings assumed by Merin, this doesn't really help him. Instead, [Merin 1999] proves and makes use of the following fact:

Fact 1. *If a and b are cells of a partition of the possibility space into propositions and $R(h, a) > R(h, b) > 0$, then $R(\neg h, {}^{<S} b) > R(\neg h, {}^{<S} a) > 0$ (and $R(h, {}^{\geq S} a) \geq R(h, {}^{\geq S} b) > R(h, b)$).*

What we want to explain is that if a is 'better' than b, 'not b' is 'better' than 'not a'. Think now of 'not a' as as expression denoting $\overline{{}^{\geq S}[a]} = {}^{<S}[a]$, and that it is made with respect to the opposite 'goal' as 'a' would be. Then it indeed follows that 'not b' is higher on the relevant scale than 'not a'.

Particularized scalar implicatures Now consider the particularized scalar implicature due to Mary's answer at her job interview to the question whether she speaks French by saying that her husband does. Obviously, the goal proposition, h, now is that Mary gets the job. Naturally, the proposition a = [Mary speaks French] has a higher relevance than the proposition b = [Mary's husband speaks French]. The net meaning of Mary's actual answer is claimed to be ${}^{\geq S}b = \bigcup\{s \in S | R(h, s) \geq R(h, b)\} \cap \bigcup\{s \in S | R(h, s) \leq R(h, b)\} = {}^{\leq S}b$. Now suppose that $b \in S$. Then, as always, $b \subseteq {}^{\geq S}b \cap {}^{\leq S}b$, and thus nothing is gained (though perhaps something is lost). Notice that if $b \in S$, the net meaning of b can only rule out that Mary herself speaks French, if this is already ruled out by the semantic meaning of b. So, if Merin assumed that $b \in S$, for the desired inference to go through, he also had to assume that the semantic meaning of b is something like 'Mary's husband speaks French and nobody else'. Given that [Merin 1999] claims that a and b are (presumably) logical independent, this cannot be what he had in mind. Perhaps Merin assumed with the rest of us that b has semantically speaking an 'at least'-meaning saying that Mary's husband speaks French and perhaps others do as well. But then, of course, it has to be ruled out that $b \in S$. This could be done if we assume that S itself *partitions* the state space (and, as we have seen above, this is what he normally assumes, although he is almost never explicit about it). Presumably, this partition is induced by a question like *Who speaks French?* On this assumption it indeed follows that the elements of the partition compatible with a = [Mary speaks French] are not compatible with the downward cone of b, and thus are ruled out correctly.

General scepticism [Merin 1999] contrasts Grice's view of conversation as cooperative information exchange with his own argumentative view on conversation. Unfortunately, Merin's view – at least on its most straightforward reading – is in contradiction with general game theoretical results. [Crawford & Sobel 1982] have shown that the amount of information that agents can transfer credibly in communication games depends on the extent to which the preferences of the conversational agents are aligned. If the preferences are diametrically opposed, as proposed by Merin to be the default case, we would predict that no credible information can be communicated at all! On the other hand, [Lipman & Seppi 1995] show that communication is possible in real debates when it cannot be excluded that the information transferred might be falsified, and that communicating false 'information' is punished. The possibility of falsification and punishment makes the communicative situations with seemingly non-aligned preferences ones where the preferences are more aligned after all, but then at a 'deeper' level. So I don't think that no sense can be made of Merin's bargaining view on communication at all, but he has to show us in what 'deeper' sense the preferences of the agents are still aligned.

But even if Merin can provide us with this missing link, his analysis of conversational implicatures I still find wanting. Although he typically can account for the phenomena, the hypothesis that conversational implicatures result from taking the intersection of the upward and downward cones is unconvincing.

First, Merin's proposal is counterintuitive in general: even if in actual conversation the goals and preferences of its participants are not always in complete alignment, we certainly do not always argue against each other. But even if we ignore this counterintuitive aspect of Merin's proposal [Merin 1999], his analysis of implicatures is still less than convincing. For numerals and course-grained temperature expressions the analysis already presupposes what should be explained,[2] and to arrive at the intuitively correct upward and downward cones we have to assume counterintuitive goal-propositions (the desire for many children and very high temperatures). For particularized conversational implicatures (and disjunction) his analysis crucially depends on the identity of the set S of relevant propositions, and he doesn't make it very clear *why* the chosen set should be chosen, nor why the upward and downward cones of a disjunctive sentence should be as he assumed.

But Merin's analysis of scalar implicatures should not be ruled out

[2]Except, of course, for 'at least'-expressions.

just because it relies on a view of communcationn which is on one reading incompatible with general game theoretical results, nor because his treatment is, intuitively speaking, unconvincing. These 'higher-order' arguments can be used only if we can account for the same empirical phenomena in a theoretically more appealing way. In the next section I will set myself to this task by providing a, perhaps, more familiar and natural analysis of the implicatures discussed in [Merin 1999, 2003].

