

8 Input and context

So when you are listening to somebody completely, attentively, then you are listening not only to the words, but also to the feeling of what is being conveyed, to the whole of it not part of it.

Jiddu Krishnamurti (1895–1986), Indian theosophist

8.1 Some terms defined

8.1.1 Decoding and meaning building

My aim in the present chapter is to provide a framework for thinking about the processes involved in listening. Much of the discussion will be shaped by the important distinction that was made in Chapters 5 and 7 between the two major operations that make up the skill:

- *decoding*: translating the speech signal into speech sounds, words and clauses, and finally into a literal meaning;
- *meaning building*: adding to the bare meaning provided by decoding and relating it to what has been said before.

We need to understand the relationship between these two operations if we are to achieve a clear picture of listening. The parts they play also have relevance to decisions made by the listening teacher. Which should we give priority to in a process approach? Which of the two – decoding or meaning building – is likely to prove the more critical in assisting an L2 listener at an early stage to crack the code of speech?

In decoding (dealt with in Chapters 9 to 11), the listener has to make sense of the speech signal. The main goal is to identify words. As soon as a native listener has formed a word match, it triggers a rapid and automatic link to the word's meaning. However, that is not the end of decoding, as the listener then has to go on to trace a grammatical pattern in the words that she has assembled.

What I have called 'meaning building' (Chapters 12 and 13) covers two general functions. The listener (a) expands on the meaning of what the speaker says, and (b) adds incoming pieces of information to her overall picture of the talk or conversation. The first function operates at both word and sentence level. The listener has to relate the speaker's words to the context and situation in which they occur. Say the listener hears the word *turn*. It will not open up a single specific meaning but a range of possibilities: one can turn a corner, turn a handle, turn over a page.

A process view of listening

One can even turn pale or turn thirty. The precise sense of *turn* that the speaker intended will only be recognised once the listener has taken full account of the words that surround it. A similar process operates at sentence level. What the listener extracts from an utterance is a bare and literal meaning. It needs to be enriched by using knowledge of the world (*what do I know about this topic?*), and by deciding its relevance to the present situation (*why did the speaker say this at this point?*).

The second function of meaning building entails making decisions about which pieces of information are important and which are not. The listener then uses the relevant ones to construct a record of the whole listening encounter.

The decoding/meaning-building distinction helps us to separate the hard evidence of the speaker's words from the conclusions that the listener forms about them. It makes sense to examine them separately because of the different processes they demand (highly automatic ones in the case of decoding and more rational ones in the case of meaning building). But we should not lose sight of the fact that they are very closely interconnected. A good example has just been given in word meaning. It is retrieved automatically as part of decoding, but the general range of meanings which the listener obtains then has to be refined according to the context in which the word occurs.

8.1.2 'Input' and 'context'

Let us now consider the types of information that supply the material for decoding and meaning building. There are three:

- **input** (also referred to here as the **speech stream** or the **signal**): the sounds reaching the ear of the listener; and the syllables, words and clauses that those sounds represent;
- **linguistic knowledge**: knowledge of the sounds, vocabulary and grammar of the language (including knowledge of word meanings);
- **context**: which includes (a) general knowledge and personal experience which the listener provides; (b) knowledge of what has been said so far in the conversation.

Decoding is closely associated with the first and the second: the listener has to use knowledge of the language, whether L1 or L2, to turn the speech stream into words, phrases and sentences. Meaning building is especially reliant upon the third, though it also draws upon linguistic knowledge in the form of word meaning (see the example above of the word *turn*).

In the rest of this chapter, we will examine the contribution that input (and with it decoding) makes to the message that the listener derives, and

compare it with the contribution made by context (and with it meaning building). The topic has quite important implications for how we view the methodology of the listening lesson. There is general agreement that listeners need to draw upon both input and context. But the heavy emphasis placed by current methodology upon ‘comprehension’ as the target of listening practice seems to have contributed to a perception that using contextual information is more central to successful L2 listening than recognising words and phrases accurately. The impression has perhaps been bolstered by experience of listening in a first language, where our recognition of words is apparently effortless, as compared with the attention which we need to devote to working out the speaker’s intended meaning.

