26

Ear-Training: developing analytic listening skills for English

Michael Ashby

Department of Phonetics and Linguistics
University College London
Gower Street
London WC1E 6BT

1 It is important to make a clear distinction between the perception of speech and the production of speech. Most people think of phonetics as being concerned with training the correct production of speech But equally important is that side of phonetics which is concerned with hearing differences in sound, and using phonetic symbols and description to record the things we perceive. In the child's learning of its first language, perception ability develops way ahead of production ability. Parents talk to their children from the moment they are born, but do not expect very much back from the child for the first year or two. This is very different from what happens when students begin on a second language in school or college. Often, they will be expected to speak, or even to construct sentences, from the beginning. In my view this is a serious mistake, because it means that incorrect pronunciation habits may be learned right from the beginning. We know from measurements and experiments that people differ in the precise ways they use their tongues, lips, jaws and vocal folds in forming particular sounds. But providing the output sounds right, none of this variation matters. There is plenty of evidence, too, that talkers listen to themselves to monitor their own output. Interfering with auditory feedback can have a big effect on the speed, loudness, and fluency of a talker's output. What this means is that a talker is not just going through a set of speech-production movements, but instead trying to match a series of sound targets.

2 Ear-training - that is, systematic and regular classes devoted to identifying and discriminating different sounds - has always been an important part of the British tradition of phonetic training. In a typical ear-training

session, the teacher dictates phrases or words which the students must take down in phonetic symbols. Each item will be repeated six or eight times until the students have completed their transcriptions. Recorded materials can be used, but a live speaker gives the opportunity for the students to see the speech movements as well as hear them, and can also interact with the students if they want to ask questions. At the end of an ear-training course, assessment is by a dictated test. Students get a numerical score which is a measure of how many sounds they have heard correctly. Ear-training is a part of all phonetics programmes at UCL. Daniel Jones, who founded the UCL department, even used to do regular ear-training with his staff.

3 One of the basic requirements for a learner of a language is to be able to make all the phonemic sound discriminations needed for the language. Many language courses and tapes use exercises on minimal pairs - pairs of word like beat/bit or feel/peel. If learners have more time to devote to ear training, they will need to learn to transcribe in phonetic symbols so that they can give answers in more advanced ear-training tasks. Students transcribe single words, then whole phrases and dictated texts, becoming more skilled at remembering how longer and longer sequences were spoken. Made up words (sometimes called "nonsense" words) are often used for part of the training, because the learner cannot use knowledge of vocabulary or context to help in guessing. Connected speech, stress, rhythm and intonation all form part of more advanced ear-training exercises. For instance, in teaching English intonation, a good way to begin is with ear-training on discriminating simple pitch movements: fall vs rise vs fall-rise.

4 Material used for ear-training must be carefully constructed and graded according to level of ability. Minimal pairs should be words that learners know and understand. Invented materials need to focus on real difficulties for the learner and should not contain too many distractions. For connected speech, whether short phrases or whole texts, it is better to use real spoken material as the basis. There are quite a lot of differences between written (literary) language and real everyday speech. Invented examples, or examples taken from books, often sound unconvincing. I collect conversations which I overhear around me

and use these as the basis for much of my material.

5 Students are often afraid of ear-training, or claim to be "tone-deaf"(strictly, this means unable to hear pitch differences properly). We have to do everything we can to overcome their fears. Often teachers too are afraid, even though they are qualified, and for this reason teachers may avoid doing ear-training. We have to find ways of teaching where mistakes are not felt as embarrassing. Where possible, we use a computer-based speech analysis system in the ear-training class to give objective measurement of what the teacher has dictated.(See our annual report at http://www.phon.ucl.ac.uk/home/PUB/ar/teaching). For some students, working with multimedia presentation on a personal computer may be a good way around this. Where certain sound distinctions are found to be particularly difficult, they must not be avoided, but instead clearly identified and tackled head-on. For students whose main concern is not phonetics but a desire to learn English, we must show the relevance of ear-training with convincing examples of how their comprehension could improve. Ear-training does bring real improvements in our students' auditory skills, even in a short intensive course of two weeks. After a year-long postgraduate training, our Masters students are able to transcribe the most difficult material I can invent or pronounce, covering English and the whole of the International Phonetic Alphabet, sometimes achieving scores of 90% or more correct.

