

Early Interpretation by Implicature in Definite Referential Descriptions

Raquel Fernández

Institute for Logic, Language & Computation
University of Amsterdam

Abstract

This paper discusses the processes by which dialogue participants incrementally compute the conveyed meaning of referential descriptions. I show that some of the phenomena observed in the psycholinguistics literature can be accounted for by theories of conversational implicature that exploit ingredients from computational models of Referring Expression Generation. The result is a system that computes inferences incrementally, from partial utterances, without need to reason with hypothesized complete descriptions.

1 Introduction

It is widely accepted that utterances in conversation often communicate information that goes beyond the conventional meaning conveyed by the linguistic expressions used. A good deal of the processes that help dialogue participants reduce the gap between the conventional meaning of utterances and the enriched meanings actually intended by speakers has to do with *inference*—a notion which is at the core of Gricean pragmatics and the theory of conversational implicature (Grice, 1975; Grice, 1989). It is also generally agreed that in conversation speakers and addressees produce and understand language incrementally, in (at least) a word-by-word fashion rather than in one go once, say, the end of an utterance has been reached. The question thus arises as to whether theories of conversational implicature (which, as most semantic and pragmatic theories, were originally designed to operate at the utterance level) can be accommodated within the *incremental turn*.

In this paper I look into how dialogue participants incrementally compute the intended meaning of referring descriptions (i.e. their intended referent). I show that a slightly modified version

of Hirschberg (1985)'s computational theory of scalar implicature that allows us to compute implicatures at the sub-utterance level can account for incremental effects observed in psycholinguistic experiments.

I start by giving an overview of previous work on the incremental processing of referential descriptions regarding both resolution and generation. In Section 3 I survey experimental results that indicate that information that goes beyond conventional semantic meaning is used incrementally at the sub-utterance level. To account for the pragmatic inferences observed, I explore an account that combines ingredients from generation models with a theory of scalar implicature. After introducing the rudiments of such a theory and some basic semantic processing in Section 4, I sketch my proposal in Section 5 and apply it to some examples.

2 Definite Referential Descriptions

Referential descriptions in the sense of Donnellan (1966)—i.e. definite descriptions that serve the purpose of letting the addressee identify a particular entity out of a set of entities assumed to be in the current focus of attention—have been studied extensively in dialogue research, specially in the context of *referential matching tasks* such as those used in the psycholinguistic experiments of Krauss and Weinheimer (1966) and Clark and Wilkes-Gibbs (1986). Work within the collaborative model of grounding put forward by Clark and colleagues (Clark, 1996) has emphasized the fact that the form and meaning of these descriptions often depends on *historic* aspects such as conceptual pacts established with specific conversational partners during the course of interaction (Brennan and Clark, 1996). Thus these approaches have mostly focused on *subsequent mentions*, i.e. descriptions that refer back to entities mentioned previously, or that refashion earlier descriptions that were not

grounded (Clark and Wilkes-Gibbs, 1986).

In the present paper I concentrate on the arguably simpler case of *first mention* referential descriptions. The main reason for this is that experiments that study incremental processing at the sub-utterance level—which is our focus here—have essentially investigated first mention descriptions in simple tasks, where the referring goal can be achieved with little interaction,¹ by means of a single referential description. Several approaches have investigated this kind of descriptions from the generation and the resolution perspective. In the remainder of this section, I summarize some of them and then give an overview of some experimental results reported in the psycholinguistics literature.

2.1 Resolution

A good deal of computational work on incremental interpretation of referring descriptions—such as the early proposals of Mellish (1985) and Haddock (1989) as well as more recent approaches such as Schuler (2003)—models incremental reference resolution as a symbolic process of constraint satisfaction, where predicates are associated with sets of constraints. The core idea is that, as a referring expression is processed from left to right, the constraints introduced by each predicate in the expression progressively narrow down the set of potential referents. Consider, for instance, the description ‘*the black wooden chair*’. Here processing ‘*black*’ would eliminate from the set of potential referents those elements in the context that are not black; processing ‘*wooden*’ would narrow down that set further to the subset of black elements that are made of wood; while finally processing ‘*chair*’ would pick up the chairs among the black wooden elements.

