Pergamon

Language Sciences, Vol. 19, No. 2, pp. 101–137, 1997

0388–0001/97 $17.00+0.00

PII: S0388–0001(96)00067–8

PHONOLOGY WITHOUT TIERS: WHY THE PHONETIC REPRESENTATION IS NOT DERIVED FROM THE PHONOLOGICAL REPRESENTATION

HELEN FRASER

This article is essentially an argument against the idea—one of the foundational tenets of generative phonology—that the phonetic representation and the phonological representation are related to one another by computational transformation. The basis of the argument is an analysis of the philosophical presuppositions of the concept 'representation'. The analysis is made from the perspective not of cognitive science, but of phenomenology. The results of the analysis suggest an alternative approach to phonology, which is outlined in the second half of the article. Finally, some implications of this new framework for practical work are explored, and contrasts drawn between phenomenological phonology and generative phonology. © 1997 Elsevier Science Ltd. All rights reserved

1. Introduction

This article has two main aims. The first is to explain my very serious concerns regarding one of the fundamental tenets of generative phonology, the idea that the phonetic representation is derived from or generated by the phonological representation, and that the relationship between them is one of computational transformation. Though this idea is the basis of vast amounts of phonological research carried out with the aim of specifying the details of that computational relationship, it is rarely questioned or justified. Here I not only question it, but find it severely wanting.

My main intention, however, is not criticism for its own sake of the work of others. The analysis reported here was carried out with the intention of discovering the key causes of my personal dissatisfaction with the generative approach to phonological phenomena. I reasoned that once I had hit upon the essential source, as opposed to peripheral niggles, it might be possible, by changing one or two key assumptions, to build up a new phonology—one which enabled me to do the sorts of things I wanted to do, rather than focusing on the technicalities of specifying a computational transformation of one representation into another.

The second half of the paper presents the framework for phonological analysis which I came up with by this method, building on the work of Fraser (1992) (henceforward, 'SSP'). The framework is based on principles of phenomenological philosophy, which are very different from the philosophical principles of cognitive science, upon which generative phonology is founded.

No doubt my framework will be as personally unsatisfying to some readers as generative phonology is to me. But perhaps there is room, nearly 30 years after the publication of *The Sound Pattern of English*, for a genuine alternative to generative phonology. Indeed, I would like to think there is room for several, and would be very pleased to find that agreement with my assessment of the flaws of generative phonology coupled with disagreement as to my proposed solutions, spurred others to develop frameworks to suit other needs.

Correspondence relating to this paper should be addressed to H. Fraser, Department of Linguistics, University of New England, Armidale, NSW, Australia (E-mail: hfraser@metz.une.edu.au).

2. A generative introduction to phonology

Recently a new, very large, very thorough and very highly regarded[1] phonology textbook has appeared (Kenstowicz, 1994, henceforward 'PGG'). It is by far the most detailed and comprehensive phonology textbook ever written. Kenstowicz begins it by describing some of the observations about speech sounds which have motivated the enterprise of phonology since the time of Sapir, and which attract new students to the discipline every year.

He quotes from Sapir's vivid description of the 'phonetic illusion' of a speaker of Sarcee, that there was a [t] on the end of certain words where no [t] objectively existed, and adds the similar more familiar example of speakers of non-rhotic dialects of English routinely claiming to hear and pronounce an [r] in words like 'soar', where, again, no [r] can objectively be observed. Based on discussion of these cases, and other puzzling phenomena of speech sound behaviour, Kenstowicz goes on to state the motivating questions of phonological research as follows:

> Why are Sarcee speakers compelled to represent *díní* 'this one' as [dìnít] and r-less English speakers [sɔ˞] 'soar' as [sɔr]? How does a child growing up in the Sarcee/English language environments where the relevant sounds are always absent in the pronunciation of these words discover them in the course of acquiring the language? A reasonable first guess is that the phantom [t] in Sarcee and the [r] in English derive from the fact that these sounds actually appear in the related suffixed words *dinit'-i* and [sɔriŋ] 'soaring'. But what precisely do we mean by a related word? How is it possible for the pronunciation of one word to influence another? Why doesn't the influence run in the opposite direction: that is, why doesn't the absence of a final consonant in the isolation form induce deletion in the suffixed forms? Providing a serious scientific answer to these kinds of questions forms an essential part of the research program of generative phonology. (PGG, p. 5)

Surely any reader coming to phonology for the first time will be intrigued by these observations, and the other 'phonetic illusions' discussed in the next few pages, and keen to find within the book answers to such questions as:

- Why is there this divergence between what is objectively there and what people think is there?
- How does the divergence work?
- What role or function does it play?
- Is this relationship between what is objectively there and what people think is there something special to phonology, or does it occur in other domains?

Kenstowicz's introduction goes on to suggest that these phonological phenomena are essentially of the same kind as the syntactic grammaticality judgements[2] that motivate generative linguistics, and that therefore a generative approach is appropriate in phonology, no less than in syntax. This involves a reformulation of the questions on page 5 (quoted above) within the generative framework, as follows:

> For any given language, what are the phonological representations and rules that have developed from UG in individual speakers as a result of exposure to the language of their environment? How do these rules apply to compute the phonetic representation? At a more theoretical level, what is a possible rule and representation? What elements are representations composed of? What principles govern their combination? Precisely how does UG [= Universal Grammar] develop into G_{Korean}, G_{French}, and so on, in response to the language of the environment? (PGG, p. 10)

Our novice reader might at this point be forgiven for feeling that some sort of swindle has been perpetrated! We started with an observation that challenges our everyday assumptions and poses deep questions about the nature of human language and even the human mind, but in the space of a few pages these have been reduced to far more banal questions about representations and rules couched in an alienating computational jargon. Such a reader might be further

disappointed by the way the book progresses. In the remaining 700 pages we are led through a tangle of technicalities, which are not just daunting in their complexity, but which bear little apparent relationship to anything linguistic or even human. Such a reader is likely to give up after some time and turn to another branch of linguistics which has a better grasp of what human beings and human language are all about.

These remarks are by no means intended as criticism of Kenstowicz's book in particular, which I agree is truly excellent as an account of generative phonology. However the situation seems to me to be a great shame, since the original observations are fascinating and challenging, and finding an account of them that does justice to our understanding of human language and human nature seems like a worthwhile goal.

The problem in my view is not with Kenstowicz's account, but with generative phonology itself, which falls a long way short of this ideal.

In this article, I would like to analyse the root causes of the dissatisfaction I find with the generative approach, and then use this analysis to present an approach to studying phonological phenomena which I find more satisfying than the generative one.

To help me explain the difficulties I have with generative phonology, I will start by briefly outlining its basic conceptual framework.[3] As well as PGG, I will also refer to Chomsky and Halle (1968), the seminal work of generative phonology, henceforth 'SPE'.

3. What is generative phonology?

The basic observation of generative phonology (GP) is that there are always two possible representations of speech sounds: the one (to stick with the Sarcee example) with the 't' and the one without the 't'. These can be called the 'phonological representation' and the 'phonetic representation', respectively. Every phonology textbook begins with examples to demonstrate the existence of these two representations, a fact of which language users are generally unaware until it is pointed out to them, since they pay attention only to the phonological representation.

In GP the phonological and phonetic representations are understood as being different representations of the same thing, differentiated by their relative abstractness. The phonetic representation is the observable pronunciation (the one without the 't', or without the 'r', in our examples); we can say it is a relatively concrete description of the sounds of speech. It can be crudely characterised as showing the 'actual pronunciation' of words.

The phonological representation[4] is relatively abstract. It cannot be directly observed, but its existence can—indeed must, according to GP[5]—be inferred from observation of 'alternations' (variant pronunciations of the same morpheme):

> If the alternations are regular, we assume that the morpheme has a unique underlying representation, such that the various phonetic shapes arise from sound changes introduced by context-sensitive phonological rules (PGG, p. 89).

The phonological representation is thus often called the underlying representation (UR). Various kinds of evidence can be used to determine its nature, which is more controversial than its existence.

Both the phonetic and the phonological representations of speech are given in GP in terms of Distinctive Features (DFs). These are a small set of phonological units, which, by being combined in a range of different ways, can represent all the speech sounds of languages. Use of DFs serves two purposes for GP: firstly it suggests that the ultimate number of elements out of which language is composed is small; secondly it gives a mechanism whereby relationships between sounds can be shown. The use of features is considered to be a natural progression from the desire to show that speech sounds form 'natural classes' in terms of their

phonological behaviour. For example, it is commonly observed that phonological rules affect classes of sounds, such as 'obstruents' or 'high front vowels', rather than random assortments of sounds. GP captures this observation theoretically by postulating that phonological rules operate in terms of features, and thus specify a whole class of sounds (PGG, p. 19ff).

The relationship between the two representations, phonetic and phonological, is conceived in terms of processes which derive the phonetic representation from the more abstract underlying or phonological representation[6] (Carr, 1993, p. 46).

In order to understand GP, we have to see it, as Kenstowicz does, as part of Chomskyan, or generative linguistics (GL). GL, in contrast to the structuralist linguistics which preceded it, is concerned not merely with the description of languages, but with fundamental questions of how speakers and hearers are able to use language, and particularly with the knowledge (competence) which underlies that ability. More generally, GL is interested in specifying the aspects of linguistic competence which are universal to all languages, and can be postulated as part of a putative innate universal grammar (UG).

In focusing on linguistic competence, GL is showing its interest not in language merely for its own sake, but in language as a window onto the human mind. Ideas about the nature of the human mind were very important in the development of GL from the previously dominant school of linguistics, Bloomfieldian structuralism. GL subscribes to the representational philosophy of mind, which, simply put, states that in order for a person to recognise something, they must have a representation of that thing in their mind. The fact that things which are physically different can be recognised as 'the same' is accounted for by the idea that representations in the mind can be transformed by rule-governed processes.

Language is thus seen as a series of levels of representation (structures) related by processes which transform one representation into another. In general, 'higher' or more abstract levels of representation are more 'efficient' than lower, in the sense that they involve fewer symbols in more parsimonious patterns. The reason that GL is called 'generative' is that it provides a mechanism to explain how the finite human mind, able to store a limited number of symbols and rules, can generate an infinite number of sentences.

The goal of GL is to make a model of linguistic knowledge. In doing this it uses a framework which shows a series of components, or modules, each of which is input or output to other components, as shown in Figure 1. The exact theoretical relationship between a linguistic description in these terms and the processes that an actual speaker or hearer goes through in producing or understanding a sentence is controversial, as we will see below.

Phonology, in GL, is concerned with the lower levels of this overall process, the phonological component,[7] which serves to make the abstract structures of language pronounceable. The input to the phonological component is the underlying representation (UR), or phonological representation. This representation is roughly equivalent to the lexical or morphemic representation that is the output of the next higher level of processing.[8]

> An underlying or phonological representation will contain all and only the unpredictable (distinctive) information for each lexical item. Predictable features are added to the underlying representation by grammatical rules and principles. These rules operate on the basis of information in the lexical item's phonological representation and the context in which it is located. For each possible word constructed by the morphology and for each possible sentence constructed by the syntax, the phonological rules will thus 'compute' or 'derive' a (surface) phonetic representation. (PGG, p. 60)

The output of the phonological component then is the phonetic representation. This represents (rather than describing in perfect detail) the actual pronunciation of speech sounds in their particular context, which is very commonly different from the UR. A standard idea about the phonetic representation is that is provides the information on which the speech production

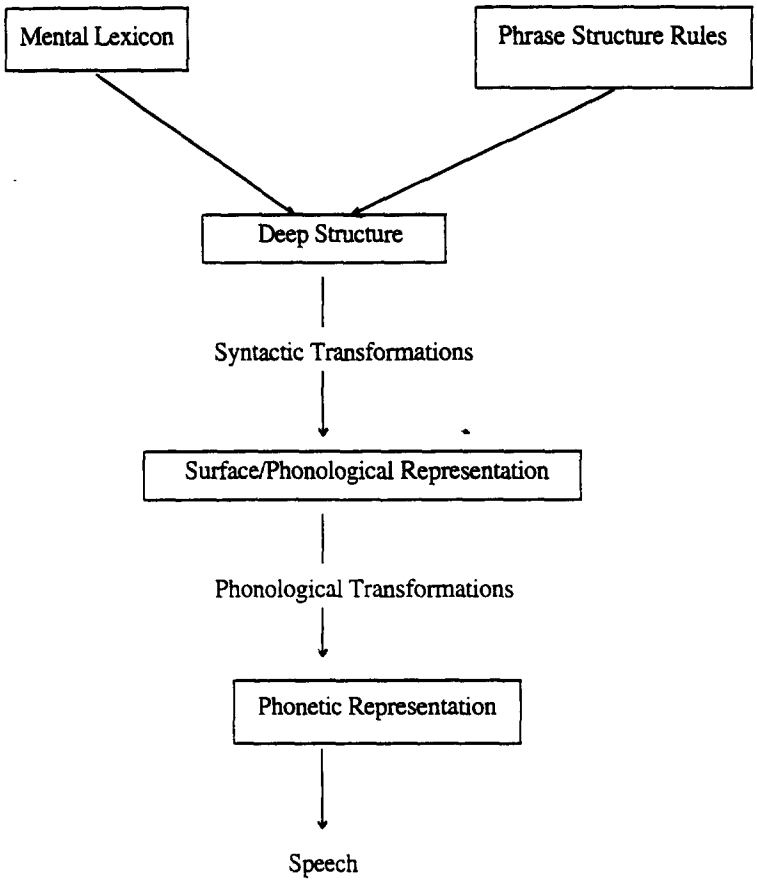

Fig. 1. Generative phonology (*cf.* Carr, 1993, pp. 46, 108; Katamba, 1989; Durand, 1990, p. 35; Goyvaerts, 1994a, p. 1385; PGG does not give a diagram, but the text supports this general model as a starting point).

system can base a sequence of muscle commands which will result in an appropriate string of speech sounds (PGG, p. 39); or according to which a linguist or speech synthesiser could build a full description of the speech, by adding knowledge of universal phonetic properties of sounds. It is seen as the task of phonetics, as opposed to phonology, to specify these detailed descriptions of speech sounds; phonology is concerned to specify the relationships between sounds, rather than their actual nature.

So, to give an English example, the URs of the sounds spelled with a 't' in 'top', 'stop', 'hot' and so on would all be the same (avoiding for present purposes the question of exactly how this sound should be represented). Rules in the phonological component change these into the various forms that occur in the phonetic representation: aspirated, unaspirated, unreleased, and so on. This representation is then acted upon by low-level implementation rules which provide the instructions to the muscles to move in a way which will produce the required word.