6 With my colleagues at UC I have recently been working on new methods of doing ear-training. We have developed a technique which we term Analytic Listening, which permits the students to give their answers on multiple-choice forms, rather than having to transcribe in full. Each question focuses on one important sound distinction at a time. For example, a question may take the form "You will hear items which are of the form VCV(vowel-consonant-vowel). Indicate whether the consonant is voiced or voiceless." This is then followed by a series of dictated or recorded items to which the listener must respond by choosing from a fixed range of answers. Our paper on the Analytic Listening technique can be read on the web at http://www.phon.ucl.ac.uk/home/shl9/ashby/ma.htm. This approach to ear-training has a number of advantages. First, no phonetic symbols need be used. Secondly, analytic listening exercises can be

graded very much more quickly than conventional transcriptions, and the choices which the student has made give us a profile of each student's auditory skills. Thirdly, this method of presentation is very suitable for interactive presentation on a personal computer. In the method we are developing, material will be presented on web-pages accompanied by sound. The students' responses will be made on a web-page form, and the system will grade performance and give automatic feedback. This system can be sampled on the web at http://holtz.phon.ucl.ac.uk/wbt.

Handout for plenary keynote speech -one page giving main headings:

Ear-Training: developing analytic listening skills for English

Michael Ashby

Department of Phonetics and Linguistics
University College London
Gower Street
London WC1E 6BT

1 Perception of speech and production of speech
- speech as movements and speech as sound
- how a child learns its first language
- what happens when students learn a second language
- what experiments tell us
- speech in linguistic theory

2 Ear training as part of phonetic training
- how ear-training is done
- live, or recordings ?
- how ear-training is assessed and examined

the Daniel Jones/UCL tradition

3 Types of ear-training
- structured questions: ship or sheep ?
- learning to transcribe in symbols
- why nonsense words (made up words) are useful
- dealing with connected speech and intonation

4 Ear-training materials
- inventing material at different levels of difficulty
- using real speech examples

5 How students react to ear training
- overcoming inhibitions
- facing difficulties head-on
- showing the relevance of ear training
- how good can our students get?

6 New ideas: Analytic Listening
- structured questions focus on important distinctions
- quick assessment and feedback
- good for multimedia and interactive presentation

Department of Phonetics and Linguistics

ANALYTIC LISTENING: A NEW APPROACH TO EAR-TRAINING

Michael ASHBY, John MAIDMENT and Evelyn ABBERTON

Abstract

In this paper we outline a new method for training and assessing auditory skills in students of phonetics, which has now been introduced into all the undergraduate programmes involving phonetics at UCL. Conventional ear-training uses dictated or recorded material of approximately phrase length, and requires students to transcribe their answers in full starting from a blank sheet of paper. The new technique, called Analytic Listening, involves multiple-choice testing. Each question focuses the students' attention upon one relevant phonetic parameter at a time and their perception is then tested with a short battery of varied stimuli. They are not asked to transcribe what they hear, but to choose from a fixed range of responses. Analytic Listening offers rapid and simple marking and feedback, leads to statistically reliable profiles of students' auditory skills, and reduces the emphasis on symbols and transcription. Because it focuses upon and tests an explicit range of relevant phonetic parameters, it permits auditory-skill benchmarks to be incorporated into syllabus design. At the same time, it encourages a realistic view of the power and limitations of humanly-made phonetic judgements alongside instrumental measurements of speech, which increasingly form a part of classroom training.

