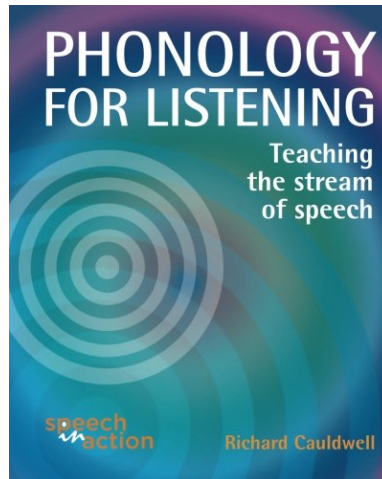**Phonology for Listening: Relishing the messy**

April 2002 (revised December 2013)

[Two publications of mine have the title
'Phonology for Listening'. This paper, and my book
*Phonology for Listening: Teaching the Stream of Speech*
which was published in April 2013, and is available here.]



Richard Cauldwell

**Abstract**

Traditional phonology has failed the student listener of English: its focus on the tidy forms of speech misrepresents the 'mess' of everyday spontaneous speech. The stream-like characteristics of everyday spontaneous speech change familiar words to such an extent that they become unrecognisable. The methodology for listening is also at fault: activities often serve other goals – discussion, writing, grammar. In addition the text-based components of the exercises distract attention from the direct encounters with recordings which contain the mess of streamed speech. In this paper I demonstrate, insofar as print allows, two pervasive types of messiness, one at lexical level, the other at syntactic level. I then use a phonological pattern from everyday speech (the double-prominence tone-unit) to introduce a tool – the word-crusher – which will help both teachers and learners practise and relish the mess of the stream of speech – so that they can improve their ability to listen to everyday spontaneous speech.

**Introduction – A Definition**

This paper addresses the relationship between language description (phonology of English) and language education (teaching listening to non-native speakers of English). Traditional phonology, which focuses on the tidy forms of speech (phonemes to citation-forms), has an adverse effect on the teaching of listening. Its findings and techniques have had value in teaching pronunciation, but have little application to listening. Indeed, representing speech sounds as a sequence of phonemes and citation forms is a misrepresentation of speech as far as listening is concerned. This is a case of 'linguistics applied', as opposed to 'applied linguistics': a linguistic discipline (phonology) has been applied to issues in language teaching, without an appropriate formulation of what language teaching requires of this discipline.

Lass (1980) argues that 'The 'ideal' phonemic system and tactic rules of [citation forms] become a kind of background against which the extremely messy products of [casual speech] can be interpreted' (p. 298). However, language pedagogy – following academic phonology – has avoided dealing with 'the extremely messy' aspects of everyday speech (Cauldwell, 1998; 2000). This situation needs to change if we are to improve the teaching of listening skills. It is necessary to devise a description – a phonology – of the 'extremely messy' features of casual speech, which starts from the premise that such speech – in all its intra- and inter-speaker diversity – is normal, not deviant.

I propose an applied linguistics definition of Phonology: A description of fast everyday spontaneous speech which aids the teaching of listening and comprises a goal, a set of items to teach, a methodology, and a technology for teaching it. This paper will focus on identifying items to teach.

**1. The listener's dilemma**

A phonology for listening has to address the issues that learners confront when listening. Goh (1997, p. 366) reports one learner's (Ying's) thoughts:

> I believe I need to learn what the word sounds like when it is used in the sentence. Because sometimes when a familiar word is used in a sentence, I couldn't catch it. Maybe it changes somewhere when it is used in a sentence.

For Ying, words which should be familiar to her become unrecognisable in the stream of speech. But interestingly Ying describes her problems in terminology that itself is an obstacle to progress: the concepts of 'word' and 'sentences' need to be deconstructed.

