Machine Translation **15:** 125–147, 2000.

Discourse Particles and Discourse Functions

MANFRED STEDE
Universität Potsdam, Institut für Linguistik, P.O. Box 601553, 14415 Potsdam, Germany
(E-mail: stede@long.uni-potsdam.de)

BIRTE SCHMITZ
KPMG Consulting Germany, Kurfürstendamm 207-208, 10719 Berlin, Germany

Abstract. Spoken language, especially spoken German, is rich in particles that do not contribute to the propositional content of utterances, but play important roles in steering the flow of the dialogue and in conveying various attitudes and expectations of the speaker. Languages differ widely in their conventions on particle usage, and therefore these words pose significant problems for translation. As a solution, we propose an inventory of "discourse functions" that characterize the pragmatic impact of particles. These functions are to be assigned to particles in the analysis phase, so that the translation step can use the abstract information to decide on the best way of rendering the same effect in the target-language utterance.

Key words: dialogue, particles, pragmatics

1. Overview

When comparing spoken to written language, one soon notices the abundance (types and tokens alike) of "particles" in speech: The many occurrences of *well*, *oh*, *let's see* and others are a typical dialogue phenomenon. At first sight, they seem to be innocent little words that contribute little to the propositional information conveyed; however, they do play important roles in steering the flow of the dialogue and in conveying various attitudes and expectations of the speaker. Furthermore, individual languages differ in their reliance on either lexical means (i.e., particles) or other prosodic and syntactic means for achieving these effects. Therefore, particles pose significant problems for automatically processing – and especially translating – spoken language. In this paper, we are concerned with spoken language translation (SLT) between German and English; the former is well known for offering a particularly large number of particles used for pragmatic purposes, and we will therefore examine German particles and their appropriate renderings in English.

In the next section, we first define the term "discourse particle" in relation to other particles and then demonstrate that discourse particles can be quite problematic in translation, drawing on extensive corpus analyses we performed in the Verbmobil project (Wahlster, 1993); our examples are from the domain of appointment scheduling, which was addressed in the first phase of the project (1993–1996). Section 3 introduces an inventory of discourse functions that we designed for

characterizing the pragmatic impact of particles, following especially the aims of translation. The task of automatically identifying the discourse function of a particle in understanding utterances is discussed in Section 4, where we outline the relevant processing steps taken in Verbmobil, and describe our rules for assigning discourse functions. Finally, we draw some conclusions on the design of SLT systems.

2. Discourse Particles

> Even without knowing the linguistic and non-linguistic context, it is clear that a speaker uttering [(1a)] conveys a different attitude towards the content of the utterance, towards the hearer, or towards other aspects of the situation than a speaker saying [(1b)]. (Bublitz, 1978: 1, our translation from German)

(1) a. *Du bleibst hier?*
YOU STAY HERE
'You're staying here?'

b. *Du bleibst doch hier?*
YOU STAY particle HERE
'You're staying here, aren't you?'

As a first example of particle usage,[1] the *doch* in (1b) "colours" the question to the effect that the speaker's hope for a positive response is made very clear. Other particles can play other roles, and moreover one and the same particle can play quite different roles in different contexts. Thus, as a first step towards providing a definition for the term "discourse particle" we investigate ambiguities.

2.1. PARTICLE AMBIGUITY

Discourse particles, in our terminology, are words that are uttered not because of their contribution to propositional content, but because of some pragmatic function for the ongoing discourse. Imagine, for instance, utterance (2) in the midst of a discussion; the *right* at the beginning serves mainly to signal turn-taking and initiating some kind of break in the conversation.

(2) Right, now let's discuss our trip to San Francisco.

The exact function of a discourse particle is often difficult to determine, though, and thus the need for disambiguation arises: in the German utterance (3) the *ja* can be a mere filler that smoothes the intonation contour, or it can mark the overall information as *given*, e.g. in a situation where the participants have just closed off the final topic.

(3) *Dann sind wir ja fertig.*
THEN ARE WE particle FINISHED
'So we're finished.'

The disambiguation problem is aggravated considerably by the fact that the vast majority of words that can be used as discourse particles also have a "standard", semantic reading.[2] In (2), the *now* can be read as a temporal adverb or as a semantically empty discourse particle. Similarly, in (3) the *dann* is most likely just an "uptake" particle, but in general it is a temporal adverb as well ('then').

We will follow the terminology of Hirschberg and Litman (1993), who distinguish between the "sentential" usage and the "discourse" usage of such words. In these terms, the problem is that many particles have one (or more than one) sentential usage(s), but it may very well be irrelevant in certain contexts, where only the discourse usage of the particle is relevant – and often that usage has no systematic relationship to the sentential usage. Then, of course, many particles also have a variety of discourse readings, and hence these words are altogether highly ambiguous.

To illustrate the kind of problems under consideration here further, we take another look at the interesting German particle *doch*. When speakers utter *Doch!* as the sole response to their partner, they deny what the partner has just said and at the same time reaffirm their own opinion. In English, one would say something like *On the contrary!* or *Yes it is!* (or use the more specific verb in question).[3] Also, *doch* can be used as an adversative conjunction roughly corresponding to the English *but*. Both these readings are sentential ones.

As for discourse readings, when *doch* is used as part of an utterance, it can have at least the following functions. In (4a), in a context where the other participant has just uttered a sentence that the speaker considers agreeable, the *doch* signals that the proposition is assumed to be shared knowledge, or self-evident. A suitable English rendering is (4b). This is the reading present in the example, which was a question; thus, the reading is independent of sentence mood.

(4) a. *Das ist doch klar.*
THAT IS particle CLEAR
b. That's clear anyway.

When a sentence such as (5a) opens an exchange, *doch* merely smoothes the utterance and should not be translated at all, as in (5b), since English does not offer a corresponding word.

(5) a. *Lassen Sie uns doch einen Termin ausmachen*
LET YOU OURSELVES particle AN APPOINTMENT ARRANGE
b. Let us arrange an appointment.

Finally, in an utterance like (6a), where *doch* is prosodically marked, it signals a return to a previous state of the negotiation: Tuesday had been suggested earlier but was rejected, and now the rejection is taken back. Again, there is no equivalent English particle; instead, a speaker can signal the reversal of their position by saying (6b) for instance.

