

MENTAL REPRESENTATION IN PHONOLOGY:
A CASE STUDY IN THE PHILOSOPHY OF COGNITIVE SCIENCE

AVNİ BURAK ERBORA

BOGAZİÇİ UNIVERSITY

2011

MENTAL REPRESENTATION IN PHONOLOGY:
A CASE STUDY IN THE PHILOSOPHY OF COGNITIVE SCIENCE

Thesis submitted to the
Institute for Graduate Studies in the Social Sciences
in partial fulfillment of the requirements for the degree of

Master of Arts
in
Cognitive Science

by
Avni Burak Erbora

Boğaziçi University

2011

Thesis Abstract

Avni Burak Erbor, “Mental Representation in Phonology: A Case Study in the
Philosophy of Cognitive Science”

Cognitive science is first and foremost an interdisciplinary field, where research is based fundamentally on the idea that minds process information - that the brain is a biological computer. Cognitive scientific *disciplines*, on the other hand, are referred to as such because their field of study relates to the notion of information processing one way or the other. Psychology, linguistics, and computer science are all such disciplines. However, even though in an interdisciplinary study, cognitive scientists from various disciplines work in collaboration and look into the work of other disciplines to compare to their own work – the disciplines themselves are not united in terms of terminology, literature, traditions and concepts they make use of. Furthermore, with the idea of information processing also comes the idea of representation – and all of these disciplines make use of the concept of a *mental representation*. On the other hand, philosophers who defend representational theories of the mind usually point out the empirical success of the cognitive sciences, and claim that the representational theories of the mind are vindicated by such success in sciences that assume mental representations. The belief here, therefore, is that the notion of representation that is being used in the cognitive sciences is in fact the same notion that the philosophers use in their articulation of the representational theories of mind. The aim of this thesis is to intensify that belief, and it will focus on the *phonology* field of linguistics for an inquiry. The central claim will be that one can find the philosophers’ conception of mental representation in the actual practice of phonology – and that a phonological representation as it is conceived of by linguists is a mental representation *precisely* in the sense that the philosophers use the term.

Tez Özeti

Avni Burak Erbor, “Mental Representation in Phonology: A Case Study in the Philosophy of Cognitive Science”

Bilişsel bilim öncelikle ve kesinlikle interdisipliner bir alandır ve çalışmalar temelde zihinlerin bilgi işledikleri ve beyinin bir biyolojik bilgisayar olduğu fikri üzerine kuruludur. Diğer taraftan, bilişsel bilimsel *disiplinler* bu şekilde isimlendirilirler, çünkü alanları bir şekilde bilgi işlemek fikri ile alakalıdır. Psikoloji, dilbilim, ve bilgisayar bilimi bu tip disiplinler arasındadır. Fakat, interdisipliner bir çalışma içerisinde bilişsel bilimciler birlikte çalışırken birbirlerinin disiplinlerine bakıp kendilerinininki ile karşılaştırıyor oldukları halde, disiplinlerin kendileri kullandıkları terminoloji, kaynaklar, gelenekler ve kavramlar açısından birleşik değildir. Ayrıca, bilgi işleme kavramı yanı sıra temsiliyet kavramını getirmektedir – ve bütün bu disiplinler bir *zihinsel temsiliyet* nosyonu kullanmaktadırlar. Diğer taraftan, temsiliyetçi zihin teorilerini savunan felsefeciler genellikle zihinsel temsiliyet varsayarak çalışan bilimlerdeki empirik başarıya işaret edip, bu başarının temsiliyetçi teorilerin teyiti olduğunu iddia etmektedirler. Buradaki inanç, o zaman, bilişsel bilimlerde kullanılan temsiliyet kavramının felsefecilerin temsiliyetçi zihin teorilerini kurarken kullandıkları kavram ile aynı olduğudur. Bu tezin amacı, bu inancı güçlendirmektir, ve burada dilbilimeki fonoloji alanına odaklanılacaktır. Bu tezin merkezi iddiası, felsefecilerin anladıkları manada zihinsel temsiliyet kavramının fonoloji pratiğinde bulunabileceği olacaktır – yani bu tez dilbilimcilerin anladıkları manada fonolojik temsiliyetin, *tam anlamıyla* felsefecilerin bahsettikleri manada bir zihinsel temsiliyet olduğunu tartışacaktır.