A belief in the power of context appears to be quite deeply established. Language teachers sometimes prove quite resistant to the idea of training learners in speech perception on the grounds that if there are local difficulties in matching sounds to words, the listener’s knowledge of the topic, the world or the prevailing circumstances will surely be sufficient to resolve any ambiguities. Similarly, any review of research into second language listening will show that, until recently, interest in how learners make use of various aspects of context (knowledge of the topic, the world, the conversation so far) greatly outweighed any work that was done on how learners handle the speech signal.

The notion that ‘context saves the day’ will be subjected to quite careful examination here in the light of what we know about first language listening. If it proves to be soundly based, then it will confirm that a process approach to listening should focus first and foremost upon meaning building. If it does not, then we might need to consider giving considerably more time than at present to training the L2 listener to identify sounds and words with confidence.

8.2 Input

Let us first examine what we mean by ‘input’. What reaches our ears is not a string of words or phrases or even a sequence of phonemes. It is a group of **acoustic features**. Clusters of these features occur together, providing evidence of the speech sounds that the speaker has made. We must not think of the words or phonemes of connected speech as transmitted from speaker to listener. It is the listener who has to turn the signal into units of language.

Decoding takes the form of a matching process. On one side of the process is the group of acoustic cues which have reached the listener’s ear; on the other is the listener’s knowledge of the language being used.

A process view of listening

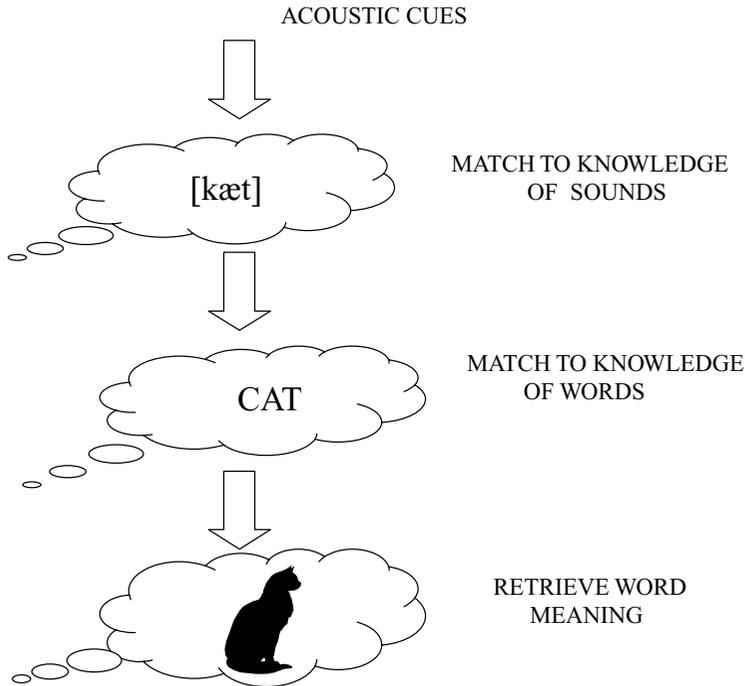


Figure 8.1 Speech perception as a process of matching.

That knowledge is stored long term in the listener's mind and consists of the spoken forms of words and maybe the individual sounds of the language as well. It seems likely that it also includes chunks of language in the form of familiar and recurrent sequences of words (*just about, do you know, should have done, anything else*). The matching process, very simplified, is represented in Figure 8.1. The thought bubbles serve as a reminder that the process is taking place *in the mind of the listener*.

It is important not to lose sight of the concept of matching. It is easy to slip into the assumption that sounds and words are present as independent units in the speech stream. They are not: it is the listener who brings form and meaning to the input by drawing upon her knowledge and experience of the language being used. In the case of the inexperienced L2 listener, the operation is complicated by the listener's limited ability to recognise the sounds of the target language or limited vocabulary against which to make a match. It may also be complicated by the listener's lack of confidence in the matches she makes.

Decoding, of course, extends beyond simple matching. The outcome of the identification process is a string of words, with meanings and intonation attached; but the listener then has to impose a grammatical pattern upon the string. The final product of decoding is a piece of information which is no longer in the form of language but has been turned into an abstract idea.