3 Scalars and exhaustive interpretation

It makes a lot of sense to assume that (truthful) speakers say as much as they can about for them desirable situations. In case the speaker is taken to be well-informed, we can conclude that what speakers do not say about desirable situations is, in fact, not true. To account for this we can formulate a 'pragmatic' interpretation rule for sceptic hearers that have to 'decode' the message following this reasoning, to hypothesize what kind of situation the speaker is in. We can call this rule one of *exhaustive interpretation*. For this interpretation rule we can assume that '$v < w$' if and only if the speaker prefers world w to world v.

Definition 2. *Exhaustive interpretation (general)*

$$exh(A) = \{w \in [A] | \neg\exists v \in [A] : v < w\}$$

Now consider the example of scalar reasoning again that was a serious problem for standard Gricean analyses: the case where Mary answers at her job-interview the question whether she speaks French by saying that her husband does. Intuitively, this gives rise to the scalar implicature that the 'better' answer that Mary herself speaks French is false. Notice that if we assume the scale to be the preference order (between states) of the speaker, we can account for this example in terms of the rule of exhaustive interpretation. All we have to assume for this analysis to work is that the state where speaker Mary speaks French herself is more preferred to one where she does not.

In Gricean pragmatics, most conversational implicaturess – the scalar ones in particular – are due to the second submaxim of quantity, which requires a speaker to say as much about an issue as she knows to be true. [Groenendijk & Stokhof 1984] formulated an exhaustivity operator applied to (term) answers to questions that implements this principle in a natural way. As it turns out [Van Benthem 1989], this operator is virtually identical to [McCarthy 1980] rule of interpretation by predicate circumscription, which can be given a natural semantic characterization in terms of a minimal-, or preference-, model analysis.[3] Assume that the (question)-predicate at issue is P ('who danced?'), and that answer A ('John danced') is given. In that case, the exhaustive interpretation of A with respect to predicate P, or the circumscription of A with respect to P, will be the set of A-worlds where P has a minimal extension (John danced and nobody else):

Definition 3. *Exhaustive interpretation with respect to a predicate*

$$exh(A,P) = \{w \in [A] | \neg\exists v \in [A] : v <_P w\}$$

Thus the exhaustive interpretation of A contains all those states w that verify A, and for which no more minimal state v exists that also verifies A. What still has to be specified is the ordering relation '$<_P$'. In standard circumscription it is assumed that $v <_P w$ if the extension of P in v is a subset of the extension of P in w, i.e. $P(v) \subset P(w)$.[4]

The rule of exhaustive interpretation can be used to account for scalar implicatures. In definition 3 it crucially relies on a particular question-predicate. In accordance with, perhaps, more standard analyses of Gricean implicatures (e.g. in the work of [Horn 1972]), we can also make exhaustification relative to a particular language in question. The standard analysis of exhaustification can be described alternatively relative to a particular language L as follows:

Definition 4. *Exhaustive interpretation with respect to a language*

$$exh(A,L) = \{w \in [A] | \neg\exists v \in [A] : \{B \in L | v \in [B]\} \subset \{B \in L | w \in [B]\}\}$$

Obviously, these definitions are equivalent if we take language L to be defined in terms of predicate P as follows: $L = \{P(a) | a \in NAME\}$, if we assume that every individual has a (unique) name.

We want to know whether in terms of exhaustive interpretation we can account for those phenomena for which [Merin 1999] claims an argumentative view on language use is required. We will take up the phenomena discussed in section 2 one by one. I would like to mention that almost none of the assumptions I make below are particularly new or surpising, though perhaps it has not been very clear how they can be used to describe the phenomena at hand.

[3] For the relation between exhaustive interpretation, knowledge, and the Gricean maxims of conversation, see [Van Rooij & Schulz 2004].

[4] It is also assumed that with respect to all other predicates P', $P'(v) = P'(w)$. This is important for the circumscription analysis of conditionals, for instance, but won't bother us in this paper.

