Collection and analysis of referring expressions used in collaboration

Philipp Spanger Yasuhara Masaaki Tokunaga Takenobu

Department of Computer Science
Tokyo Institute of Technology

Abstract: In order to pursue research of referring expressions occurring in interaction between humans and agents, we first set up a data-collection experiment recording human-human interactions. Adopting a collaboration task of solving the Tangram puzzle by a pair of participants, every utterance in interaction was recorded in synchronysation with the position of every Tangram piece and the operations carried out on them by the participants. Referring expressions in this data were annotated with their referents to build a referring expression corpus. We provide preliminary results of the corpus analysis from various standpoints. Finally, we draw some preliminary conclusions and then discuss ideas on how to extend this research direction.

1 Introduction

Recent progress in various fields related to agent technology has encouraged more concentrated research into human-agent interaction from a variety of standpoints. This research has focused not simply on "entertainment robots", but in particular on interaction in situations where humans and agents carry out a task in collaboration. Extending linguistic capabilities of agents (both understanding and generation of natural language) is a critical part of this research. Referring expressions have been studied for a long time, since they are linguistic expressions often used to refer to a certain object in interaction. Furthermore, they are a crucial device for enabling smooth collaboration between humans and agents in the achievement of a common task.

Initial research into understanding and generation of referring expressions [Dale 89, Dale 95], focussed largely on studying isolated referring expressions in a specific environment and how those expressions were able to distinguish the target object from the given distractors. Often, as an initial stage and in order to address the complexities of dealing with full referring expressions, research focussed simply on the selection of attributes of the target object. The assumption here was that no other factor outside the current situation (e.g. the context of interaction) would have an impact on the formulation of the referring expression.

The seminal work in this field was the Incremental algorithm (IA) proposed in [Dale 95]. In recent years, there have been various extensions to this algorithm improving several major limitations of the IA (e.g. [Krahmer 02, van Deemter 02, Krahmer 03]). The need for a systematic approach in order to allow a unified evaluation of such algorithms provided the motivation for the creation of the TUNA corpus that was developed at Aberdeen University as part of the TUNA project [van Deemter 07]. Work has begun to use this corpus for evaluating different algorithms for attribute selection, e.g. [Gatt 07]. The TUNA corpus is the most extensive collection of referring expressions to date. At the same time, the TUNA corpus has the limitation of only taking into account individual expressions in an interaction-free setting.

Within this context of seeking a unified and comparative evaluation of corpus-based algorithms, the TUNA corpus laid the basis for holding The Attribute Selection for Generating Referring Expressions (ASGRE). It was organised as a Shared Task Evaluation Challenge (STEC) and held as part of the UCNLG+MT workshop in 2007[1]. While STECs have been common in many areas of NLP, as the organisers noted, this ASGRE was organised as a pilot STEC in the GRE area [Belz 07], since they had not yet been tried in the area of Natural Language Generation (NLG).

[Viethen 08] notes that "NLG researchers have tended towards data gathering exercises that explore some specific aspect of referring expression generation, focussing on hypotheses relevant to algorithm development." [Viethen 08] concentrates on the question of the use of spatial relations in referring expressions. Their data collection experiment, while slightly less artificial and less simple than in the TUNA corpus was also exclusively concerned with individual referring expressions outside the framework of any collaboration.

While more intensive research in this direction can lead to increasing our understanding of isolated referring expressions (including attribute selection), at the same time there are important inherent weaknesses in this type of approach, which put a strict limit on the type of knowledge we can gain with this kind of research. In fact, referring expressions to a very large degree exactly *do not* occur in isolated instances but *as part of an interaction*. In particular, looking exclusively at individual expressions – outside the framework of the actual interaction (in which they are usually uttered) and without taking into account neither the linguistic nor task context – deprives us of potentially critical clues to understanding referring expressions in interaction. In addition, studying simply individual expressions also contradicts the experience, that referents are often identifiable thanks to the collaboration context. This in

[1] A further one was held in 2008 and now planned for 2009

turn indicates the necessity to study referring expressions in the complex environment in which they are uttered, in particular in task-oriented dialogue where linguistic activity goes along with actually carrying out a task in collaboration.

