

Utterance Units in Spoken Dialogue

David R. Traum¹ and Peter A. Heeman²

Abstract. In order to make spoken dialogue systems more sophisticated, designers need to better understand the conventions that people use in structuring their speech and in interacting with their fellow conversants. In particular, it is crucial to discriminate the basic building blocks of dialogue and how they affect the way people process language. Many researchers have proposed the *utterance unit* as the primary object of study, but defining exactly what this is has remained a difficult issue. To shed light on this question, we consider grounding behavior in dialogue, and examine co-occurrences between turn-initial grounding acts and utterance unit signals that have been proposed in the literal, namely prosodic boundary tones and pauses. Preliminary results indicate high correlation between grounding and boundary tones, with a secondary correlation for longer pauses. We also consider some of the dialogue processing issues which are impacted by a definition of utterance unit.

1 INTRODUCTION

Current spoken dialogue systems tend to restrict themselves to simple exchanges with the user. Unlike natural human conversation, turn-taking is generally fixed to the expression of a single sentence or speech act, with little flexibility about when the turn can change. To move beyond this will require a better understanding of how spoken dialogue works. An important starting point is a clear formulation of the basic units of language production and comprehension. Our opinion is that the best units of dialogue for an artificial system to attune to are the very same ones that humans use. Since spoken dialogue systems are meant to interact with humans, they must be able to use appropriate conventions in responding in a timely and appropriate fashion.

It has often been claimed that for spoken dialogue, *utterances* rather than *sentences* are the primary object of study [3, 4]. But just what *are* utterances? How they are built up from more primitive bits of speech, and how do they cohere with other utterances by both the same and other speakers? Following Bloomfield [2], the term *utterance* has often been vaguely defined as “an act of speech.” The problem is that action comes in many different types and sizes. As discourse analysis of written text concerns the relationships between different sentences rather than sentence internal relationships, discourse analysis of spoken dialogue should concern the relationships between utterances. Finding an appropriate definition of *utterance units* is thus an important starting point for distinguishing utterance-internal language processes (e.g., phonology, speech repairs) from those that operate at a discourse level, (e.g., turn-taking, grounding, rhetorical relations).

Analysts have proposed many different definitions of utterances and *utterance units*. The *turn* is the unit of dialogue that has most often been proposed for study as a basic utterance unit. Fries [8], for example, uses the term *utterance unit* to denote those chunks of talk that are marked off by a shift of speaker. Some authors (e.g., [20]) also distinguish *speaking-turns* in which new information is conveyed from *backchannel items*, which are short responses such as *ok*, *right*, *yeah*, and *mm-hm*. While the turn has the great advantage of having easily recognized boundaries,³ there are several difficulties with treating it as a basic unit of spoken language. First of all, the turn is a multi-party achievement that is not under the control of any one conversant. Since the turn ends only when another conversant speaks, a speaker’s turn will have only an indirect relation to any basic units of language production. If the new speaker starts earlier than expected, this may cut off the first speaker in midstream. Likewise, if the new speaker does not come in right away, the first speaker may produce several basic contributions within the span of a single turn.

From a purely functional point of view, many analysts have found the turn too large a unit for convenient analysis. Fries, for example, noted that his utterance units could contain multiple sentences. Sinclair and Coulthard [27], found that their basic unit of interaction, the *exchange*, cut across individual turns. Instead, they use *moves* and *acts* as the basic single-speaker components of exchanges. A single turn might consist of several different moves, which might be part of different exchanges. In the next section, we review several proposals for utterance units smaller than a single turn.

Even though we believe that the basic utterance units can be smaller than individual turns, turn boundaries are still the easiest place to recognize utterance unit boundaries. If we assume that the turn is composed of utterance units that are smaller than turns, then the end of a turn should also be an end of an utterance unit. When a new conversant starts to speak, this is partial evidence that he believes that the speech by the previous speaker contained a complete interactional unit. The way he responds can give even more evidence about whether he thinks it is complete. If the new speech builds on or responds to the previous speech, this supports a hypothesis that the previous speech was adequate, whereas if the previous contribution was ignored, this is evidence that it might have been deficient in some manner and not obliging the same type of response.

By examining the relationships of features of the end of the old turn and the beginning of the new turn, we can hope to see which features signal the end of utterance units. We looked at the *grounding* behavior displayed by the new turn with respect to previous units by the old speaker, contrasting the distribution of types of relationships based on whether boundary tones and pauses were present at the end of the old speaker’s turn.

¹ TECFA, FPSE, Université de Genève, 9 Rte de Drize Bat. D, 1227 Carouge, Switzerland. Email: David.Traum@tecfa.unige.ch

² Computer Science Dept., University of Rochester, Rochester, NY 14627, USA. Email: heeman@cs.rochester.edu

³ Difficulties still remain, such as how to count turns when more than one conversant is speaking, and in determining whether a particular utterance counts as a backchannel item.