Traditional ear-training and its limitations

Traditionally, the training and testing of auditory skills has been done by Ear-Training Dictation. The teacher repeats aloud certain speech sequences (alternatively, of course, recorded materials may be used). The students write down what they hear in phonetic symbols. The students' versions are compared with what was intended, and a score is worked out which is essentially a measure of how far what the student has written

corresponds with the teacher's original version. Two types of material have generally been used. (1) connected speech in a specified accent of the principal language being studied. For this a systematic (more-or-less phonemic) transcription is usually required; (2) invented or `nonsense' words, either drawing sounds from the phonological system of a specific language or else from a more general repertoire. For items drawing on a wide range of sounds, an impressionistic (non-systematic) transcription is appropriate. The obvious advantage of nonsense materials is to test sound identification and discrimination without assistance from context.

The marking of dictation scripts is labour-intensive. A single script can easily contain scores - even hundreds - of "errors". Each error must be found, and its severity ranked. Commonly, elaborate schemes are worked out for costing anticipated types of error. But answers can depart from what is intended in numerous unforseen ways, creating difficulties for those marking and moderating the test. Similarly, feedback to students needs to be tailored to their individual patterns of error. Detailed comments must be given in one-to-one tuition, or written individually on submitted work. To a limited extent, students may correct their own (or each others') work against a correct answer, but this encourages the belief that there is one correct answer. However, since transcription is to some extent subjective, more than one answer may in fact have some merit, once the most elementary level of transcription has been passed. For instance, transcribers can differ - quite reasonably - over such apparently basic issues as how many segments a given sequence seems to contain.

Analytic Listening

We have attempted to devise an approach to auditory training and assessment which will both reduce the labour involved in marking and assessment and also place more emphasis on critical, analytic listening rather than the process of getting down the "right" symbols. We call the technique Analytic Listening. It seems to go a considerable way towards overcoming both the theoretical and practical difficulties associated with the traditional approach, and is now running as part of training and assessment in all UCL undergraduate programmes involving basic practical phonetics.

Analytic listening is done in ear-training class in the regular way, and assessed as before with a dictation examination. What is different is that the students do not sit down with a

blank sheet of paper, but with a multiple-choice form. Each question in the multiple choice test explicitly focuses attention upon some relevant property or parameter, and the students' ability to detect that property, or make discriminations along the parameter, is then tested several times over with changing material.

Basic form of an Analytic Listening question

A short sample of Analytic Listening material accompanies this paper. A teaching or examination session would involve a number of questions (perhaps 10 or 20) taking two to three minutes each in a live presentation. Each question deals with a different phonetic point. In Q1 of the sample given here, the focus is upon the presence or absence of a glottal stop at the beginning of a short sequence - whether or not there is a "hard attack". (This is, in fact, a realistic example of the sort of phonetic judgement that has to be made routinely by a Speech and Language Therapist in a voice clinic). Each question in turn comprises a number of items - we have settled on five in the exercises we have designed so far. The five items within the question all test the same perceptual judgement, but with a different dictated sequence. Each item is repeated typically three times before the students mark their responses and move on to the next. The dictated material may be of various types - sometimes English-like, sometimes nonsense - and can be constructed so as to contain potential distractors (e.g. in this case a glottal stop at the end instead of the beginning). The similarity of the technique with forced-choice perception testing will be obvious. Of course, the listener may sometimes guess, but the overall scoring is arranged to take account of this (as with any Multiple-Choice test). Each question alone has some statistical utility anyway: a listener who makes the right choice five times out of five has only a three percent chance of doing this by guessing. When all five items have been dealt with, the students move on to the next question, which will generally focus upon a different phonetic distinction.

For training, rather than testing, the teacher can of course use a question of this sort to introduce the discrimination in question. One can begin by giving feedback after each item and one can repeat contrasting items in sequence. In a way this merely gives some structure to the kind of improvised repetition and comparison that has always formed part of every good ear-training class in response to students' questions, but it is worth emphasising some differences from the traditional method which this simple example illustrates. First a single clearly-defined phonetic difference is at issue and the

expectations for the task clearly set out. Second, the discrimination is tested both thoroughly and reliably (it is unclear what statistical reliability - if any - may attach to results obtained in conventional ear-training dictation). Thirdly, although a ready-made answer form is required (either on paper, or computer-presented), this may be re-used indefinitely with changing dictation material, leading to long-term economies in preparation time for teachers.