In general, phonological processes operate to transform a smaller number of URs into a larger number of actual pronunciations, by adding or removing aspiration, voicing, and so on, or by changing the place and manner of articulation of segments in particular contexts.[9]

Given a basic model such as this, the main task of GP phonologists is to specify in detail

how this transformation of phonological representation or UR into phonetic representation (which has turned out to be very complex) takes place, both in specific instances, and in general. In order to do this, the nature of the UR has to be determined. The UR, as we have seen, is unobservable. Though its existence is accepted as a cornerstone of GP, its nature is controversial. In order to achieve the goal of specifying the processes that transform the UR into the phonetic representation, it is necessary also to specify the nature of the UR.

> The crucial task of the phonologist ... will be to set up phonological (underlying) representations and to discover the rules that map them onto phonetic surface forms. (Goyvaerts, 1994b, p. 3560)

What this means in practice is that phonologists observe the actual pronunciation of words in a language, take this as evidence of the nature of the phonetic representation, and then deduce both the UR and the rules that convert the UR into the phonetic representation.

This practice leaves quite a lot of room for manoeuvre. Phonologists recognise the co-dependence of postulation of the UR and of the processes that are needed to transform the UR into the phonetic representation. Over the years there have been many changes in GP, with regard to the way the UR is conceived, and with regard to the type of processes that are envisaged as converting the UR into the phonetic representation. Some of these developments will be discussed in the next section.

4. Theoretical concerns within GP

In this section I would like to outline some general concerns about the exact formulation of the task that GP has set itself. It should be emphasised that here I am bringing together concerns which have been raised on occasion from within a GP framework, rather than raising my own original criticisms. However, where they have mostly been raised and addressed in the literature as individual concerns,[10] I believe it is important to see these individual concerns together; I hope that analysis of them will help me to motivate and develop the alternative theoretical perspective I will offer below.

4.1. The nature of the phonological representation

Specification of the phonological representation, both in universal terms and in terms of the grammars of specific languages, is a main focus of work in GP. There is general agreement about several characteristics of the phonological representation, and considerable debate about others. The phonological representation is universally agreed, as we have seen, to be more abstract than the phonetic representation, and to be a more efficient encoding of linguistic information, containing only the distinctive information about speech sounds.

However there are serious questions around the issue of exactly how abstract the phonological representation is, and debate about this has a long history in phonology. The criterion that the UR should be a maximally efficient representation of morphemes led, famously, to postulation of forms which many found impossible to credit as related to the way they appeared in real speech.

The attempt to solve this problem raised even more serious issues regarding what evidence or criteria should be considered relevant in deciding how abstract the phonological representation is. The issue became one of how to *constrain* the UR.

For some, it is important that the UR should not be allowed to diverge too far from the phonetic reality of its realisation. This constraint pursued rigorously led to the development of Natural Generative phonology (Anderson, 1985, p. 339ff), but is also in a more general way a constraint on many generative phonologists.

For others, it seems obvious the nature of the UR must be decided with reference to psychological evidence. From this point of view it seems unlikely that URs are always the maximally simple and efficient representations suggested by standard GP: there is much evidence from many quarters that the assumption that the human mind operates with minimum storage and maximal computation is not always justified. Again there have been some attempts to reform GP to take account of these concerns, notably Ohala's 'Experimental Phonology' (e.g. Ohala and Jaeger, 1986). These concerns (Wheeler, 1980, raises many of them) have had some impact on GP in general.

It remains true however that in mainstream GP, the main method for deciding on the nature of the UR is examination of a corpus of data, with less reference to either detailed phonetic evidence or psychological experimentation than would satisfy many critics.

> Although other more 'experimental' approaches have been proposed from time to time (see Ohala, 1986) the study of ... 'corpus-internal' generalisations ... continues to be the major avenue of research into the speaker's internalised grammar. (PGG, p. 57)

What this means in practice is that theorists bring in psychological evidence, and indeed more detailed phonetic evidence, as and when they wish, rather than in a principled way.

It also means that a range of assumptions about the nature of the human mind (such as the simplicity, efficiency, etc. of its operation[11]) are brought into play as 'constraints' on the specification of URs, often with very little explanation or justification, and in defiance of large bodies of psychological evidence (e.g. Margolis, 1987) that the human mind does not operate in such a formal way.

We still lack general discussion of the extent to which phonetic, psychological or other criteria should influence postulation of URs. The result is that issues of abstractness remain a major problem in GP, with no very clear path towards their solution opening up.

> While shifts in the field have significantly altered the terms of the debate, the question of abstractness [of the UR] unquestionably remains central in phonological theory. (Broselow, 1994, p. 1398)

4.2. The status of the phonetic representation

The phonetic representation has certainly not been a focal concern for GP (e.g. Diehl, 1991). It is however a concern of phoneticians working within the GP framework. We have seen above that the division of labour between phonetics and phonology has been considered one in which phonology investigates the relationships between sounds, or their formal properties, whereas phonetics investigates the actual nature of the sounds, or the content of the formal categories defined by the phonological relationships.

Phoneticians have often observed that if the phonetic representation of GP is intended as a record of what speech is 'really like', it is a very poor one (e.g. Rischel, 1990). Speech after all is continuous, rather than a string of discrete units as given in the phonetic representation; and the differences between elements are gradient rather than categorial. Even a narrow phonetic transcription represents speech only incompletely.

To this, theoretically-minded phonologists have responded that it is not the intention of the phonetic representation to give an accurate record of 'real speech': the phonetic representation itself is abstract from the point of view of the real nature of speech—though not as abstract as the phonological representation.

> ... phonetic transcription is understood ... not as a direct record of the speech signal, but rather as a representation of what the speaker of a language takes to be the phonetic properties of an utterance, given his hypothesis as to its surface structure and his knowledge of the rules of the phonological component. ... Since the phonetic transcription ... represents the speaker-hearer's interpretation rather than directly observable

> properties of the signal, the existence of certain discrepancies between the transcription and the signal can be understood. Thus it is no longer a problem that the transcription is composed of discrete symbols whereas the signal is quasi-continuous, or that the transcription provides information only about some properties of the signal and not about others ... The phonetic transcription can therefore be taken to be a two-dimensional matrix in which the columns stand for consecutive units and the rows stand for different features. (SPE, p. 294)

These attitudes led to considerable bad feeling between the disciplines of phonetics and phonology in the decades following publication of SPE, despite the fact that they were ostensibly dealing with the same phenomenon, namely speech sounds (*cf.* Keating, 1988). Phoneticians derided phonologists' cavalier attitude to data, and accused them of simply dreaming up a phonetic representation which suited the rules they wanted to postulate as having produced it; phonologists derided phoneticians' prosaic attitude, which focused attention on mere description, and away from the arena of theoretical interest.

More recently there has been some rapprochement between the disciplines. Some phoneticians devised models of speech that both took account of the nature of speech and understood the systematic nature of language in the Chomskyan framework—postulating for example, multi-valued rather than binary features (e.g. Ladefoged, 1975), or motivating features from cross-linguistic evidence (e.g. Kohler, 1984). Phonologists for their part paid more attention to the phonetic detail they were accounting for theoretically. PGG for example includes large sections on descriptive phonetics.

This increased attention to the phonetic nature of speech was part of the motivation for the development of non-linear or autosegmental approaches in GP (Clements, 1994; Goldsmith, 1989). The representation of speech in tiers, with features structured as hierarchical trees rather than unordered lists of distinctive features in a sequence of segments, allows a nearer approximation to the continuous nature of real speech.

On the other hand, such a representation is still essentially a discrete and categorial one, and as such needs translation into the continuous and gradient form of actual muscle movements or acoustic waves. The assumption that this translation can be achieved by low level implementation rules has been seriously challenged in several quarters. Though this problem has not yet gained central status in phonology (e.g. Kenstowicz's textbook barely touches on it), there is a growing interest in the 'phonetics–phonology interface' problem. For these theorists, the question of the relationship between these two components, and the question of where exactly the transformation from discrete to continuous representation occurs, takes on key importance (Keating, 1991; Blumstein, 1991; Cohn, 1993; Rischel, 1991).

This highlights a serious issue for GP—the fact that there is no general agreement on exactly what the phonetic representation actually is in the model. Is it the last of the linguistic (categorial) representations, or the first of the physical (continuous) ones? Is it acoustic or articulatory? What general characteristics does it have: what units of representation are appropriate to it? Given that it is abstract,[12] how abstract is it? Indeed how should this be determined? What constraints ought to be placed on its abstractness?

Very few phonologists address these general questions explicitly. On the other hand, all make implicit assumptions as to their answers. Theorists use the concept of phonetic representation with a wide range of different connotations and presuppositions, as has been dramatically illustrated by Nolan (1982), who compared the ways six speech theorists understood the phonetic representation, and found five widely divergent ways of conceiving it.[13] The hope that the differences are not really important since all the representations are transformable into each other is certainly unfounded.

In short, few phonologists act according to Pierrehumbert's conclusion that 'Phonetic

representation is one of the most difficult problems in linguistics.' (Pierrehumbert, 1990, p. 391). Indeed the truth is that most GP phonologists still take remarkably little interest in the phonetic representation. It is not unfair to say that the subtle theoretical distinction between a description of 'real speech' and a somewhat abstract phonetic representation is rarely at the forefront in phonological work.[14] It could be argued that practical work is guided by a much more simplistic idea that the phonetic representation is the 'obvious' description of the observable physical nature of speech, and as such, frankly uninteresting. The main interest is still in the work of specifying the UR and the processes that operate on it. The phonetic representation is simply the transcription used as the starting point for analysis of what these constructs must be like: whether the particular transcription used is strictly accurate or not is of less interest than the processes that must be postulated as generating it.

This has allowed an interesting ambiguity to persist in the term 'phonetic representation': it can mean either the level of representation in the model of the language-user's competence, or the transcription on which phonological analysis is based. In this regard it is worth pointing out that in practice, even where a detailed phonetic description of the data is carried out, and even when consideration is given to the formal characteristics of the phonetic representation used in the model, universally the transcription given to motivate the postulation of particular rules and underlying representations is given in a rough segmental transcription. Which is why there have been continuing calls for phonological theories to be more accountable to their data, despite the improved apparatus for representing the non-linear nature of speech (e.g. Diehl, 1991; Kingston and Diehl, 1994; Docherty, 1992), and also for more attention from phonology to 'real speech' as opposed to the idealised citation forms generally used as the corpus from which phonological theorising begins (e.g. Dressler and Moosmuller, 1991).

Even Kenstowicz, when displaying data for discussion and analysis, uses none of the sophisticated phonetic representations he presents in Chapter 4. Rather he gives it in the bizarre combination of orthography and phonetic symbols—typical of many phonologists—which already presupposes most of the analysis he is about to argue for. The impact on theory should not be underestimated.

4.3. The nature of the transformational processes

Specification of the processes by which the phonological representation is transformed into the phonetic representation is another key task of GP. The strong interaction between the two tasks, postulation of the UR and of the processes needed to transform that UR into the phonetic representation, has been noted (see especially Anderson, 1985). Many of the same issues arise in considering the nature of the transformational processes as have arisen for the nature of the UR.

Again there is wide agreement on the general necessity for such processes, but considerable disagreement about their more specific nature. It is interesting to note however that this necessity is rarely justified or explicated to nearly the same extent as the need to postulate an underlying representation. Once the distinction between UR and phonetic representation has been determined, the relationship between them seems obvious in GP. Even in SPE there is no serious discussion or justification of the assumption that the UR is transformed into the phonetic representation.

The transformational processes are generally assumed to be computational, i.e. to make as little reference as possible to the meaning or other non-phonological characteristics of the units being transformed.[15] In other words, they should operate so far as possible on the *form*, not

the *content*, of the sounds—though there is some disagreement as to the degree to which this is possible.

In the early days the assumption was made that the transformation should be conceived in terms of rules, and rules postulated which maximised the efficiency and simplicity of the process. As with the abstractness of the UR, however, application of this assumption led to problems: the proliferation of rules (and intermediate levels of representation), and the need for a complex system of ordering for their application, which seemed intuitively unlikely to be part of normal linguistic processing.

Here too, the reaction to such problems was a call for constraints on the rules that could be postulated. The constraints invoked were general considerations about the nature of the human mind, and, to a lesser extent, psycholinguistic evidence (e.g. Olson and Clark, 1976), suggesting for example that speakers producing speech do not always go through the same levels of transformation as linguists postulated in their theories, or that perception and production are not simple mirror images of one another, as implied by the GP model. One well-known development was Natural Phonology (Anderson, 1985, p. 343).

More recently, there have been challenges to the idea that the transformation involves rules of the traditional kind at all. In tune with developments in cognitive science and neural network computing, non-linear phonology has suggested that rather than explicit rules operating on a traditional UR, the process can be achieved by a system of constraint statements that must be satisfied. These ideas have been taken further in, for example, Goldsmith (1993), aptly titled 'The Last Phonological Rule'.

These ideas, linked with developments in thinking about the UR, have led to proposals for 'monostratal phonology', in which the concept of two levels linked by processes is abandoned (Bird, 1995). This approach concentrates wholly on the phonological representation, stating constraints on its form rather than rules for translating it into a phonetic representation. Mention is however made of a phonetic representation, conceived in articulatory terms along the lines of the 'gestural score' representation suggested by Browman and Goldstein (e.g. Browman and Goldstein, 1992). It is not clear therefore that this theory really is 'monostratal', since it does recognise a phonetic representation separate from the phonological representation; presumably the intention is that the phonetic representation is much less abstract than is usual in GP, with muscular movements being determined directly by the phonological representation, rather than being mediated through the stage of a symbolic phonetic representation. But the issue of how the phonological representation is transformed into the phonetic representation remains unsolved, so far as I can see.

5. An outsider's perspective

Taking together all these issues that have arisen for GP allows me to point out my concerns with this model as a way of understanding the sound systems of languages.

The situation as I see it can be summarised as follows: the basic framework of GP sets the task of specifying a transformation of the phonological representation into the phonetic representation, a task to which large amounts of energy are directed by GP theorists. But when one probes with questions of what exactly the phonological and phonetic representations are, and why language should involve the transformation of one into the other, no very convincing answers are forthcoming. In particular, one looks in vain for explication of the terms 'representation',[16] and 'abstract', or justification for the central role of representation in the model, or the hierarchical relationship between them.