**2. Words 1**

A pronunciation key in a dictionary will give the citation form of a word; for example the pronunciation of 'where' is given as |weə|. However, the citation form of a word (within a given variety of English) is fixed and bricklike: when

read aloud, speakers customarily precede and follow the citation form with a pause and it occurs in a tone-unit of its own – it is prominent (stressed) and – generally – occurs with a falling tone. It is bricklike in that its boundaries (initial and final consonants) and its internal structure (syllables, vowel quality) are reproduced with clarity.

However, if we look at the fate of word 'where' in a recording ('In a strange town', taken from Brazil 1994) we find that the citation form is rare. Of the four occurrences of 'where' in the recording only one resembles the citation form. Table 1 shows the tone-units in which the four instances of 'where' occur. The tone-units are shown in the left-hand column: syllables in lower case letters are non-prominent, syllables with upper-case letters are prominent, syllables with upper-case letters and underlining are both prominent and tonic (i.e. they are the location of the tone of the speech unit).

*Table 1 The fate of 'where' in the stream of speech*

| Tone-Unit | Duration of 'where' | Intonation choice on 'where' |
|---|---|---|
| 1. \|\| where there were STREET LIGHTS \|\| | 0.1 | non-prominent |
| 2. \|\| where she'd SAID \|\| | 0.1 | non-prominent |
| 3. \|\| WHERE MARket street was \|\| | 0.2 | prominent |
| 4. \|\| but i WASn't sure WHERE \|\| | 0.3 | tonic |

You can hear these examples [here].

In both tone-units 1 and 2, *where* is non-prominent, and lasts about one tenth of a second, it sounds like an abrupt truncated bleat of a lamb and is spoken at a speed of 600 words per minute (wpm). In tone-unit 3 *where* |we| is prominent, lasts twice as long (spoken at a speed of 300 wpm) but still falls short of the citation form. Only in tone-unit 4, in which *where* occurs before a pause does the citation form occur, at a speed of 200 wpm.

In using the recording in which these tone-units come, one of my groups of advanced learners reported that they could not hear the word *where* in tone-units 1 and 2. When the non-prominent sections *where there were* and *where she'd* were isolated, they still had difficulty reconciling the short sharp bleats with their knowledge of the citation form |weə|.

These learners (themselves teachers of English) were being hampered by their knowledge of the citation form. The citation form (the brick-like word) was an obstacle to improvement of their listening skills. This is, it seems to me, a natural consequence of the way in which the spoken language is represented and taught. Following traditional phonology, it is represented as a set of citation forms. But in natural language use, the words are streamed: the citation form (as can be seen from our examples above) only occurs under special conditions (tonic prominence before a pause). It does occur, but it is a rare beast: it is a useful reference point (Lass's 'background'), but it is misrepresentation of the essential streamlike nature of all speech.

As a consequence of the focus on the citation form, students expect to hear citation forms in speech, and when they fail to hear them, they blame themselves for their failure. This is emphatically <u>not </u>a failure on their part, it is a failure of pedagogy which (following traditional phonology) does not mediate between the 'background' of the citation form, and the 'extremely messy products' (Lass again) of the stream of speech.

## 3. Dissolving words 2

One could argue at this point that words such as *where* belongs to a category of 'function words', or 'weak forms' and are thus far more subject to being messed up by the stream of speech than non-function words or non-weak forms. Not so. Ying's dilemma expresses a truth about all words. Consider the following tone unit from Cauldwell (1997) which was spoken at a rate of 408 wpm:

> *Example 1*
> ‖ this is ONE I'm going to be looking at in slightly more DEtail in fact ‖

Table 2 shows the fifteen words of this tone unit both orthographically and in citation form. Notice that the speaker's pronunciation of the word *one* is |wɒn| not |wʌn| (column 3).