An initial important attempt to focus more on studying natural language in dialogue and in particular in a task-oriented setting was the construction of the HCRC Map Task corpus. This corpus was collected in Glasgow and Edinburgh in the early 90s ([Anderson 91]). For building this corpus, an orientation task on a map is employed. However, the exploration of dialogue in this context is carried out in a very controlled way which leads to quite an artificial set-up (e.g. participants maps don't correspond). Furthermore, researchers from a wide range of interests participated in the design of the data collection experiment. While providing thus potentially a resource for studies of natural dialogue from many different perspectives, this also led to a lack of focus on the specific type of expressions we attempt to deal with: referring expressions. To our knowledge, no systematic attempt has been made to carry out an in-depth analysis of the referring expressions occurring in this corpus, except [Bard 00] dealing with the dependency between "intelligibility" of referring expressions and the knowledge of hearers.

More recently, there has been research carried out by [Gergle 07]. They note how groups benefit considerably from access to shared visual information in collaboration and provide some initial experiments that demonstrate this claim in their chosen environment. Their stated aim is "developing a computational model that integrates linguistic cues with visual cues". In their experiment a Helper guides a Worker in constructing an online jigsaw puzzle, while both share a view of the work area rendered on each of their computer screens. Since they focus on the contribution of the visual information, they isolate the impact of visual context (use of "delayed" set-up). While their research is certainly thought-provoking, they only deal with resolving pronouns without paying much attention to other types of referring expressions.

[Foster 08] in contrast concentrates on a specific type of referring expression in task-oriented dialogue which they call "haptic-ostensive" references, which involve manipulating an object. They point out that previous research in the field of generating multi-modal referring expressions has assumed only a very little amount of "shared knowledge" and note that in a collaborative context, "objects can be made accessible simply by manipulating them as part of the task". Hence, they pose the necessity of dealing with the context of jointly manipulating objects ("mutual task in a shared environment"). The setting of their experiment is quite realistic compared to the Map Task in that both participants can intervene into the task. At the same time their experiment seems quite (possibly overly) simple as both participants have access to all necessary information for completion of the experiment. This is also indicated by the fact that participants needed less than 3 referring expressions on average for completion of the task. While an interesting implementation is realised using "haptic-ostensive" references, no further work is done is analyzing other types of referring expressions. Thus their work provides new insight into an important type of referring expression in interaction, however at the same time it is limited by exclusively looking at this one type of expression. They do not deal with other types or how these different expressions interact in the process of reference (e.g. in which situations, which types of expressions are used?).

Our work is carried out in the overall framework of researching referring expressions in a collaborative task. In order to develop a model for understanding and generating referring expressions in this environment, we need a better understanding not only of particular types of expressions – which in itself is important – but of the whole range of expressions used in interaction and their role in the collaborative process. Hence we have been working on collecting data and have begun to analyze the collected referring expressions from various standpoints to get a broader understanding of the referring process in collaboration. This is a necessary part of developing algorithms for generating referring expressions that not only deal with various individual questions (or types of referring expressions) but that seek to deal with the overall complexity and the variety of expressions as used by humans.