Cueing the listeners in

We believe that a very important feature of Analytic Listening is the way in which the question first cues the listener in to the point under consideration. When students are asked to transcribe items in full in the traditional manner, they must generally do so against a background of essentially arbitrary conventions, established for the course or even by individual teachers. For instance, it may be established in advance that some property (aspiration, nasality, length, etc) is/is not required to be marked, that all vowel qualities are to be shown as Cardinal Vowels but without diacritics, and so on. Without such conventions, students' versions would differ unpredictably and chaotically so that marking and feedback would be virtually impossible. Learning to do well in an ear-training class is thus partly a matter of acquiring and falling-in with conventions. But because Analytic Listening directs attention to a specific phonetic property in each question, and gives a direct indication of the type of answer expected, these arbitrary conventions are largely eliminated. Q2 gives another illustration of how the cueing-in process works. The wording of the question tells students to make the assumption that these items are VCV, then make a judgement about the C part. This is a very different matter from simply expecting the students to assume that the sequences are VCV. The orientation suggested in the rubric is to be taken a sort of working hypothesis. It permits the use of quite ambiguous sequences as dictation items, for instance a sequence such as [awa], (where it is questionable whether the portion represented [w] is a "consonant") or items containing affricates (where the number of segments is problematic). Further, the question does not ask "is the consonant a voiced fricative?" but rather "does it have voiced friction?" That friction might indeed be within a fricative proper, but it might be in an affricate, or might be added to a type generally familiar only as a sonorant, such as a nasal. We are thus able to elicit fine phonetic judgements about specific parts of sequences without reinforcing simple assumptions about the nature of those sequences.

Applicability of the technique

The two examples given so far have been based on single segments of consonantal type, but of course the technique is applicable to a wide range of segmental and non-segmental effects. Q3 is an illustration of the technique applied to vocalic sounds.. We can ask specifically about frontness, rounding, and so on, in particular vowels, or ask questions which require the comparison of two vowel tokens (same/different, longer/shorter, etc). Equally, stress or pitch may be what is focused on; for instance, listeners may be asked which of two identified syllables within a phrase has the higher pitch, where the highest pitch within a sequence is located, and so on. The technique thus offers a way to ask quite specific questions about prosodic properties, without requiring that students approach the material with a ready-made ("phonological") framework for the analysis of intonation.

Since in general there are no phonetic symbols used in the questions or required in the answers, we can employ sound-types which the students have not explicitly met before, or for which there are no generally agreed symbols. Learning of symbols is not infrequently seen as a laborious and unattractive task. Some speech professionals (eg Speech and Language Therapists) who have qualified in phonetics and still require to use phonetics tell us that the forgetting of symbols, or the failure to keep up with changes in symbolisation, are major embarrassments to them. Analytic Listening will therefore have an immediate application in refresher courses, and should help to get across the message that forgetting symbols is not the same as forgetting one's phonetic skills.

Eventually, of course, most of our students do learn a wide range of phonetic symbols, and as they progress students begin to include symbols alongside their answer choices and to attempt here and there complete transcription of the dictated items. This in our view puts symbols in their rightful place: very much secondary to the essential phonetic judgements. Our hope is that we can wean students from the idea that there is a fixed repertoire of sounds, and that everything they hear must be an example of one of them. We do not want our students to be able to cope only with what they have "done" previously. We would prefer them to think of speech as a complex of numerous parameters, that can combine in new and possibly surprising ways. Human attempts to represent the complexity of speech (by transcription) are limited and relatively unreliable; where human listeners excel is in making judgements concerning specific features - even

tiny details - within that complexity.