> *Table 2*

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|
| this | is | **ONE** | i'm | going | to | be | looking | at | in | slightly | more | **DE**tail | in | fact |
| ðɪs | ɪz | wɒn | aɪm | gəʊɪŋ | tuː | biː | lʊkɪŋ | æt | ɪn | slaɪtli | mɔː | diːteɪl | ɪn | fækt |

This is a double-prominence tone-unit, with prominences occurring in columns 3 (ONE) and 13 (DE). All the words between columns four and twelve (inclusive) are non-prominent: they include the 'content' and supposedly 'non-weak' forms *going*, *looking at* and *slightly* and *more.* The very fact that these words are occurring non-prominently illustrates an extremely important fact about the stream of speech: all words – content words included – have weak forms.

Table 3 attempts to represent in graphic form what can only be properly appreciated by listening to sound-files. It shows, in the right hand column, the messy products of this tone-unit, aligned with the citation form and the speed of each part of the tone unit in words per minute. Although at this point I should have used a greater range of phonetic symbols to represent the detail of what is happening, I have, in the interests of clarity and readability, used symbols with which readers will be familiar.

Table 3 From words to messy products in a fifteen word tone-unit

|   | Words | Speed | Citation forms | Messy products |
|---|-------|-------|----------------|----------------|
| 1 | This | 588 | ðɪs | ʔɪs |
| 2 | is | 600 | ɪz | z |
| 3 | ONE | 465 | wɒn | wɒn |
| 4 | i'm | 938 | aɪm | aɪ ŋ |
| 5 | going | 438 | gəʊɪŋ | gəʊ |
| 6 | to be looking at | 1218 | tuː biː lʊkɪŋ æt | lʊkɪŋæt |
| 7 | in slightly more | 450 | ɪn slaɪtli mɔː | ɪslaɪʔimuːt |
| 8 | DEtail | 143 | diːteɪl | diːteɪl |
| 9 | in fact | 130 | ɪn fækt | ɪn fækt |

Row 7 indicates the fate of the words 'in slightly more'. The citation form sound shapes of are thoroughly messed up: the |n| of 'in' is omitted; the consonant cluster |tl| is reduced to a glottal stop; the vowel quality of 'more' is shifted from low-back to high-back. Extreme changes indeed, compounded by the speed of utterance: – three words spoken in four tenths of a second at a speed of 450 words per minute.

So how are we to prepare students for such speech? A common argument at this point goes: 'It is impossible to teach every phonetic form that a word might take'. Rost, for example, states:

> it is doubtful that 'fast speech rules' can be learned deductively and consciously applied in real time whenever one encounters an unfamiliar blur of sound (1990, p. 57)

Rost goes on to argue that only immersion in 'contexts of actual use' will allow learners to develop 'gradually a phonological sensitivity to the new language' ... leading to the adoption of 'language specific principles in decoding speech' (p. 57). Rost's pessimism is widely felt, and is representative of the view that the ability to cope with the 'extremely messy products' of the stream of speech can only be learned by osmosis – by extensive exposure. However, it is my belief that we can, using the structures of the stream of speech help learners much more than we do at present. But to do so, we need to abandon a cherished friend, the sentence.

**4 The fate of the sentence**

The sentence is the second obstacle to effective listening that we noted in Ying's diary above. Traditional phonology's approach to the sentence is that phenomena such as prominence, tone, and pausing occur at grammatical boundaries. But the facts of everyday spontaneous speech most often confound these expectations. Take the sentence 'While I was at university, I was very involved with the students arts society which was called the Arts Umbrella'. Traditional Phonology would predict that each of the three clauses

would co-occur with a tone-unit, 'content' words would be prominent, function words would be non-prominent, and tones would occur on the last lexical item in the tone-unit.

*Example 2*
1 ‖ ⬈ WHILE i was at uni<u>VER</u>sity ‖
2 ‖ ⬈ i was very in<u>VOL</u>VED with the STUdents' ARTS so<u>CI</u>ety ‖
3 ‖ ⬊ which was CALLED the ARTS um<u>BREL</u>la ‖

Conventions are as above for Table 1, with the addition that the tones (two rising and one falling) occur on the underlined syllable – the final prominence in the tone unit.