(6) a. *Dann nehmen wir d o c h Dienstag*
THEN TAKE WE particle TUESDAY
b. All right, so we *do* take Tuesday.

If words of this kind were rare, one could argue that SLT systems simply have more pressing problems to attend to. In German, however, the number of particles is quite significant. Among the vocabulary of the Verbmobil system, we have identified 49 particles that cause problems of the kind just described.[4] For illustration, (7) shows a short excerpt from a dialogue of the Verbmobil corpus, with the potential discourse particles underlined.

(7) **A:** *Ja, Herr Helfer, dann sagen Sie mir nochmal, wann haben Sie im März denn noch Zeit?*
'Well, Mr Helfer, so tell me again, when do you have time in March?'

B: *Der März wäre bei mir eigentlich, äh, die letzte Märzwoche, von Montag dem fünfundzwanzigsten März, bis Freitag den neunundzwanzigsten, da hätte ich frei. Wäre Ihnen das recht?*
'March is for me, uhm, last week in March, Monday 25th to Friday 29th, then I'm available. Is that all right for you?'

A: *Ja, prima, da hab' ich auch noch keinen Termin. Also, das könnte mir ganz recht sein. Hm, vom Montag, dem fünfundzwanzigsten bis zum Freitag den neunundzwanzigsten. Das muss ich mir gleich mal aufschreiben.*
'Yes, fine, I don't have any appointment at that time. So that could be all right for me. Hm, Monday 25th to Friday 29th. I should make a note right away.'

The most frequent and difficult discourse particles are the 20 words listed in Table I. Considering only this set, when randomly taking 1,000 turns from the German Verbmobil corpus, we found 1.75 particles per dialogue turn on average. This figure supports our earlier informal counts, according to which a typical Verbmobil dialogue (which consists of about 15 turns) contains 20 to 30 occurrences of such particles. Therefore, the particle-disambiguation problem for an SLT system is not only difficult but also important.

2.2. TOWARDS A DEFINITION

Having initially described discourse particles as particles used in a discourse reading rather than a sentential reading, we can now take steps towards a more precise definition. Several studies of English discourse particles have suggested that utterance-initial position is the central criterion to identify a discourse reading (e.g.

Reichman, 1985), and certainly, adopting such a criterion makes it much easier to gather and evaluate the data (cf. Byron & Heeman, 1997). The position criterion, however, holds for a certain class of discourse particles at best; Hirschberg and Litman (1993) give a number of counterexamples. Furthermore, it does not apply to the German language, where most particles can occur almost anywhere in the utterance.

For the particles found in the Verbmobil corpus, Bos and Schiehlen (1999) provide a classification into classes such as "focus particle", "sentential adverb" or "discourse relation adverb". These classes are sometimes relevant for disambiguation and translation, in particular for distinguishing sentential usages from discourse usages. For selecting one of the candidate discourse readings, however, we found that positional or structural criteria help only occasionally (see Section 4).

Thus, the class of discourse particles is a heterogeneous one and very difficult to demarcate in syntactic or other formal ways. Schiffrin (1987: 31ff) develops a lengthy definition of "units of talk" and then treats discourse particles ("markers" in her terminology) as bracketing such units. Fischer (1998) characterizes the class of German discourse particles as comprising both "interjections" and "segmentation markers". Interjections "signal the spontaneous expression of a cognitive state" and subsume the classes of modal particles (occuring utterance medially, integrated in the phrasal intonation contour) and hesitation markers. For segmentation markers, Fischer refers back to Schiffrin (1987), who lists these criteria: they bear no grammatical relationship to other elements in the sentence; they are not inflectable; they may be phonologically ill-formed; they connect utterances as a kind of "discourse glue".

Finally, a useful criterion for separating sentential usage from discourse usage is the deletion test: if a particle is removed from the utterance, and the translation is still a truthful rendering of the source utterance (even if it possibly differs in pragmatic force), then a discourse particle is at hand.

2.3. TRANSLATION PROBLEMS

For the task of translating utterances containing discourse particles, the monolingual ambiguity problem is extended by that of finding an appropriate translation, and different languages have developed quite different conventions for using particles. German is known to offer an especially wide range of discourse particles, and therefore it is not surprising that many of them do not have any straightforward English translation at all. Instead, their function in discourse needs to be signalled by different means. And in many cases a particle is best dropped from the translation process altogether, if its existence is due merely to certain conventions in spoken German, which do not carry over to the target language.[5] The problem is, given a particular utterance containing a particle, to tell which case is at hand.

König and Stark (1991) note with respect to the general category of "function words":

> In only a few cases do function words in German have a straightforward counterpart in English. Despite the close genetic relationship between these two West Germanic languages, their lexical structures seem to be very different as far as the appropriate closed classes are concerned. The typical situation is that for each use of a function word in German a choice has to be made from among several English "equivalents". Moreover, equivalence often has to be established on the level of the next larger unit, the phrase, the clause, or the whole sentence. (König & Stark, 1991: 304; our translation)

In support of this observation, Fischer (1998) analyzed particles in bilingual data produced by professional translators and concluded (p. 136) that the traditional word classes interjection, modal particle, segmentation marker, tag question, etc. cannot be maintained in translation, since items from different categories are very often regarded as functional equivalents by the translators. In addition, a lexeme-based translation, even if possible, is quite likely to do harm to the stylistic colour of the utterance. Fillmore (1984: 133) noted that "these expressions, if used in English as their counterparts are used in German, would produce very mannered speech."

Therefore, with discourse particles, the need to set lexical items in correspondence with syntactic constructions or, in the case of speech processing, even with prosodic features, is the rule rather than the exception. Consider the following list of English translations for the German *doch* in its various readings and functions, as compiled from the dictionary entries in König et al. (1990): *but*, *(and) yet* (formal style), *though* (also postponed), *nevertheless*, *nonetheless*, *however*, *but ... still*, *because*, *since*, V-*ing as ...*, emphatic *do*, *anyway*, *after all*, *really*, tag-question, *surely*, *I (dó) wish*, *if (...) only*, *why don't you ...*, *yes*, *of course*, *naturally*. Even when abstracting from the selection task, the sheer range of the *kinds* of translations underlines the difficulty of providing any formal system of translation rules.