Numerals Just like standard Gricean pragmatics, also the standard analysis of exhaustive interpretation [Groenendijk & Stokhof 1984] has problems to account for the fact that the exactly-interpretation of sentences like *John has three children* is missing for the version modified by 'at least'. But then, like [Merin 1999, 2003], we can simply assume that *three* means something different than *at least three*. Indeed, let us just assume (temporarily) that [three] = 3. Moreover, let us assume that 'John has three children' is true in state w if 3 is an element of the extension of property λn[John has n children] in state w. In general, we assume that $P(3)$ is true in w iff $3 \in P(w)$. Of course, we want to assume that if John has 3 children, he has 2 children as well. This can be accounted for by assuming that property λn[John has n children] is monotone *downward entailing*. A property P is downward entailing with respect to numerals iff it holds for all natural numbers n and m and states w that if $n \in P(w)$ and $m < n$, then also $m \in P(w)$. Thus, for downward entailing P, $P(3)$ entails $P(2)$, but the truth of $P(2)$ does not rule out that also $P(3)$ is true. In this sense numerals are still predicted to have an 'at least' reading when used with a downward-entailing predicate. Obviously, we also have predicates that denote monotone upward entailing properties, and properties that don't behave monotonically at all. An example of the former kind of predicate is 'Can run the 100 meters in n seconds', while mathematical predicates like '$3 + n = 7$' are prime examples of expressions denoting non-monotonic properties.

To account for the fact that *John has <u>at least</u> three children* doesn't give rise to the implicature that he has no more than three children, we have to slightly change the above analysis: we don't say anymore that 'three' means 3 and that P(three) is true in w iff $3 \in P(w)$, but rather with Kadmon (1987) that 'three' means $\{x|card(x) = 3\}$ and that P(three) is true in w iff $\exists x \in$ [three] : $x \subseteq P(w)$. Similarly, we say that 'at least three' means $\{x|card(x) \geq 3\}$ and that P(at least three) is true in w iff $\exists x \in$ [at least three] : $x \subseteq P(w)$. Van Rooij & Schulz (submitted) show that in combination with some standard assumptions in dynamic semantics, this analysis leads to the correct prediction that while the exhaustive reading of P(three) gives rise to an exactly reading, this is not so for P(at least three).

Temperature expressions [Horn 1972] accounts for the implicatures of temperature terms using two scales: $\langle$boiling, hot, warm, lukewarm$\rangle$ and $\langle$freezing, cold, cool, tepid$\rangle$, while Merin showed that these scales could, in principle, be derived. However, it is not made clear which particular propositions *should* be at issue to derive those orderings, nor why *only* these propositions *could* be relevant.

Though [Merin 1999] bases his analysis on Ducrot's general argumentative view on language, Ducrot himself thinks of this argumentative orientation somewhat differently from Merin. While Merin takes the argumentative orientation with respect to a *contextually given* proposition, Ducrot assumes that the argumentative orientation is *inherent* to the *language* itself, and accordingly is context independent. Thus, it is expressions themselves that have, according to Ducrot, already a certain argumentative orientation. To take one of Ducrot's favorite examples, whereas *little* has a negative argumentative orientation and can be used to argue for 'negative' conclusions, *a little* has a positive one. The same can be said for *warm* versus *cold*, they have *as part of their meaning* already an argumentative orientation. In logical terms this is normally expressed by saying that whereas the one has a monotone *increasing* meaning, the other is monotone *decreasing*. Suppose that also *warm* and *cold* have a conventional, or 'semanticized', argumentative orientation: the former 'goes to' high temperatures and the latter to low ones. But then we can state the meanings of the two expressions by [warm] $= \{n|n \geq 15^0C\}$ and [cold] $= \{n|n \leq 5^0C\}$. The scalar inferences can now be accounted for by standard exhaustive reasoning.

Let us assume that a statement like *It is warm/cold* is normally made if the temperature is at issue. So, to account for the implicatures it seems natural to assume that the relevant language L consists of coarse-grained temperature expressions like *freezing*, *cold*, *cool*, *tepid*, *warm*, *hot*, and *boiling*. We will assume the following meanings of these expressions: [freezing] $= \{n|n \leq 0^0C\}$, [cold] $= \{n|n \leq 5^0C\}$, [cool] $= \{n|n < 10^0C\}$, [tepid] $= \{n|n \geq 10^0C\}$, [warm] $= \{n|n \geq 15^0C\}$, [hot] $= \{n|n \geq 25^0C\}$, and [boiling] $= \{n|n \geq 100^0C\}$. Notice that from these meanings we can derive the following emtailment relations: boiling $\models$ hot $\models$ warm, and freezing $\models$ cold $\models$ cool, but that no inference relations exist between warm and cold, for instance. In fact, if two of the above expressions do not stand in an inference relation to one another, they are incompatible.

Applying now the exhaustivity operator with respect to a language,[5] and assuming that the language L contains of the coarse-grained temperature expressions mentioned above, we immediately can infer from *It is warm* that it is not hot, and thus that it is between 15 ^0C and 25 ^0C. Similarly, we can infer from *It is cold* that it is not freezing, and thus that it is between 5 ^0C and 0 ^0C.

[5] Of course, doing things in terms of exhaustivity with respect to a predicate would be equivalent.