After reviewing some proposals for sub-turn utterance units in the next section, we describe some of the reasons why a fixed notion of such units is important, particularly for the development and functioning of spoken dialogue systems. Then, in Section 4 we review some previous work on grounding. In Section 5, we present our scheme for labeling relatedness between utterances, which captures the most immediate aspect of grounding – how a new speaker initially responds to prior speech. We then describe how we used this scheme, along with prosodic markings to label a dialogue corpus (Section 6) and the results of this study (Section 7). We conclude with some observations on how to extend this study and some implications for spoken dialogue systems.

2 PROPOSALS FOR SUB-TURN UTTERANCE UNITS

Sacks et. al, [24] present a theory of the organization of turns as composed of *turn-constructional units* (TCUs). At the conclusion of each TCU there occurs a *transition-relevance place* (TRP), at which time it is appropriate for a new speaker to take over (or the current speaker may extend her turn with a subsequent TCU). TCUs may consist of differing types of syntactic contributions, including lexical, phrasal, clausal, and sentential constructions. Much subsequent work on turn-taking (e.g., [6, 20, 7]) has tried to analyze what features are used to signal a TRP. The features that were examined included syntactic completions, pauses, and various prosodic features including boundary tones.

When people speak, they tend not to speak in a monotone. Rather, the pitch of their voice, as well as other characteristics, such as speech rate and loudness, varies as they speak. (Pitch is also referred to as the fundamental frequency, or f_0 for short.) The study of intonation is concerned with describing this phenomenon and determining its communicative meaning. For instance, as most speakers of English implicitly know, a statement can be turned into a question by ending it with a rising pitch.

Pierrehumbert [21] presented a model of intonation patterns that later was taken as part of the basis for the ToBI (Tones and Break Indices) annotation scheme [26]. This model describes English intonation as a series of highs (**H**) and lows (**L**). The lowest level of analysis deals with stressed words, which are marked with either a high or low *pitch accent*. The next level is the *intermediate phrase*, which consists of at least one stressed word, plus a high or low tone at the end of the phrase. This *phrasal tone* controls the pitch contour between the last pitch accent and the end of the phrase. The highest level of analysis is the *intonational phrase* (IP), which is made up of one or more intermediate phrases and ends with an additional high or low tone, the *boundary tone*, which controls how the pitch contour ends. Pierrehumbert and Hirschberg [22] theorize that boundary tones “play a considerable role in the conveyance and perception of discourse segmentation”. They claim that an **H** tone signifies that the intonational phrase is forward-looking, which would account for its use both on yes/no questions, and for its use to signal that the next utterance will continue the current one.

In addition to marking intonational phrasing, the ToBI annotation scheme also uses break indices to segment spoken dialogue. In its current incarnation, the break indices have a direct mapping to the intonational phrasing, but work by Wightman et al. [32] uses two additional levels above the intonational phrase. The first level (level 5) represents a superior major tone [14], and the next level (level 6) is “reserved for marking sentence boundaries.” In this work, Wightman has not focused on spoken dialogue, but prepared text spoken by

newscasters.

Previous spoken dialogue research efforts in the TRAINS project have needed to take a stand on the issue of utterance segmentation. In looking at the relationship of prosody to discourse structure, Nakajima and Allen [18] used four principles to segment turns into utterance units: utterance units correspond to sentences of text, they can correspond to basic speech acts, they are at most a single turn, and they can be marked with a pause of at least 750 msec. Gross et al. [10], in segmenting the TRAINS-91 dialogues used slightly smaller units, dropping the syntactic and speech act principles, and using one of the following as a signal of an utterance unit boundary: a boundary tone, a pause in speech longer than a single beat, or a resetting of the pitch level, starting a new intonational phrase. The identified utterance units thus correspond roughly to Pierrehumbert’s IPs, although they also split IPs at points of disfluency.

There have also been some more radical proposals for smaller basic units. For example, Poesio [23] proposed that each word or morpheme be a basic unit, accessible to all aspects of the discourse processing, including semantics, reference resolution, and dialogue management. This position has some attractions because there is psycholinguistic evidence that people do in fact use context even at such basic levels and are able to act appropriately (e.g., by gazing at an object) as soon as the necessary words to disambiguate the meaning are heard [28]. Additionally, this position allows a uniform style of processing and representation for phenomena like pronoun resolution which have both intra- and inter-sentential components.

3 USES FOR UTTERANCE UNITS

We feel, however, that such a step as word by word discourse processing goes too far — there is a legitimate role for sub-turn utterance units in distinguishing local from discourse phenomena in dialogue processing. While certain aspects of the local context will be necessary for an agent to attend to, many issues can better be dealt with by attending to the utterance unit and performing some kinds of processing only within UU boundaries, or between utterances. We briefly review some of these issues in this section.