We conclude that particle translation cannot be performed directly by lexical correspondence rules. Instead, we opt for first providing an abstract characterization of their pragmatic impact in discourse. Then, the translation process can decide on the means suitable for conveying this impact in the target language, given the specific target-language context. To this end, we first need to classify the specific roles that discourse particles play in conversation.

3. Discourse Functions

To investigate the variety of purposes that particles are employed for, we analyzed 30 German dialogues from the Verbmobil corpus, isolated 49 particles, and checked their relevance for translation. With the deletion test mentioned above, we first separated sentential usages from discourse usages, and for the latter decided

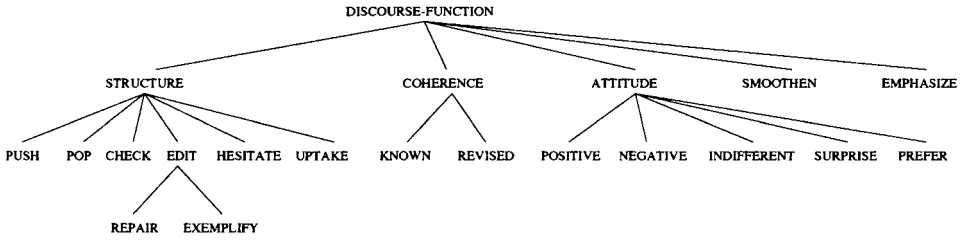

Figure 1. Taxonomy of discourse functions.

on the most appropriate English translation of the clause containing the particle (or several particles). In this way, we determined the different readings for each particle; accordingly, the resulting classification does not amount to a general survey of particle functions (such as that of Fischer, 1998), but was primarily determined by the specific requirements of translation into English. In a second step, we compared the results of the corpus analyses to some classifications proposed in the research literature, which led to a few modifications. The resulting hierarchy of discourse functions is shown in Figure 1. Table I gives the assignment of discourse functions to the 20 most frequent particles in our corpus. (Recall that each particle also has one or more sentential usages, so that a one-to-one correspondence in the table does not represent an unambiguous particle.) We now explain these functions in turn and illustrate them with examples from the Verbmobil data.

3.1. STRUCTURE

3.1.1. *Push, Pop*

These particles mark the opening of a new subtopic (or a brief insertion) and the return to the previous topic, respectively (see for example Grosz & Sidner, 1986). Typical English markers for these purposes are *by the way*, *but*, and *anyway*. Example (8) illustrates PUSH, (9) POP.

(8) *Das ist bei mir sehr schlecht. Mein Name ist übrigens Gürtner. Also, zwanzigster Januar ist sehr ungünstig.*
'That's very bad for me. My name is, by the way, Gürtner. Anyway, January 20th is quite bad.'

(9) *Das ist bei mir sehr schlecht. Mein Name ist übrigens Gürtner. Also, zwanzigster Januar ist sehr ungünstig.*
'That's very bad for me. My name is, by the way, Gürtner. Anyway, January 20th is quite bad.'

Table I. Frequent German particles and their discourse functions

Particle	Function(s)
aber	SMOOTH, PUSH, HESITATE
allerdings	SMOOTH, EMPHASIZE
also	UPTAKE, EDIT, POP, HESITATE
auch	SMOOTH
bloß	SMOOTH, EMPHASIZE
denn	SMOOTH
doch	SMOOTH, REVISED, KNOWN
eben	SMOOTH
ehrlich	CHECK
eigentlich	SMOOTH
erstmal	SMOOTH
etwa	SURPRISE, EXEMPLIFY
gern	SMOOTH
irgendwie	SMOOTH
ja	SMOOTH, EMPHASIZE, UPTAKE, CHECK, KNOWN, HESITATE
jedenfalls	SMOOTH
nämlich	SMOOTH, EDIT
nicht	SMOOTH, CHECK
oder	CHECK
vielleicht	SMOOTH

3.1.2. *Check*

Particles with this function are turn-yielding signals, prompting the dialogue partner to respond. Specifically, by using a CHECK, the speaker often seeks approving feedback from the hearer (example: *isn't it?*).

(10) *Morgen ist Freitag, oder?*
'Tomorrow is a Friday, isn't it?'

3.1.3. *Edit*

This function indicates a modification of the previous utterance segment in the sense of Gülich and Kotschi (1995): a preceding segment can be modified, specialized or reformulated. An example for a general EDIT is (11).

(11) *Da haben wir wieder einen gemeinsamen Termin, nämlich die Tagung in Wien.*
'We have another joint activity then, namely the convention in Vienna.'

We distinguish two sub-cases of EDIT, namely REPAIR and EXEMPLIFY.

With self-repair markers, a previous segment can be explicitly overwritten. According to Gülich and Kotschi (1986), self-repairs consist of an element that gets overwritten, an element that replaces the "old" one, and a reformulation indicator, which we here label with the discourse function REPAIR. Typical English markers are *I mean* and *sorry*.

(12) *Ja, mit den Ferien werden wir uns natürlich etwas in die Quere kommen, weil Sie den ganzen August, nein, ab 15. August nicht können.*
'Well, we're going to have some trouble with the holidays, since you won't be available for the whole of August, no, from August 15th on.'

The second subtype of EDIT marks the addition of more specific, example-like information to the previous segment.

(13) *Ich könnte nächste Woche, etwa am Montag.*
'I could make it next week, for instance on Monday.'

3.1.4. *Hesitate*

This function characterizes particles employed to fill pauses; the speaker signals problems with planning or formulating the utterance and wishes to hold the turn (see for example Schiffrin, 1987).

(14) *Ja, bei mir wäre es günstig, der siebte, Sonntag der siebte, oder, ja, auch der einundzwanzigste wäre möglich, und achtundzwanzigste ebenso.*
'Yes, for me the 7th would be OK, Sunday the 7th, or, well, the 21st would also be OK, and the 28th as well.'

3.1.5. *Uptake*

Foremost, an UPTAKE marks turn-taking: the speaker signals that the partner's preceding utterance has been understood, and that the exchange will go on. UPTAKE particles can also occur in the midst of a turn in order to hold the turn and to connect a new sentence to the previous ones, when no hesitation is involved.

(15) *Also, ähm ich wollte gerade sagen, der Mittwoch, der ist paßt mir ausgezeichnet.*
'So, uhm I just wanted to say that Wednesday is is fine with me.'