One way of representing the whole decoding operation is in terms of the kind of sequential process shown in Figure 7.1 (p. 114): a series of stages in which smaller units are progressively built into bigger ones with the help of knowledge of the language. From bundles of acoustic cues in the speech signal, the listener manages to identify phonemes, the sounds of the language. Then the phonemes are built into syllables, the syllables into words, the words into phrases and the phrases into clauses or sentences. Finally, the sentences have to be converted from language into ideas.

This analysis seems plausible at first glance, but there are a number of problems with it. Firstly, it is by no means sure that listeners do employ all these **levels of representation**. For example, the sounds of a language vary greatly according to the syllable in which they appear, so they are a very unreliable unit to make a match to. Commentators have suggested that listeners may not use phonemes at all when analysing the speech signal. They might use syllables instead, or might go straight to a match with a word.

Most importantly, it has been shown that a listener does not wait until the end of a clause or a sentence before deciding what a speaker is saying. Listening seems to be very much an **online** activity, with the listener decoding the sounds of speech at a delay behind the speaker of as little as a quarter of a second. A quarter of a second is about the length of a syllable, so this adds support to the idea that the syllable is an important unit of processing.

If listening is so immediate, then it cannot proceed in neat steps (syllables into words, words into phrases and so on). It seems likely that the listener forms an idea about what the speaker is saying quite early on in the utterance, but constantly revises it as she hears more and more. So decoding is not the simple sequential operation that it is sometimes said to be.

8.3 Context

The final outcome of decoding is no longer in word form but is an abstract idea (sometimes called a **proposition**) which contains the literal

A process view of listening

meaning of what has been decoded. But a proposition does not mark the end of the listening process. Consider the literal meaning that a listener might extract from the sentence *It's going to rain*. We could represent it like this (bearing in mind that it is no longer in the form of words):



That is the literal meaning. But the final message that the listener derives will vary enormously depending upon the situation in which the sentence has occurred. Consider the different responses that would be called for if:

- a. the speaker is a keen gardener and there has been a drought;
- b. the speaker has tickets to watch some open-air tennis;
- c. the speaker and listener are having a picnic;
- d. the speaker points at some dark clouds;
- e. the listener knows that British people often make relatively meaningless statements about the weather;
- f. the comment happens during a conversation about climate change.

Whereas in the case of (a), the listener might respond *Great!*, an appropriate response in the case of (b) would be *Oh dear!*. In both instances, the listener's understanding of the message is enriched by personal knowledge about the speaker. Examples (c) to (f) suggest other possible sources of evidence: (c) draws upon the immediate situation, (d) upon world knowledge that is shared by speaker and listener (dark clouds presage rain), (e) upon cultural knowledge, and (f) upon knowledge of the topic (the likelihood that future climate change will reduce rainfall). All of these tend to be referred to very generally as constituting 'context'. However, the word has become rather a catch-all, and it is greatly preferable to specify exactly what kind of information the listener is bringing to bear.

We might envisage a further situation in which the utterance *It's going to rain* occurs during a conversation about a current water shortage. Here, the comment draws its relevance from what has been said earlier. This, too, is sometimes loosely referred to as 'context', but it is a very different type in that the information comes not from the listener's long-term knowledge but from short-term recall of this particular conversation. It is useful to distinguish this type of information by calling it 'co-text' (Brown and Yule, 1983a: 46) or 'text-so-far'.

The main point illustrated by the examples just discussed is that a listener draws upon multiple sources of evidence, which go beyond the raw information in the speech input and give depth and relevance to

the message that is finally extracted. We have treated this evidence as instrumental in meaning building. But it does, in fact, contribute to decoding as well. A listener might need to draw upon co-text in order to select the appropriate sense for a word: compare *right* in *Turn right at the traffic lights* with *right* in *She got three questions right*. In addition, co-text and context assist in correcting possible misperceptions. A listener who identifies the word *knickers* during a radio discussion about the Church might conclude that the word she heard was actually *vicars*. A listener who hears the word *dessert* in a talk about camels might assume that the speaker mispronounced *desert*.

So we should note that context and co-text fulfil two distinct functions, which are sometimes confused by commentators on L2 listening:

- They enrich the raw meaning of the utterance and make it relevant to the current situation.
- They provide extra evidence that assists the decoding process.