3.1 Speech Repairs

Heeman and Allen [11] propose that speech repairs should be resolved very early in speech processing. Speech repairs are repairs where the speaker repairs or repeats something she just said, as in this example from [13].

we’ll pick up a tank of uh the tanker of oranges
reparandum alteration

In the above example the speaker intends that the text marked as the *reparandum* be replaced by the *alteration*. Speakers tend to give strong local clues to both mark the occurrence of speech repairs [16] and to help the hearer determine the intended correction [15].

One type of speech repair of particular interest is the *fresh start*, in which the speaker abandons what she is saying, and starts over. Exactly how much is abandoned is not explicitly signaled. Consider the following example from the TRAINS corpus [12] (d93-13.3 utt7), with intonation phrase endings marked as ‘%’.

two hours %
um okay %
so um I guess I’d like to take um
and tankers only take orange juice right %

In this example a fresh start occurs starting on the fourth line. From the context, it seems that the speaker is canceling the text starting the word “okay” – this is just the current utterance, not material communicated in previous IPs.

Other examples of fresh starts are not as clear as to how much material should be cancelled. Consider the following example (d93-24.3 utt10):

and pick up one tanker %
now how long d-
it takes an hour between Corning and Dansville % (*with question intonation*)

In this example, the question is whether the text “and pick up one tanker” is canceled along with the text “now how long d-”. If it is, then it might be a counterexample to the hypothesis that fresh starts just cancel the current utterance. Although the hearer does not take up this suggestion, a satisfactory alternate explanation is that it is his discourse obligation to first respond to the speaker’s question [30].

Not all same-turn repairs are speech repairs. Consider the following example (d93-26.2 utt41):

oh that wouldn’t work apparently %
wait wait %
let’s see %
maybe it would %
yeah it would %
right %
nineteen hours did you say %

In the above example the speaker changes what she is saying, yet due to the complete intonational phrase endings, it doesn’t seem as if a speech repair is taking place. The repair is at a more interactional level, since the speaker needs the participation of the hearer to complete it. For true speech repairs, the processing is much more localized, and perhaps strictly to the current utterance unit. A firm definition for utterance units would thus provide a good metric for deciding which repairs could be done locally, and which repairs need more information and have a more global effect.

3.2 Speech Act Identification

While there is a variety of action that proceeds at different levels throughout spoken dialogue, determination of utterance unit sizes can be helpful for distinguishing certain types of action. As an example, the theory of Sinclair and Coulthard [27], mentioned above, is composed of hierarchical actions, some of which occur across speaker turns, and some of which occur within the turn of a single speaker. The notion of an utterance unit could be useful in isolating which span of speech represents particular tank of action. Also, the multi-level conversation act theory proposed in [31] used an utterance unit distinction. In that proposal, turn-taking acts occurred within utterance units, grounding acts (see Section 4) occurred at the utterance unit level, and traditional core speech acts were composed of multiple utterance units.

Obviously, for such a theory, the nature of an utterance unit will affect how some acts are classified. For example, the set of grounding acts from [31] included acts for *repair* and *continue* functions. Trivially, smaller utterance units would lead to more continue acts, but consider the case in which a repair (e.g., a speech repair) occurs *within* an utterance unit. In this case, the function of the entire unit with respect to the previous context would be considered, and the unit might be seen as performing a continue function. If, however,

such a unit were split up so that the repair was in one unit and the reparandum is in a previous unit, the new unit might be termed a repair.

While this might would not pose a serious problem for processing within a system, it does make it much harder to judge validity of such a coding scheme. Although there is a general agreement that it is important for dialogue systems to attend to action, there is still not much agreement on what the proper set of actions should be. For action types which are sensitive to an utterance unit distinction, performing any evaluation becomes next to impossible if the proper set of units are not held constant for different coders.

3.3 Dialogue Management

There are a number of ways in which a good definition of utterance units would facilitate the tasks of a dialogue manager. For one thing, the basic decision of when and how to respond to input from a user can be informed by awareness of utterance units. Even if full discourse processing of language proceeds within an utterance unit, as suggested by [23], the decision of when to respond still mainly follows regular patterns, occurring at regular places [24]. Awareness of the markers for turn-endings and TRPs can help with the decision of when to respond and when to wait for further input.

More importantly, the nature of the previous utterance unit can be a good signal of how to respond. Suppose, for instance, the speaker projects a desire to continue speaking. In this case, unless there is some major problem with the previous input, the system can issue a backchannel type response, without detailed processing (yet) of the speaker’s implications of the contribution.