3.2. COHERENCE

These functions indicate pragmatic presuppositions on speakers' knowledge, attitudes, or the flow of discourse; thus, they are an important means of giving coherence to the dialogue. We distinguish two types of coherence-oriented functions, KNOWN and REVISED.

3.2.1. *Known*

The speaker signals that the utterance relates to and builds upon knowledge shared with the hearer. In example (16), the *ja* points out that the information presented is assumed to be already given.

(16) *Es handelt sich ja um eine zweitägige Reise.*
'It is, as you know, a two-day trip.'

3.2.2. *Revised*

The speaker revises an assumption or a statement made earlier, or moves back to an older assumption; "the proposition *q* holds unexpectedly, given the contextual factors" (König et al., 1990: 59; our translation). Content is actively being re-planned, which distinguishes this function from REPAIR.

(17) *Da kann ich d o c h.*
'I *can* do it, after all.'

3.3. ATTITUDE

With particles in these functions, speakers signal their attitude toward the propositional content of the utterance.[6] We distinguish five cases: POSITIVE (18), NEGATIVE (19), INDIFFERENT (20), SURPRISE (21) and PREFER (22).

(18) *Oh, können wir ruhig jetzt hier besprechen, also, ich habe jetzt auch nicht so viel Zeit.*
'Oh, we can talk about it right now, it's just that I don't have so much time.'

(19) *Ansonsten bin ich leider schon ausgebucht.*
'Otherwise I'm unfortunately booked already.'

(20) *Machen wir es vielleicht meinetwegen am Nachmittag um drei.*
'As far as I am concerned, we can do it in the afternoon at three.'

(21) *Hast Du den Termin etwa vergessen?*
'Did you actually forget the appointment?'

(22) *Ich würde eher dann den Dienstag danach vorschlagen.*
'I would rather suggest the following Tuesday.'

3.4. SMOOTH

This class corresponds to the German *Abtönungspartikeln* (toning-down particles), which label the propositional content as more suggestive than matter-of-fact, and hence also express politeness. They often serve to smooth the intonation contour of the utterance. More than the other particles, they are typical for spoken language; they "comment upon the sentence from a meta-level and anchor it in the dialogue context" (Hentschel & Weydt, 1989: 14; our translation).

(23) *Wann hätten Sie denn dafür mal Zeit?*
'When would you have time for that?'

3.5. EMPHASIZE

Particles with this function add extra emphasis to a statement; they "characterize a predicate as ranking high on a scale" (König, 1991: 18).

(24) *Ja, das ist eine sehr gute Idee. Das ist wirklich prima.*
'Yes, that's a very good idea. That is really great.'

4. Using Discourse Functions in Verbmobil

With the inventory of discourse functions in place, we can now turn to the task of automatically assigning these functions to particles, i.e. to particle disambiguation. We first sketch the information flow in the relevant portion of the Verbmobil system, and then describe our recognition rules for discourse functions. Thereafter, we explain how the information is used in the translation procedure.

4.1. ANALYSIS AND TRANSLATION IN VERBMOBIL

In the Verbmobil system, several distinct processing lines work on the translation task in parallel, employing different MT paradigms. In one of these lines, "traditional" syntactic and semantic analysis methods are combined with a transfer-based translation module (Dorna & Emele, 1996). Here, the central data structure is a syntactic/semantic representation called the Verbmobil Interface Term (VIT) (Bos et al., 1998). When the analysis components have constructed a VIT for an utterance segment, they pass it on to both the transfer module (which converts it to a target-language VIT) and to the context evaluation module (ConEval), which is in charge of deeper semantic and pragmatic analyses. Figure 2 illustrates the module structure for this particular part of Verbmobil. The ConEval module performs disambiguations, resolves anaphors, determines the dialogue act of the utterance, and extracts the "kernel message" from the utterance (which is later used for a dialogue protocol). For these tasks, we map the VIT to a conceptual representation

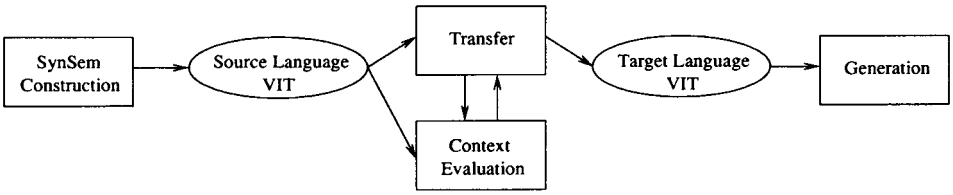

Figure 2. Excerpt from Verbmobil system architecture.

implemented in the description logic LOOM (MacGregor, 1991); the sequence of utterance representations is our context model.

When the transfer module is working on the source VIT, it calls ConEval in case it needs contextual information to decide on the best mapping for a particular predicate. This happens whenever an ambiguous lexeme cannot be disambiguated by inspecting only the utterance segment in which it occurred (i.e. the current VIT), or when information about an anaphor antecedent is required. One class of requests from transfer to ConEval concerns particles: when transfer notices a translation-relevant ambiguity, it asks ConEval for the discourse function of the particle. Answering this request includes making the decision whether a discourse usage is present at all; if it is not, ConEval responds with the pseudo-discourse-function NONPRAGMATIC.

4.2. PARTICLE DISAMBIGUATION IN CONEVAL

Analyzing corpus data recorded from human speakers and coding disambiguation rules for an implemented SLT system are, of course, different enterprises. Restrictions on what can be done result from two sources. First, the "deep analysis" modules can only process those utterances that have successfully passed through the earlier phases, in particular through syntax (assuming a pipe-line architecture among these modules). This constraint is the same as in understanding written language, but when dealing with German discourse particles, grammatical analysis often becomes especially difficult. Second, prosodic information, which yields highly useful disambiguation clues for human listeners, is available only to a limited extent. Even though Verbmobil has made much progress in supplying prosodic features to subsequent analysis modules, not all desirable information can be extracted, and it is not always reliable.

Accordingly, our implemented disambiguation rules account only for a subset of the phenomena we encountered in our corpus analyses. The rules are tailored to the output of the syntax and semantics modules; in particular, they account for the *segmentation* performed by these modules. Dividing the utterance of a speaker into segments suitable for further analysis and translation is one of the most crucial decisions made in the overall system. Since particle usage often does not conform with grammar rules, its treatment happens to be particularly affected by segmentation decisions. Extensive testing with these components of Verbmobil led to sets

of sample VITs with utterances containing the 20 particles listed in Table I. By generalizing from these VIT representations to similar contexts of the particles, we developed a set of rules for each of them. In the following, we first briefly characterize our general approach to disambiguation, then describe the kinds of information (features) used in the particle rules, and finally illustrate the framework with the example *ja*.