8.4 Context and the less skilled listener

There has been disagreement about the extent to which language learners with limited vocabulary and grammar are able to make use of context and co-text. One well-established view (Oakeshott-Taylor, 1977; Osada, 2001) has it that their attention is so focused upon the effort of decoding unfamiliar sounds and words that they have little left to spare for wider considerations. Evidence from language testing (Hansen and Jensen, 1994: 265) shows that lower-level learners report much less successfully on the global meaning of a listening text than do more advanced ones.

However, contradictory findings have been reported in relation to unskilled L1 *readers*. They have been shown to make quite heavy use of context because they find it easier than decoding what is on the page. In a much-quoted experiment, Perfetti and Roth (1981) showed that it is weak readers, not good ones, whose ability to recognise words is most assisted by a clear context. There are similar findings within L2 listening. The most extensive investigation of the 'input/context' issue to date was undertaken by Tsui and Fullilove (1998), who analysed answers given by 20,000 Hong Kong examination candidates to different types of listening question. The difference between successful and less successful listeners was found to be that the successful ones were much better at answering test items where the candidates could not fall back on world knowledge for support. They could get by without it because their decoding skills were so good. This suggested that it is the *unskilled* listener who is more dependent upon context.

A process view of listening

So there is evidence supporting the view that less experienced second language listeners rely heavily upon contextual and co-textual information. Yet a contrary view persists that the need to focus upon decoding the input distracts these individuals from using context/co-text to build larger patterns of meaning.

There is an easy way to resolve this apparent contradiction. The truth is that both unskilled and skilled listeners make use of context, but that they do so for different purposes.

- Skilled listeners and readers (whether in L1 or L2) make use of context to *enrich their understanding of the message*. Less skilled listeners are not always able to achieve this wider understanding because their attention is so heavily focused upon details of the signal.
- Less skilled listeners and readers (whether in L1 or L2) make greater use of context and co-text to *compensate for parts of the message that they have not understood*. In second language listening, the failure might be due to problems of decoding, problems of word and grammar knowledge or problems in recognising the relationships that link ideas.

8.5 'Bottom-up' and 'top-down'

It may seem strange that we have so far avoided the terms '**bottom-up**' and '**top-down**', which are very often used in discussions of second language listening. This is intentional because they have come to be used rather confusingly. Strictly speaking, the terms refer to *directions of processing*: distinguishing between, on the one hand, building phonemes into words and words into phrases and, on the other, using context and co-text to help identify words that are unclear (Field, 1999). It is in this sense that they will be used in this book: 'bottom-up' referring to building small units into larger and 'top down' to the influence of larger units when identifying smaller ones. Even here, as we have seen, there are complications. Because listening is online, we cannot assume that there is an easy 'bottom-up' progression from sounds to syllables to words to phrases. And the 'top-down' uses of context can serve two very different purposes: to compensate for gaps in understanding or to enrich a fully decoded message.

It is when we move on to consider how a listener manages to combine information from the input with information from context that the terms 'bottom-up' and 'top-down' cause particular problems. They are sometimes associated loosely with 'decoding' and 'meaning building'; they are also sometimes treated as if they were synonyms of 'input' and 'context'. Commentators might contrast a 'bottom-up view' of

processing (meaning a view that listening relies upon input) with a ‘top-down view’ (a view that listening relies upon context).

This is unhelpful for two reasons. Firstly, when terms such as ‘view’ or ‘model’ are used, the impression is given that we are dealing with contrasting theories of listening. But it should be obvious that a listener has need of both. We could not identify the topic of an utterance without some minimal decoding, and we could not appreciate the relevance of the utterance without some minimal use of context and co-text. For this reason, psychological descriptions of listening assume that the processes involved are highly interdependent.

Secondly, if we think of ‘bottom-up’ and ‘top-down’ as representing types of processing (one small-to-large, one large-to-small), it is misleading to imply that top-down processes involve only context and co-text. Here is an example. Let us suppose that you hear somebody say the word *veshtables* [ˈveʃtəblz] and succeed in matching it to a word in your vocabulary. How did you do it? If you had proceeded in a bottom-up way, you would have been fazed by the presence of the unexpected sound [ʃ]. To resolve the issue, you might have drawn upon co-text, for example: *cabbages, carrots and other . . .* Or you might have drawn upon context (the fact that you are in a greengrocer’s or ordering a meal in a restaurant). But you might equally well have drawn upon your knowledge of a familiar chunk of language (*fruit and vegetables*) or just of the word *vegetables*. You might even have drawn upon the knowledge that the syllable *vesh* is not used in English. The point at issue is that all of these are examples of top-down processing in that all of them involve using larger units to resolve a decoding problem that concerns a smaller one in the form of the unorthodox sound [ʃ]. As well as context and co-text, the listener has other (linguistic) means available for resolving decoding problems in a top-down way.