On the other hand, suppose there is some problem in understanding the current input. In this case, there are two main options left to the system: (1) wait for the user to say more and hopefully clarify or repair the problem, (2) notify the user of the problems, requesting some sort of repair. Attending to utterance unit production can help in deciding between these two options – in the midst of a unit, it will be preferable to wait for the user, using speech-repair detection techniques, such as those presented in [11]. At an utterance unit boundary, the dialogue manager can shift strategies, dependent on the anticipated continuation by the user. Given a signal that the user would like to continue her turn, if the problem is some sort of ambiguity or under-specification that is likely to be cleared up by future speech, the system can continue to wait. If, however, the problem is an error in understanding (or disagreement about) the previous utterance, the system can decide to break in with a repair.

An attunement to utterance units can also help a dialogue manager decide how to structure the speech output of the system. By using regular utterance units, a system could provide the same resources to a user for providing feedback of understanding. Also, a system could structure the output in order to request help from the user, such as by ending an utterance unit at a point where the user could fill in a referring expression that the system is having difficulty calculating. In general, the division of contributions into more manageable sized utterance units, which can build up turns at speech, will be a more efficient mechanism for flexibly coordinating the contributions of the dialogue to reach a mutual understanding of what has been said.

4 GROUNDING

Clark and Schaefer [5] call the process of adding to common ground between conversants *grounding*. They present a model of grounding in conversation, in which *contributions* are composed of two phases,

presentations and *acceptances*. In the presentation phase, the first speaker specifies the content of his contribution and the partners try to register that content. In the acceptance phase, the contributor and partners try to reach the *grounding criterion*: “the contributor and the partners mutually believe that the partners have understood what the contributor meant to a criterion sufficient for the current purpose.” Clark and Schaefer describe several different methods that are used by the partners to accept the presentation of the contributor. These include feedback words such as *ok*, *right*, and *mm-hm*, repetition of the previous content, and initiation of a next relevant contribution. Completions and repairs of presentations of the contributor also play a role in the grounding process.

Traum and Allen [29] built on this work, presenting a *speech acts* approach to grounding, in which utterances are seen as actions affecting the state of grounding of contributed material. In addition to some acts which present new material, there are *acknowledgment* acts which signal that the current speaker has understood previous material presented by the other speaker, *repairs* and *requests for repair*. Acknowledgment acts include three types, *explicit* acknowledgments which are one of the feedback words, whether they appeared as a backchannel or not, *paraphrases* of material presented by the other speaker, and *implicit* acknowledgments, which display understanding through conditional relevance. Novick and Sutton [19] looked specifically at explicit acknowledgments, cataloging the interaction sequences that they play a part in.

5 GROUNDING AND RELATEDNESS

For the purposes of determining an appropriate utterance unit boundary, we are not as concerned with whether something is ultimately grounded, but rather whether the responding agent *starts* the grounding process (with a repair or some component of the acceptance phase), and *which* previous utterances are being grounded. We thus use a less detailed labeling of utterances than that of [29], with a slightly different focus. We lump the repair, request for repair, and the paraphrase and implicit categories of acknowledgment together into one category we call *related*. While, for a particular utterance, it can be difficult to judge *which* one of these functions is being performed, it is usually very straightforward to determine whether or not it performs some one of them. We also separate out the *explicit* acknowledgments, since, while they generally perform some sort of acknowledgment, it is not possible to tell with certainty *what* they are acknowledging. Likewise, for utterances which follow backchannels and other turns which consist solely of these signals, there is no real content for the new turn to be related to. We also allow categories for *unrelated* utterances, which either introduce new topics, or cohere only with previous speech by the same speaker and do not play a grounding role towards presentations by the other speaker. Our final category is for those utterances for which it is *uncertain* whether they are related or not.

We used single character annotations to indicate relatedness of utterances. These characters are shown on the line above the text which they label. The coding scheme is explained below, including some examples.

5.1 Category: Explicit Ack

Label: e

Description: This category includes all utterances by a new speaker which begin with an **explicit** acknowledgment – the first word of the turn by the new speaker is a word like “okay”, “right”, “yeah”,

“well”, or “mm-hm” (also “no” and “nope” in a negative polarity context). These markers are also called *backchannel items* [20], or simply acknowledgments [5, 19]. This category shows relatedness to *something*, although it is usually not really possible to tell what exactly (if anything) it is related to. This marking is used both for true *backchannels*, which do not claim a turn but overlap speech by the other, isolated acknowledgments between turns by the other speaker, and longer turns beginning with one of these items. This category is also used for repair requests for which it is not clear exactly what should be repaired (e.g., an utterance consisting solely of “what?”).