GP practitioners are unlikely to accept this remark. To them, these issues are philosophical and not phonological. GP looks to the philosophy of cognitive science for its 'big picture'. It is from here that GP takes the basic framework of the representational nature of mind and the transformational nature of language outlined above.

But turning to the more philosophical discussions of cognitive science brings little comfort for someone questioning GP's use of the concept of representation, since its justification there is little better. Many problems have been raised, by both linguists and non-linguists, about the representational view of mind, based as it is on the untenable 'conduit metaphor' of communication and representational/computational view of human thought (see for example, Shanon, 1993; Harris, 1981; Grace, 1987; Carello *et al.*, 1984; see SSP for my own criticisms).[17] But despite the fact that representation is a central concept of cognitive science, little notice is taken of the many problems that have been pointed out, from a wide variety of perspectives, with the representational view of the human mind; indeed the whole concept of representation is remarkably little discussed within that discipline, despite its central status in the theory.

I have to say that it strikes me as odd that GP practitioners are not more concerned than they are by this situation. Presumably it all makes sense to them, and there is no doubt that much useful work has been done within the framework by linguists who do not seem to feel as troubled as I do by the lack of clarity regarding exactly what it is they are doing.

Personally however I find this lack of clarity about 'big picture' questions very disturbing. It is immensely frustrating to be confronted with bookfuls of machinery for transforming something into something else, with no very clear explanation of what either entity is, or of why one of them should be transformed into the other. It is also frustrating to find the discussion always at such a removal from 'real speech', 'real people', and the interesting questions Kenstowicz starts with. When problems are pointed out, GP moves to implement 'constraints', but my frustration continues with the lack of discussion of principles for deciding upon relevant constraints and implementing them consistently.

To me, in short, the movement in PGG (or any of the other introductions to phonology) from identifying the phenomena we wish to investigate, to recasting the questions in the framework of GP is far too fast. Before we rush into creating machinery for transforming one representation into another, we need first to address questions such as the following in depth:

- What *is* a representation?
- What does it mean for a representation to be abstract?[18]
- Why should there be several representations of speech?
- Why should one representation be transformed into another?

Of course I am not the only one with concerns about GP, as has been seen. There are many criticisms, and various alternative theories have been proffered, some of which have been mentioned above. Many of these, frankly, do not impress with the depth of their theoretical analysis: there is often a sense that the main concern is to get that over with so as to return quickly to practical[19] work.

Two analyses which I do find very impressive are those of Fowler and Linell (e.g. Fowler, 1985; Linell, 1979, 1982).

Fowler is primarily interested in speech production theory, and motivated perhaps most strongly by dissatisfaction with the lack of attention GP gives to phonetic evidence, and the relationship between the continuous and discrete representations of speech. She criticises the division of labour between speech science/phonetics on the one hand, studying speech in its continuous aspect, and phonology on the other studying speech in its discrete categorical

linguistic aspect. The assumption each makes, that either of these views is easily translatable into the other, she considers to be called into serious question by the philosophical problems of the mind–body relationship. This undermines the value of theoretical work in both disciplines.

Her recommended solution is to see the two aspects of speech as *both* being physical instantiations of abstract phonological form, one in the physical medium of muscle movement, and the other in the physical medium of neural activity. The change from one form to another is replication of a pattern in a different physical medium, not translation of a mental pattern into a physical medium. On this view, the assumption must be made that phonemes or other linguistic units should be observable in the physical (articulatory or acoustic) realm. The fact that they are not obvious in the form that we expect to find them should make us question our assumptions about what they are like: perhaps they are not so static, discrete and context-free as we think? These ideas lead Fowler to much interesting work seeking more insightful phonetic representations of speech.

Linell's view is somewhat different, though motivated by similar frustrations with the constant focus of GP on the UR and postulation of rules to transform it into a phonetic representation. The motivation for his work seems to be dissatisfaction with the lack of respect GP shows for psychological issues regarding the nature of the mind of the language user. He considers it important to constrain the abstractness of the postulated UR in a non-arbitrary way. The view he develops is that phonological representations are plans for the production of speech, which makes them concrete, rather than abstract, entities. The focus, in his view, should be on observation of linguistic contrast at the phonetic level, not postulation of contrast at some unobservable abstract level, for whose existence there is little real evidence. These ideas result in a range of interesting phonological analyses of real connected speech data.

While each of these arguments is admirable in its own way, and each begins from dissatisfactions with GP similar to my own, neither locates the problems with GP in quite the same way as I do. To me, the key issue is GP's use of the concept of 'representation', and I would like to explore that issue more fully, and follow its implications for the conception of the language user.

There have been analyses of similar theoretical problems in other branches of linguistics which bear some affinity to my views (e.g. Davis and Taylor, 1990; Haiman, 1985) but to my knowledge nothing in phonology that takes the same perspective as the analysis I offer—though Bybee's 'functional phonology' (Bybee, 1994) makes some congenial points, and I will be interested to explore the relationship between her ideas and my own at a later date. I also have some sympathy with Kelly and Local's (1989) pragmatic approach (see further discussion below), in which they eschew theoretical analysis in favour of detailed description of speech sounds. However, I am sceptical as to the ultimate separability of description and theory, and I would like to think that my approach here is complementary to theirs.

In what follows I offer my own perspective on the phonological issues raised in the first part, based on insights of phenomenological philosophy. This gives one way, which I happen to find appealing, of answering the questions 'What is a representation?' and 'What does it mean for a representation to be abstract?' No doubt there are other ways, which would lead to different suggestions for theoretical development from the one offered here. I hope my discussion so far has raised them as serious questions in need of discussion, whether or not the solution I offer is accepted.

6. What is phenomenology?

Phenomenology is a much misunderstood philosophy. Many of its key insights are rather

subtle, and difficult to understand without a real change in the attitude with which one regards the world. This has left it open to two kinds of abuse. On the one hand, it has been commandeered as ammunition for a particular social and artistic movement, which suggests that any description is as good as another, and the only thing that differentiates 'good' from 'bad' is the power of those who make the judgement. On the other hand it has been ridiculed as either stating the obvious in the most obtuse way, or (often in the same breath) denying the obvious.

I believe that despite this there is a strong core of extremely valuable insight in the best phenomenological work, and that it is this insight which can help us to understand the nature of the phonological problems we are facing.

It also needs to be said that phenomenology is a very diverse discipline; indeed it is not so much a single school of thought as a range of loosely related methods sharing a few central commitments. In the present discussion I am giving a personal account of the phenomenological ideas that seem most relevant to the concerns of phonology, making no attempt to be an impartial or comprehensive coverage of the field (see Spiegelberg, 1982; Schmidt, 1985; Schmitt, 1969; Spurling, 1977; Roche, 1973; King, 1964; Llewelyn, 1985; Ricoeur, 1981 as recommended readings among a wide range of often less accessible texts; my own particular perspective is defended more fully in SSP).

The starting point for my presentation of phenomenological ideas is the observation that there is a distinction between a thing and its description, or between reality-in-itself and our perception or interpretation of reality.

This observation is of course common to many philosophies, but the use phenomenology makes of it is somewhat different. In phenomenology, attention is directed away from the many possible descriptions or perceptions of reality, and turned towards the maker of those descriptions (the Subject), and the processes by which descriptions are made. This shift of focus gives rise to many interesting philosophical insights—in my belief, it actually provides the way out from the age-old conflict between realism and idealism—and a large range of philosophical and theoretical questions which need not detain us here. It also has important implications for various kinds of scientific endeavours, most notably those of the human sciences, among which linguistics can be numbered.

Phenomenology takes very seriously the insight that any description is Subjective,[20] in the sense of being the product of a process performed by a Subject, and pursues rigorously the implication that no description—not even those of the hard sciences—is more objectively valid than any other.

This does not mean, as it is often parodied, that we can never say anything about anything, or that any description is as valid as any other. It means that, for any description to be accepted as valid in a particular context, the process whereby that description has arisen must be specifiable. For any endeavour, a foundational level of description must be defined and defended. There is no expectation that a framework suitable for one endeavour will necessarily be relevant to another.

Since no representation or framework is seen as objectively 'accurate', there is considerable flexibility in choice of a relevant framework for any project. On the other hand there is the responsibility to justify that framework in terms of the project being undertaken. For example, I will argue below that while it is useful for linguists to postulate the two representations of speech we have been considering and to specify the relationship between them in terms of formal rules when the project is one of coming to a linguist's understanding of what languages are like, this is not valid when the project is one of understanding the internal workings of the language-user's mind.

These phenomenological ideas, in my view, give an approach that can be taken to some of the fundamental questions raised above about the nature and status of the various representations of speech, and their relationships to one another, and which can be used to define frameworks for various kinds of research, and justify their applicability to the project in hand. That approach involves asking questions about the descriptions used in the framework, regarding the Subject who makes the descriptions, and the process of their constitution. In the following section, I will use this approach to unravel some of the issues that trouble me about the use of the term 'representation' in phonological theory.

7. What is a representation?

A straightforward way of understanding the nature of representation[21] is to see a representation as a product of a three-way[22] relationship, in which *someone* represents *something* with *something* else. A prototypical example of representation is a colonel representing a battalion with a salt-cellar. The salt-cellar in this case, while retaining its existence as a salt-cellar, also takes on another existence as a representation: it stands for the battalion. This existence as a representation of a battalion however presupposes the existence[23] of two other entities: someone doing the representing (in this case, the colonel) and something being represented (in this case, the battalion). For a representation to work, these two other roles must be performed by some other entities. I argue that this is the case with any representation. Nothing can exist as a representation without someone using it as a representation.

Use of the term 'representation' commits one to being able to specify who is doing the representing and what is being represented. That is easy enough to do in 'salt-cellar' cases, but what about other uses of the term 'representation'? It is common for example to say that 'language represents reality'. Who or what plays the three roles in this case? Framing this question makes it immediately obvious that a more precise locution would be to say that 'people represent reality with language'. Language is the representation; the representer is some language-using being.

The things being represented with language are however quite different in this case than in the salt-cellar case. When the Colonel used the salt-cellar to represent the batallion, it was clear that the salt-cellar already existed as a salt-cellar, and continued to do so even while it was being used as a representation of a battalion. What happens however when someone represents the salt-cellar as a salt-cellar in the first place? It is not already represented or described as anything prior to its representation in language as a salt-cellar. When we use language to represent reality the thing we are representing is some not-already-represented aspect of reality. There isn't something we can point to that plays the same role as the salt-cellar. In this respect, representation using language is quite different to the prototypical salt-cellar case.

I believe it is most important to acknowledge this. It is also rather difficult, since we language users develop, during the course of learning to use a language, a very powerful intuition that using language is a matter of attaching labels to aspects of reality. This feeling however is erroneous. Part of using language is constituting aspects of reality to which words and sentences can be attached.

Linguists are in general more sensitive to the need to credit this level of processing than other people. This is one of the main legacies of Saussure's work, which emphasises that both the signified and the signifier are language-specific constructs, rather than objective labels for bits of reality. The work of Sapir, Whorf and followers too has made it clear that even such apparently objective 'bits of reality' as colours are constructed differently by different languages, as are cultural objects, and social entities such as emotions (Wierzbicka, 1992, and others).

Another example where the constructed nature of the reality we represent with language is commonly brought to awareness is in psychological counselling, where it can be helpful to recognise that the way one represents a situation is to some extent under one's own control. For example, one approach to giving up smoking (Carr, 1994) is to change one's representation of the experience of nicotine withdrawal from 'I feel awful because I have denied myself the pleasure of smoking' to 'I feel a little uncomfortable because I recently ingested some poison'. In both representations, the undescribed sensation is the same, but the way it is represented or conceptualised is different—with, hopefully, significant differences in the action one bases on the representation.

These examples call attention to the three-way relationship which I argue is relevant in any use of language to represent reality. There is always a not-yet-represented reality, and someone doing the representing, even though in the vast majority of cases we notice only the representation and pay no attention to the other two elements. The thing being represented by language is some not-already-named aspect of reality. Unless we assume that reality comes already prepackaged so that words can simply be attached to it as labels, we must acknowledge some process of constitution of the aspect of reality to be named—in all cases, but most especially when we are dealing with bits of reality which are not physical entities or natural kinds.

Enough has now been said that some more precise terminology becomes necessary.

The one doing the representing is some language-using being, for which we can use the term *Subject*. A Subject is a being for whom (or which) representations can exist. 'The' Subject is a generalised version of this concept. The Subject must therefore be a being of a kind which can create representations of previously unrepresented reality. This ability requires a faculty that has been called 'background understanding' (Dreyfus, 1979; Varela *et al.*, 1991; Winograd, 1981), which involves understanding the general meaningfulness of a situation.

The not-yet-represented reality is a difficult concept to discuss, for the obvious reason that it is necessary to represent it to talk about it, and in doing so it is changed from being not-yet-represented to being represented. We can use the term 'Something'—or 'reality-in-itself', or 'World'—to speak generally about not-yet-represented reality. The not-yet-represented World certainly exists (phenomenology is not idealism), but there is no description of it that can count as 'objective'. There is no way of describing or representing a Something except from the point of view of some Subject. We can point to Somethings as existing in the World, and try to gain experience or background understanding of what they are like before we represent or describe them. Further exemplification of this concept will be given below; see also Fraser (1995).

Talking about the representation itself is also tricky when we are in the domain of language, rather than of salt-cellars. The word 'representation' is not ideal: it gives too big a temptation to think that using language is like using a salt-cellar to represent a battalion, and that something in the world is playing the same role as the salt-cellar.

There is a non-technical term for the concept I am interested in here: description. *Description* is in many ways a more useful characterisation of what language users do than *representation*; and it certainly avoids some of the unnecessary theoretical issues thrown up by the term representation, while still calling attention to the three-way relationship presupposed by its use. I will use the term 'description' commonly in the following discussion, intending it to emphasise the fundamental point that there is a difference between a thing and its description, with which I began my argument, and to highlight the point that the Something without a description given by a Subject is *undescribed*, not in the way of the salt-cellar before the colonel uses it to represent a battalion, but in the way of the salt-cellar before it has been called a salt-cellar (or anything else).

I will use the word 'description' commonly in the following discussion. There is also a more technical term that I will use when needed: Object. In order for a description to be made and used, relevant characteristics of the Something have to be selected. These characteristics can be *constituted* into an Object, in a sense that will become clearer as the discussion progresses. Perhaps the nearest ordinary word is 'concept', but this word too has too much theoretical baggage for comfort, and I prefer 'Object' for its complementarity with 'Subject'.