4.2.1. *Cascaded Disambiguation*

When considering the kinds of information needed to determine the discourse function of a particle, the task turns out to be very similar to the other contextual evaluation problems tackled in ConEval. For recognizing dialogue acts and disambiguating content words and particles alike, we use a cascaded procedure that operates in three stages.

The first stage employs a set of "strict" rules that allow safe assignment of the intended reading. These rules operate solely on the VIT representations of the current and the preceding utterance, and on the recorded dialogue act of the preceding utterance: no further contextual knowledge is available. An example of a situation that can be handled at this stage is the particle *doch* constituting a complete utterance (or utterance segment): its three pragmatic discourse functions (cf. Table I) all require it to appear in the middle of a segment, and hence it can only assume its semantic reading when appearing in isolation.

When a decision cannot be made in the first stage, a second rule set is invoked. These are not strict implications but " weighted default rules": disambiguation (and also anaphora resolution) is typically not a clear-cut decision but involves weighing up conflicting evidence, resulting in the "most likely" answer. The apparatus of weighted default rules has been developed in our group as an extension of a description logic; its application to the task of dialogue-act recognition is explained in Schmitz and Quantz (1995). For lexical disambiguation, the rules accumulate weights for the candidate readings; the final decision, however, is deferred to the third stage of interpretation.[7]

In the third stage, the VIT is mapped to a conceptual representation in a description logic (LOOM), as mentioned above. Now, the sequence of utterance representations of the preceding dialogue is available, as well as the general knowledge and inference rules encoded in the domain model. Nouns and verbs from the VIT representation become LOOM objects that instantiate concepts from the domain model taxonomy. Furthermore, all predicates relating to temporal information are collected and mapped to a separate representation, for which we have developed a specific reasoning scheme (see Stede et al., 1998) that allows equivalence and inclusion relations between temporal expressions to be computed. These operations are used to track the progress in negotiating an appointment.

At the level of LOOM representations, a third set of default rules can modify the weights that were assigned to candidate interpretations in the second stage. When all these rules have been checked, the highest-ranked reading wins.

There are two motivations for using the cascaded approach. The first stems from the observation that the various disambiguation tasks often interact; for instance, it is sometimes necessary to resolve an anaphor before the dialogue act can be assigned, and disambiguating a particle is sometimes simplified if the dialogue act is known. Therefore, any information inferred by the strict rules in the first phase can be utilized by the default rules in the second phase. The second reason is efficiency: since Verbmobil aims at operating almost in real time, we try to make decisions as soon as possible and without invoking unnecessary and costly reasoning procedures.[8]

4.2.2. *Information Used in the Disambiguation Rules*

The dichotomy between strict rules and default rules reflects differences in the status of the information used. On the one hand, some features are necessary for a particular reading, whereas others are only vaguely indicative. On the other hand, some features used in the rules come from more reliable sources than others; for instance, while information from word order can be taken as correct, prosodic labels in the VIT should not lead directly to a conclusion without possible counterevidence being checked for, since prosodic analysis in the current state of the art is not perfect.

We obtained the features used in our rules by first analyzing the corpus data and then mapping the resulting insights to the shape of the actual VITs produced by Verbmobil. The features can be grouped into the following categories.

Collocations
Supplied by the VIT representations, this is a reliable source of information. Therefore, it is employed by several strict rules in the first stage of disambiguation. We can illustrate this with two examples:

- *aber* is assigned the function SMOOTH when it occurs in phrases like *aber gern*, *aber klar*, *aber ja* and others (which all correspond to the English *sure*);
- when *nicht etwa* occurs in a question (which can also be read off the VIT), *etwa* is assigned the function SURPRISE.

Position in utterance
Under the assumption that utterance segmentation worked well, this is also a reliable feature. Some particles can be quite reliably disambiguated when they constitute a complete utterance (e.g. *doch*, *ja*), or there are at least strong tendencies (e.g. *eben* in its temporal reading occurs only as a response to a *when*-question). Further, we found that when *doch* occurs at the end of the segment, it usually signals the function KNOWN (with the exception of some idoms like *sag' ich doch* 'that's what I say'). Segment-initial position is a necessary but not sufficient condition for UPTAKE. *Also*, for instance, can be classified as such when followed by a complete main clause, as opposed to (25), where it is a conjunctive signalling a causal relationship.

(25) *Also treffen wir uns morgen.*
SO MEET WE US TOMORROW
'So we'll meet tomorrow.'

Syntactic features
These are a reliable source of information, but they serve only as heuristics for disambiguating particles; hence, they are not used in strict rules. Three examples follow:

- *eben* occuring in a clause with past tense usually signifies its temporal, i.e. non-pragmatic, reading (corresponding to *just*);
- when *doch*, unless prosodically marked, occurs in a clause expressing a proposal (e.g. in imperative mood), its function is usually SMOOTH;
- the same holds for *gern* when occuring in the scope of the modal verb *können* 'can', as in (26a) which is a polite way of saying (26b).

(26) a. *Wir können uns gern nächste Woche treffen.*
WE CAN US particle NEXT WEEK MEET
b. I don't mind meeting next week.

Dialogue history
Given the accumulated possibilities of recognition and parsing errors in an SLT system, the dialogue history recorded by our module in LOOM is not always complete and accurate; therefore, we use it only in rules with relatively low weights. One example that was already mentioned is *doch*, which – when turn-initial – is sensitive to the feature of an explicit negation being present in the preceding utterance by the other speaker. The particle *bitte* is ambiguous between 'you're welcome' and 'please say that again'; we check whether the preceding dialogue act was THANK and in that case assign the first reading. A more complex case is *noch*, which can either merely SMOOTH the utterance or mean 'another'. Given our conceptual representation of the preceding turns, we try to determine the event that *noch* has scope over in the present utterance and then attempt to find an event of the same type in the dialogue history. If this query succeeds, we assign the non-pragmatic reading; but if not, we cannot reliably conclude the opposite, due to the incompleteness of domain model and reasoning.