These resolution models thus focus on computing semantic denotation in a model-theoretical way, largely ignoring any aspects related to pragmatics and implicature. However, as we shall see in Section 3, hearers can incrementally use pragmatic information that goes beyond conventional interpretations to identify referents at stages where there is semantic ambiguity, thus speeding up the process of establishing the speaker’s meaning. Constraint-based models can naturally be enriched with pragmatic constraints, as done, for in-

¹Although see Brown-Schmidt et al. (2005) for a promising attempt to use eyetracking methods in interactive settings.

stance, by DeVault and Stone (2003), who complement the conventional content with goal and and plan-related inferences. I shall follow a similar approach and extend a constraint-based semantic model with scalar implicatures.

A different trend of approaches (e.g. Roy (2002) and Schlangen et al. (2009)) explore probabilistic models of reference. In this case, the referential potential of linguistic expressions is learned from data by exploiting statistical correlations between the linguistic expressions used and particular referential configurations (i.e. the context of utterance with its set of potential referents and knowledge about which referent is the intended one). Models of this sort thus do not make a clear distinction between semantics and pragmatics: They capture expectations about linguistic meaning and cooperative behaviour implicitly and with the same mechanisms.

2.2 Generation

The Generation of Referring Expressions (GRE) is one of the key areas within the field of Natural Language Generation. Traditionally, researches in this area have attempted to generate *minimal descriptions*, i.e. the shortest possible descriptions that succeed in uniquely identifying the intended target referent in a given context. For instance, if the context includes one black chair and one brown table, the description ‘*the black chair*’ is not minimal, while the description ‘*the chair*’ is. The idea is that minimal descriptions are consistent with Grice’s Maxim of Manner, in particular with the Sub-maxim of Brevity: “be brief; avoid unnecessary prolixity” (Grice, 1975). Since the shorter description ‘*the chair*’ succeeds in identifying the referent, the presence of the predicate ‘*black*’ is considered *redundant* and hence susceptible of generating false implicatures by violating the maxim of Brevity.

The problem with this approach is that it does not take into account the fact that interpretation is a continuous process and ignores the possibility of reasoning incrementally: whether a predicate is considered redundant or not is determined by reasoning with *complete* descriptions. The predicate ‘*black*’ is considered redundant in the complete description ‘*the black chair*’ because there is an alternative, shorter, complete description that does without it. This view thus misses the point that what may count as redundant upon completion of

an utterance can be informative during incremental processing.

The seminal work of Dale and Reiter (1995) addresses precisely this issue. The Incremental Algorithm proposed by these authors uses properties in a predefined *preference order* incrementally, as long as they have *discriminatory power*, i.e. as long as they rule out some distractors—elements that are not the intended referent—from the context set. For instance, in our earlier example, ‘*black*’ has incremental discriminatory power because at the point when the modifier is uttered it rules out the table, which is not black. Thus, if colour is a particularly salient property (it is ranked high in the preference order), the use of a colour predicate such as ‘*black*’ can help the hearer to more easily identify the intended referent—a point also made by Grosz (1981): “A speaker should be redundant only to the degree that redundancy reduced the total time involved in identifying the referent.” The optimization of property preference orderings and the use of properties that are redundant *a posteriori* but that do have incremental discriminatory power are issues that are actively being investigated in GRE research (Viethen et al., 2008; Krahmer et al., 2008).