Levels of analysis

It may perhaps be objected that by cueing the listeners in, and giving a fixed range of choices for responses we are producing a task that is too easy compared with unrestricted transcription. Our response to this is that the technique removes largely irrelevant distractions and pressures, but does not remove or in any way dilute the essential auditory discriminations to be acquired. It is certain that if students cannot handle the important sound discriminations under the controlled conditions of Analytic Listening, they will not be able to do so in the rough-and-tumble of unrestricted transcription.

A further advantage of cueing the listeners in, is that the expected level of `narrowness' can be varied from question to question. Most introductory practical phonetics courses, for instance, would not expect as routine the marking of aspiration differences, or of a wide range of secondary articulations. But specific Analytic Listening questions can focus on such matters, as Q4 and Q5 illustrate. We can thus teach and test finer discriminations than is easily possible with conventional dictation material.

Materials - development and future

The first Analytic Listening materials were constructed in 1991 and introduced into course P101 (Introduction to Phonetics and Phonology A) which forms part of the BA programme in Linguistics at UCL. In 1994 the technique was introduced into course B104 of the BSc (Speech Sciences) programme, and in summer 1995 assessment by means of Analytic Listening replaced the nonsense-word component of the examination for that course. Samples of Analytic Listening materials were distributed nationally to accompany presentations at a NetPhon workshop in 1995 and to the British Association of Academic Phoneticians at the 1996 Colloquium in York.

Quite extensive recordings of Analytic Listening material were completed in 1995 for use in the Department's Listening Centre and for sale to students. For these materials, students record their responses in printed answer books. Printed answer sheets and books are also used in regular classes and for examinations. This simple method of presentation has the advantage that students can listen to the material whenever they wish, and the printed answer books give a permanent record for feedback. But there will be real advantages also in linking the Analytic Listening technique to a system for Computer-Aided Learning.

Such a system could deliver training and test materials, give rapid feedback, keep long-term records of progress, and perhaps work adaptively to match the different needs of users.

One promising system is LAPT (London Agreed Protocol for Teaching), developed originally at UCL with input from other London Medical Schools. This is a flexible PC-based system for delivering computer-based teaching. It has so far been used principally by medical and biomedical science students. As the exercises are simply written on a wordprocessor, new subject matter from a range of contributing colleagues is easily incorporated. A unique feature of LAPT is the introduction of confidence assessment into multiple-choice testing. After each answer, students are asked to declare a confidence level of 1, 2, or 3. For a correct answer, the marks gained are 1, 2, or 3. For an incorrect answer, the marks are 0, -2 or -6. The intention is encourage students to think realistically about the confidence they feel in their answers. We believe this scheme will be especially valuable when extended to auditory judgements. In our view, it is very important that those trained in phonetics should be able to form some estimate of the reliability of their perceptions (in terms of repeatability, likely consensus with other observers, and agreement with what instrumental analyses might be likely to show).

At present LAPT does not have facilities for outputting sound, and our intention is to marry it with the system developed by our colleague M. Huckvale which he has called CD-Mike. In this, a PC becomes a sophisticated player for an audio CD, any part of which can be randomly accessed using information in a stored `script' giving identification and timing of the contents of the disk. Advantages of this system are that the audio CD can be a relatively straightforward transfer from existing recorded material, while any new CD will also be usable as a stand-alone recording.

Conclusion

We have outlined the background to the development of Analytic Listening, described some of its main features and indicated one line of future work. It remains to put the technique into a broader context.

Because it involves questions each with a clear perceptual focus, Analytic Listening opens the way to designing the practical phonetics syllabus around an explicit hierarchy of phonetic properties and parameters. Instead of being based upon a certain type of material

(nonsense, RP, etc), or upon a fixed number of "sounds" to be covered , our syllabus can be based directly upon auditory skills (ability to make judgements concerning voicing, nasality, pitch, aspiration, and so on). For each programme, and for each course within the programmes, we must specify just what phonetic properties it is relevant to teach and to test at each stage. The Analytic Listening technique than gives us the means to train awareness of these properties, and to test them reliably and directly. There are possibilities for national and international comparisons between courses and examinations and even for standardisation of recognised tests.