2 Building the corpus

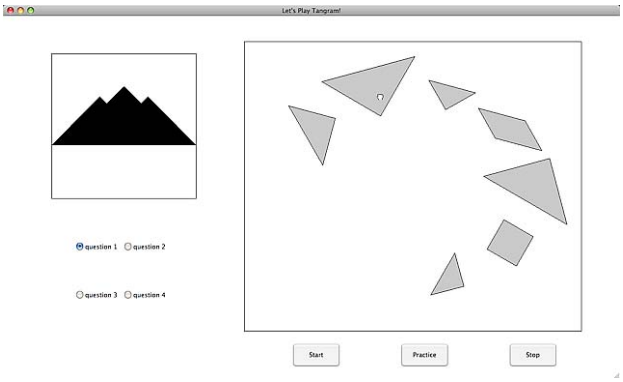

Figure 1: Screenshot of the Tangram simulator

2.1 Experimental setting

We recruited 12 graduate students of the Department of Cognitive Science, four females and eight males, and split them into 6 pairs. Each pair was instructed to solve the Tangram puzzle cooperatively. The Tangram puzzle is a kind of geometrical puzzle that originated in ancient China. The goal of Tangram is to construct a given shape by arranging seven pieces of simple figures as shown in Figure 1. The pieces includes two large triangles, a medium-size triangle, two small triangles, a parallelogram and a square.

In order to record the precise position of every piece during the solving process, and every action the participants made, we implemented a Tangram simulator in which the pieces on the computer display can be moved, rotated and turned over with simple mouse operations. The simulator displays two areas: a goal shape area (the left side

of Figure 1) and a working area (the right side of Figure 1 where pieces are shown and can be manipulated.

We assigned a different role to each participant of a pair: a *solver* and an *operator*. Given a goal shape, the solver thinks of the necessary arrangement of the pieces and gives instructions to the operator how to move them. The solver can not directly manipulate the pieces. The operator manipulates the pieces with the mouse according to the solver's instructions. The operator does not see the goal shape. Sharing the simulator's working area, the solver and operator can communicate only through exchange of linguistic expressions. In this interaction, we can expect many occurrences of referring expressions, particularly of expressions referring to the pieces of the puzzle.

[Foster 08] conducted a similar experiment where a pair of participants cooperatively arrange several pieces as instructed. In their task, however, both participants can equally manipulate pieces. They collected 512 dialogues and found only 1,333 referring expressions in those dialogues. This means the participants used less than three referring expressions in a dialogue. Although their experimental setting might be natural, it is inefficient from a viewpoint of collecting referring expressions. In order to collect referring expressions more efficiently, we assigned asymmetric roles to the participants as described above. In our setting, the solver is obliged to refer to the pieces by using linguistic expressions since she can not directly manipulate the pieces and has to ask the operator to move them. We expected such an asymmetric setting to elicit comparatively more referring expressions than [Foster 08], and in fact it did.

In our Tangram simulator, all pieces are the same colour, thus colour is useless in identifying a specific piece. Unlike in the TUNA corpus, here only size and shape are useful object-intrinsic attributes. Rather, we can expect other attributes such as spatial relations, deictic reference and actions performed on pieces to be used more often in the dialogues.

Figure 2: Experimental setting

The participants of a pair sit side by side as shown in Figure 2. Each participant has her own computer display that share the working area. Only the operator has a mouse as well for manipulation of pieces, but does not have a goal shape on the screen. The solver has a goal shape on the screen but not a mouse. A shield screen was set between the solver and operator to prevent the operator from seeing the goal shape on the solver's screen, and to restrict their interaction to only linguistic expressions.

Each participant pair is assigned four exercises: to form two symmetric and two asymmetric shapes. The participants exchanged their roles after two exercises, i.e. a participant first solves a symmetric and an asymmetric one as the solver and then she does the same as the operator, and vice versa. Before starting the first exercise as the operator, each participant had a short training exercise in order to learn the manipulation of pieces with the mouse. The initial arrangement of the pieces was randomised every time. We set a time limit of 15 minutes for an exercise. In order to prevent the solver from getting into deep thought and keeping silent, the simulator is designed to give a hint every five minutes by showing a correct piece position in the goal shape area. After 10 minutes have passed, a second hint is provided, while the previous hint disappears. A dialogue ends when the goal shape is constructed or the time is up. Utterances by the participants are recorded separately in stereo through headset microphones in synchronisation with the position of the pieces and the mouse actions. We collected 24 dialogues (4 exercises by 6 pairs) of about four hours in total. The average length of a dialogue was 10 minutes 43 seconds.