Prosody
When the VIT lists a word as prosodically marked, we take this information as reliable, but we do not draw conclusions from the absence of the feature, due to the limitations of prosodic analysis. The *noch*, mentioned above, can be safely assigned its non-pragmatic reading when it is marked, because no particle merely SMOOTHes the utterance when it is stressed.[9] Similarly we treat a stress mark as sufficient for disambiguating *doch*, *allerdings* and *jedenfalls*.

World knowledge

As pointed out above, reasoning with a domain model is necessarily incomplete. For the subdomain of date and time expressions, which we represent in a specific framework external to LOOM, we perform quite reliable deductions. Particle disambiguation uses these to check whether one date/time expression is more specific than another one, which is an indicator for the EDIT and EXEMPLIFY functions performed by *also*, *etwa*, and *nämlich*. In the LOOM domain model, we have implemented a "contrast check" that determines whether two objects have a common ancestor that is not void of information (e.g. not as abstract as "Thing") and at the same time are not identical. For example, this yields the information that *Hamburg* and *Berlin* are potentially contrastive when they are arguments of the same predicate (e.g. *drive to Hamburg/drive to Berlin*), which can lead to assigning *aber* its non-pragmatic reading as a contrastive connective (which is not obvious from syntax in German). Still, when the contrast predicate returns "false", this can either be a correct answer or a mistake due to an incomplete model.

4.2.3. *Example: Disambiguating* ja

To illustrate our rule mechanism, we now explain the disambiguation procedure for the particle *ja*, which is the most ambiguous and at the same time one of the most frequent discourse particles in German. Adding its NONPRAGMATIC reading to the six discourse functions listed in Table I, the choice is to be made among seven possible readings. After collecting occurences of each reading from the corpus, we related them to actual VIT representations in Verbmobil and isolated a number of features that characterize the usages. As outlined above, features can be classified along two dimensions: (a) reliability of the information source, characterized by the annotations (R+) and (R-); (b) descriptive power of the feature. Here, we distinguish necessary (N) and heuristic (H) features, where (N) indicates that the feature is present in all occurences of the reading in the corpus, whereas (H) means that the feature is present in most of the cases. The features that turned out most useful for describing *ja* are the following:

- turn-initial: *ja* is the first word of a turn
- turn-final
- after pause: *ja* occurs after speaker has paused
- medial: *ja* within a segment, not at beginning or end
- with acceptance: *ja* before or after a segment uttered by the same speaker, with dialogue act ACCEPT or FEEDBACK-POSITIVE
- with statement: *ja* after a segment with declarative sentence mood
- with request
- complete response: *ja* is a complete utterance, following a question by the other speaker
- with known proposition: *ja* within a segment conveying a proposition the hearer knows already
- stress

- rising intonation
- extended duration

The readings, here illustrated with examples, are characterized by the features as follows:

- NONPRAGMATIC either complete response (R+) (N), or with acceptance (R-) (H) and not medial (R+) (N) (27)
- CHECK with statement (R-) (H) and rising intonation (R-) (N) and turn-final (R+) (N) (28)
- EMPHASIZE medial (R+) (N) and stressed (R-) (H) and with request (R-) (H) (29)
- HESITATE extended duration (R-) (H) and not medial (R+) (N) (30)
- KNOWN medial (R+) (N) and with known proposition (R-) (N) (31)
- UPTAKE turn-initial or after pause (R+) (N), and not NONPRAGMATIC (32)
- SMOOTH medial (R+) (N) and not stressed (R-) (N) and not KNOWN (33)

(27) *Ja.*
'Yes.'

(28) *Dann legen wir das fest, ja?*
'So that's definite, isn't it?'

(29) *Kommen Sie ja pünktlich!*
'Make sure you are on time!'

(30) *Ja, ich glaube schon.*
'Well, I think so.'

(31) *Das ist ja ein Wochenende.*
'That is a weekend, after all.'

(32) *Ja, da muss ich mal nachschauen.*
'Well, I have to check.'

(33) *Da haben wir ja etwas wenig Zeit.*
'We have only a little time.'

Only features annotated with both (R+) and (N) can be used in strict rules in the first disambiguation stage. Thus, we have one rule that strictly assigns NONPRAGMATIC when *ja* is a complete response, and two rules that eliminate several candidate readings when the feature "medial" is present or absent, respectively. The remaining decisions have to be made by the default rules, which suggests that the meaning of *ja* is highly context-dependent. The default rules increase weights of the discourse functions as stated above, and they are divided among stages 2 and 3 according to the information they need; for *ja*, only the rule using the feature "with known proposition" is deferred to stage 3 (where domain model and discourse history are available). Since adjusting the numeric weights is an ongoing task that keeps up with changes in the VITs and the segmentation, we leave these details aside here.

4.3. TRANSFER BASED ON DISCOURSE FUNCTIONS

When Verbmobil's transfer module receives the answer to a discourse function request, it decides on the most appropriate translation. In summary, there are four cases for the target-language correspondent:

(a) In sentential usage, the particle has a "literal" lexical translation (i.e. one that is found in a bilingual dictionary), as in (34).

(b) In discourse usage, the particle can have a "non-literal" lexical translation, e.g. *doch* with function KNOWN translated as *after all* (35).

(c) In discourse usage, the particle can have a non-lexical translation, i.e. it is rendered by a syntactic construction or intonation feature, for example *ja* with function CHECK translated as a question tag (36), or stressed *doch* with function REVISED translated as a stressed verb (37).

(d) In discourse usage, the particle can have a zero translation, such as *denn* with function SMOOTH (38).

(34) *Montag ist gut, ich habe aber nur eine Stunde Zeit.*
'Monday is fine, but I have only one hour.'

(35) *Samstag ist schlecht, das ist doch ein Wochenende.*
'Saturday is bad; it's a weekend, after all.'

(36) *Dann treffen wir uns morgen, ja?*
'So we meet tomorrow, don't we?'

(37) *Also treffen wir uns d o c h morgen.*
'So we *do* meet tomorrow.'

(38) *Wieviel Zeit haben Sie denn?*
'How much time do you have?'

Mapping the particle to an intonation feature (c) is currently not implemented in Verbmobil, but a thorough treatment of target-language intonation, which is a rather long-term goal for SLT, should include these cases.