3 Psycholinguistic Evidence

In a series of experiments, Sedivy and colleagues (Sedivy et al., 1999; Sedivy, 2003) have used the eye-tracking paradigm to investigate how humans interpret instructions that contain referential descriptions with different types of modifiers, including colour (e.g. *black*, *red*), material (e.g. *plastic*, *wooden*), and scalar adjectives (e.g. *tall*, *big*). In these experiments, subjects wearing a head-mounted eye-tracker are shown an array of objects and are asked to pick up one of them with instructions such as ‘*Pick up the plastic spoon*’. Since subjects direct their gaze towards potentially referred objects, the precise information provided by the eye-tracker offers direct evidence about the alternative referents that are being considered by a subject at precise points in time and about the point at which a commitment to an interpretation is made.

The displays used in the experiments were such that upon hearing the adjective, more than one referent was possible while some referents could be discarded. For instance, one sample scenario included two objects that were made of plastic—

the target spoon and a comb—plus a tie and a bulb as distractors. As expected, upon hearing the adjective subjects looked at the two objects that matched its semantic content (the spoon and the comb). The more interesting result was that in scenarios where a contrasting object was also in the display (e.g. a metal spoon) a preference for the target plastic object could be observed before the head noun was uttered, regardless of the semantic indeterminacy. That is, in these scenarios the identification of the target object took place *earlier*, at a point when the ongoing utterance was still semantically ambiguous. Experiments also show that this effect, let us call it “*contrastive bias*”, does not obtain with all kinds of modifiers: hearers did not show this kind of bias when interpreting colour adjectives, while they did for descriptions with material and scalar adjectives.²

With regard to an experiment that included also production, Sedivy (2003) reports that what distinguished modifiers that gave rise to a contrastive bias from those that did not was the frequency with which that kind of modifier was spontaneously generated to describe an object in a context where modification was not required for unique identification. That is, modifiers such as colour adjectives that are often produced “redundantly”—redundantly *a posteriori*, but that may still help to reduce the search space for the hearer incrementally—do not lead to a contrastive bias, while those that are typically used only when they are needed for unique identification of the referent are understood as such and hence give rise to a contrastive inference.

This latter result seems to indicate that the observed contrastive bias could be successfully modelled by a probabilistic approach that is sensitive to the statistical correlations in the data. However, a follow-up experiment demonstrated that this is not entirely trivial. Grodner and Sedivy (forthcoming) found that when subjects were explicitly told beforehand that the speaker suffered from an impairment leading to linguistic deficits, they did not show evidence of a contrastive bias regardless of the statistical patterns in the data.³ This seems

²Throughout the paper, I will ignore scalar adjectives and exemplify my points with colour and material modifiers. The special features of gradable adjectives, which I have addressed elsewhere (Fernández, ms), are not critical for the approach presented here.

³Results of a similar nature regarding disfluencies are reported by Arnold et al. (2007): subjects who hear disfluent descriptions infer that the referent is difficult to describe;

to indicate that the explanatory power of statistical regularities is limited.⁴

In her discussion of the experimental results I have surveyed, Sedivy (2007) seems to appeal to an idea of redundancy that—in line with the point made in Section 2.2—relies on reasoning with hypothesized complete descriptions: “*Thus, when the display contained a referent for which the use of a modifier was communicatively motivated, people showed a preference for this referent compared to displays in which there was no clear reason to refer to the same target referent using a modifier*”. Here I shall adopt a different perspective and show that these results, as well as the differences observed for different types of modifiers, can be accounted for by combining elements from computational models of Referring Expression Generation with ingredients from standard theories of scalar implicature to arrive at a system that operates incrementally, without the need to reason with complete descriptions.

4 Preliminary Notions

Before moving on to sketch an account of the effects I have described, I shall first introduce the rudiments of scalar implicature theories, the notation I will use, and the basic process of incremental semantic interpretation I assume.

4.1 Scalar Implicature

Scalar implicature is a kind of conversational implicature (Grice, 1975) whose computation is dependent upon the identification of some salient relation that orders a concept referred to in an utterance with other concepts of the same type. The idea is that the use of an expression in the scale (i.e. the ordering) implicates that (the speaker believes that) the other expressions in the scale (often considered stronger) do not apply.