2.2 Annotation

Recorded dialogues were transcribed with a time code attached to each utterance. Since our main concern is collecting referring expressions, we defined an utterance to be a sentence which is as complete as possible. In this respect, our "utterance" tends to be longer than those of ordinary dialogue corpora in which an utterance is usually a segment between pauses with a certain length. In the transcribed text, referring expressions were annotated together with their referents by using a multi-purpose annotation tool SLAT [Noguchi 08]. A quick look at the text reveals that there are various types of referents of referring expressions: one piece or several pieces, a part of a piece and a location. In the first step, we focus on those referring expressions whose referent is a piece or a set of pieces in the working area. Because representation of the referent is easy for these cases, i.e. a (set of) piece identifier(s).

Two annotators (two of the authors) annotated four dialogue texts separately and decided on the following criteria to identify a referring expression after discussing discrepancies between the two.

- A minimum span of a noun phrase with necessary information including repairs with their reparandum and disfluency [Nakatani 93] are marked. e.g. "*ôkii sankaku iya sikaku* ([the] big triangle, no, rectangle)" → "*ôkii sankaku iya sikaku*" is marked.
- Demonstrative adjectives are included in expressions. e.g. "*sono sankaku* (that triangle)" → "*sono sankaku*" is marked.

- Obviously wrong expressions are marked with a comment. e.g. "*tiisai sankaku* (small triangle)" referring to a square is marked with a comment "lexical error".
- An expression without a definite referent is marked without its referent identifier.
- All expressions appearing in muttering to oneself by any of the participants are excluded.

According to the above criteria, we annotated all 24 dialogue texts and corrected discrepancies by discussion between the annotators. We collected referring expressions of 449 types and 1,509 tokens out of 24 dialogues. Comparing to Foster's experiment mentioned above (1,333 out of 512 dialogues), we succeeded in collecting referring expressions very efficiently.

3 Analysis of the corpus

In this section we report some initial results from analysing the collected corpus. We first provide some figures characterising the overall structure of the corpus and then discuss some specific characteristics.

3.1 Overall structure

As noted, we collected a total of 1,509 tokens and 503 types of referring expressions over 24 dialogues. The experimental setting tended to encourage referring expressions from the solver, while the operator is constrained to confirming his understanding of the utterances of the solver. This is reflected in the number of referring expressions by the solver (1,287) largely outnumbering those of the operator (222). In addition, those expressions uttered by the solver and the operator matched in 18 types of expressions. This leads to the fact that the number of types counted for each solver (470) and operator (51) simply added give a number higher than the total number of types. Table 1 summarises the statistics of the corpus.

Table 1: Number of collected referring expressions

	solver	*operator*	*total*
types	470	51	503
tokens	1,287	222	1,509

In the annotation process, we distinguished the referents of the referring expressions according to the number of pieces an expression refers to. We noted there are a number of expressions (215 expressions or about 15% of the total) whose referent is more than 1 piece (2 pieces or more). For now, we excluded them from our analysis, noting that in future we intend to deal with them. In the present work, we exclusively deal with expressions that refer to either 0 or 1 piece. In this context, an expression that has as a referent "0 pieces" is an expression whose referent cannot be determined without ambiguity. For example if there are two large triangles in the puzzle, and if the solver uses the expression "*ôkii sankaku* (a large triangle)" this can refer to either one of them (i.e. there is not a determined one referent). Table 2 summarises the expressions we excluded in this study. We note that there are 9 types of expressions that were used both by solver and operator (thus the total number of types/tokens is not the simple sum of solver and operator).