Unfortunately, the word 'Object' is potentially misleading due to confusion with the ordinary word 'object'. The capitalised word denotes any aspect of reality that has been described or talked about by a Subject. It can be a material object, or an action, or a characteristic of either an object or an action. It is not intended to imply that a representation must always be an object in the ordinary sense, nor that it must always be a representation of an object.

The important thing is that the existence of any Object (including those of scientific or mathematical description) presupposes that there has been an activity (which we can call Objectification, or constitution) by a Subject. Though this is often forgotten or ignored in everyday life (quite rightly), it should not be lost sight of in theoretical discussion. The representation is the product of this process, and is fundamentally different in kind from the not-yet-represented reality. Whether this process—call it 'constitution', 'Objectification', 'conceptualisation' or whatever—occurs prior to or as part of language use, and exactly what the relationship is between a concept and the words used of it, though an important issue, is complex and need not detain us here.

An application where these distinctions become crucial is in building models of language and language users. We should not be allowed to put a representation into a model without giving an account of its presuppositions. Any representation used in the model must relate to the right kind of Subject and the right kind of Something in the model. The phenomenological framework is useful in making sure that this happens properly. I believe that it is the key to sorting out some of the problems of phonological theory pointed out above.

A particular problem arises in using the framework to discuss phonology however: we use language to represent reality, but we also need to use language to represent language; in other words we need to use language as a metalanguage. Unfortunately the vocabulary of the language and of the metalanguage are often very similar, so confusion commonly arises as to which level one is speaking on. Language is itself a Something that needs to be described or Objectified by a Subject. Thus the undescribed sounds of speech have to be objectified as meaningful words and sentences in order to perform their function as language. This is the process we are interested in when discussing phenomenological phonology, and will be the topic of the rest of this paper.

Before leaving this section, it is appropriate to mention a third common use of the term 'representation'. This is not specific to linguistics but is important in understanding about the general characterisation of the human mind as representational. It is common in Western philosophy to think of knowledge as representation: thus it is assumed that, in order to recognise something, the mind must contain a representation of that something, and rules for relating that representation to others. This view is not only a fundamental tenet of cognitive science and generative linguistics, but a fundamental component of 'common sense' in Western culture.

The phenomenological approach opposes it however, on the grounds that it invokes a homunculus (recall note 17 above). As may be clear following the discussion of representation above, for a representation to work, there must be someone to make it, and to interpret it. Thus the idea that there are representations in the mind suggests that there is 'someone' inside the mind, creating and (more importantly) interpreting the representations. Similarly with the rules

for relating the representations—the rules may exist, but someone is needed to apply them (Wittgenstein is famous for having pointed this out). See SSP for more background; also Ellis (1993) and Shanon (1993) for similar arguments.

The fact that the representational view of mind is part of Western 'common sense' makes it difficult to bring it to attention, and even more difficult to oppose. However, it has been opposed in various circles, and its limits are becoming more widely accepted. Fowler (1994) gives a useful overview of some alternative conceptions of 'representation' which are now thoroughly accepted in many of the arts and social sciences, noting Critical Sociolinguistics and Halliday's Systemic Linguistics as two of the few branches of linguistics to have taken notice of these philosophical developments.[24] The phenomenological view presented here is far more compatible with modern ideas of representation as construction. The challenge is to apply these ideas of representation to a discipline such as phonology.

8. The generative enterprise from the perspective of phenomenology

It is interesting to look from the perspective of the framework just outlined at the generative enterprise of finding a computational relationship between the phonological and phonetic representations. The framework provides a way of understanding the nature of representation, and sets clear constraints on how phonetic, phonological and other representations in a model of human language use should be conceptualised: it must always be possible to specify who is doing the representing and what is being represented.

Asking the relevant questions about the phonetic and phonological representations of GP soon makes it clear that each is made by a different Subject: the phonetic representation is made by the linguist doing the modelling, and the phonological representation is made by the language-using Subject being modelled.[25] This is particularly evident in Sapir's descriptions of his fieldwork, quoted by Kenstowicz (see above).

With this insight, there is no reason in the phenomenological view for these two representations to be transformable into one another. Even if such a transformation of phonetic and phonological representations could be done simply, there would not be the temptation to ascribe it to the language-using Subject, since the task is obviously performed by the linguist, not the Subject.

In fact, of course, this computational transformation cannot be done simply. The entire history of GP, including the reasons for its divergence from Bloomfieldian structuralist phonology, demonstrates this. The self-imposed constraint that the transformation must be done computationally, with as little as possible reference to meaning, has resulted in the proliferation of levels of representation (*cf.* the similar observation in Fowler, 1985), and the complication of the computation—surely beyond anything that could possibly be ascribed to human subconscious mental processing, as a range of critics have pointed out.

It is easy to see the development of non-linear phonology, with its tiers and trees of representation, as the latest stage in the same pattern of development. With the understanding of representation that I have presented here, non-linear theory is no closer to being an account of human knowledge about speech sounds than any of the other versions of the theory that attempts to capture the relationship between phonetic and phonological relationships by computation. This of course is not to deny the value of any particular phonological theory when it is used for a particular purpose. As I will argue below, all kinds of models can be useful for increasing one's understanding of speech, so long as they are not pushed beyond the limits of their application.

However, to think that in finding a computational relationship between these two representations we are accounting for human phonological knowledge is to seriously misrepresent the nature of the language-using Subject, and what is happening when the Subject makes representations of the World, including representations of speech sounds.

Perhaps it will one day be possible to transform phonological representations into phonetic ones with a better degree of precision than it is now. But such success would be on a level with that of a grand theory that could, with no reference to their evolutionary ancestry, transform a representation of any bird, Escher-like, into that of a corresponding reptile. Perhaps such a theory can be made, but it would fall a long way short of being a satisfactory account of evolution.

To suggest that what the language-using Subject is doing is transforming one representation into another, even if we could show how it is done, cannot account for the really interesting thing Subjects do, which is to *create* representations of speech sounds from something previously *un*represented.

And indeed this highlights precisely what I think GP does: it misunderstands the nature of the Subject. In the sections following, I will clarify this claim that in the phenomenological view GP is misguided in spending so much energy specifying the nature of the transformation between the phonetic and phonological representations.

9. Speech sounds in the phenomenological framework

In this section I will flesh out the general framework given above with more specific reference to speech sounds, developing the theory of phenomenological phonology. We will be looking at what is involved in calling elements of phonological theory 'representations', and, more generally, what is involved in talking about and describing speech sounds.

9.1. The Something: speech sounds before they are described

The Something we are dealing with in phonology is undescribed speech sounds. Of course it is not possible in this framework to describe speech sounds objectively. All that is possible is to point to the Something whose description from different points of view we are interested in phonology. In this case the Something is the sounds of speech, the noises that people make with their vocal tracts in order to communicate linguistically. These can be described or Objectified in numerous ways, which we will discuss in more detail below.

A particular speech sound Something can be isolated without being described by being captured on tape or computer. This can then be inserted into different contexts and changes to its description noted. For example, it is possible to isolate the portion of the word 'spat' which is represented in spelling with '_pat', and play it with and without the preceding 's' sound. Doing this will make English speakers want to describe it as beginning with a 'p' in one case and as beginning with a 'b' in the other case. It is tempting but inaccurate to suggest that the 'p' has changed into a 'b', or that any other transformation of one representation into another, has taken place. To do this is to make a judgement that one of these descriptions is more accurate or 'real' than the other.

The PP framework encourages us rather to see the situation as one in which an undescribed Something is described in one case as a 'p' and in the other as a 'b'. The relationship between the 'p' and the 'b' is that they are both derived from the same Something, not that one is transformed into the other.

9.2. The Subject: the person describing the sounds

The Subject is the person who makes a description of the undescribed Something, or Objectifies it in a certain way. Any time there is a described speech sound, PP encourages us to ask which Subject has made that description, and to look at the process of its constitution.

A key insight is that in phonology there are *two* Subjects creating representations of speech sounds. The first is the language-using Subject (LUS) being modelled by linguistic theory. The second is the linguist *making the model*. A linguist in this sense is someone taking a detached and scientific attitude to language, as opposed to actually using language to communicate. The linguist obviously is also an LUS; and any LUS can take a detached attitude to language from time to time (though of course specialised training is needed to make the LUS into a linguist in the full sense of the term). At any one time however, we are dealing theoretically with the viewpoint of one Subject or the other—and the two are very different. It is crucial to keep them distinct in all theorising about speech sounds—though doing this takes a major effort, and makes us as linguists do something we are not used to doing as Subjects, namely consider our own Subjective contribution to the creation of the Objects with which we operate.

The linguist is modelling or theorising the language behaviour of the LUS. The goal is to make a model of a certain kind, and the background understanding has to do with what a model is, what the purpose of this model is, what 'science' is, and so on.

The LUS is using the language to understand not-already-described reality and to communicate. The goal is the linguistic meaning of the sounds, and the background understanding is general knowledge of the language, culture and situation.

The LUS is different from the Subject as conceived in GP. This must be so, because the Subject assumed by GP is engaged only in transforming one representation into another by means of some form of computation. The Subject in PP is engaged in creating representations or descriptions of not-already-described-reality, and understanding the relationships between the representations and between the aspects of reality being represented. These issues are explored in greater detail in SSP.

One of the factors in determining representations of reality is the knowledge the Subject has about the world, the language and the context in which the representation is being made. In PP as in GP, a major research question is the specification of the knowledge Subjects have that enables them to constitute linguistic Objects. However in PP the nature of this knowledge is conceived quite differently from the representational view of GP. Much of the knowledge is 'knowledge how', rather 'knowledge that', and there are limitations on the extent to which it is possible to think of knowledge as a catalogue of representations and rules.

Even to the extent that it is useful to think of the Subject's knowledge as a catalogue, the type of catalogue is very different in PP: the information is not organised like the books in a public library, according to arbitrary call numbers based on a carefully designed system of categories representing the content of each book. Finding information is more like finding a book at home, where location depends on a personally meaningful system, based on the way books are experienced in the life of the Subject: books are located sometimes by size, sometimes by order of acquisition, sometimes by where they were last used, sometimes by the order in which the owner intends to read them, or has read them, sometimes by the place they are most likely to be needed, and so on: a mixture of many criteria are needed to understand the organisation, and yet (mostly) the owner is able to retrieve a particular book very effectively.

9.3. The description

The Objects with which we are concerned in PP are the descriptions, or representations, of speech sounds made by each of these two Subjects in the contexts in which they operate. There are many such descriptions. Some relate to the LUS; others relate to the linguist; and they are at many different levels of abstraction. In PP these two kinds of description are fundamentally different, and there is little temptation to confuse them or to consider them more similar than they are. It is worth noting that this type of confusion is all but inevitable in GP, given the well-known 'systematic ambiguity' GP deliberately encourages between 'grammar' as a linguist's construct and 'grammar' as a native speaker's competence (e.g. SPE, p. 3). This type of ambiguity is anathema to PP.

In this section, I will look at both linguist's descriptions and LUS's descriptions. To some extent, the former can be equated with the phonetic representation of GP, and the latter to GP's phonological representation. However, the viewpoint from which these representations are made gets very confused in GP—due to the 'systematic ambiguity' just mentioned. I will therefore call PP's version of the phonological representation the 'psychophonological' description, to keep it distinct from GP's phonological representation, or UR.

9.3.1. Psychophonological description. A psychophonological description (PsD) is a description of speech sounds made by the LUS, and PP is committed to demonstrating that any description so called in a model is indeed one made by the LUS. To do this, it is necessary to observe not just speech sounds but Subjects, and the way Subjects use speech sounds.

In order to determine what psychophonological descriptions are like, we have to ask, 'As what kinds of Object does the LUS Objectify speech sounds?', or in simpler words, 'How does the LUS describe speech sounds?' It becomes evident immediately that there is not just one psychophonological description, but many.

It seems clear that the first and most basic description of speech sounds by the LUS is as units like those known in ordinary language as 'words'. For the time being we can define[26] a 'word' as 'a linguistically meaningful unit of speech sound'—the sort of sound that native speakers of a language can interpret, use and refer to.

It is important to emphasise that the linguistically-meaningful-sound Object is the first and most basic Object that the Subject constitutes from the speech sound 'Something'. It is very tempting to ask, as GP does, what *components* of the sound make up this Object. In PP this question does not make sense. Components of Objects are themselves Objects in PP, so questions arise as to the constitution of the components as much as the constitution of the Object. Explaining the perception of words in terms of the prior perception of smaller units (perhaps acoustic features, or segments) is considered question-begging in PP, since there is no assumption in PP that a detailed acoustic or articulatory description is a 'more basic' representation than one in terms of meanings[27] (*cf.* further discussion below of the role in PP of features of the 'Something', as opposed to the features of GP).

It may be worth reminding readers that I am talking at this point about the description of speech sounds as words, not about the description of the world using words. This distinction is important to maintain in both GP and PP (*cf.* Dell, 1980, p. 21, fn 14), though doing so is not always easy, since we have to use language as a metalanguage.

I am claiming then that the word is the most basic description of sound to the LUS. The word is the basic 'building block' of PP. (Recall that in GP the Distinctive Feature plays this role; the difference is very significant.) The rationale for this claim has several parts, which for lack of space I will merely rehearse briefly here: in language acquisition, it is this kind of

meaningful unit that is learned first;[28] in speech perception and production, these are the units that people are most conscious of and into which they most obviously Objectify speech sounds; if we think about the bigger picture, language is about meaning and its communication.

The word is not however the only kind of Object the LUS makes of speech sounds. The LUS can also describe parts of words. The first partition of this kind is the division of words into a sound-part and a meaning-part, so that we can think of the sound of the word 'dog' separately from the meaning 'dog'. This is certainly not something that Subjects do automatically, or that we as linguists can presuppose. Learning to do it is a similar process to that involved in learning to abstract the concept of 'red' from observation of red things. It takes children a long time to learn to do it; and there is evidence that speakers in non-literate cultures might do it to a lesser degree than we take for granted in our culture.

It is traditional in linguistics to think of words as a pairing of sound and meaning (in an arbitrary relationship).[29] For PP, while this is a useful way of talking for many purposes *for linguists*, it cannot be literally true *for Subjects*. Rather PP sees both sound and meaning as being derived from understanding of an unpartitioned whole comprising both sound and meaning in a 'linguistically meaningful sound'. The traditional idea of an arbitrary relationship is seen in PP as analogous to considering a coloured object to be a pairing of 'redness' and 'ballness', an idea which may appeal to componential semanticists but seems preposterous to many others.