Example: d93-18.3 utt11-12

u: so we have to start in Avon
e
s: okay

5.2 Category: related

Description: This category is for those utterances that *demonstrate* a relationship to a previous utterance. This can be a demonstration style acknowledgment (i.e., repetition or paraphrase of a previous utterance by the other speaker), a repair or repair request, or some kind of continuation. Any second part to an *adjacency pair* [25] (such as an answer to a question) would fit in this category. In addition, the recency of the relationship is marked as follows:

Label	Description
0	utterances related to the most recent utterance unit by the previous speaker.
1	related to the utterance unit one previous to the most recent but <i>not</i> related to the most recent.
2	related to utterance unit two previous to the last one (and not to anything more recent).
etc.	higher numbers for related to utterances further back.
,	related to previous material by the other speaker, but it is unclear to the coder whether they are related to the immediately previous utterance unit (which would be marked 0), or to an utterance unit further back (which would be marked 1, or 2, etc.).

5.3 Category: unrelated

Label: u

Description: these utterances are **unrelated** to previous speech by the old speaker. This could be either the introduction of a new topic, or could show a relation to previous speech by the same speaker, while ignoring speech by the other speaker in the prior turn.

Example: d93-18.3 utt13-18

U: how long does it take to bring engine one to Dansville
0
S: three hours
e
U: okay <sil> and then <sil> back to Avon to get the bananas
0
S: three more hours si(x) - six in all
u
U: how long does it take to load the bananas

5.4 Category: uncertain

Label: ?

Description: The coder is **uncertain** whether or not these utterances relate to previous speech by the other speaker.

5.5 Combination Categories

In addition, we have a special subclass of markings for those cases in which the main category labeling is misleading or uninteresting.

Category: After explicit

Label: X-e

Description: A suffix of **-e** is used for those utterances which are **after** turns which consist *only* of **explicit** acknowledgments (or backchannels). Since these utterances do not have any content that new material *could* be related to, the categories for unrelated to last (**u**, **1**, etc.), are not very significant. Therefore, we append the marking **-e** to the other category in these cases.

Subcategories:

Label	Description
u-e	material which is unrelated to previous material by the other speaker, but for which the last utterance unit by the other speaker was just an explicit acknowledgment.
1-e	material which is related to the penultimate utterance unit by the other speaker, but for which the last utterance unit by the other speaker contained just an explicit acknowledgment. This signal is often used for popping sub-dialogues, such as in the following example, in which the last turn by speaker S refers to the first utterance by speaker U.

Example:d93-18.3: utt51-54

```
U: how long does that take <sil> again
   0
S: that ta(ke)- <sil> so just to go from Bath to Corning
   e
U: mm-hm
   1-e
S: two hours
```

6 Data

As part of the TRAINS project [1], a corpus of problem-solving dialogs have been collecting in which a user collaborates with a planning assistant to accomplish some task involving manufacturing and shipping goods in a railroad freight system. Since this corpus contains dialogues in which the conversants work together in solving the task, it provides natural examples of dialogue usage that computer systems will need to handle in order to carry on a dialogue with a user. For instance, our corpus contains instances of overlapping speech, backchannel responses, and turn-taking: phenomena that do not occur in collections of single speaker utterances, such as ATIS [17]. For our current study, we looked at 26 dialogues from the TRAINS-93 corpus [12]. This totaled over 6000 seconds of spoken dialogue comprising 1411 total turn transitions.

6.1 Prosodic markings

All turn-initial utterances are marked with respect to their direct relatedness to the previous full or partial intonation phrase (IP) by

the previous speaker. Full IP's are terminated by a boundary tone (labeled %). Partial phrases are anything from the last complete IP to the beginning of speech by the new speaker. If there was no previous IP by the previous speaker than the entire turn so far is counted as a partial IP. The amount of silence between turns is also noted, labeled below in seconds within square brackets (e.g., [.42]).

6.2 Relatedness Markings

Each turn-transition was marked as to how the initial installment of the new turn related to the last few completed or uncompleted IPs produced by the other speaker. For cases of overlapping speech, the speech of the new speaker was marked for how it relates to the current ongoing installment by the continuing speaker. For simultaneous starts of IPs, only the speech by the new speaker was marked for how it related to the previous speech by the current speaker. The continuing speech of one speaker was not marked for how it related to embedded speech by the other speaker, including true backchannels,

Table 1 summarizes this coding scheme described in Section 5, as used to mark turn-initial relatedness to utterance units from the previous turn.

Label	Description
e	explicit acknowledgment (e.g., "okay", "right", "yeah", "well", or "mm-hm")
0	related to the most recent utterance by the previous speaker
1	related to the utterance one previous to the most recent but <i>not</i> related to the most recent
2	related to utterance two previous to the last one (and not to anything more recent)
,	related to previous material by the other speaker, but it is unclear to the coder whether they are related to the immediately previous utterance unit (which would be marked 0), or to an utterance unit further back (which would be marked 1 , or 2 , etc.)
u	unrelated to previous speech by the old speaker
?	uncertain whether these utterances relate to previous speech by the other speaker
u-e 1-e	the same meaning for the first item, but follows a turn by the other speaker consisting only of an item marked e

Table 1. Relatedness Markings

6.3 Examples

Below we show some examples of labeled dialogue fragments. The prosodic and relatedness markings are shown above the line they correspond to. The first example shows a simple sequence of **0** and **e** relations, all following boundary tones, with clean transitions. The first response by S shows a relationship to the first contribution by U. Then the last two start with explicit acknowledgments.