Table 2: Number of expressions (referents ≥ 2)

	types	*tokens*
solver	103	161
operator	28	54
total	122	215

We analysed the remaining 1,294 tokens that referred to either 0 or 1 piece (Table 3). It is notable that in this case as well, there were exactly 9 types of expressions that both solver and operator used. In addition, we noted that among those expressions employed by both solver and operator (all expressions including those excluded from analysis), in both cases at least a third were demonstratives (demonstrative adjectives or demonstrative pronouns). This seems to indicate that in our task setting, demonstratives are one of the most "general" kinds of expressions, while the use of other kinds of expressions rather depends on the specific context.

Table 3: Number of expressions (referents < 2)

	types	*tokens*
solver	367	1,126
operator	23	168
total	381	1,294

A more detailed analysis of the collected expressions (referring to 0 or 1 piece) led us to distinguish them according to whether they included the following syntactic/semantic elements: i) demonstratives (adjectives and pronouns), ii) object attribute-values, iii) spatial relations, iv) actions on an object and v) others. If for example an expression contained both a demonstrative adjective and mentioned an action on an object, the expression was characterised by both elements. Thus, the numbers of expressions listed in Table 4 do not simply add up to the total number of expressions in the corpus (1,294). We divided the spatial relations again according to whether they were "projective" (including expressions such as "*migi ni aru sankaku* (the triangle on the right)") or "topological" (expressions that include topological relations such as "*ima ichiban chikaku ni aru sankakkei* (the triangle that is the closest right now")).

The elements listed in Table 4 are of a quite different kind; with the "demonstrative adjective" or "pronoun" being a syntactic category while categorizing an expression according to whether it includes a mention of an action, is related to the expression's semantics. In order to provide a more concrete understanding of our analysis, Table 5 shows an example utterance for each element that we defined. We show the example in the Japanese original with its English translation. The element that we identified in the whole expression is underlined.

Table 4: Elements in collected referring expressions

	Element	*types*	*tokens*
i)	*demonstrative*	118	745
	adjective	100	196
	pronoun	19	551
ii)	*attribute*	303	641
	size	165	267
	shape	271	605
	direction	6	6
iii)	*spatial relations*	129	148
	projective	125	144
	topological	2	2
	overlapping	2	2
iv)	*action*	89	96
v)	*others*	29	30
	remaining	15	15
	similarity	14	15

Overall, we note here a tendency to employ object attributes, in particular the attribute "shape" (605 tokens or slightly over 45%) and the attribute "size" (267 tokens or about 20%). We also observe a tendency to use expressions that include demonstrative pronouns (551 tokens or slightly over 40%). These kinds of referring expressions are quite general and appear in a variety of other non-collaborative settings as well. A very small number of collected expressions use characteristics of the pieces or the task ("overlapping" of pieces, "remaining" pieces and "direction" of pieces). Another kind of expression referred to the "same" type of object (the "same" as previously handled or referred to). In addition, we found another kind of expression not usually employed by humans outside of collaborating on a task; a referring expression mentioning an "action on an object".

96 expressions (slightly over 7% of the total) of this type occurred in our corpus. [Foster 08] note that "In particular, when conversational partners cooperate on a mutual task in a shared environment, objects can be made accessible simply by manipulating them as part of the task." Thus they define "*haptic - ostensive reference*, that is, reference which involves manipulating an object". However, they are here only dealing with references that contain both a linguistic expression as well as a physical action carried out by the participant. In our experiment, we found this kind of ("*haptic-ostensive*") reference, but also much more prominent linguistic referring expressions that included a mention of an action. In this sense, the expressions we deal with include a broader variety of referring expressions. These are the expressions we focus on in the next section.

3.2 More on expressions mentioning actions on objects

Out of the different kinds of expressions employed by the participants in our experiment, we selected those expressions that include mention of an action on an object, and further analysed them. Since this kind of expression does not generally occur outside of a collaborative task environment, we expect that focussing on a further analysis in this area can help to understand the referring process in collaboration.