In PP, sound and meaning are not originally separate entities, paired by a syntax, as they are in GP. They are separate Objects derived from the same Something. This view gives a very different perspective on the relationship between sound and meaning. In GP this is conceived as an arbitrary relationship. In PP, to see it as an arbitrary relationship requires the detached and reflective view of the linguist, comparing a range of languages. In the experience of the LUS, few relationships are as non-arbitrary as that between the sound and the meaning of the basic words of that Subject's own language.

Although PP cannot simply assume that sound and meaning are separate from the outset, it does seem to be the case that part of being a fully competent user of language involves being able to break words down into a sound part and a meaning part. There is an empirical prediction here: it seems likely that adult speakers of any language, whether literate or not, will have some way of referring to the sound as opposed to the meaning of words, for example in commenting on different accents or ways of speaking, or in using words in ways that involve attention to rhythms, rhymes, etc. Exploration of this prediction would be very interesting in view of the common observation that in non-literate cultures people are far less inclined than in Western academia to conceptually separate the sounds and meanings of words (e.g. Hutton, 1992).

Further, it seems that the LUS also subdivides the sound-part of words into sublexical units of various kinds. This is certainly true of literate Subjects, and PP predicts that it will be true, though undoubtedly to a lesser extent, even among non-literates. Evidence for this too comes from several sources. People can talk about speech in terms of component units of words; and their behaviour in classifying and comparing words—as pointed out in GP—seems to indicate some kind of awareness of words as having an internal structure. The types of units that seem to be relevant include syllables, onsets and rimes of syllables, and segments. It would also seem that Subjects are able to divide up the sound-aspect of words in a non-linear way, abstracting tone, intonation and other Objects. This idea is discussed further below.

We can assert then that the LUS does have some kind of representation of sublexical speech sound, though this representation is more, not less, abstract, than that of 'word'. This makes for a significant shift from a well-known concept of linguistics: the 'double articulation' of language, i.e. the view that languages are composed of small meaningless units which can be

combined in order to make larger units and then associated with a meaning. The concept of double articulation is a useful one for practical purposes of analysis. However it has limitations; in particular it is not a view which can form the basis of a model of a human language-user—for reasons which will be discussed in more detail below.

9.3.2. Phonetic description. A phonetic description (PhD) is also a description of speech, but a linguist's one. The role the phonetic description plays in PP theory is quite different from the role of the phonetic representation in GP. However it is not correct to attach the prefix 'psycho-', since the phonetic representation does not play any role in the Subject's mental processing during speech perception or production. The PhD is a linguists' description, not an LUS description; there is no temptation in PP, as there apparently is in GP, for linguists to make a phonetic representation and then ascribe it to the LUS.

In PP, the phonetic description is *more abstract* than the psychophonological description (notice the reversal of status from that given to the two representations in GP). The reason for this is clear from earlier discussion. In order to make a PhD it is necessary to understand a lot about language in general, and, usually, something about the particular language you are transcribing. The idea of making a PhD would only occur to someone who was able to take an attitude of detachment in observing language (most likely someone familiar with more than one language or dialect). A phonetic description could not arise for a LUS without this attitude of detachment; a psychophonological description on the other hand is necessary to the use of language in communication. We have seen how the least abstract level of psychophonological description is the word, with sublexical units of sound only arising from several levels of abstraction. A phonetic description involves processes of abstraction still beyond those. The making of a phonetic description relies upon prior understanding of psychophonological Objects; this makes it a more abstract description than any of the psychophonological ones.

As evidence for this point of view, consider the following well-attested observations: There is no sense in which the phonetic representation is anything like an automatic transduction of 'real speech', as is clearly demonstrated by the difficulties of making any kind of 'low level' transcription by machine, and indeed by the difficulty even highly trained phoneticians face in making a phonetic transcription of a language they do not know, or of a known language where 'higher level' information is artificially withheld. To be able to make a phonetic transcription, a linguist has to know the meanings of the sounds being transcribed, or at least have access to an informant to say whether they are indeed meaningful words of the language.

There is thus no sense in ascribing the linguist's phonetic description to a model of the LUS as a low-level representation. It is certainly not an Object that would arise for a Subject in the course of normal life, as is clearly demonstrated by the fact that phonetic descriptions of sounds have a very limited role in ordinary conversation, and the difficulty Subjects have in focusing on the phonetic aspects of speech sounds when asked to by linguists.

This idea that the PhD has no role for the LUS undoubtedly seems extremely odd to anyone who is used to thinking of the PhD as the level of representation which drives the motor control of speech production, or as the low-level acoustic representation on which speech perception is based. Here again it is necessary to keep in mind the distinction between the Something of speech, and its description. The Something certainly exists and has certain characteristics; but it can be described in many different ways from different points of view. The description of speech as a series of motor commands or acoustic parameters, while a valid one, can only be made from the linguist's perspective.

In PP then, the ambiguity of the term 'phonetic representation' that arises in GP does not

apply. In GP, it will be recalled, 'phonetic representation' could refer to either the element of the model which was the output of the processing done in the phonological component, or the transcription on which postulation of the UR and transformational processes is based. In PP, only the latter meaning is relevant.

However, while phonetic description is not part of the PP model of the LUS, it is an important part of PP, and PP sets great store by detailed and accurate description and transcription of speech, as will be discussed in more detail below.

So given that the phonetic description is a linguist's description, what can be said about it? The first observation is that there are many kinds of phonetic description, as was the case with the psychophonological description. These different phonetic descriptions can be exemplified by the many different kinds of phonetic transcription that are possible.

As there can be many different goals and points of view for the linguist, so there can be many different kinds of phonetic transcription. To ask 'What is *the* phonetic representation of this particular stretch of speech?' is not very helpful. The answer of course is 'It depends'. It could be articulatory, acoustic or auditory, in units of segments, features or parameters, it could be relatively 'broad' or relatively 'narrow', etc. (see Laver, 1994, for some useful discussions). It depends on the context in which the transcription is to be used, what it is to be compared with, and the purpose for which it is being used, as well as on the skill of the phonetician making it, and their knowledge of the language being transcribed. It is important to consider the context of use, to choose justifiably from among the large range of possible transcription systems the one that best fits the purpose, and to stick to the chosen level throughout the whole transcription. The PP framework can help achieve the most appropriate level and detail of transcription, and consistency in applying it. A little more will be said on this below.

9.4. Summary

The PP framework presented in this section can be summarised with a diagram. This has enough parallels with the GP framework presented earlier to allow comparison of the two theories.

10. A programme for phenomenological phonology

The main value of the analysis presented so far is that it encourages a rethinking of the need to spend a large proportion of phonological research energy on specifying a transformational relationship between the phonological and phonetic representations.

However, having developed the framework of PP for this purpose, it seems worth exploring whether it might constitute an interesting way of doing phonology in its own right. It certainly raises questions, and provides methods and criteria for answering them. Though these are all significantly different from the questions, methods and criteria of GP, they have some validity in their own right and might lead to a fruitful research programme.

Some questions which arise in this framework are:

- As what kinds of Object do Subjects represent or describe speech sounds?
- What are the descriptions of speech sounds which are most useful to linguists?
- What are the processes of constitution of an Object from something not already described like?
- What generalisations can we make about the sound systems of languages?
- What is the best description of a particular language for a particular purpose?

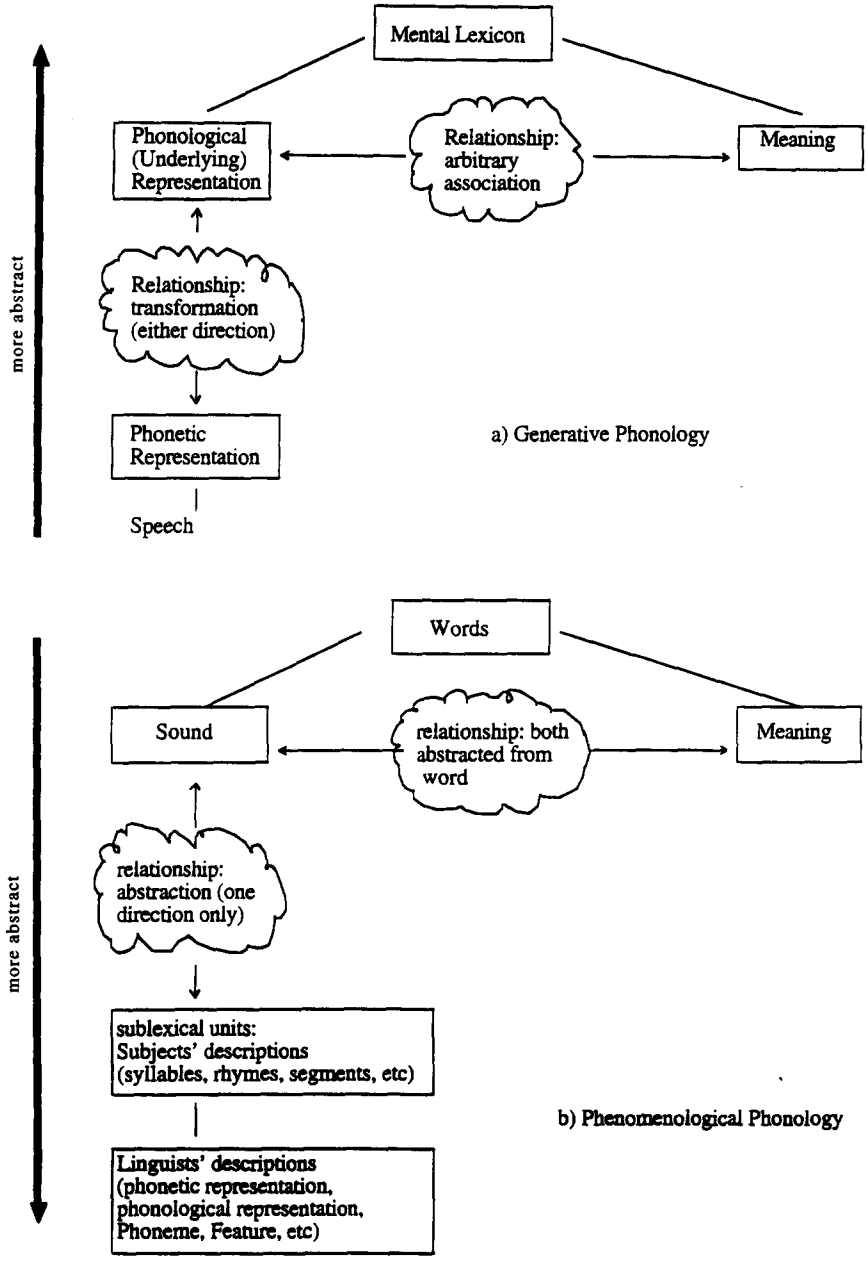

Fig. 2. A comparison of generative phonology and phenomenological phonology.

In answering these questions, rigour in PP is achieved by constant reference of phonological experience and observation to the big picture of the PP framework—as opposed to allowing theory to 'run away with itself', or the formalisms to become more of an object of study than the reality, criticisms often levelled at GP (e.g. Coleman and Local, 1991). This is in fact

exactly the purpose and value of having a framework such as the one presented above: it allows a more hermeneutic theory development (see SSP).

So, as an example, while it is of course useful and necessary to develop 'ways of talking' about phonological matters (linguists' Objects, in fact), it is important not to allow these ways of talking to wander too far from an understanding of the Something, the Objects and the characteristics of speech as they are experienced by language users. For example, we can talk about a 'sound system' of a 'language', and we can use the concepts of 'double articulation' and segments being combined according to rules of co-articulation. But we should not lose sight of the fact that in the Subject's experience, as opposed to linguists' theories, there is no such thing as a sound system, or even, arguably, a language: rather what we have in experience is 'people talking', and linguistically meaningful sounds. It is from understanding of these that we constitute the more abstract Objects of speech sounds and sound systems. A little more will be said about these processes of abstraction and the relative abstractness of Objects below; see also Fraser (1995, 1996a).

There is no doubt that work in the PP framework is in its early stages, or that its eventual credibility will depend as much on its practical achievements as its theoretical cohesion. In the meantime, the following sections give a little more detail on what PP might look like. It is perhaps worth pointing out that it is clear that there are many things that GP can do that PP will never be able to do, and is not intended to do. Criteria for 'success' in PP are the attainment of insights and descriptions which are useful for practical tasks, and which help illuminate and tie together observations made in a variety of different domains (e.g. related areas of psychology, philosophy and linguistics), not an ability to do what can be done in GP.

It will be observed that PP focuses more on Subjects' perception and understanding of speech and speech sounds than on speech production. In PP the two processes are seen as being significantly different—though no doubt sharing common knowledge to some extent. Future work will look in more detail at the processes of speech production in a PP framework. (Linell's work by contrast focuses more on production than perception, which may be one reason for differences between his approach and PP.)

10.1. Psychophonological description

A psychophonological description (PsD) is a description of speech sounds made by the language-using Subject, as opposed to descriptions of speech made by the linguist in a reflective mode. A PsD can be defined as 'any description of speech sounds made by a LUS'.

The first observation is that for any individual LUS there are many psychophonological descriptions. For example there is description in terms of words, letters of the alphabet, intonation patterns, rhythmic structures, various sublexical units, and so on. The PP framework allows us to place these many psychophonological descriptions in a hierarchy of abstractness according to the processes of abstraction which are needed to constitute them (see Fraser, 1995, 1996a).

As well as the descriptions of speech sounds that are mentioned spontaneously by the LUS, it seems clear that there is an underlying knowledge of speech sounds and their relationships which is rarely brought to consciousness by the LUS. The evidence commonly adduced by GP to justify the kind of rule-governed linguistic competence it postulates is certainly valid as evidence that Subjects know how to use speech sounds. For example, given a new word, someone who 'knows the language' can form the correct plural, find the correct relationships with other words, and so on. However from the perspective of PP, this evidence does not

immediately and unarguably imply the existence of GP-style rules, or validate the formal computational 'competence' of GP. The choice is not a forced one between a knowledge-less Subject and a computational one, as it often portrayed in introductions to GP.