However, this 'sentence' was not uttered as in Example 2: it is a stretch of attested unscripted, spontaneous, speech from Cauldwell (1997) which occurred as follows:

*Example 3*
1 ‖ ⬈ WHILE i was at uni<u>VER</u>sity ‖
2 ‖ → i was VEry in<u>VOL</u>VED ‖ [pause]
3 ‖ → with <u>THE</u> ‖ [pause]
4 ‖ → <u>STU</u>dents ‖ [pause]
5 ‖ ⬊ <u>ARTS</u> society ‖
6 ‖ ⬊ which was CALLED the ARTS um<u>BREL</u>la ‖

Although the first and third clauses occur in tone units much as predicted, the 'main clause' is broken up into four tone-units (2-5). The speaker uses level tone in speech units 2, 3 and 4 tones to give herself time to decide what to say next; she uses <u>two</u> falling tones (5 & 6) instead of the 'predicted' <u>one</u>. The non-co-occurrence of clause and tone-unit, and such use of level and falling tones are very common phenomena in spontaneous speech – they are symptomatic of the structural messiness of everyday spontaneous speech.

From the point of view of the tidy concepts of words (as citation forms) and sentences (as units of speech with predictable characteristics) everyday speech is 'messy'. The tidy concepts of word and sentence contribute to the dilemma Ying feels herself to be in, to my advanced students' inability to perceive *where*.

## 5. A goal for listening

In order to improve this situation we need a goal for applied phonology, a goal which treats the characteristics of everyday speech not as messy deviant phenomena, but as natural and normal. It should be a goal that respects the word-dissolving, non-clause-boundary-respecting nature of the stream of speech.

We thus need a goal for listening work:

to make students familiar and comfortable with the real-time acoustic blur of the stream of speech, and the way in which this stream is shaped by speakers to communicate meanings in all contexts.

This goal has a number of features worth highlighting. First, *make ... familiar and comfortable* means that teachers have to devote time to helping students with those parts of the stream of speech that are most likely to seem strange, difficult, and frustrating. Second, by *acoustic blur* I mean the normal stream-like state of everyday spontaneous speech, where the boundaries of words are blurred together and the sound-shape of the word-bodies are altered by speaker decisions (as with *where* above) and by speaker characteristics (such as gender, accent, and voice quality). Third, there is mention of the *real-time* status of speech – that it typically occurs once, quickly, and is unavailable for inspection other than in the short/medium term memory – crucially it is unsupported by written transcription. Fourth there is mention of both the <u>form</u> (*shaping the stream*) and <u>meaning</u> of speech: all work relating to this goal should relate to both the perception, and the interpretation of speech. Lastly there is mention of *contexts*: most of the textbook rules relating to pronunciation and listening come from an examination of decontextualised scripted language – most recordings for listening are done in studio surroundings, producing ideal, clear speech – this has resulted in students such as Ying being led into false expectations concerning the acoustic clarity of what they will hear.

Before I suggest how we might begin to address this goal in specific actvities, it is necessary to explore further why a goal specific to listening is necessary.

### 6. A problem with methodology

One of the problems with contemporary approaches to listening is that the focus is on listening-as-an-activity rather than listening-as-a-goal.

White (1987) asked a range of teachers of differing nationalities why they valued listening materials. The left hand column of Table 3 lists the replies as published in Anderson & Lynch (1988, p. 66) and the right-hand column (added by myself) relates these activities to their implied goals.