Example: d93-13.2: utt18-22

```
U: how long is it from Elmira to Dansville % [.42]
   0
S: Elmira to Dansville is three hours % [1.23]
   e %
U: okay um so why don't uh % [1.42]
   I send engine two with two boxcars to Corning
   e %
S: okay
```

The following example shows some of the other categories. The second turn by U shows the **u-e** category, since the previous turn was just the utterance of “okay”. This turn also ends without a boundary tone. S’s second turn seems to be some sort of clarification attempt, but it is not clear if it is at all related to the content of U’s previous utterance. U’s final utterance is merely a continuation of his own previous utterance and is unrelated to the last installment by S (which also has no final boundary tone).

Example: d93-16.2 utt27-31

```

% [.73]
U: then do that
e % [.38]
S: okay
u-e
U: um what is tank- what is engine okay engine two is [1.54]
? [.26]
S: you have you told
u
U: it's picking up the oranges

```

7 RESULTS

7.1 Prevalence of Grounding Behavior

Tabulating the markings on the beginning of each stretch of single speaker speech yields the results shown in Table 2. This shows how the next utterance is related to what the other speaker has previously said, and so gives statistics about how much grounding is going on. Of all turns, 51% start with an explicit acknowledgment (category **e**); 29% are related to previous speech of the other speaker (one of the following categories: **0 1 2**, **1-e 2-e**, **-e**); 15% are unrelated to previous speech of the other speaker, but follow an acknowledgment (**u-e**); 2% are possibly related or possibly unrelated (category **?**), and only 3% are clearly unrelated and do not follow an acknowledgment.

Category	#	%
Explicit	696	51%
Related	400	29%
Unrelated after Explicit	199	15%
Unrelated	42	3%
Uncertain	29	2%
Total	1366	100%

Table 2. Prevalence of Grounding Behavior

These results give strong evidence of grounding behavior at turn transitions. Fully 80% of utterance display grounding behavior, while another 15% occur in positions in which (according to the theory of [29, 31]) further grounding is unnecessary. It is only in 3-5% of turn transitions in which a lack of orientation to the contributions of the other speaker is displayed.

7.2 Boundary Tones

Table 3 shows how relatedness correlates with the presence of a boundary tone on the end of the preceding speech of the other speaker. Here, we have subdivided all of the markings into two groups, those that occur at a smooth transition between speaker turns (*clean transitions*), and those in which the subsequent speech overlaps the previous speech (*overlap*). For the overlap cases, we looked for a boundary tone on the last complete word before the overlapping speech occurred.

The distribution of the overlaps into tone and no-tone categories is still somewhat problematic, due to the potential projectability of IP boundaries [9]: a new speaker may judge that the end of a unit is coming up and merely anticipate (perhaps incorrectly) the occurrence of a tone. Thus for some of the entries in the second to last column, there is a boundary tone which occurs after the onset of the new speaker’s speech.

Type	Clean Transitions			Overlaps		
	Tone	No Tone	%	Tone	No Tone	%
e	501	24	95%	77	94	45%
0	267	17	95%	16	41	28%
1,2	7	4	64%	7	11	39%
,	9	4	69%	1	6	14%
1,2-e	7	0	100%	3	0	100%
u	18	7	72%	2	15	12%
u-e	186	2	99%	5	6	45%
?	17	3	85%	6	3	67%
Total	1012	61	94%	117	176	40%

Table 3. Boundary Tones and Relatedness

For the clean transitions, we see that more than 94% of them follow a boundary tone. Of more interest is the correlation between the relatedness markings and the presence of a boundary tone. For explicit acknowledgments and utterances that are related to the last utterance, we see that 95% of them follow a boundary tone. For transitions in which the next utterance relates to an utterance prior to the last utterance, or is simply unrelated, we see that only 64% and 72% of them, respectively, follow a boundary tone.

7.3 Silences

We next looked at how the length of silence between speaker turns (for clean transitions) correlates with boundary tones and relatedness markings. The relatedness markings that we looked at were related-to-last (**0**), and unrelated-to-last (**1 2 u**). Due to the sparseness of data, we clustered silences into two groups, silences less than a half a second in length, (labeled *short*), and silences longer than a half a second, (*long*). The results are given in Table 4.