Following the seminal work of Horn (1972) and others such as Gazdar (1979), Hirschberg (1985) proposes a theory of scalar implicature that specifies the conditions under which a speaker may license a scalar implicature and a hearer may infer it. In her theory, scalar implicatures are calculated from surface semantic representations of complete utterances by (1) identifying a potential

such inferences are however cancelled when subjects are told speakers suffer from a linguistic impairment.

⁴Or perhaps that an adequate probabilistic model should also take into account the cooperativity of the speaker.

scalar sub-formula in the logical form, (2) identifying the scale or scales that this subformula belongs to, and (3) inferring negative implicatures for alternate and higher values in the scale(s).

I shall essentially follow this approach,⁵ but allow for the possibility of incrementally computing scalar implicatures from sub-formulas as they become available during incremental processing, and for the possibility of computing the implicatures from formulas other than those strictly corresponding to the logical forms of utterances (this should become clearer in Section 5.2).

As extensively discussed by Hirschberg (1985), scales can be of several types. I consider two types of scales: alternate scales $\{\sigma, \sigma', \dots\}$ containing scalar expressions that are not ordered but simply contrast with each other, and linear scales $\langle \sigma, \sigma', \dots \rangle$ containing expressions that are linearly ordered according to some suitable relation.

The inference rule for scalar implicature I assume is the following, where ψ and $\psi[\sigma/\sigma']$ are identical except for the fact that all occurrences of expression σ in ψ have been substituted by σ' in $\psi[\sigma/\sigma']$:

- (1) SCALAR IMPLICATURE INFERENCE RULE
 Given a formula ψ , a scalar expression σ in ψ , and a scale S that includes σ :
 $\forall \sigma'. ((\sigma' \in S \wedge \sigma' > \neq \sigma) \rightarrow \neg \psi[\sigma/\sigma'])$

We can use the rule in (1) to compute, for instance, the scalar implicatures inferred from the utterances in (2) and (3) (indicated by \rightsquigarrow) by considering, in the standard way, that ψ corresponds to the semantics of the whole utterance.

- (2) Some people left the party early.
 a. $S : \langle \text{all, some} \rangle$
 b. \rightsquigarrow Not all people left the party early.
- (3) A: Do you have apple juice?
 B: I have grape, tomato or bloody mary mix.
 a. $S : \{ \text{grape, tomato, bloody, apple} \}$
 b. \rightsquigarrow I don't have apple juice.

4.2 Domain Representation

I model the domain in a way similar to how input databases are modelled in GRE systems, i.e. characterising entities in terms of attributes and values.

⁵I employ a simplified version of the original approaches, which amongst other things ignores epistemic operators, but which suffices to illustrate the points that occupy us here.

I assume a first-order-logic system with all the usual logical symbols, including equality = and a top \top symbol; a domain of entities U ; a set of relational symbols A corresponding to attributes (such as colour, material, type); and a set of constant symbols V corresponding values (such as blue, red, plastic, metal, spoon, comb). A function $Val : A \rightarrow \mathcal{P}(V)$ assigns to each attribute a set of appropriate values. Variables e, e' range over elements in U , variables att, att' over elements in A , and variables val, val' over elements in V . The interpretation function assigns to each $att \in A$ a relation $U \times Val(att)$. I write $att(e) = val$ (instead of the more standard notation for relations $att(e, val)$) to express that element e is related to value val by attribute att .

4.3 Incremental Interpretation

I shall use formulas $att(e) = val$ as logical forms of adjectives and nouns in definite referring descriptions. For instance, ‘red’ will be interpreted as $colour(e) = red$; the extension of this expression is then the set of red elements in the context.