*Table 4 Reasons for valuing listening materials, and their implied goals*

| Reason | Implied Goal |
|---|---|
| good for starting discussions | speaking skills |
| can be used for self-access learning | learner autonomy |
| contains a variety of tasks | motivation |
| entertaining | motivation |
| easy to use | teachability |
| practises guessing from context | general comprehension |
| amusing | motivation |
| uses authentic material | relevance |
| consolidates language | revision of vocab & structure |

None of the reasons listed in the left-hand column relate directly to the listening skill. These reasons imply a focus on the goals listed in the right-hand column: so that (for example) a listening activity which is 'good for starting discussions' is serving a speaking goal. Listening is thus relegated to the status of a servant with speaking as the master. We have got so used to using listening activities in this way that we fail to use the properties of the recordings (its messy words and syntax) to improve students' abilities to cope with everyday listening.

Consider for the moment the (increasingly common) case of an authentic unscripted natural recording in a textbook for which there are the following components: pictures, introductory text, the recording itself, questions to answer while listening, and (probably in the back of the book) an answer key plus an orthographic transcript of the recording.

The acoustic blur (Brown, 1990) – a politer term than 'mess' from – of everyday spontaneous speech <u>only occurs on the recording</u>: one cannot tell from the other components, exactly which words will be the most blurred, or which the least; one cannot be sure how the speaker will pattern their stream of speech in relation to syntactic categories. One certainly cannot tell this from the orthographic transcript: although predictions can be made, they are made on the basis of rather primitive notions of what the speaker would be likely to 'stress' in a particular context. Thus given a knowledge of the context (a lecturer deferring detailed comment on a particular topic until later) and the 'rules' of stress assignment from traditional phonology, one would predict Example 1 to go as follows:

*Example 1a*
THIS is ONE i'm GOing to be LOOking at in SLIGHTly MORE DEtail in fact

instead of (cf. Table 3)

*Example 1*
this is ONE i'm going to be looking at in slightly more DEtail in fact

The essential characteristics of everyday spontaneous speech therefore reside (in pedagogic situations) in the acoustic blur <u>of the recording</u>: they do not reside in the transcript (even where this contains every 'um' and 'ah'), they do not reside in the questions and answers, nor in the contextualising matter, however beautifully they are designed and produced. The accompanying components perform many useful functions (not least, motivating students) but we allow ourselves to be seduced by them – we are tempted by them to spend insufficient time with the recording itself (cf. Cauldwell, 1998; 2000).

Research by Field (1998a, 1998b) has shown that the least time in a listening lesson is spent in the post-listening feedback phase: large amounts of time are spent three other phases – contextualising, 'while listening', and follow-up activities (e.g. discussion, writing).

It is a pedagogic imperative is that we adjust our listening methodology so that we <u>spend more time with the recording in the post-listening phase</u>. But doing what exactly? In the following section, I propose a simple technique for this phase in listening which relates listening <u>activities</u> more closely to listening <u>goals</u>.

**7. The word-crusher**

It is necessary to get the students to mimic, or recreate, the features of the stream of speech in their own pronunciation. They should attempt to mouth the features of the acoustic blur that they can hear, that they have noticed, or had pointed out to them. They should, in other words, relish the acoustic blur of speech with their own physical mechanisms. They should do this not with the goal of being able to mimic the acoustic blur in their own speech (though this could be a goal if they so wish) but with the goal of becoming *familiar and comfortable* with the acoustic blur of normal everyday spontaneous speech of other people.

In other papers (Cauldwell 1996; forthcoming) I have described how electronic technology can help teachers and students identify and focus on the fastest (therefore potentially the most problematic in terms of listening) meaning-bearing stretches of speech. These can be made the target of listening comprehension questions.

There is not space to reproduce all this description in this paper, so let's assume that Example 1 is just such a target, and that it becomes the focus of work in the post-listening phase of the lesson.

> *Example 1*
> this is ONE i'm going to be looking at in slightly more DEtail in fact

The teacher should demonstrate, with both his|her own voice, and with the use of repeated playback of the recording, the relationship between the citation form version and the original version. The teacher's role is to demonstrate intermediate stages of blurring between the extremes of slow speech and the normal speech, by gradually increasing speed, removing pausing, increasing the extent of the deletion and linking of sounds. The role of the original is to provide continual reminders of the natural features of the stream of speech (variations in speed, pitch height, etc.) – which are also worth relishing. It is also vitally important that <u>the students mimic </u>this process of incremental blurring and speeding up – it is of course their abilities in listening that we are trying to improve.