For some particles, straightforward correspondence rules between discourse function and corresponding English translation can be given. In general, though, there is still a choice task to be solved. Returning to the example of *doch*, when the function SMOOTH has been assigned, sometimes deletion is the best option (39). But when a speaker asks the partner to perform some action, where the *doch* mitigates the command-like effect of an imperative, a non-literal translation with a *Why don't you ...* phrase is appropriate (40). Therefore, transfer rules should combine the discourse function of particles with other information about the target clause and then select the best translation.

(39) *Da muss ich doch nochmal nachsehen.*
'I have to check again.'

(40) *Schauen Sie doch einmal nach!*
'Why don't you check that?'

To what extent particles and other utterance segments are deleted from the translation depends on the perspective taken on the translation task: an automated interpreter can be "defensive" and translate as much source information as possible, or it can actively filter the information processed by the system; an ideal system would parameterize this decision and have the users decide. The speech recognition in Verbmobil can detect and remove certain self-corrections. For example, (41a) is not translated literally as (41b) but simply as (41c). Similarly, ConEval can determine a number of cases where a unit of information overrides another, and this is compressed for the purposes of protocoling the dialogue. For example, (42a) literally means (42b), but with the help of calendar knowledge can be identified as a self-correction and reduced to (42c).

(41) a. *Wie wäre es Mon- äh Dienstag?*
b. How about Mon- uhm Tuesday?
c. How about Tuesday?

(42) a. *Ich schlage vor Dienstag den achten, ich meine den neunten.*
b. I suggest Tuesday the eighth, I mean the ninth.
c. I suggest Tuesday the ninth.

5. Summary and Conclusions

In spoken language, discourse particles are employed for pragmatic reasons rather than for their contribution to propositional content. As such, they cannot be translated on a simple lexeme-to-lexeme basis for two reasons: first, individual languages differ widely in the inventory of particles they offer; second, even if a "corresponding" particle is available in the target language, the resulting stylistic colour of the translation can be very different from that of the source utterance.

Therefore, we have proposed a taxonomy of abstract discourse functions to represent the pragmatic impact of discourse particles. There is, still, no one-to-one mapping between particles and discourse functions in analysis, nor between discourse functions and their realizations in the target language. It is thus important to separate the task of assigning the intended discourse function (in the analysis stage) from deciding on the most appropriate translation. For the first step, different kinds of contextual knowledge are required; the second step can select the best translation in the light of the overall target utterance and its options for achieving the pragmatic effect represented by the discourse function of the source particle.

In dialogue, the discourse functions we proposed can be accomplished not only with particles or with prosodic means, but also with "routine formulas": conventionalized phrasal expressions that cannot be translated literally (very similar to idioms). To give just two examples here, the German hesitation marker (43a) should not be translated literally as (43b); rather, an equivalent formula common in English is (43c).

(43) a. *Da muss ich mal gucken ...*
b. There I have to look ...
c. Let me see ...

The discourse function CHECK can be realized by phrases like (44a) which also should not be translated word for word (44b) but with a conventional phrase such as (44c).

(44) a. *Sehe ich das richtig?*
b. Do I see that correctly?
c. Am I right?

Parsing routine formulas and any other phrasal expressions is difficult enough when dealing with written language,[10] but the problem is even more complex in SLT, where the additional uncertainties of speech recognition have to be reckoned with. Our corpus analyses showed that routine formulas play an important role in task-oriented dialogues, and in designing our taxonomy of discourse functions, we tried to incorporate their typical usages. At this point, though, we have to leave this issue to future research.

For the general task of SLT, the issues of particles and routine formulas demonstrate that the methods needed can be quite different from those used in standard MT. Spoken language contains many elements that go beyond conveying propositional content, and these elements can occasionally be even more important than the propositions. Accordingly, SLT has a greater need for *pragmatics-oriented* approaches that pay attention to goals, beliefs, and attitudes of the dialogue participants. This information can be expressed with quite different conventionalized means in different languages. For automatic translation, a strictly lexeme-based transfer strategy (as used in Verbmobil) should be augmented with interlingual elements that abstract over language-specific realizations – and our taxonomy of discourse functions is meant as one such interlingual resource.

One of the most important questions for SLT, in our experience, is that of segmentation, i.e. finding the right units of analysis and translation. In systems processing written language, the "right" unit is taken to be the grammatically well-formed sentence. In spoken language quite often there are no well-formed sentences. Choosing the segment boundaries in the analysis has enormous influence on what the subsequent processing modules are able to do, and ultimately on the quality of the translation. Our problem of finding the discourse function of a

particle is just one of the tasks that depend on the segmentation decision. At the same time, particles very often are an important indicator for the presence of a segment boundary, so that the decisions on segmentation and on particle disambiguation should be made congruently. The Verbmobil system architecture does not allow for feedback from the ConEval module to the segmentation phase, which of course simplifies processing, but a consideration of pragmatic factors in the recognition and segmentation phase would be desirable.

Finally, another reason for giving prominence to abstract, interlingual representations in SLT is the role of prosody in both the source and target utterances. Pragmatic effects are highly dependent on the appropriate intonation, and realizing it in the synthesis phase of SLT (e.g. to produce an effect that was rendered by a discourse particle in the source utterance) requires input that is more abstract than a mere sequence of words: a "concept-to-speech" component.

Acknowledgements

Peter Pollmanns helped with the design and implementation of the particle-disambiguation rules. For their valuable suggestions for improving earlier versions of this paper, we thank Dagmar Barth, Brigitte Grote, and three anonymous reviewers.

Notes

[1] For obvious reasons, we do not attempt to provide a literal gloss for particles in the German examples, and the English translations can only approximate the meaning of the German phrases, due to the very problems under discussion here.

[2] Exceptions are, for instance, the English *oh* or the German *ach*, which are always discourse particles.

[3] *On the contrary* is, however, more general than this use of *doch*, which specifically requires an explicit negation in the preceding utterance, as in (i).

(i) A: *Wir gehen nicht zum Fussball.*
WE GO NOT TO-THE FOOTBALL
'We're not going to the football.'
B: *Doch!*
'Yes, we are!'

[4] This figure does not include modal and other particles that can also pose problems for translation but do not have a discourse usage.

[5] Furthermore, the problem is amplified by the fact that German particles can be combined to form conglomerates, as in (ii).