The semantic interpretation of a description such as ‘the red plastic cup’ then proceeds as shown in (4). I use the symbol \top to initialise the existentially quantified formula introduced by the definite article.⁶

- (4) t_0 The t_1 red t_2 plastic t_3 cup t_4
 $t_1 : \exists e. \top$
 $t_2 : \exists e. colour(e) = red$
 $t_3 : \exists e. colour(e) = red \wedge$
 $\quad\quad\quad material(e) = plastic$
 $t_4 : \exists e. colour(e) = red \wedge$
 $\quad\quad\quad material(e) = plastic \wedge$
 $\quad\quad\quad\quad\quad type(e) = cup$

This incremental process is in line with the symbolic constraint-based approaches to incremental reference resolution described in Section 2.1. Expressions add constraints that incrementally narrow down the set of potential referents. My aim is to complement this semantic process with default pragmatic inferences that can be computed using the ingredients of scalar implicature theories. I turn to this in the next section.

⁶I do not include the presupposition of unique existence in the semantic representation.

5 Early Interpretation by Implicature

In this section I show how we can use the main elements of the theory of scalar implicature I have sketched to account for the early interpretation effects observed in the psycholinguistic experiments described in Section 3.

5.1 Scales

Determining what is a possible salient scale in a given situation can be a tricky issue. In general expressions within an ordering share a common type (e.g. they are quantifier expressions, of juice types). But even so, it is not trivial to determine which elements of the relevant type can be considered part of a scale that is salient for both speaker and hearer. This point is emphasized by Benotti and Traum (2009) in their account of comparative implicatures, where they opt for deriving the simplest possible scale ⟨no, yes⟩ for scalar adjectives such as ‘safe’ in comparative constructions.

I take *attributes* and *values* to be scalar expressions, i.e. expressions that can be associated with a scale of related concepts. Values give rise to alternate scales $\{val, val', \dots\}$ containing different values relevant for one attribute, while attributes take part in ordered scales $\langle dim, dim', \dots \rangle$ that rank several attributes according to salience. Here I shall directly borrow the notion of “property preference ordering” (or “list of preferred properties” (Dale and Reiter, 1995)) from REG systems and assume that the hearer’s representation of the contextual domain (like the generator’s) includes such a scale of attributes. I will however make minimal assumptions about the elements that belong to a particular scale. By definition, salient scales include the triggering scalar expressions that have been overtly uttered. For instance, an utterance of say ‘plastic’ with logical form $material(e) = plastic$ can give rise to two scales:

- (5) a. $S_1 : \{\dots, plastic, \dots\}$
b. $S_2 : \{\dots, material, \dots\}$

The question is then how these scales that are assumed to be part of the common ground of speaker and hearer are further populated. Since we are concerned with a visually shared situation, I will define default rules for including additional expressions into a scale that rely only on properties of the shared visual context.⁷ In particular,

⁷Of course other aspects may render expressions salient

we assume that if an expression is *witnessed* in the shared visual context it can enter a contextual scale. Intuitively, a value such as *red* is witnessed if the visual context includes an entity that is red. Similarly, an attribute such as *colour* is witnessed if the context includes an entity for which that attribute is applicable. As mentioned, attributes take part in ordered scales—preference orderings. Preferred attributes are those that are more salient. Here I equate salience with ease of perception: in a shared visual context, an attribute is salient if it can be easily perceived (its values easily discriminated) by both speaker and hearer. I use $>_{sal}$ to denote the preference ordering.

The following scale construction inference rules make more precise what has just been explained informally. Note that for ordered scales, the only attributes that are relevant are those that are more salient than the attribute evoked by the utterance itself.⁸

- (6) SCALE CONSTRUCTION DEFAULT RULES
 Given a (sub-)utterance u with semantics $\text{att}(e_i) = \text{val}$, and potential scales $S_1 : \{\dots, \text{val}, \dots\}$ and $S_2 : \langle \dots, \text{att}, \dots \rangle$:
- a. $\forall e_j \text{ val}' . \text{att}(e_j) = \text{val}' \wedge \text{val}' \neq \text{val} \wedge e_i \neq e_j \rightarrow \text{val}' \in S_1$
 - b. $\forall e_j \text{ att}' . \text{att}'(e_j) = \text{val} \wedge \text{att}' >_{sal} \text{att} \rightarrow \text{att}' \in S_2$