PP is very careful about the danger of confusing linguists' Objects with Subjects' Objects. The Subject knows about linguistically meaningful sounds and their relationships, but may or may not create Objects based on comparison of these linguistically meaningful sounds. The Subject's knowledge is usefully summarised by the linguist with 'sound system' terminology, but it remains an open question to what extent the Subject 'has' a sound system in this sense. To the extent that it does, it is very unlikely, according to PP, to be a logically regular formalistic sort of system (*cf.* the library analogy above).

It has been argued already that the word is the basic, most 'concrete' description from the LUS's point of view, with sublexical units derived from understanding and comparison of words. This change of perspective has important consequences for the concept of co-articulation as understood in the theory.

Consider again the 'double articulation' view of language that underpins GP. This view is useful to linguistics, but from the PP perspective very limited, raising as it does questions about how the Subject could ever apprehend these small meaningless units of sound and constitute them as Objects.[30] In PP, these small meaningless Objects are abstractions from meaningful words. This difference between GP and PP makes for some interesting shifts of perspective in research.

The 'double articulation' view of language raises questions about how the small meaningless units of language are joined together to create larger meaningful units, leading to the vast research on co-articulation. In PP by contrast, the obvious question is 'How do smaller units get created from the basic linguistically meaningful units?' Considering the rules that would be necessary to join phonemes together into a smooth flow of continuous speech is still useful in PP, for practical tasks such as speech synthesis. But there is no temptation to look for the rules devised by linguists undertaking such a practical task in the minds of Subjects.

When we come to investigate the LUS's use of speech from a PP perspective, there is more focus on questions like those raised by Kenstowicz at the beginning of PGG: questions like how Subjects ever come to consider the 't' at the beginning of 'top' and the 't' in 'stop' to be the 'same sound'. Many of the hairy theoretical issues around co-articulation in GP-based phonetics and phonology (*cf.* Beckman, 1988; Daniloff and Hammarberg, 1973; Fowler, 1980; Hammarberg, 1976; Kent and Minifie, 1977; Ohala, 1993, p. 229ff) could be readily sorted out, I claim, by distinguishing these two points of view—that of the linguist and that of the Subject—as suggested by the PP framework.

As well as looking at the psychophonological descriptions that are created by an individual generalised Subject, it is also possible for PP to compare the *characteristics* of PsDs across a range of LUSs. While it seems likely that there are aspects of any LUS's PsDs which are idiosyncratic, it also seems to be the case that significant generalisations can be made about Subjects who use a particular language. It is here that the concepts of 'languages' and their 'sound systems' are particularly useful to the linguist, though with all the cautions on use of these concepts mentioned above.

Many of the well-known observations of linguistics can be set in the PP framework in a way which invites further investigation. Speakers of some languages (such as English) seem predisposed to constitute speech sound Objects through a linear analysis of words. Speakers of other languages do things differently; for example, speakers of Semitic languages constitute 'consonant frames'; speakers of tone languages focus on pitch contour in a certain way, and so on.

Similarly when it comes to counting the units in a linguistically meaningful sound, speakers of different languages behave differently. In some cases, the most natural unit is the syllable, in others the segment. In some cases, glottal stop is not counted as a unit, in others it is; similarly with glides. For example, the same phonetic sequence, characterisable by the linguist as 'stop-glide-vowel', can be given quite different psychophonological descriptions by speakers of English (two consonants and a vowel), Korean (a consonant and a special kind of vowel) or Irish (a special kind of consonant and a vowel).

There is a large quantity of information like this about many languages in the personal knowledge of many linguists. Only a small fraction of this information gets to be focused on as part of phonological theory if the main energy is directed towards technical questions, and attention is focused on a very few sets of very small fragments of the sound systems of languages (*cf.* Copeland, 1994; Love, 1992). Gathering together such information and formulating it into a PP framework, distinguishing the phonetic Something from the particular description it is given, provides a useful basis for many tasks in language teaching and the understanding of foreign accents. In turn, information gained from Subjects through applying this method will be very valuable to PP in refining its conception of abstraction and constitution.

10.1.1. Methods for investigating psychophonological description. In method, PP has some affinity with the various experimental movements in GP mentioned above. In determining the nature of a PsD it is not sufficient to use 'corpus-internal generalisations'. It is essential to observe the linguistic behaviour of Subjects, and more importantly to discuss with them what they think about the speech sounds being investigated.

In this respect PP is very different from GP, which is wary of 'introspection' and suspicious of speakers' own ways of talking about language.[31] In PP, certainly Subjects' ways of talking are not accepted at face value as an accurate reflection of the nature of the PsD. However, the Subject's conceptualisations of speech sounds are respected as a starting point for discussion, rather than being dismissed as 'illusions' or errors.

Another important avenue towards understanding the PsD is study of the phonetic Something from which Subjects are creating their speech sound Objects. This can be done by attention to the phonetic description.

10.2. Phonetic description

The phonetic description (PhD) can be defined in PP as any description of speech made from the linguist's point of view—that is, with an attitude of detachment (observing the language, rather than being actively engaged at that moment in using language to communicate linguistic meaning), and with an interest in systematising the units of speech sound and their relationships. All the while the linguist also of course has some understanding of the units as linguistically meaningful sounds, but this understanding is in the background, and the more scientific attitude is in the foreground.

Phonetic description, in PP, is not considered simple and uninteresting; there is no sense that phonetic description is 'mere' observation. In fact, for PP the issues involved in providing a relevant phonetic description are if anything more tricky that those of the PsD, where there is always the possibility of returning to the 'base level' of the Subject's own naming of the linguistically meaningful units of the language.

This might be a good place to consider briefly the relationship between PP and the 'concrete phonology' often used as a 'straw man' with which to contrast to GP (PGG, p. 72ff; Katamba,

1989, p. 137): an approach which recognises only phonetic representations in the GP sense described above, and gives no place to a UR or transformational rules—suggesting to GP theorists a lack of appreciation of the knowledge and mental processing necessary to using language.

Such a theory, if anyone holds it, is quite different from PP. PP sees the phonetic description as a highly abstract linguist's description which has no status in the mental processing of language by the LUS. The 'concrete' level in PP (though of course this is only relatively less abstract within the framework; PP does not believe in the idea of 'concrete description' in general) is the psychophonological description,[32] which is the one that has reality in the immediate experience of the LUS (compare the abstract nature of the phonetic description: it takes work by a linguist even to persuade a naïve LUS of the existence of a phonetic description separate from the PsD).

We have seen then that in PP, the phonetic description is a linguist's construct with no status as part of the 'processing' Subjects carry out in using speech. Nevertheless, the PhD is an important element of PP, since it helps the linguist gain an understanding of the phonetic Something from which the Subject creates the various Objects of Psychophonological Description.

This is not to contradict the view that the phonetic description is not a direct representation of 'real speech'. Certainly the phonetic description is an abstraction, created from a particular Subjective point of view. But there can be many phonetic descriptions, and PP is interested in them all. Working with them can increase the linguist's understanding of the phonetic Something that lies behind them all and also of the presuppositions and other aspects of the Subjective viewpoint from which they are constituted. Becoming a good phonetician then is a matter of taking various descriptions of speech and considering the contribution to each made by the Something and the Subject—especially one's own Subjective viewpoint as a linguist.

What is not considered useful in PP is the 'solidification' of phonetic description into units or symbols, to the extent that one is working more with the symbols or formal descriptions than with the Something. This is an accusation which has been levelled at GP, with its focus on formalisation and symbolisation, and its reduction of all the phonetic richness of language to a few Distinctive Features. In PP it is considered useful to transcribe or formalise speech, but important to return often to hear the actual speech being represented by the symbols. It would be very unlikely that a PP conference presentation, for example, would be given without a recording or detailed phonetic transcription of the speech being discussed. GP talks on the other hand often present an entire argument based on highly abstract representations of speech—often of a language which has never been heard by members of the audience. In PP we are interested in gaining understanding of the phonetic Something of speech, and use symbols and descriptions to this end (here there is considerable affinity between PP and Kelly and Local, 1989; *cf.* also Coleman and Local, 1991).

There is little place in PP for the Distinctive Features of GP, whether arranged as matrices or in other geometries, and indeed the whole notion of the 'features' of speech is quite different in PP. The Something certainly has characteristics or features which can be observed, compared, abstracted and Objectified. Once the characteristics or features have been Objectified in this way however they are Objects, not features of a Something. In PP it is considered important to give the correct ontological status to the elements of the theory: features are features, rather than entities (*cf.* similar point in Linell, 1979). PP is thus in strong disagreement with the GP view expressed succinctly by Basbøll:

> As soon as we operate with natural classes of segments, we in fact also operate with distinctive features characterising these classes, at least implicitly. (Basbøll, 1988, p. 196)

Two very useful analogies from the GP literature allow me to explain the difference in attitude to features between GP and PP. In the first, Anderson (1974) likens the Distinctive Features to the members of an orchestra, who, when they play together, create the music—which is thus conceptualised as a whole made up of many parts. Such an analogy does not hold good in PP however. In speech, the basic experienced unit is the word, and the sub-lexical units are derived from this by abstraction. In an orchestra, the basic experienced unit is the person, and an orchestra is made up by combining these units.

Bird (1995) makes a very imaginative analogy of features with the primary colours, which, when combined, make up a range of the actual colours of experience. While this is an extremely useful analogy for explaining GP's view of how DFs work as the basic units of sounds, it also highlights the quite different understanding in GP and PP of the relationship between things encountered in experience, and scientific analysis of those things. In the phenomenological view, primary and secondary colours are equal in experience; the knowledge that some colours can be created by combining other colours is the product of a detached or scientific appraisal. It would be rather unusual to suggest that in order for a Subject to experience the colour orange they must at some 'lower' level make a representation of the colours red and yellow.

Overall then PP has a strong commitment to detailed phonetic observation and description, and is particularly interested in issues of transcription in various styles and at different levels. In this respect PP has some affinity with those branches of GP which consider 'accountability to the data' to have a high priority. As Kelly and Local (1989) point out, because of the focus in GP on formalisation and on analysis of small and often unrealistic bits of speech, there is still much descriptive data to gather, even for the best-studied languages, like English.

PP's commitment involves not just detailed observation of speech but also attention to the question of the appropriate level of transcription of the data. It is generally not appropriate to include all the observable detail in a transcription. What is important is that the Subjective viewpoint, the level of abstraction and the degree of detail should be consistent throughout a transcription. It would be very unlikely for example that a PP account of speech sounds would include transcriptions in a mixture of orthographic and phonetic symbols. It may be interesting to look at some issues of traditional phonetic terminology from this perspective; its use often involves a mixture of viewpoints (e.g. articulatory versus acoustic) which may not be the most appropriate (*cf.* Docherty and Fraser, 1992).

10.2.1. Methods for giving phonetic descriptions. The aim of creating phonetic representations of speech is to increase one's understanding of the nature of speech sounds as a Something. This has value not just in PP, in helping to understand the way the Subject Objectifies this Something, but also as technical expertise in a range of areas such as speech technology and language teaching.

To do this it is necessary to learn to take yourself and your own linguistic prejudices out of the picture as much as possible, so that you are working as much as possible with a *phonetic description* and as little as possible with your own particular *psychophonological description*.[33] This is difficult to do; in fact it is arguably impossible to attain this goal. However it is possible to improve one's skill in making phonetic transcriptions. The starting point is to come to realise the contribution being made by your own background understanding in making any representation of speech. The use of laboratory techniques is very helpful in developing this skill.

The methods used in PP for achieving this understanding involve firstly direct observation

of speech sounds collected from a wide range of sources and analysed with a wide range of laboratory techniques; and secondly consideration of a wide range of transcription conventions and issues, and how they can be related in the PP framework. Phonetic transcription is a focal interest in its own right, not just a side issue or preliminary.

10.3. The process of constituting the psychophonological description

The emphasis in PP is on the constitution of phonological Objects from the undescribed 'Something' of speech, rather than on the transformation of one Object (described by the linguist in physical terms) into another (described by the linguist in 'meaning' terms).

The basic process involved in the constitution of Objects is seen in PP as one of abstraction. In abstraction, the Subject attends to certain characteristics of the Something, and has those more clearly in focus than other characteristics of the Something. These attended characteristics can then be Objectified by the Subject, or, in other words, they can be constituted into an Object. This is the only way that Objects can be created, through the agency of a Subject, and it is the reason why in PP we say that the existence of any Object presupposes some prior activity on the part of a Subject. It is also the reason why we say that any Object is abstract; none is a direct or 'concrete' representation of the Something.

However, while it is true that all Objects are abstract, some Objects are more abstract than others.

Objects created by the process just described can be compared and contrasted by the Subject. In doing this, other characteristics are brought to light, and these in turn can be Objectified. This process can be iterated many times (see Fraser, 1995; 1996a, for more detail; also Ellis, 1993, for a similar argument). Thus we have the basis of determining a hierarchy of abstractness for Objects: those closest to the level of experience are least abstract; those that require understanding and comparison of many layers of Objects and a very particular Subjective point of view are more abstract.

In determining the nature of the PsD in PP, we are more interested in the lower levels of abstraction of speech sounds, in the constitution of phonological Objects from the Something of speech. If two bits of speech 'Something' are Objectified in the same way, there is no assumption that there must be a third representation underlying them both from which they are both derived. Rather the questions are asked: 'What causes these two distinct "Somethings" to be classified as the same Object?'; 'What factors influence the creation of one PsD rather than another?'

The PP framework suggests a three-way classification of the factors that can influence the constitution of the PsD:

- The nature of the 'Something' that is being named.
- The context in which that 'Something' is viewed.
- The nature of the Subject making the representation, particularly, in the case of phonology, the Subject's knowledge of the language.

Of these, the first two, while crucial in the sense that without them there would be no Object, seem, interestingly, to be of less influence than the third. PP focuses a lot of attention on the Subjective qualities that influence the constitution of speech sound Objects, and particularly on the question 'What does the Subject need to know in order to be able to understand speech?' The similarities and differences of this PP question to the analogous one in GP have already been addressed.

One piece of knowledge that seems particularly important to the Subject in constituting sub-lexical Objects is knowledge of a writing system.[34] The writing system seems for example often to be a major factor in determining the way the flow of speech is segmented into units. This poses a large number of interesting theoretical questions, and is certainly an area in need of far more research than it has so far received. It seems the PP framework might be a useful one for this type of work.

It is interesting to observe here that very often it seems that what a person is 'told' about speech sounds by the writing system given by the language is of more influence in the process of Objectification than what is 'experienced' by that Subject. This seems to be a general feature of much of our everyday knowledge; indeed it is this characteristic which motivated the early phenomenologists to press strongly for a return to 'the phenomena', which were more reliable than the 'facts' of science and schooling (see SSP and references therein).