Table 5: Examples of elements in referring expressions

Element	*Example*
demo.-adj.	"*ano migigawa no sankakkei*" (that triangle at the right side)
demo.-pron.	"*kore*" (this)
attr.-size	"*tittyai sankakkei*" (the small triangle)
attr.-shape	"*ôkii sankakkei*" (the large triangle)
attr.-dir.	"*ano sita muiteru dekai sankakkei*" (that large triangle facing to the bottom)
spa.-pro.	"*hidari no okkii sankakkei*" (the small triangle on the left)
spa.-top.	"*ôkii hanareteiru yatu*" (the big distant one)
spa.-over.	"*sono sita ni aru sankakkei*" (the triangle underneath it)
action	"*migi ue ni doketa sankakkei*" (the triangle you put away to the top right)
remaining	"*nokotteiru ôkii sankakkei*" (the remaining large triangle)
similarity	"*sore to onazi katati no*" (the one of the same shape as that one)

Expressions including mention of actions on objects were about 7% of the total amount of expressions considered here, which is not a negligible figure but is neither comparable in number to other type of expressions, like expressions including demonstrative pronouns. A closer look at the different types of expressions including "actions" showed some interesting points. We observed that those expressions can be again divided up into three categories: i) use of temporal adverbials without explicit mention of an action (about 22% of this kind of expression), ii) combination of a temporal adverbial with a verb indicating an action ("turned", "put", "moved", etc) (about 77%) and iii) expressions exclusively containing an action - verb (about 12%). Table 6 summarises data on this type of expression.

Table 6: Expressions mentioning actions on an object

	types	*tokens*
temporal adverbials w/o verb	19	22
temporal adverbials w/ verb	59	62
verb only	11	12
total	89	96

In terms of those expressions only including a "temporal adverbial", about half used the Japanese "*sakki no* ... (the ...from [just] before)". As in English, the Japanese is ambiguous in terms of referring to a previous *linguistic* (an utterance) or *physical* act (e.g. moving an object). This ambiguity had to be resolved by looking at each of the collected conversations as well as the videos of the experiments and then resolve this ambiguity. The other half of expressions was "*ima no* ... (the current .../the ...you are *verb*-ing/ the ...from just before)" which indicates a current time, i.e. an object the operator is handling (holding, moving, etc) or which both focus on at the current

time. However, at the same time, the expression *ima no* in Japanese can also indicate a just passed point in time and is thus ambiguous. Further analysis should look at these cases in more detail.

A large majority of expressions included a verb explicitly referring to an action on the object.We note that only slightly over 10% overall of this kind of expression did not include any type of temporal expression. This indicates that when mentioning an action in collaboration, humans tend to at least seek to order actions according to "previously done" or "currently doing" and thus this needs to be taken into account in any algorithms dealing with this type of expression.

3.3 Discussion

In this section, we presented some basic findings from an analysis of the corpus and then focussed on expressions mentioning an action on an object. In our experimental setting, about 7% of expressions were of this type. While in [Foster 08], about 36% of initial linguistic expressions were of this kind, we can assume that the particular weight of this way of referring to objects is largely dependent on the experimental setting. In an initial detailed look at these expressions, we noted that they are often combined with temporal adverbials. In particular participants make a distinction between what was done before and what is being done right now. This indicates the need for further study into this type of expressions and its relation to time expressions.

4 Conclusion and future work

This study deals with referring expressions in collaboration. As a first step towards developing efficient algorithms for understanding and generating of referring expressions in this context, we collected a corpus of referring expressions used by humans working on a collaborative task. We employed a Tangram puzzle simulator as an effective way of collecting this type of expression and carried out an initial analysis of the collected expressions.

The referring expressions we collected were of various kinds; incorporating demonstrative adjectives, pronouns or whole expressions mentioning actions on objects. As a type of referring expression that generally does not appear in utterances outside a collaborative framework, we focussed on those expressions including a mention of an action on an object. The analysis we have carried out on the collected expressions is an initial analysis, in so far as we have not yet made progress on a more detailed analysis of any other type of expression mentioned.