This account of the listening process is important to bear in mind. It illustrates how we deal with the issues that arise when we decode the input of a non-native speaker or of somebody with an unfamiliar accent. Above all, it illustrates the way in which a second language listener behaves when confronted with input that is difficult to follow, reminding us that multiple sources of top-down information (not all of them contextual) are available to assist her.

So, modern interactive accounts of decoding see it as a kind of negotiation, where the listener weighs a number of pieces of evidence in order to decide what is in the input. The issues that concern researchers today are not whether listening is ‘bottom-up’ or ‘top-down’ – since it is clearly both – but which source of knowledge an unskilled listener is most reliant upon. Which source prevails if information from ‘bottom-up’ processing conflicts with information from ‘top-down’?

8.6 Compensatory processing

A tentative answer to these last questions will now be suggested. Let us consider the case of a native listener. For such a listener, input is decoded in a way that is highly automatic. Decoding is fast and accurate and makes few demands upon the listener's mind. By comparison, having to draw upon the evidence provided by context or co-text is much slower. So it makes sense in L1 to rely upon input, and to use other sources of information as a fallback in cases of ambiguity, inconsistency or lack of clarity in the signal.

Now compare the case of somebody listening in a foreign language, especially a less experienced listener. She is likely quite often to feel the need to make checks upon the accuracy of her decoding. We can represent the situation in terms of her *level of confidence*. If a listener feels that she (a) has succeeded in decoding a sufficiently large proportion of the input and (b) is confident about the accuracy of what has been decoded, then there will be less need to rely compensatorily upon information provided by context and co-text. On the other hand, if the listener feels unsure of what she has made of the input, then she will rely more heavily on external information.

This perspective owes much to Stanovich's (1980) **Interactive Compensatory Hypothesis**, which aims to account for the way in which weak L1 readers handle text. The trade-off is shown in Figure 8.2. Where confidence in the input is high, the role of compensatory 'top-down' information (the dark part of the column) is relatively small. But, where confidence is low, 'top-down' information makes a much greater contribution.

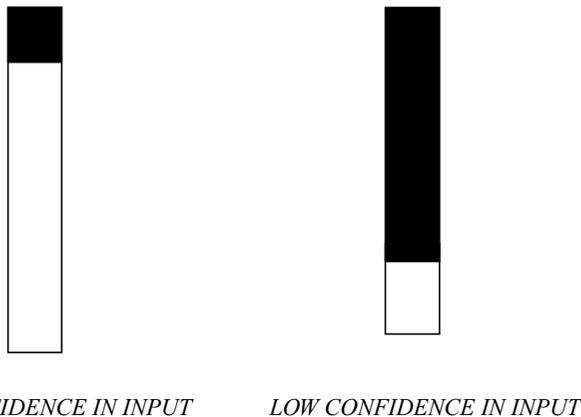


Figure 8.2 Stanovich's Interactive Compensatory Hypothesis.

Although we have assumed that first language listeners can generally rely upon their decoding skills, the model also accounts for situations in L1 listening. Compare the kind of listening that takes place in a lecture theatre with the kind that takes place in a crowded pub. The absence of noise in the first situation allows listeners to feel a high level of confidence in their ability to decode what is said. Though external evidence will be used to enrich meaning, there will be little need to rely compensatorily upon it to supply missing words – except perhaps when an unknown piece of terminology occurs. By contrast, the high level of noise in the pub situation means that listeners cannot trust the input so completely; they need to draw much more heavily upon context and co-text.

This suggests that there is no constant relationship between input and context but that the way in which they influence each other varies from one situation to another. It also suggests that listeners approaching a second language are already well practised in the compensatory process. Of course, the ‘noise’ in an L2 context is rather different: it is created by sections of the text that the listener cannot decode because of problems of recognition or lack of linguistic knowledge. But the process still involves striking a fine balance between confidence in the input and the need to draw upon external information. If the use of context to plug gaps in understanding is such a feature of second language listening, it must be because it is already a familiar experience in L1 situations.