Now consider what we predict for negative statements like *It is not hot*. This, of course, depends on what we take to be the relevant language. We make the following two constraints on a language with respect to which we have to apply exhaustification: First, all elements of L have to be mutually compatible. This means that the relevant temperature expressions should either be all monotone upward, or monotone downward. It allows {boiling, hot, warm, lukewarm} to be a suitable language L. The second assumption is that for negative statements we switch to the following language: $\bar{L} \stackrel{\text{def}}{=} \{\bar{a}|a \in L\}$ (where $\bar{a}$ denotes the complement of a). It is easy to check that now we indeed predict that negation gives rise to the desired scale-reversion.

In conclusion: we don't have to associate an argumentative orientation with temperature expressions to account for the desired conversational inferences. It is enough to assume that some expressions have a 'at least' and others an 'at most' meaning, and that temperature statements are interpreted exhaustively w.r.t. a suitable language L.

Particularized scalar implicatures Earlier in this section we accounted for the implicature arising from Mary's answer at her job-interview in terms of the speaker's preference relation among states. [Merin 1999] proposed that this example should be accounted for in terms of his relevance relation $R(\cdot,\cdot)$, while we have shown above how the use of a language, or a set of alternatives, can be useful for the analysis of implicatures. Here I want to suggest that we can also account for Mary's answer by taking relevance and a particular language into account.

In definitions 3 and 4 we defined our exhaustivity operator (implicitly or explicitly) in terms of the standard notion of entailment. However, we can generalize this operator by taking a notion of relevance into account. In general we can use the following exhaustivity operator, which depends not only on a language L, but also a goal proposition h and a relevance function R:

Definition 5. *Exhaustive interpretation with respect to relevance*

$$exh(A, L, h) = \{w \in [A] | \neg\exists v \in [A] : v <_h^L w\}$$

Here we assume the following ordering relation between states:

$$v <_h^L w \quad \text{iff} \quad R(h, \bigcap\{B \in L | v \in [B]\}) < R(h, \bigcap\{B \in L | w \in [B]\})$$

Assuming h normally to be $\bigcap L$,[6] we see that this new exhaustivity operator is really a generalization of our earlier ones. Consider now the set $\bigcap\{B \in L|t \models B\}$ in case predicate P is at issue. It denotes the proposition that *at least* all individuals that have property P in state w actually have property P. That is, it denotes the following proposition: $\lambda i[P(w) \subseteq P(i)]$. Van Rooij & Schulz (submitted) show that the resulting ordering relation $v < w$ between states mirrors an entailment relation exactly when all information is relevant. In particular, this shows that in our special case $V(h, \bigcap\{B \in L|w \in [B]\}) < V(h, \bigcap\{B \in L|w \in [B]\})$ if and only if $\{B \in L|w \in [B]\} \subset \{B \in L|w \in [B]\}$. Thus, in the special case under discussion, definition 5 reduces to definitions 3 and 4.

However, by making exhaustive interpretation dependent on relevance, we can account for more phenomena than we could until now. In particular, for the scalar inference resulting from Mary's answer at her job-interview.

In this paper I have argued against Merin's analysis of scalar implicatures in terms of his upward and downward cones. But, by taking relevance into account, for many examples an analysis in terms of exhaustive interpretation is in fact very similar to what Merin proposed. By a clever – though sometimes unnatural – choice of goal proposition h, it is in most (if not all) cases possible that what we take to be the semantic meaning of an expression with a monotone increasing meaning is the same as what he takes to be the upward cone of the proposition expressed by a phrase with a (semantically) non-monotonic meaning. Let denote this proposition by A. In those cases, Merin's intersection of this resulting upward cone with a proposition's downward cone will be the same as our exhaustive interpretation of A.

4 Conclusion

In this paper I have contrasted the standard cooperative view on communication with an argumentative one. According to the former view, we communicate information that is good for all participants of a conversation, while according to the latter, we communicate always to argue for a particular hypothesis and do this always against an opponent. In this paper I have put some doubts on the universal applicability of the latter

[6] Notice that this results in the empty set if the elements of L are not mutually compatible, something that we actually ruled out. But even then this goal proposition need not be natural, but I just want to show that we could do things in terms of goals as well.

view. Though I agree with Ducrot and Merin that an argumentative perspective is useful for the analysis of (at least) adversary connectives, adopting this perspective for the analysis of all scalar implicatures (or at least in the way proposed by Merin) was argued to be unconvincing. I have argued that a more natural explanation is possible when we assume speakers to say as much as they can, and hearers to interpret exhaustively. Although this assumption seems rather standard, it does not require perfect alignment of preferences as normally presupposed in Gricean pragmatics. Thus, even if neither the Gricean cooperative view on language use, nor the alternative argumentative view has universal applicability, this doesn't mean that conversational implicatures cannot still be accounted for by means of a general rule of interpretation. Obviously, I take the principle of exhaustive interpretation to be such a general rule.