Type	Tone			No Tone		
	Short	Long	%	Short	Long	%
0	160	107	40%	6	11	65%
u,1,2	15	10	40%	8	3	27%

Table 4. Silences

We find that when there is a boundary tone that precedes the new utterance, there is no correlation between relatedness and length of silence (a weighted t-test found the difference in distributions for related-to-last and unrelated-to-last to not be significant, with $p=0.851$). This suggests that the boundary tone is sufficient as an utterance unit marker and its presence makes the amount of silence unimportant.

In the case where there is no boundary tone, we see that there is a correlation between length of silence and relatedness markings. Only 27% of unrelated transitions follow a long pause (the mean silence length was 0.421 seconds, with a standard deviation of 0.411), while

65% of the related transitions follow a long pause (the mean silence length was 1.072 seconds, with a standard deviation of 0.746). Although there are few data points in these two categories, a weighted means t-test found the difference in the distributions to be significant ($p=0.014$). So, long pauses are positively correlated with the previous utterance being grounded, and thus long silences seem to be a secondary indicator of utterance unit completion.

8 DISCUSSION

Our results are still preliminary, due to the small sample of the relevant categories. However, they do show several things convincingly. First, although grounding behavior is prevalent throughout these problem solving dialogues, there are different degrees to which the speech is grounded. Since adding to the common ground is a prime purpose of conversation, grounding should prove a useful tool for further investigating utterance units and other dialogue phenomena. Second, the claim that utterance units are at least partially defined by the presence of an intonational boundary seems well supported by the conversants' grounding behavior: in addition to serving as a signal for turn-taking, boundary tones also play a role in guiding grounding behavior. Finally, the grounding behavior suggests that pauses play a role mostly in the absence of boundary tones.

There are several ways in which to follow-up this study. First would be just to examine more dialogues and thus get a larger sample of the crucial *related to prior* category, which indicates that the last (partial) unit was deficient but a previous unit was worth grounding. Another possibility would be to look closer at the overlapped category to gain a handle on projectability, and look closely at the contrast in distributions from the clean transition case. Even our preliminary numbers from Table 3 show that the overlap category contains more cases of no boundary tone, which is what we would expect in a case of disagreement over turn-taking.

8.1 Implications for Spoken Language Systems

In Section 3, we described some of the benefits for spoken language systems that could be derived from having a clear notion of utterance units. The results of this study show that the use of prosodic information will play a key role in the ability to make use of such units. Recognition of prosodic features, especially boundary tones will be crucial to the determination of utterance units. Also, production of appropriate prosodic features will be essential to achieving useful contributions. Since silence can be seen as the termination of an utterance unit in the absence of a boundary tone, it is also important for a production system to avoid pausing where it would not like a user to make such a judgment of unit finality. This could be achieved by using a continuation tune, or perhaps a filled pause.

One of the important features of utterance units is due to the role they play in grounding. They seem to be the basic unit by which human speakers coordinate to ground their contributions. Hence, utterance unit boundaries define appropriate places for the machine to offer backchannel responses, or to check its space of interpretations to determine whether it is appropriate to give positive (e.g., an acknowledgment) or negative (e.g., a repair) feedback.

ACKNOWLEDGEMENTS

Funding for the second author was gratefully received from NSF under Grant IRI-90-13160 and from ONR/DARPA under Grant N00014-92-J-1512. We would also like to thank James Allen.