8.7 The importance of decoding

Listening teachers might interpret the Stanovich model in one of two ways. They might conclude that L2 listeners are not at all dependent upon their decoding skills because they have another resource to hand in the availability of contextual and co-textual knowledge. Or they might conclude that L2 listeners need to give priority to developing their decoding skills so as to reduce their dependence upon outside information.

As noted at the outset, there has been a tendency in recent years to downgrade the part played by input in L2 listening, and to assume that, if a listener has perceptual problems, they can readily be resolved by drawing upon contextual evidence. From a pedagogical perspective, the tendency has been reinforced by the high importance which the comprehension approach gives to the outcomes of listening rather than to the processes which give rise to them. As a result, much recent research and comment on L2 listening has focused upon the contribution made by world knowledge (e.g. Long, 1989) or topic knowledge (Long, 1990).

However, the evidence from Tsui and Fullilove (1998) cited above paints a very different picture. It suggests that what differentiates skilled

A process view of listening

from unskilled listeners is that the latter have to rely more heavily upon contextual and co-textual evidence to supplement their decoding. They need to compensate for gaps in their understanding where decoding has failed and for their lack of confidence in the accuracy of the word matches they have made. If a teacher asks what characterises a skilled listener, the answer would seem to be *accurate and automatic decoding*, not the ability to make use of context.

Alongside the research evidence, there are other objections to the notion of 'context saves the day'. We need to look closely at what commentators mean when they mention 'context' in this way. There is no doubt that world knowledge and knowledge of topic, situation and speaker provide useful support to a decoding process that is running into trouble. But much of the discussion of compensatory processing in L2 assumes that listeners draw heavily upon co-text, i.e. upon their recall of what the conversation has been about so far. This leads to a very circular argument. The recall can only be of value *if the listener was able to decode accurately what the speaker said previously*. Far from offering an alternative to poor decoding, co-text depends entirely for its reliability upon whether the listener's decoding skills are adequate or not!

To make this clearer, consider what happens when decoding fails the L2 listener. At one extreme, the listener might have such a lack of decoded information that it is impossible to construct any clear notion of what the utterance is about. There would then be no co-text available to help resolve later problems of decoding. Alternatively, and perhaps worse, inaccurate decoding at an early stage of listening might have a 'knock-on' effect as far as later understanding was concerned. A listener who mis-segmented a sequence like *I went to assist her* might come to assume that the remainder of the text would be about a female sibling or a nun (*I went to a sister*).

A further reason for stressing the importance of input draws upon what we know about memory. There are two major components in our memory store: one that holds long-term knowledge and one that holds and operates upon short-term information. What is clear about the latter, termed **working memory**, is that it is very limited in what it can contain. This has important implications for listening. If a listener is able to decode the input effortlessly, the result is to leave a great deal of working memory free for thinking about larger issues such as the overall meaning of the text. If (as with a novice L2 listener) decoding is uncertain and makes heavy demands upon attention, then it leaves no memory resources spare for interpreting what has been heard or carrying forward a recall of what was said earlier.

The wrong message thus seems to have been getting through to practitioners. If early learners emerge as too dependent upon bottom-up

processing, teachers should not conclude that the solution is to switch the focus mainly or entirely to top-down alternatives. Instead, they need to direct their efforts towards ensuring that learners' decoding becomes more expert and thus demands less effort.¹ In this way, they ensure that there are fewer demands upon the learners' attention, enabling them to devote some of that attention to wider meaning rather than to compensating for understanding.

Quite apart from the freeing of working memory, there are also benefits so far as confidence is concerned. It is no coincidence that, when questioned on their concerns about listening, language learners frequently cite decoding difficulties. Evidence that phoneme and word recognition are a major source of concern for low-level L2 listeners comes from a study by Goh (2000). Of ten problems reported by second-language listeners in interviews, five were connected with perceptual processing. Low-level learners were found to have markedly more difficulties of this type than more advanced ones. Here is an explanation for the listening anxiety that often arises as a result of the transitory nature of the speech signal.