APPENDIX

Examples of Analytic Listening material

The first few questions in this short sample are referred to in the text by question number (Q1, etc) and illustrate specific points mentioned there. The remaining questions are included as additional illustrations. The set of ten questions is not intended to illustrate a balanced lesson plan or assessment test, but merely to exemplify the technique. The phonetic transcription given at right shows suitable dictation material to accompany the items and would not of course be visible on students' copies of answer books.

Q1	Indicate whether the dictated item begins with a glottal stop.		
1	with	without	[aʔ]
2	with	without	[ʔoʔ]
3	with	without	[ɛ]
4	with	without	[ʔɛ]
5	with	without	[ʔha]

Q2 You will hear items of the form VCV. Indicate whether the consonant has voiced friction or not.

1	yes	no	[avə]
2	yes	no	[arə]
3	yes	no	[aɣə]
4	yes	no	[adʒə]
5	yes	no	[aɱə]

Q3 Indicate whether the item you hear contains a monophthong or a diphthong.

1	monophthong	diphthong	[peɪd]
2	monophthong	diphthong	[pɛːd]
3	monophthong	diphthong	[puːd]
4	monophthong	diphthong	[pɜːd]
5	monophthong	diphthong	[pɛyd]

Q4 The items you will hear are of varying structures, but each contains one voiceless plosive. Indicate whether that voiceless plosive is aspirated or not.

1	aspirated	unaspirated	[əˈkʰɜː]
2	aspirated	unaspirated	[ata]
3	aspirated	unaspirated	[ˈwɔːtʼə]
4	aspirated	unaspirated	[stʰeɪ]
5	aspirated	unaspirated	[ˈwɔːʔə]

Q5 You will hear items of the form VCV. Indicate whether the consonant is velarised or not.

1	yes	no	[oɬo]
2	yes	no	[olʲo]
3	yes	no	[oɬo]
4	yes	no	[oɮo]
5	yes	no	[oʎo]

Q6 You will hear items of the form VCV. Indicate whether the consonant you hear is voiced or voiceless.

1	voiced	voiceless	[ɛʁɛ]
2	voiced	voiceless	[ɛɳɛ]
3	voiced	voiceless	[aɮa]
4	voiced	voiceless	[aɻa]
5	voiced	voiceless	[ɛj̥ɛ]

Q7 You will hear forms with two vowels.Compare the second vowel with the first and indicate whether the second vowel is more open, about the same height, or closer than the first.

1	more open	same height	closer	[bobœ]
2	more open	same height	closer	[bebi]
3	more open	same height	closer	[bybo]
4	more open	same height	closer	[dɔde]
5	more open	same height	closer	[dɯdy]

Q8 In each item, listen for the consonant with the manner specified, then identify its place of articulation.

1	plosive:	retroflex	palatal	uvular	[hɛce]
2	fricative:	palatal	velar	uvular	[jɛχɛ]
3	nasal:	palatal	velar	uvular	[çoɲa]
4	lateral:	dental	alveolar	retroflex	[maɻa]
5	fricative:	alveolar	postalveolar	retroflex	[quʂu]

Q9 You will hear versions of the English word *written*. Indicate whether there is a schwa vowel in the second syllable, or whether the second syllable consists only of a syllabic nasal.

1	schwa	syllabic nasal	(schwa)
2	schwa	syllabic nasal	(syllabic nasal)
3	schwa	syllabic nasal	(syllabic nasal)
4	schwa	syllabic nasal	(schwa)
5	schwa	syllabic nasal	(schwa)

Q10 Each item will contain two vowels.Indicate whether the first vowel is nasalised or oral.

1	nasalised	oral	[sɛsɛ̃]
2	nasalised	oral	[sɛ̃sɛ̃]
3	nasalised	oral	[fũfu]
4	nasalised	oral	[fufũ]
5	nasalised	oral	[fũfũ]