Our data collection experiment was limited, since we only allowed for linguistic interaction, excluding other data, available even in this setting (e.g. eye movements) or that was artificially excluded (e.g. gestures to each other). In future, this information should also be incorporated into an extended analysis of referring expressions in collaboration.

References

[Anderson 91] Anderson, A., Bader, M., Bard, E., Boyle, E., Doherty, G. M., Garrod, S., Isard, S., Kowtko, J., McAllister, J., Miller, J., Sotillo, C., Thompson, H. S., and Weinert, R.: The HCRC Map Task Corpus, *Language and Speech*, Vol. 34, No. 4, pp. 351–366 (1991)

[Bard 00] Bard, E. G., Anderson, A. H., Sotillo, C., Aylett, M., Doherty-Sneddon, G., and Newlands, A.: Controlling the intelligibility of referring expressions, *Journal of Memory and Language*, Vol. 42, pp. 1–22 (2000)

[Belz 07] Belz, A. and Gatt, A.: The Attribute Selection for GRE Challenge: Overview and Evaluation Results, in *Proceedings of the MT Summit XI Workshop Using Corpora for Natural Language Generation: Language Generation and Machine Translation (UCNLG+MT)*, pp. 75–83 (2007)

[Dale 89] Dale, R.: Cooking up referring expressions, in *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, pp. 68–75 (1989)

[Dale 95] Dale, R. and Reiter, E.: Computational interpretation of the Gricean maxims in the generation of referring expressions, *Cognitive Science*, Vol. 19, No. 2, pp. 233–263 (1995)

[Foster 08] Foster, M. E., Bard, E. G., Guhe, M., Hill, R. L., Oberlander, J., and Knoll, A.: The roles of haptic-ostensive referring expressions in cooperative, task-based human-robot dialogue, in *Proceedings of 3rd Human-Robot Interaction*, pp. 295–302 (2008)

[Gatt 07] Gatt, A., van der Sluis, I., and van Deemter, K.: Evaluating algorithms for the generation of referring expressions using a balanced corpus, in *Proceedings of the 11th European Workshop on Natural Language Generation (ENLG-07)* (2007)

[Gergle 07] Gergle, D., Rosé, C. P., and Kraut, R. E.: Modeling the impact of shared visual information on collaborative reference, in *Proceedings of 25th Computer/Human Interaction Conference*, pp. 1543–1552 (2007)

[Krahmer 02] Krahmer, E. and Theune, M.: Efficient context-sensitive generation of referring expressions, in van Deemter, K. and Kibble, R. eds., *Information Sharing: Reference and Presupposition in Language Generation and Interpretation*, pp. 223–264, CLSI (2002)

[Krahmer 03] Krahmer, E. and van Erk, S.: Graph-based generation of referring expressions, *Computational Linguistics*, Vol. 29, No. 1, pp. 53–72 (2003)

[Nakatani 93] Nakatani, C. and Hirschberg, J.: A speech-first model for repair identification and correction, in *Proceedings of 31th Annual Meeting of ACL*, pp. 200–207 (1993)

[Noguchi 08] Noguchi, M., Miyoshi, K., Tokunaga, T., Iida, R., Komachi, M., and Inui, K.: Multiple purpose annotation using SLAT – Segment and link-based annotation tool, in *Proceedings of 2nd Linguistic Annotation Workshop*, pp. 61–64 (2008)

[van Deemter 02] van Deemter, K.: Generating Referring Expressions: Boolean Extensions of the Incremental Algorithm, *Computational Linguistics*, Vol. 28, No. 1, pp. 37–52 (2002)

[van Deemter 07] van Deemter, K.: TUNA: Towards a UNified Algorithm for the generation of referring expressions, Technical report, Aberdeen University (2007), www.csd.abdn.ac.uk/research/tuna/pubs/TUNA-final-report.pdf

[Viethen 08] Viethen, J. and Dale, R.: The use of spatial relations in referring expression generation, in *Proceesings of 5th International Natural Language Generation Conference*, pp. 59–67 (2008)