(ii) *Wir sollten wohl doch noch mal einen Termin ausmachen.*
'We probably should arrange an appointment'.

To what extent these can be analyzed and translated compositionally is an open question. In this paper, though, we deal only with individual particles.

[6] For the attitudinal discourse functions, as well as for EMPHASIZE (see below), it is difficult to provide clear criteria for distinguishing sentential and discourse usage; it might be argued that an attitude towards a proposition is a matter of semantics rather than pragmatics. Still, the deletion test

provides some evidence: when the particle reiterates information already present in the utterance, we see it as semantically redundant, which indicates a discourse usage. Consider, for example, the redundant *meinetwegen* 'as far as I'm concerned' in (20) and contrast it with its being used non-redundantly as a complete utterance in which case it would be translated as 'fine by me'.

[7] At present, the weights used in the particle rules are set manually and adjusted by evaluating the results. An automatic learning procedure would require a large quantity of VITs annotated with discourse functions, which so far is not available.

[8] The mapping to the LOOM representation is performed for every VIT, so that our context model is kept up-to-date. But if a disambiguation is performed in the first stage, the request from the transfer module is answered immediately, so that the translation process can proceed in parallel with our further computations.

[9] Hirschberg and Litman (1993) also point out that particles in their discourse reading are unstressed. There are some exceptions, though. A stressed German *also*, for example, is a good POP marker, similar to the English *so*.

[10] See for example Levin and Nirenburg (1994) for a treatment of idioms in MT.

References

Abraham, Werner (ed.): 1991, *Discourse Particles*, John Benjamins, Amsterdam.

Bos, Johan, C. J. Rupp, Bianca Buschbeck-Wolf, and Michael Dorna: 1998, 'Managing Information at Linguistic Interfaces', *COLING-ACL '98: 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics*, Montreal, Quebec, pp. 160–166.

Bos, J. and M. Schiehlen: 1999, 'Klassifikation der deutschen Partikeln in Verbmobil' [Classification of German particles in Verbmobil], Verbmobil Memo 141, Universität des Saarlandes/Universität Stuttgart.

Bublitz, Wolfram: 1978, *Ausdrucksweisen der Sprechereinstellung im Deutschen und Englischen* [Modes of expression of speakers' attitudes in English and German]. Niemeyer, Tübingen.

Byron, D. and P. Heeman: 1997, 'Discourse Marker Use in Task-Oriented Spoken Dialog', *Proceedings of the 5th European Conference on Speech Communication and Technology (Eurospeech)*, Rhodes, Greece.

Dorna, Michael and Martin C. Emele: 1996, 'Semantics-Based Transfer', *COLING-96: The 16th International Conference on Computational Linguistics*, Copenhagen, pp. 316–321.

Fillmore, Charles: 1984, 'Remarks on Contrastive Pragmatics', in J. Fisiak (ed.), *Contrastive Linguistics. Prospects and Problems*, Mouton, Berlin.

Fischer, K.: 1998, 'A Cognitive Lexical Pragmatic Approach to the Functional Polysemy of Discourse Particles', Doctoral dissertation, Universität Bielefeld.

Grosz, Barabara J. and Candace L. Sidner: 1986, 'Attention, Intentions, and the Structure of Discourse', *Computational Linguistics* **12**, 175–204.

Gülich, E. and Th. Kotschi: 1986, 'Reformulierungshandlungen als Mittel der Textkonstitution. Untersuchungen zu französischen Texten aus mündlicher Kommunikation' [Reformulation as a Means of Text Constitution. Research on French Texts as Oral Communication], in W. Motsch (ed.), *Satz, Text, Sprachliche Handlung*, Akademie Verlag, Berlin.

Gülich, E. and Th. Kotschi: 1995, 'Discourse Production in Oral Communication', in U. Quasthoff (ed.): *Aspects of Oral Communication*, deGruyter, Berlin.

Hentschel, R. and H. Weydt: 1989, 'Wortartenprobleme bei Partikeln' [Category Problems with Particles], in H. Weydt (ed.), *Sprechen mit Partikeln*, deGruyter, Berlin.

Hirschberg, Julia and Diane Litman: 1993, 'Empirical Studies on the Disambiguation of Cue Phrases', *Computational Linguistics* **19**, 501–530.

König, Ekkehard: 1991, 'Identical Values in Conflicting Roles: The Use of German *ausgerechnet, eben, genau* and *gerade* as Focus Particles', in Abraham (1991).

König, Ekkehard and Detlef Stark: 1991, 'The Treatment of Function Words in a New Bilingual German-English Dictionary', in Abraham (1991).

König, Ekkehard, Detlef Stark, and Susanne Requardt: 1990, *Adverbien und Partikeln: Ein deutsch-englisches Wörterbuch* [Adverbs and Particles: A German–English Dictionary], Julius Groos, Heidelberg.

Levin, Lori and Sergei Nirenburg: 1994, 'Construction-Based MT Lexicons', in A. Zampolli, N. Calzolari, and M. Palmer (eds), *Current Issues in Computational Linguistics. In Honour of Don Walker*, Kluwer Academic Publishers, Dordrecht.

MacGregor, R.: 1991, 'Using a Description Classifier to Enhance Deductive Inference', *Proceedings of the Seventh IEEE Conference on AI Applications*, pp. 141–147.

Reichman, Rachel: 1985: *Getting Computers to Talk Like You and Me: Discourse Context, Focus, and Semantics (an ATN Model)*, MIT Press, Cambridge, MA.

Schiffrin, Deborah: 1987, *Discourse Markers*, Cambridge University Press, Cambridge.

Schmitz, Birte and J. Joachim Quantz: 1995, 'Dialogue Acts in Automatic Dialogue Interpreting', *Proceedings of the Sixth International Conference on Theoretical and Methodological Issues in Machine Translation TMI 95*, Leuven, pp. 33–47.

Stede, M., S. Haas, and U. Küssner: 1998, 'Understanding and Tracking Temporal Descriptions in Dialogue', in B. Schröder, W. Lenders, W. Hess, and T. Portele (eds), *Computers, Linguistics, and Phonetics between Language and Speech*, Peter Lang, Frankfurt.

Wahlster, Wolfgang: 1993, 'Verbmobil: Translation of Face-to-Face Dialogues', *Proceedings of the Third European Conference on Speech Communication and Technology*, Berlin.