8.8 Some implications for the teacher

The main argument of this chapter has been that both input and context play an important part in second language listening, regardless of the level of the learner. The important difference is that a novice listener is likely to make use of context to compensate for inadequate decoding skills, while a more experienced one employs context to enrich understanding of a message. The argument has been put that expertise in listening is assisted by the ability to decode connected speech in a way that is automatic and accurate. This (a) gives the listener confidence in her ability to shape sounds into words and (b) releases attention that can be switched from basic processing to deeper issues of meaning.

There are clear implications for the listening teacher wishing to adopt a process approach to the skill. Firstly, it is useful in the early stages of listening instruction to dedicate time and effort to building up the learners' decoding processes. The goal is to ensure that their listening comes to approximate more closely that of a native listener who enjoys the benefit of being able to decode automatically. Secondly, the way in which 'context' is treated by the teacher needs to vary according to the level of the learners. At lower levels of English, the emphasis will be on

¹ Though this should by no means rule out a parallel 'top-down' approach by the teacher, which equips learners with strategies that enable them to make the most of the little they succeed in decoding.

A process view of listening

encouraging learners to make use of world knowledge, topic knowledge, etc. to compensate for gaps in understanding. At higher levels, learners can be encouraged to expand on what they hear by relating it to background knowledge just as they would in L1.

This raises issues about the convention of a pre-listening phase in a lesson, when learners reflect on the topic of the recording. At either level, it makes perfect sense; but the purposes for which it is employed should differ significantly. For the novice, it provides a framework which may assist in interpreting some imperfectly grasped fragments of speech; it also alerts the learner to certain words associated with the topic which may occur in the listening passage. In other words, it assists in decoding. For the more advanced listener, it creates a mindset which brings forward knowledge of the topic available from L1 and possibly raises questions and expectations about the content to be heard. In other words, it assists in building meaning.

When thinking about the less skilled listener, it is worth bearing in mind the distinction made between context (= external knowledge) and co-text (= 'text-so-far'). Teachers need to recognise that the latter only slowly becomes a resource on which the L2 listener can depend, since its reliability depends upon how much of the input the listener is able to decode. This means that there will be a progression from a situation where the learner supports decoding almost entirely with outside knowledge to one where co-textual information can increasingly be relied upon.

For the process teacher, then, the relationship between input and context/co-text is not a constant one, but one that evolves. This has consequences for the types of task that are set within the listening lesson, and in particular for the ways in which the teacher encourages the learner to make sense of what has been heard. Early work on compensatory strategies should first draw on *contextual* cues and only gradually involve the use of *co-textual* ones.

We should also consider moderating the meaning-building demands imposed upon learners in the early stages of listening. Here, two considerations come together: the limited attention that the novice listener can allocate to wider meaning and the unreliability of the listener's understanding of 'co-text'. It is wise to restrict the number of questions which require the listener to report on the meaning of the recording as a whole, or to interpret or evaluate it. It is also wise to employ texts which are narrative or instructional, and where one assertion leads to another chronologically ('then') or in an additive way ('and'). This spares listeners from having to build complicated meaning relationships at the same time as dealing with the problems posed by inexperienced decoding.

We have now identified some very general guidelines for a developmental element in a process-based programme:

- early priority accorded to decoding skills;
- early attention to the development of strategies that make use of context;
- a carefully moderated increase in reliance upon co-text;
- minimal use in the early stages of questions relating to wider meaning or interpretation of what has been heard;
- a shift over time in the purposes for which we use the pre-listening stage;
- an early preference for texts where the meaning construction is additive or chronological rather than entailing more complex meaning structures.

Further reading

- Field, J. (2004) 'An insight into listeners' problems: too much bottom-up or too much top-down?' *System*, 32: 363–77.
- Field, J. (1999) 'Key concepts in ELT: bottom up and top down'. *ELT Journal*, 53/4: 338–9.
- Lynch, T. (2006) 'Academic listening: marrying top and bottom'. In A. Martinez-Flor and E. Usó-Juan (eds.), *Current Trends in Learning and Teaching the Four Skills within a Communicative Framework*. Amsterdam: Mouton, pp. 99–101.
- Tsui, A. and Fullilove, J. (1998) 'Bottom-up or top-down processing as a discriminator of L2 listening performance'. *Applied Linguistics*, 19: 432–51.