10.4. The relationship between psychophonological and phonetic descriptions

I have been at pains to argue throughout that the psychophonological and phonetic descriptions are not related by transformation from the point of view of the language-using Subject. However this is not to say that there is no relationship at all between them. Far from it. They are after all Objects derived from the same 'Something', and specifying a relationship between them is an important goal of PP.

Of course, it is important to see this relationship for what it is. The two representations are related, not directly but through the 'Something'. To translate one into the other is not a matter of applying a computational process. It is a genuine translation in the sense of translating of one language to another. Just as, to translate (effectively) from Japanese to English it is not sufficient to keep each Japanese Object and simply transform it into a corresponding English Object, so the same strategy fails as a method of translating the PsD into the phonetic representation. In language, it is necessary to *understand* the Japanese sufficiently to go back to the meaning behind the words and reconstitute that meaning in English words. Similarly, to translate one description of speech into another, it is necessary to get behind the description and come to some understanding of the 'Something' from which it is constituted; from there it is possible to constitute another description of the same Something, by taking a different Subjective point of view, or seeing it in a different context.

The ability to do this is valued by the phenomenological phonologist, and has many practical uses. Not least of these is the attainment of a set of cross-linguistic generalisations about the nature of the sound systems of languages. Despite the fact that a major goal of GP is the specification of linguistic universals, it is remarkable how little general knowledge about the nature of sound systems of languages is represented in textbooks or the theoretical literature[35] (Hammond, 1994a, 1994b; Maddieson, 1984; Ringen, 1994; Kaisse, 1990; Greenberg, 1978;[36] Ferguson, 1978).

The PP framework suggests a way of systematising the notion of allophonic and phonemic descriptions, with which all students start studying linguistics, in a relatively theory-neutral way, and considering the relationships between these two levels of description across a wide range of languages. Because the framework gives a very clear understanding of these two descriptions as linguists' Objects, there is no need, or temptation, to start formalising the relationship between them in a computational way, as if they were representations in the Subject's competence, related by transformation. The focus can remain clearly on using these two descriptions as the basis of comparison across languages. One could hope that the result

might be as useful as Ladefoged and Maddieson (1996), in which the concept of 'segment' is used so successfully as a linguists' unit for the practical comparison of languages, without entering into discussion as to its status as an Object for the LUS.

Such a work (*cf.* Keating, 1988, p. 282, who calls for a similar study) would provide data from which a wide range of extremely interesting questions might be addressed, and which could form the basis of theory-development in a wide range of directions. The following are a selection of such questions, to which answers are currently unavailable:

- What proportion of allophony can be accounted for as phonetic assimilation?
- What characteristics do *exceptions* to phonological rules show?
- Are different kinds of phonological relationships associated statistically with different kinds of semantic or morphological relationships? For example, there is evidence in the phonological systems of Tamil, Korean and some other languages that fortition is associated with intensification and transitivity—to what extent is this association general across languages?
- What are the very *uncommon* phonological rules, as opposed to the impossible or unattested ones?
- To what extent do phonological systems relate phonetically similar sounds as opposed to more dissimilar ones?
- What is the 'setting' of a particular process within the phonological system of whole languages? Chomskyan phonologists most often discuss individual processes in isolation from the system as a whole, whereas PP is interested to see how that process fits into the whole system.

Such a database would also form a useful extension of the admirable work by Maddieson (1984), which used the 'common-sense' linguist's Object, the phoneme, much in the way advocated above, as the basis of comparison of the phoneme inventories of a wide range of languages.

Interestingly, this work threw up quite a number of observations that run somewhat counter to the expectations prompted by a GP understanding of the sound systems of languages. A generally GP attitude to language would suggest that sound systems would be fairly regular at the level of the underlying description of the segment or feature. Maddieson's work confirms the well-known fact that this is true only up to a point.

A roughly GP perspective also suggests that the characteristics of sound systems might reflect an 'efficient information-transfer' aspect of language—say an optimal number of basic elements, or certain kinds of trading relationships among the units. However, such expectations are not confirmed: 'Languages really do differ in phonological complexity' (Maddieson, 1984, p. 23).

The PP perspective on the other hand does not raise these unfulfilled expectations in the first place, since the units of sound systems are seen as abstract descriptions produced by the LUS, on the basis of the factors outlined above. There is no temptation to think of them as 'building blocks' in a linguistic system which is 'efficient' in this information-theoretic sense.

11. Conclusion

Let me be the first to acknowledge the two most likely criticisms of this work.

First, there is no doubt that the PP framework contains many gaps and much speculation. It is presented as a starting point for future work, rather than a completed project. I am sympathetic to those self-styled pragmatists who reserve judgement on the theoretical framework

I have offered until they see 'what can be done with it'. Personally however I find value in conceptual clarity for its own sake, and observe that much can be lost by rushing too quickly to 'practical applications'. I espouse Einstein's dictum, 'There is nothing so practical as a good theory'. That is the reason I find it important to build up a broad framework before filling in the details—in the hermeneutic style—rather than concentrating first on details and practicalities.

Second, it is likely that some will find my criticisms of GP to be less well-founded than they might be. Surely there are defences that could be made against my critique of the GP project of transforming an underlying representation into a phonetic one. In this respect it should be remembered that the criticisms of GP raised here are not intended to help GP practitioners lift their game, but to help me base my explanation of my own perspective on phonological issues in a set of ideas familiar to most readers. Those who have not felt my dissatisfactions with the GP project are unlikely to be persuaded by my criticisms, but they may be better able to understand PP because of them.

On the positive side, it would seem important to judge PP in the light of the claims made on its behalf: that it gives a new perspective on old issues, that it raises questions that are worth pursuing, and that it offers a genuine alternative to those people who, like myself, are simultaneously fascinated by the subject matter of phonetics/phonology and repelled by the Chomskyan view of Subjectivity. Work on the PP framework continues in Fraser (forthcoming).

PP does not and cannot improve upon GP in the sense of doing the same thing better. It can do something well that is different but worthwhile.

Acknowledgements—This paper has been several years in the writing, and has input from many people, for which I am grateful. I would like to thank the following people particularly: Paul Foulkes, Cliff Goddard, Brian Byrne, Bruce Stevenson, John Hajek, Nick Reid. I would also like to acknowledge the support I have received from my colleagues at UNE, and financial assistance from two internal Faculty of Arts grants and a small ARC grant.

NOTES

[1]See, for example, Dinnsen (1995) and Szpyra (1995).

[2]That we know without ever being told what the missing information is in sentences like 'John is too stubborn to talk to', versus 'John is too stubborn to talk', and so on.

[3]I am of course using the term 'Generative Phonology' rather loosely here to mean any phonological theorising which takes SPE as its point of departure, including some which may differ from SPE in quite a few respects, and even some which is more commonly called phonetics (*cf.* Keating, 1988, p. 299; Goyvaerts, 1994a, p. 1388; Dinnsen, 1995, p. 132).

[4]I avoid use of the common term 'surface representation' because it is ambiguous, being used sometimes to refer to the phonological representation ('surface' in relation to the 'Deep Structure', or equivalent), and sometimes for the phonetic representation ('surface' in relation to the 'Underlying Representation').

[5]'... an approach which only recognises phonetic representations ... fails to capture significant generalisations concerning relationships among allomorphs of a morpheme.' (Katamba, 1989, p. 137; see also the similar point in PGG, p. 72ff).

[6]In principle these rules are intended to be able to transform either representation into the other, but in practice emphasis is given to the 'production' direction. In the present discussion I will focus on the production direction, as is common in GP.

[7]'Phonology is the component of our linguistic knowledge that is concerned with the physical realization of language. Possession of this knowledge permits us to realize words and the sentences they compose as speech ... and to recover them from the acoustic signal.' (PGG, p. 2)

[8]'The phonological representation is more abstract. ... Phonologists believe that it is essentially the form in which the lexical item is stored in memory.' (PGG, p. 7)

[9]'The two representations are systematically related by phonological rules that delete, insert, or change sounds in precise contexts.' (PGG, p. 7)

[10]It is impossible in the space available to make detailed reference to the historical development of GP. Rather I have used general 'state-of-art' discussions, such as Carr (1993), Brame (1972), Goyvaerts (1981), Asher (1994), Anderson (1985) and Newmeyer (1988).

[11]'. . . lexical representations and a system of phonological rules are chosen in such a way as to maximise a certain property that we may call the 'value' of the grammar, a property that is sometimes called 'simplicity' . . .' (SPE, p. 296)

[12]'. . . data and measurement only exist by virtue of decisions about what to represent and how to represent it. Decisions about the most concrete levels of representation can have consequences for our understanding of entire aspects of sound structure. Careful thought about representation at the lowest level clarifies the relationship between phonology and phonetics by clarifying the character of the rules and intermediate representations needed.' (Pierrehumbert, 1990, p. 378)

[13]'. . . as a practical tool for linguistic analysis (Ladefoged); as a cognitive representation in speech perception (Nearey); as a gestural plan underlying speech production (Browman and Goldstein); as one or more levels at the level where grammars interface to the physical world (Keating, Rischel); and as a statement of purely non-cognitive aspects of the speech event in the physical world (Pierrehumbert).' (Nolan, 1990, p. 453)

[14]The fact that it is necessary to give reminders about the abstract nature of the Phonetic Representation (e.g. Goyvaerts, 1994b, p. 3561) is evidence for this state of affairs.

[15]'The phonological component of the grammar assigns a phonetic interpretation to the syntactic description, making reference only to properties of the surface structure [i.e. the UR], so far as we know.' (SPE, p. 7)

[16]Most phonology textbooks do not define the term 'representation' at all. Trask (1996) gives only 'Any conventional rendering of a piece of speech with a set of symbols or linguistic objects appropriate to some particular level of linguistic analysis', and his treatment of 'abstract' and 'underlying' are hopelessly circular. Goyvaerts (1994b) gives a useful description of the various levels of representation used in phonology, but no general discussion of what a representation is, or why these particular levels are chosen by GP. Rischel (1990) is the only writer in Phonetics/Phonology I have come across who actually raises the questions of what a phonetic representation is and how it relates to phonological representation—but he unfortunately remains content simply to pose the questions, without attempting to answering them.

[17]Briefly foreshadowing later discussion here, for those who are familiar with the issues, my concern is that all such models inevitably presuppose the activity of an unacknowledged homunculus. A homunculus is a 'little man' presupposed but not explicitly recognised by theories such as the one that suggests that vision works by means of an image of the outside world being projected onto the retina. Such a theory is said to beg the question of 'who is looking at the image on the retina'. The computational view of mind is far more sophisticated than this of course, but is still susceptible to the same criticism: 'If the mind is a computer program, who reads the printout?' SSP also explains why Dennett's notion of 'discharged homunculi' does not allay my worries.

[18]Chomsky and Halle's assurance that '. . . the lexical representation is abstract in a very clear sense; it relates to the signal only indirectly, through the medium of the rules of phonological interpretation that apply to it as determined by its intrinsic abstract representation and the surface structures in which it appears.' (SPE, p. 12) is hardly a satisfactory treatment.

[19]Though it is sometimes hard not to question the sense in which 'practical' is used in GP theory.

[20]'Subjective' means 'relating intrinsically to the viewpoint or activity of a Subject' (see further discussion below). It should not be confused with the everyday use of the word 'subjective', which means 'relating to the viewpoint of a particular person'.

[21]This discussion is not an attempt to define 'representation'; it if were it would be a very bad one, as it uses the word 'represent' three times as part of the definition. Rather it is an attempt to unpack some of the logical implications and presuppositions of using the concept of representation in a theory of phonology. Defining what the word 'represent' *means* is a very difficult task best left to a semanticist or lexicographer, and certainly well beyond the scope of this paper.

[22]A fourth role is also relevant: the representation can represent something *to someone*. We will not be concerned with this aspect here—though it is worth exploring—but simply assume that the representer is representing to him or herself.

[23]It is not necessary to claim actual existence for the battalion: the colonel could be representing a hypothetical battalion that he merely wishes he had. This is not the place for a discussion of general ontological issues; we will limit discussion to the ontological status of representations.

[24]'Until recently, realism of representation has not been much challenged in linguistics, partly because of an uncritical acceptance of the communication or channel metaphor in which language is a conduit for conveying ideas unaltered from one brain to another'. Compare other disciplines which have an alternative conception of representation as a process of construction: 'Representation, which is always from some specific point of view, is not just an innocent process, but a social practice, conscious or not.' (Fowler, 1994, p. 3358)

[25]At least, this is the intention. We will point out later the confusion that arises from a situation in which a linguist is making a description of a Subject's representation.

[26]I have often been criticised for the lack of a formal definition for this central concept of my theory, and no doubt some readers will feel the proffered definition to be inadequate. In response I make three points: (a) most of the key units of GP do not have adequate formal definitions either; (b) most of the key units of GP actually rely on prior understanding of meaningful words for their definition; just try to define phoneme, feature, syllable or any other unit without making any reference to the word; (c) I am not convinced that a formal definition is the most appropriate

for this domain of discussion; clarity of expression is the important thing, and formal definitions are useful only insofar as they enhance communication. I am using the word 'word' here as an informal way of talking about such units of language in ordinary familiar English. The difficulty of determining the precise boundary between what is a word and what is not—cases such as the contracted form of 'is' in 'It's raining'—is not of major concern. There is an empirical claim here that all languages will have some way of referring to linguistically meaningful units of sound. Obviously in many cases this will not be an exact equivalent of the English word 'word', since many languages are organised differently. However I would be interested to hear of any language which has no way at all for native speakers to refer to any linguistically meaningful unit of sound.

[27]This is not to suggest that the meaning is basic and is somehow converted into the sound, as it has sometimes been misunderstood! The first representation of the sound-Something made by the Subject is as linguistically meaningful sound. The sound and the meaning are not separated until later—see below.

[28]See Vihman (1996) for an extended discussion of these ideas.

[29]It is interesting that Saussure, who showed us the fallacy of thinking that language is a set of pre-existing sounds attached to a set of pre-existing meanings, was also a strong proponent of the arbitrary relationship between them.

[30]It is a fundamental tenet of Phenomenology that the Subject can only apprehend aspects of the world that are meaningful in experience. It is possible for the Subject to constitute 'meaningless' Objects, such as phonemes and Distinctive Features, but only by abstraction from meaningful ones.