An important point to realise here relates to the earlier discussion of goals. The goal is related to the <u>improvement of listening skills</u>, the activity – mouthing, relishing the words – is a pronunciation activity which is serving the goal of listening.

The criticism might come: 'There is no phonology here: this is an act of mimicry of unique instances of phonetic phenomena'. The argument might continue that it is not worth doing because it teaches nothing which can be

generally applied to other extracts, other recordings, in other listening situations.

Not so. What is being practised is the partial dissolving – blurring – of words into units of speech (tone-units) which themselves have a structure. Example 1 is an instance of this structure, which is the double-prominence tone unit. (Brazil, 1997). This can be represented as a five-part structure as in Table 6:

*Table 6 The structure of a double-prominence tone-unit*

| 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|
| this is | ONE | i'm going to be looking at in slightly more | <u>DE</u> | tail in fact |

The structure of a double-prominence tone-unit is in five parts, labelled 1 to 5 in the first row. The two prominent segments are in columns 2 and 4 (shaded); the non-prominent segments are in 1, 3 & 5 (unshaded).

The double-prominence tone-unit, in having the potentially large middle non-prominent segment, is a useful tool in listening pedagogy. – when a number of words are spoken non-prominently, at speed, between two prominences (as in Example 1) they are likely to be subjected to the extremes of blurring that are normal and common in everyday spontaneous speech, and which present Ying with her difficulties. We can therefore think of the double prominence tone unit as a 'word-crusher'.

The word-crusher can be used observationally, in the design of listening comprehension activities to identify likely sources of difficulty, it can – indeed should – be used to help students become familiar and comfortable with the acoustic blur of speech. I have described above how it can be used in conjunction with recordings of everyday spontaneous speech. It can also be used without recourse to recordings, it is a tool which can be used by both teacher and students.

Here's how. You can take any utterance that has occurred in a recording, or which occurs in the textbook, or which you yourself have invented. Example 4 might be just such a sentence.

*Example 4*
He wasn't happy with you.

There are a number of ways in which this sentence might co-occur with, the double-prominence tone-unit, two of which are shown in Table 7.

*Table 7 The structure of a double-prominence tone-unit*

| 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|
| he | WAS | n't | <u>HA</u> | ppy with you |
|  | HE | wasn't happy with | <u>YOU</u> |  |

The second of the two versions has the longer non-prominent segment, and is therefore more useful for pursuing our listening goal: 'making students familiar and comfortable with the acoustic blur of speech'. Notice that there is a triple-prominence version of this sentence (*he WASn't HAppy with YOU*) but this structure is less useful for our goal, as the non-prominent segments are not sufficiently long to be challenging.

Having identified this sentence, and presented it in some form (on paper, blackboard, OHT) The teacher should model ('listen and repeat after me'), and the students mimic the tone-unit at increasing speeds. Both teacher and students should start slow, exaggerating the relationship between the prominent and non-prominent syllables by making the former extra loud, and the latter extra soft:

> *Example 4a*
>
> HE wasn't happy with YOU.

Even at the slower starting speed, however, a deliberate attempt should be made to make an acoustic blur of the pronunciation of the non-prominent segment. In other words, in pursuit of the goal of listening, both the teacher and student should aim for messy inaccuracy, for a dissolving of the words in the non-prominent segments. They should relish the mess they are creating, because in so doing, they are in direct engagement with the essential characteristics of the stream of speech. It is important to avoid, at all costs, introducing an extra prominence into the non-prominent segment.