These two types of scales—alternate value scales and scales of preferred attributes—are of a rather different nature and I shall assume that they are used in different ways by the SCALAR IMPLICATURE INFERENCE RULE given in (1) above. The implicatures inferred from value scales are the classic scalar implicatures computed from alternate scales such as that shown in (3). These scalar implicatures are standard in the sense that they can be computed from input formulas ψ that directly correspond to logical forms of (sub-)utterances without any further assumptions.⁹

5.2 Contrastive Inferences as Scalar Implicatures

We will now look into how the machinery we have in place allows us to infer scalar implicatures that

for both speaker and hearer.

⁸Note that these are only *default* rules. Certainly elements can be part of a salient scale for less overt and immediately accessible reasons.

⁹Besides those that allow us to construct the relevant scale.

can account for the contrastive bias that helps the hearer identify the intended referent at an early stage, regardless of the semantic ambiguity. The intuitive idea I want to explore is that these scalar inferences are only drawn if the attribute evoked by a sub-utterance has *incremental discriminatory power*, as introduced in Section 2.2—that is, if the context includes another element that has a different value for that attribute (and that hence would get eliminated from the set of potential referents). More formally, an expression such as ‘*plastic*’ in a partial utterance such as ‘*the plastic...*’ with semantics (7a) has discriminatory power if (7b) is supported by the context (where $e \neq e'$):

- (7) a. $\exists e . \text{material}(e) = \text{plastic}$
 b. $\exists e' . \text{material}(e') \neq \text{material}(e)$

The SCALAR IMPLICATURE INFERENCE RULE (repeated below for convenience) can then be exploited to enrich (7b).

- (8) SCALAR IMPLICATURE INFERENCE RULE
 Given a formula ψ , a scalar expression σ in ψ , and a scale S that includes σ :
 $\forall \sigma' . ((\sigma' \in S \wedge \sigma' > \neq \sigma) \rightarrow \neg \psi[\sigma/\sigma'])$

The clause in the matrix of (7b) ($\text{material}(e') \neq \text{material}(e)$) acts as input formula ψ for the rule, while the attribute *material* instantiates σ . Since the inferred implicatures ($\neg \psi[\sigma/\sigma']$) characterise both the intended referent and the additional element required for the expression to have discriminatory power, they fall under the scope of both quantifiers introducing these elements:

- (9) $\exists e . \text{material}(e) = \text{plastic} \wedge$
 $\exists e' . \text{material}(e') \neq \text{material}(e) \wedge$
 $\rightsquigarrow \neg [\sigma'(e') \neq \sigma'(e)] \wedge$
 $\rightsquigarrow \neg [\sigma''(e') \neq \sigma''(e)] \wedge$
 \dots

In order to illustrate how this works with a concrete example, let us consider the sample scenario described in Section 3. Recall that the visual context contains at least three elements: a plastic spoon, a plastic comb, and a tie—let’s assume the latter is made of silk. Let’s also assume that in this context, object *type* is a more salient attribute (e.g. can be more easily perceived) than *material* and that therefore upon processing the fragment ‘*the plastic...*’ the preference scale $\langle \text{type}, \text{material} \rangle$ is evoked. The conventional meaning of this fragment is shown in (10a). There

are two witnesses that make this formula true in the current context—the spoon and the comb—and hence there is semantic ambiguity. However the modifier has incremental discriminatory power since the context is consistent with (10b). Now, if the hearer does not have reasons to believe that the speaker is not cooperative (or capable of being so), the SCALAR IMPLICATURE INFERENCE RULE can be used to enrich (10b) with further inferences, as long as they are supported by the context. Given the preference scale $\langle \text{type}, \text{material} \rangle$, the rule can generate the implicature in (10c), which is equivalent to $\text{type}(e') = \text{type}(e)$.