REFERENCES

- [1] James F. Allen, L. K. Schubert, G. Ferguson, P. Heeman, C. H. Hwang, T. Kato, M. Light, N. Martin, B. Miller, M. Poesio, and D. R. Traum, 'The TRAINS project: a case study in building a conversational planning agent', *Journal of Experimental and Theoretical Artificial Intelligence*, 7, 7–48, (1995).
- [2] Leonard Bloomfield, 'A set of postulates for the science of language', *Language*, 2, 153–164, (1926).
- [3] Gillian Brown and George Yule, *Discourse Analysis*, Cambridge University Press, 1983.
- [4] Herbert H. Clark, *Arenas of Language Use*, University of Chicago Press, 1992.
- [5] Herbert H. Clark and Edward F. Schaefer, 'Contributing to discourse', *Cognitive Science*, 13, 259–294, (1989). Also appears as Chapter 5 in [4].
- [6] Starkey Duncan, Jr. and George Niederehe, 'On signalling that it's your turn to speak', *Journal of Experimental Social Psychology*, 10, 234–47, (1974).
- [7] Cecilia Ford and Sandra Thompson, 'On projectability in conversation: Grammar, intonation, and semantics'. presented at the Second International Cognitive Linguistics Association Conference, August, 1991.
- [8] Charles Carpenter Fries, *The structure of English; an introduction to the construction of English sentences.*, Harcourt, Brace, 1952.
- [9] Francois Grosjean, 'How long is the sentence? Predicting and prosody in the on-line processing of language', *Linguistics*, 21(3), 501–529, (1983).
- [10] Derek Gross, James Allen, and David Traum, 'The TRAINS 91 dialogues', TRAINS Technical Note 92-1, Department of Computer Science, University of Rochester, (July 1993).
- [11] Peter Heeman and James Allen, 'Detecting and correcting speech repairs', in *Proceedings of the 32th Annual Meeting of the Association for Computational Linguistics*, pp. 295–302, Las Cruces, New Mexico, (June 1994).
- [12] Peter A. Heeman and James F. Allen, 'The Trains spoken dialog corpus', CD-ROM, Linguistics Data Consortium, (April 1995).
- [13] Peter A. Heeman, Kyung-ho Loken-Kim, and James F. Allen, 'Combining the detection and correction of speech repairs', in *Proceedings of the 4rd International Conference on Spoken Language Processing (ICSLP-96)*, (1996). To appear.
- [14] D. R. Ladd and N. Campbell, 'Theories of prosodic structure: Evidence from syllable duration', in *Proceedings of the XII International Congress of Phonetic Sciences*, Aix-en-Provence, France, (1991).
- [15] Willem J. M. Levelt, 'Monitoring and self-repair in speech', *Cognition*, 14, 41–104, (1983).
- [16] R. J. Lickley and E. G. Bard, 'Processing disfluent speech: Recognizing disfluency before lexical access', in *Proceedings of the 2nd International Conference on Spoken Language Processing (ICSLP-92)*, pp. 935–938, (October 1992).
- [17] MADCOW, 'Multi-site data collection for a spoken language corpus', in *Proceedings of the DARPA Workshop on Speech and Natural Language Processing*, pp. 7–14, (1992).
- [18] Shin'ya Nakajima and James F. Allen, 'Prosody as a cue for discourse structure', in *Proceedings 2nd International Conference on Spoken Language Processing (ICSLP-92)*, pp. 425–428, (October 1992).
- [19] David Novick and Stephen Sutton, 'An empirical model of acknowledgement for spoken-language systems', in *Proceedings ACL-94*, pp. 96–101, Las Cruces, New Mexico, (June 1994).
- [20] Bengt Orestrom, *Turn-Taking in English Conversation*, Lund Studies in English: Number 66, CWK Gleerup, 1983.
- [21] J. B. Pierrehumbert, 'The phonology and phonetics of english intonation', Doctoral dissertation, Massachusetts Institute of Technology, (1980).
- [22] Janet Pierrehumbert and Julia Hirschberg, 'The meaning of intonational contours in the interpretation of discourse', in *Intentions in Communication*, eds., Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, SDF Benchmark Series, 271–311, MIT Press, (1990).

- [23] Massimo Poesio, 'A model of conversation processing based on micro conversational events.', in *Proceedings of the Annual Meeting of the Cognitive Science Society*, (1995).
- [24] H. Sacks, E. A. Schegloff, and G. Jefferson, 'A simplest systematics for the organization of turn-taking for conversation', *Language*, **50**, 696–735, (1974).
- [25] Emmanuel A. Schegloff and H. Sacks, 'Opening up closings', *Semiotica*, **7**, 289–327, (1973).
- [26] K. Silverman, M. Beckman, J. Pitrelli, M. Ostendorf, C. Wightman, P. Price, J. Pierrehumbert, and J. Hirschberg, 'ToBI: A standard for labelling English prosody', in *Proceedings of the 2nd International Conference on Spoken Language Processing (ICSLP-92)*, pp. 867–870, (1992).
- [27] J. M. Sinclair and R. M. Coulthard, *Towards an analysis of Discourse: The English used by teachers and pupils.*, Oxford University Press, 1975.
- [28] M. K. Tanenhaus, M. J. Spivey-Knowlton, K.M. Eberhard, and J. C. Sedivy, 'Integration of visual and linguistic information in spoken language comprehension.', *Science*, **268**(3), 1632–34, (1995).
- [29] David R. Traum and James F. Allen, 'A speech acts approach to grounding in conversation', in *Proceedings 2nd International Conference on Spoken Language Processing (ICSLP-92)*, pp. 137–40, (October 1992).
- [30] David R. Traum and James F. Allen, 'Discourse obligations in dialogue processing', in *Proceedings of the 32th Annual Meeting of the Association for Computational Linguistics*, pp. 1–8, (1994).
- [31] David R. Traum and Elizabeth A. Hinkelman, 'Conversation acts in task-oriented spoken dialogue', *Computational Intelligence*, **8**(3), 575–599, (1992). Special Issue on Non-literal language.
- [32] Colin W. Wightman, Stefanie Shattuck-Hufnagel, Mari Ostendorf, and Patti J. Price, 'Segmental durations in the vicinity of prosodic phrase boundaries', *Journal of the Acoustical Society of America*, **91**(3), 1707–1717, (March 1992).