This tool, can be used by both teachers and students anytime, anywhere, and with any text. It should be used often: students can practise with this tool in class (out loud), or while travelling home (in their heads). Every new vocabulary item should be subjected, by the learner, to at least equal amounts of acoustic blurring as to citation-form practice. Learners could even compare the ways in which they make the acoustic blurs (this might be especially useful in multi-lingual groups).

This type of exercise is similar to types already in existence in textbooks, especially those designed to demonstrate the so-called stress-timed nature of English speech. Example 6 (from Underhill 1994, p. 71) is such a sentence:

> *Example 6*
>
> YOU and then it's        ME and then it's        HIM and then it's        HER

There are, however, important differences in the goals of and assumptions behind the two types of exercise. The stress-timing exercise is aimed at improving pronunciation, and presents speech as more tidy than normally is (English is not stress-timed cf. Cauldwell 2000).

The word-crusher, however, is a tool to help towards the listening goal – its starting assumption is that speech is a patterned acoustic blur, that all words (not just function words) participate in the blurring. Its role is to help teachers

and learners relish the mess that is the reality of everyday spontaneous speech.

## 8. Conclusion

I have focussed on the double prominence tone-unit (the word-crusher) as a tool, but speech does not consist solely of such units. There are single, triple, and incomplete tone-units. The syllabus component of phonology for listening would have to include these, as well as many other features of speech. These include investigation and acclimatisation to a wide variety of accents, and the voice styles associated with different age-groups, the two genders, and differing speech situations.

It is time to engage directly with the features of the stream of speech. To do this we need to spend more time, with the recordings in a listening lesson observing and imitating the features of everyday spontaneous speech. Frequent use of the word-crusher will free learners such as Ying from the dilemma of believing they 'know words' yet not being able to perceive them when they hear them in speech.

## References

Brazil, D. (1994). Pronunciation for advanced learners of English. Cambridge: CUP.

Brazil, D. (1997). *The communicative value of intonation in English*. [Second Edition]. Cambridge: CUP

Brown, G. (1990). *Listening to Spoken English*. [Second Edition]. Harlow, Longman.

Cauldwell, R.T. (1996). Direct encounters with fast speech on CD Audio to teach listening. *System* 24|4, 521-528

Cauldwell, R.T. (1997). Voices in the university. [Cassette and booklet]. Birmingham: The University of Birmingham

Cauldwell, R.T. (1998) Faith, hope, and charity: the vices of listening comprehension. *The Language Teacher July 1998*

Cauldwell, R.T. (2000) Falling out of context. In C. Heffer & H. Sauntson (Eds.) *Words in context: A Tribute to John Sinclair on his retirement.*

Cauldwell, R.T. (2000) Grasping the nettle: The importance of perception work in listening comprehension. [Published on the internet at http:||www.speechinaction.com ]

Cauldwell, R.T. (2001) The Functional irrhythmicality of spontaneous speech: A discourse view of speech rhythms. [Published on the internet at http:||www.speechinaction.com ]

Cauldwell, R.T. (forthcoming). Speech in action: Teaching listening with the help of ICT. In A. Chambers, J. Conacher, and J. Littlemore (Eds.) *Teaching and Learning Languages Effectively with ICT*

Field, J. (1998a). Skills and strategies: towards a new methodology for listening. *ELT Journal* 52(2), 110-118.

Field, J. (1998b). The changing face of listening. *English Teaching Professional* 6, 12-14

Goh, C. (1997). Metacognitive awareness and second language listeners. *ELT Journal 51*(4), 361-369

Lass, R. (1984) *Phonology: An Introduction to Basic Concepts.* Cambridge: Cambridge University Press.

Rost, M. (1990). *Listening in Language Learning*. Harlow: Longman.

Underhill, A. (1994). *Sound foundations.* London: Heineman

White, G. (1987) 'The teaching of listening comprehension to learners of English as a Foreign Language: a survey'. M. Litt. dissertation, University of Edinburgh.