[31]In fact, as has often been pointed out, GP seems to have a rather ambivalent attitude to introspection, disallowing most of the descriptions of speech made by Subjects, and yet placing much faith in certain kinds of 'grammaticality judgements'. See Fraser (1996a).

[32]It is interesting that by assuming the Phonetic Representation to be the obvious or concrete one, GP subscribes to a particular ontological hierarchy in which description in physical or logical terms is considered more concrete than description in the terms of Subjective experience. This is quite antithetical to Phenomenology.

[33]It is well known that the language one learns as a child has an enormous influence on the way speech sounds are perceived. It takes a major effort to lessen the prejudices about speech sounds given to one by one's language.

[34]This observation is the basis of may work on dictionary pronunciation guides—see Fraser (1996b) and Fraser (in press).

[35]By far the greatest amount of knowledge in this area falls into the category of received wisdom, handed down from teacher to student and propagated through the literature. Sometimes this information is set out more systematically in a textbook (e.g. Katamba, 1989, chapter 5) for the benefit of students; but to serious researchers such treatments raise more questions than they answer, and serve mainly to emphasise the piecemeal and ad hoc nature of the knowledge we have about phonological systems.

[36]This book presents the work of the Stanford Universals Project, which, while admirable, highlights the fact that most of the issues need far more detailed treatment than can be given in a chapter of a book. For example, Ferguson (1978), despite its encouraging title *Phonological Processes* investigates just two processes—alveolar stop to dental fricative and the reverse—in six languages, and concludes that the question of whether to classify them as lenition or fortition is very unclear; Bhat (1978) similarly emphasises the multifarious nature of the process of palatalisation and the difficulty of stating its defining characteristics.

REFERENCES

Anderson, S. (1974) *The Organisation of Phonology*. Academic Press, New York.

Anderson, S. (1985) *Phonology in the Twentieth Century*. University of Chicago Press, Chicago.

Asher, R. (Ed.) (1994) *Encyclopedia of Language and Linguistics*. Pergamon Press, Oxford.

Basbøll, H. (1988) Phonological theory. In Newmeyer, F. (Ed.), *Linguistics: The Cambridge Survey, Volume 1, Linguistic Theory: Foundations*, pp. 192–215. Cambridge University Press, Cambridge.

Beckman, M. (1988) Phonetic theory. In Newmeyer, F. (Ed.) *Linguistics: The Cambridge Survey, Volume 1, Linguistic Theory: Foundations*, pp. 216–238. Cambridge University Press, Cambridge.

Bhat, D. (1978) A general study of palatalisation. In Greenberg, J. (Ed.), *Universals of Language*, pp. 47–92. Stanford University press, Stanford.

Bird, S. (1995) *Computational Phonology*. Cambridge University Press, Cambridge.

Blumstein, S. E. (1991) The relation between phonetics and phonology. *Phonetica* **48**, 108–119.

Brame, M. (Ed.) (1972) *Contributions to Generative Phonology*. University of Texas, Austin.

Broselow, E. (1994) Generative phonology: the abstractness controversy. In Asher, R. (Ed.), *Encyclopedia of Language and Linguistics*. Pergamon Press, Oxford.

Browman, C. P. and Goldstein, L. (1992) Articulatory phonology: an overview. *Phonetica* **49**, 155–180.

Bybee, J. (1994) A functional approach to some phonological phenomena. Australian Linguistic Institute, La Trobe University, Melbourne, Australia.

Carello, C. *et al.* (1984) Inadequacies of the computer metaphor. In Gazzaniga, M. S. (Ed.), *Handbook of Cognitive Neuroscience*. Plenum Press.

Carr, A. (1994) *The Only Way to Stop Smoking Permanently*. Penguin, Harmondsworth.
Carr, P. (1993) *Phonology*. Macmillan.
Chomsky, N. and Halle, M. (1968) *The Sound Pattern of English*. Harper and Row, New York.
Clements, G. N. (1994) Non-linear phonology. In Asher, R. (Ed.), *Encyclopedia of Language and Linguistics*. Pergamon Press, Oxford.
Cohn, A. (1993) Nasalisation in English: phonetics or phonology. *Phonology* **10,** 43–81.
Coleman, J. and Local, J. (1991) The "no crossing constraint" in autosegmental phonology. *Linguistics and Philosophy* **14,** 295–337.
Copeland, J. (1994) Unmotivated free alternation in Tarahumara: the principle of emergence in phonology. *Language Sciences* **16,** 213–227.
Daniloff, R. G. and Hammarberg, R. E. (1973) On defining coarticulation. *Journal of Phonetics* **1,** 239–248.
Davis, H. and Taylor, T. (Eds) (1990) *Redefining Linguistics*. Routledge.
Dell, F. (1980) *Generative Phonology*. Cambridge University Press, Cambridge.
Diehl, R. L. (1991) The role of phonetics within the study of language. *Phonetica* **48,** 120–134.
Dinnsen, D. (1995) Review of Michael Kenstowicz: Phonology in Generative Grammar. *Phonology* **12,** 131–134.
Docherty, G. (1992) *The Timing of Voicing in British English Obstruents*. Foris, Berlin.
Docherty, G. and Fraser, H. (1993) On the relationship between acoustic and electropalatographic representations of speech. Work in progress. 7. University of Reading, Speech Research Laboratory, p. 8–25.
Dressler, W. U. and Moosmuller, S. (1991) Phonetics and phonology: a sociopsycholinguistic framework. *Phonetica* **48,** 135–148.
Dreyfus, H. (1979) *What Computers Can't Do: The Limits of Artificial Intelligence*. Harper and Row, New York.
Durand, J. (1990) *Generative and Non-Linear Phonology*. Longman, New York.
Ellis, J. (1993) *Language, Thought, and Logic*. Northwestern University Press, Evanston, IL.
Ferguson, C. (1978) Phonological processes. In Greenberg, J. (Ed.), *Universals of Language*, pp. 403–442. Stanford University Press, Stanford.
Fowler, C. (1980) Coarticulation and theories of extrinsic timing. *Journal of Phonetics* **8,** 113–133.
Fowler, C. (1985) Current perspectives on language and speech production: a critical overview. In Daniloff, R. (Ed.), *Speech Science: Recent Advances*, pp. 193–278. Taylor and Francis, London.
Fowler, R. (1994) Representation. In Asher, R. (Ed.), *Encyclopedia of Language and Linguistics*. Pergamon Press, Oxford.
Fraser, H. (1992) *The Subject of Speech Perception: An Analysis of the Philosophical Foundations of the Information-Processing Model of Cognition*. Macmillan, London.
Fraser, H. (1995) What would phonology be like if it were based on connected speech?—An introduction to phenomenological phonology. *Work in Progress*, Speech Research Laboratory, University of Reading, pp. 13–25.
Fraser, H. (1996a) The Subject in Linguistics. In Simms, K. (Ed.), *Language and the Subject*, pp. 115–125. Rodopi.
Fraser, H. (1996b) Guy-dance with pro-nun-see-ay-shun. *English Today* **12,** pp. 28–37.
Fraser, H. (in press) Dictionary pronunciation guides. *International Journal of Lexicography*.
Fraser, H. (forthcoming) *Analysis of Speech Sounds: A Phenomenological Introduction to Phonetics and Phonology*. Oxford University Press, Oxford.
Goldsmith, J. (1989) *Autosegmental and Metrical Phonology*. Basil Blackwell, Oxford.
Goldsmith, J. (1993) *The Last Phonological Rule: Reflections on Constraints and Derivations*. Chicago University Press, Chicago.
Goyvaerts, D. (Ed.) (1981) *Phonology in the 1980s*. E. Story-Scientia, Ghent.
Goyvaerts, D. (1994a) Generative phonology. In Asher, R. (Ed.), *Encyclopedia of Language and Linguistics*. Pergamon Press, Oxford.
Goyvaerts, D. (1994b) Representation in Phonology. In Asher, R. (Ed.), *Encyclopedia of Language and Linguistics*. Pergamon Press, Oxford.
Grace, G. (1987) *The Linguistic Construction of Reality*. Croom Helm.
Greenberg, J. (1978) *Universals of Language, Vol. 2: Phonology*. Stanford University Press, Stanford.
Haiman, J. (1985) *Natural Syntax: Iconicity and Erosion*. Cambridge University Press, Cambridge.
Hammarberg, R. (1976) The metaphysics of coarticulation. *Journal of Phonetics* **4,** 353–363.
Hammond, M. (1994a) Phonological typology. In Asher, R. (Ed.), *Encyclopedia of Language and Linguistics*. Pergamon Press, Oxford.
Hammond, M. (1994b) Phonological universals. In Asher, R. (Ed.), *Encyclopedia of Language and Linguistics*. Pergamon Press, Oxford.
Harris, R. (1981) *The Language Myth*. Duckworth.
Hutton, C. (1992) Arbitrariness and rational signs. In Wolf, G. (Ed.), *New Departures in Linguistics*, pp. 250–259. Garland, New York.
Kaisse, E. (1990) Towards a typology of post-lexical rules. In Inkelas, S. and Zec, D. (Eds), *The Phonology/Syntax Connection*. University of Chicago Press, Chicago.
Katamba, F. (1989) *An Introduction to Phonology*. Longman, London.
Keating, P. (1988) The phonology–phonetics interface. In Newmeyer, F. (Ed.), *Linguistics: The Cambridge Survey, Volume 1, Linguistic Theory: Foundations*, pp. 281–302. Cambridge University Press, Cambridge.

Keating, P. A. (1991) On phonetics/phonology interaction. *Phonetica* **48,** 221–22.
Kelly, J. and Local, J. (1989) *Doing Phonology: Observing, Recording, Interpreting*. Manchester University Press, Manchester.
Kenstowicz, M. (1994) *Phonology in Generative Grammar*. Basil Blackwell, Cambridge, MA.
Kent, R. D. and Minifie, F. D. (1977) Coarticulation in recent speech production models. *Journal of Phonetics* **5,** 113–133.
King, M. (1964) *Heidegger's Philosophy: A Guide to his Basic Thought*. Basil Blackwell, Oxford.
Kingston, J. and Diehl, R. (1994) Phonetic knowledge. *Language* **70,** 419–454.
Kohler, K. J. (1984) Phonetic explanation in phonology: the feature fortis/lenis. *Phonetica* **41,** 150–174.
Ladefoged, P. (1975) *A Course in Phonetics*. Harcourt Brace Jovanovich, New York.
Ladefoged, P. and Maddieson, I. (1996) *Sounds of the World's Languages*. Basil Blackwell, Oxford.
Laver, J. (1994) *Principles of Phonetics*. Cambridge University Press, Cambridge.
Linell, P. (1979) *Psychological Reality in Phonology*. Cambridge University Press, Cambridge.
Linell, P. (1982) The concept of phonological form and the activities of speech production and speech perception. *Journal of Phonetics* **10,** 37–72.
Llewelyn, J. (1985) *Beyond Metaphysics?: The Hermeneutic Circle in Contemporary Continental Philosophy*. Humanities Press, New Jersey.
Love, N. (1992) On the need for a new departure in Phonology. In Wolf, G. (Ed.), *New Departures in Phonology*. Garland, New York.
Maddieson, I. (1984) *Patterns of Sounds*. Cambridge University Press, Cambridge.
Margolis, H. (1987) *Patterns, Thinking and Cognition: a theory of judgment*. Chicago University Press, Chicago.
Newmeyer, F. (Ed.) (1988) *Linguistics: The Cambridge Survey, Volume 1, Linguistic Theory: Foundations*. Cambridge University Press, Cambridge.
Nolan, F. (1982) The nature of phonetic representations. *Cambridge Papers in Phonetics and Experimental Linguistics*.
Ohala, J. (1986) Consumer's guide to evidence in Phonology. *Phonology Yearbook* **3,** 3–26.
Ohala, J. and Jaeger, J. (Eds) (1986) *Experimental Phonology*. Academic Press.
Ohala, J. J. (1993) Coarticulation and phonology. *Language and Speech* **36,** 155–170.
Olson, G. and Clark, H. (1976) Research methods in psycholinguistics. In Carterette, E. and Friedman, M. (Eds), *Handbook of Perception, Vol. 7: Language and Speech*, pp. 25–74. Academic Press.
Pierrehumbert, J. (1990) Phonological and phonetic representation. *Journal of Phonetics* **18,** 375–394.
Ricoeur, P. (1981) *Hermeneutics and the Human Sciences*. Cambridge University Press, Cambridge.
Ringen, C. (1994) Phonology—rule types. In R. Asher (ed) *Encyclopedia of Language and Linguistics*. Pergamon, Oxford.
Rischel, J. (1990) What is phonetic representation? *Journal of Phonetics* **18,** 395–410.
Rischel, J. (1991) The relevance of phonetics for phonology: a commentary. *Phonetica* **48,** 233–262.
Roche, M. (1973) *Phenomenology, Language and the Social Sciences*. Routledge and Kegan, London.
Schmidt, J. (1985) *Maurice Merleau-Ponty: Between Phenomenology and Structuralism*. Macmillan, London.
Schmitt, R. (1969) *Martin Heidegger on Being Human: An Introduction to 'Sein und Zeit'*. Random House, New York.
Shanon, B. (1993) *The Representational and the Presentational: An Essay on Cognition and the Study of Mind*. Harvester-Wheatsheaf, New York.
Spiegelberg, H. (1982) *The Phenomenological Movement: A Historical Introduction*. Martinus Nijhoff, The Hague.
Spurling, L. (1977) *Phenomenology and the Social World: The Phenomenology of Merleau-Ponty and its Relation to the Social Sciences*. Routledge and Kegan Paul, London.
Szpyra, J. (1995) Review of Michael Kenstowicz: Phonology in Generative Grammar. *Lingua* **96,** 189–195.
Trask, R. (1996) *A Dictionary of Phonetics and Phonology*. Routledge, London.
Varela, F. *et al.* (1991) *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, Cambridge, MA.
Vihman, M. (1996) *Phonological Development*. Basil Blackwell, Oxford.
Wheeler, C. J. (1980) On the relationship between phonology and psychology. *International Journal of Human Communication* **13,** 51–100.
Wierzbicka, A. (1992) *Semantics, Culture and Cognition*. Oxford University Press, Oxford.
Winograd, T. (1981) What does it mean to understand language? In Norman, D. (Ed.), *Perspectives in Cognitive Science*, pp. 231–263. Norwood, New Jersey.