- (10) The plastic ...
- a. $\exists e. \text{material}(e) = \text{plastic} \wedge$
 - b. $\exists e'. \text{material}(e') \neq \text{material}(e) \wedge$
 - c. $\sim \neg [\text{type}(e') \neq \text{type}(e)]$

A context with an additional, contrasting element such as a metal spoon would support the implicature in (10c) and thus would allow the hearer to disambiguate the semantics in favour of the plastic spoon (since the only assignment that makes (10) consistent with the context is one where the plastic spoon is assigned to e). Thus, we are able to predict the *contrastive bias* reported by Sedivy and colleagues and to account for the fact that hearers are able to make predictions about potential referents incrementally, at a point when the ongoing utterance is still semantically ambiguous, without need to reason with hypothesized complete descriptions.

Resorting to scales akin to the property preference orderings typically used in GRE models also allows us to account for the differences observed with different kinds of modifiers. We can explain this by appealing to the relative position of different types of modifiers within the preference scale, specially with respect to the attribute type.¹⁰ If an attribute is highly prominent and there is no other attribute higher up in the preference scale, then it will not give rise to the contrastive implicature, even though it may still have discriminatory power. It seems reasonable to assume that in shared visual situations often colour is very salient, arguably even more salient than the object type (11c). Thus, as observed by Sedivy, in a context with, say, a red plate, a red cup, and a blue cup, upon hearing ‘*Pick up the red...*’ hear-

¹⁰At least judging from the limited experimental conditions tested by Sedivy and colleagues

ers do not exhibit any contrastive bias (i.e. no preference for the red cup is observed).

- (11) ...the red ...
- a. $\exists e. \text{colour}(e) = \text{red}$
 - b. $\exists e'. \text{colour}(e') \neq \text{colour}(e)$
 - c. $S : \langle \text{colour}, \text{type} \rangle$

In this case, the scalar implicature is not inferred since there is no higher-ranked attribute that would license the application of the SCALAR IMPLICATURE INFERENCE RULE.

6 Conclusions

In this paper I have discussed how pragmatic inferences can be exploited during incremental interpretation in the resolution of first-mention referential descriptions. I have concentrated on contrastive effects observed in the psycholinguistics literature and have sketched a proposal that combines elements from computational models of Referring Expression Generation with ingredients from standard theories of scalar implicature. The result is a system that operates incrementally on partial utterances, without need to reason with complete descriptions. Clearly, the main burden of the approach is the determination of the preference scale—a problem that REG models also face. Although the present paper offers only a preliminary account, it hopefully contributes to opening the door for investigating further how pragmatic theories can meet the challenges imposed by the incremental nature of language use in dialogue.

Acknowledgements

I am grateful to the anonymous SemDial reviewers for their comments. This research has been supported by a VENI grant funded by the Netherlands Research Council (NWO).

References

- Jennifer E. Arnold, Carla L. Hudson Kam, and Michael K. Tanenhaus. 2007. If you say thee uh you are describing something hard: The on-line attribution of disfluency during reference comprehension. *Journal of Experimental Psychology Learning Memory And Cognition*, 33(5):914.
- Luciana Benotti and David Traum. 2009. A computational account of comparative implicatures for a spoken dialogue agent. In *Proceedings of the Eight International Conference on Computational Semantics*, pages 4–17, Tilburg, The Netherlands, January.

- Susan Brennan and Herbert Clark. 1996. Conceptual Pacts and Lexical Choice in Conversation. *Journal of Experimental Psychology*, 22(6):1482–1493.
- S. Brown-Schmidt, E. Campana, and MK Tanenhaus. 2005. Real-time reference resolution in a referential communication task. In JC Trueswell and MK Tanenhaus, editors, *Processing world-situated language: bridging the language-as-action and language-as-product traditions*, pages 153–172. MIT Press.
- Herbert Clark and Donna Wilkes-Gibbs. 1986. Referring as a collaborative process. *Cognition*, 22:1–39.
- Herbert Clark. 1996. *Using language*. Cambridge University Press.
- Robert Dale and Ehud Reiter. 1995. Computational interpretations of the Gricean Maxims in the Generation of Referring Expressions. *Cognitive Science*, 18:233–266.
- David DeVault and Matthew Stone. 2003. Domain inference in incremental interpretation. In *Proceedings of ICoS 4: Workshop on Inference in Computational Semantics*, pages 73–87, Nancy, France. INRIA Lorraine.
- Keith S. Donnellan. 1966. Reference and definite descriptions. *The philosophical review*, 75(3):281–304.
- Raquel Fernández. ms. Semantics and Pragmatics in Incremental Reference Resolution: The Case of Relative Gradable Adjectives. Unpublished manuscript.
- Gerald Gazdar. 1979. *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic Press.
- H. Paul Grice. 1975. Logic and conversation. In D. Davidson and G. Harman, editors, *The Logic of Grammar*, pages 64–75. Dickenson, Encino, California.
- Paul Grice. 1989. *Studies in the Way of Words*. Harvard University Press.
- David Grodner and Julie Sedivy. forthcoming. The effect of speaker-specific information on pragmatic inferences. In N. Pearlmuter and E. Gibson, editors, *The Processing and Acquisition of Reference*. MIT Press, Cambridge, MA.
- Barbara Grosz. 1981. Focusing and description in natural language dialogues. In A. K. Joshi, I. Sag, and B. Webber, editors, *Elements of Discourse Understanding*. Cambridge University Press.
- Nicholas J. Haddock. 1989. Computational models of incremental semantic interpretation. *Language and Cognitive Processes*, 4(3):337–368.
- Julia Hirschberg. 1985. *A Theory of Scalar Implicature*. Ph.D. thesis, University of Pennsylvania, Computing and Information Sciences.
- Laurence Horn. 1972. *On the Semantic Properties of Logical Operators in English*. Ph.D. thesis, University of California, Los Angeles.
- Emiel Krahmer, Mariët Theune, Jette Viethen, and Iris Hendrickx. 2008. GRAPH: The costs of redundancy in referring expressions. In *Proceedings of the 5th International Conference on Natural Language Generation*, Salt Fork OH, USA.
- R.M. Krauss and S. Weinheimer. 1966. Concurrent feedback, confirmation, and the encoding of referents in verbal communication. *Journal of Personality and Social Psychology*, 4(3):343–346.
- Christopher S Mellish. 1985. *Computer Interpretation of Natural Language Descriptions*. Chichester: Ellis Horwood.
- Deb Roy. 2002. Learning visually grounded words and syntax for a scene description task. *Computer speech and language*, 16(3):353–385.
- David Schlangen, Timo Baumann, and Michaela Atterer. 2009. Incremental reference resolution: The task, metrics for evaluation, and a bayesian filtering model that is sensitive to disfluencies. In *Proceedings of the 10th Annual SIGDIAL Meeting on Discourse and Dialogue*, London, UK, September.
- William Schuler. 2003. Using model-theoretic semantic interpretation to guide statistical parsing and word recognition in a spoken language interface. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics (ACL)*, pages 529–536, Sapporo, Japan.
- Julie Sedivy, Michael Tanenhaus, Craig Chambers, and Gregory Carlson. 1999. Achieving incremental semantic interpretation through contextual representation. *Cognition*, 71:109–147.
- Julie Sedivy. 2003. Pragmatic versus form-based accounts of referential contrast: Evidence for effects of informativity expectations. *Journal of psycholinguistic research*, 32(1):3–23.
- Julie Sedivy. 2007. Implicature during real time conversation. *Philosophical Compass*, 2/3:475–496.
- Jette Viethen, Robert Dale, Emiel Krahmer, Mariët Theune, and Pascal Touset. 2008. Controlling redundancy in referring expressions. In *Proceedings of the Sixth International Language Resources and Evaluation (LREC' 08)*, pages 950–957.