ACKNOWLEDGEMENTS

First and foremost, I would like to express my gratitude for Assist. Prof. William Wringe and Assist. Prof. Markus Pöchtrager – who helped me throughout the duration of writing this thesis with incredible insight and invaluable advices. I am grateful that I will always have the academic knowledge I gained from them, and I can only hope that in the future I can be as helpful to students as they were to me.

I am indebted to Prof. Stephen Voss, Prof. Berna Kılınç and Assist. Prof. Lucas Thorpe for the valuable comments and criticisms during my defense – with their input, my perspective is strengthened.

I deeply appreciate the unending support of my mother, Prof. Dr. Ayşe Dilek Erbora and my father Teoman Erbora – without their support and understanding, writing this thesis would simply be impossible.

I would also like to thank Gökçenur Hazinedar and Güneş Artam for the encouragement – they were always there when I needed support the most, and it would be impossible to show my appreciation for the excellent discussions we had over the years in a short space like this.

Last but not the least, I would like to thank Gürhan Hazinedar; he always kept the mood high and never let confusion get the best of me when writing this thesis.

To my mother,
Ayşe Dilek Erbora

CONTENTS

CHAPTER I: INTRODUCTION	1
CHAPTER II: PHONETIC AND PHONOLOGICAL REPRESENTATION	7
Phonetics	9
Phonology	18
CHAPTER III: MENTAL REPRESENTATION	39
Representational Theory of Mind	42
The Nature and Content of Mental Representations	56
CHAPTER IV: THE PHILOSOPHICAL NATURE OF PHONOLOGICAL REPRESENTATION	72
Psychological Realism	73
Phonological Representation in The Mind	84
CHAPTER V: CONCLUSION	97

CHAPTER I

INTRODUCTION

Cognitive science is first and foremost an interdisciplinary field, where research is based fundamentally on the idea that minds process information - that the brain is a biological computer. Cognitive scientific *disciplines*, on the other hand, are referred to as such because their field of study relates to the notion of information processing one way or the other. Psychology, linguistics, and computer science are all such disciplines. However, even though in an interdisciplinary study cognitive scientists from various disciplines work in collaboration and look into the work of other disciplines to compare to their own work – the disciplines themselves are not united in terms of terminology, literature, traditions and concepts they make use of. That is, for instance, in practice a psychologist does not follow the literature on linguistics and likewise a linguist does not necessarily know what a computer scientist is working on (and so on and so forth). Thus, all cognitive scientific disciplines have their own terminology, concepts and traditions. On the other hand, with the idea of information processing also comes the idea of representation – and all of these disciplines make use of the concept of a *mental representation*. And since it is not in the practice of contemporary empirical sciences to contemplate on the nature of the theoretical concepts they make use of, their concept of representation is usually defined operationally. Furthermore, the general assumption among scientists is that when other disciplines use this notion, they mean more or less the same thing as in their own discipline.

Now on the other hand, philosophers contemplate on the nature of things. Philosophers with an interest in how the mind works have postulated representational theories of the mind, and there is a vast literature on this. The mainstream view in the philosophy of cognitive science literature claims that, analogous to a computer, the mind works through operations of transformation and storage of mental representations that are mainly information-bearing structures with the brain as their physical basis. These mental representations, thus conceived, have ontological reality. Now, philosophers who defend this view, usually point to the empirical success of the cognitive sciences, and claim that the representational theories of the mind are vindicated by such success in sciences that assume mental representations. The belief here, therefore, is that the notion of representation that is being used in the cognitive sciences is in fact the same notion that the philosophers use in their articulation of the representational theories of mind.

The aim of the present inquiry will be to intensify that belief, to the extent that this is possible in a master's thesis. I will focus on the phonology branch of linguistics, and analyze the concept of a phonological representation – and I will argue that a phonological representation is *precisely* the same kind of representation that the philosophers have in mind when they talk about mental representations. In other words, this thesis will aim to show that one can find the philosophical concept of mental representation reflected in the actual practice of phonology. I believe such an inquiry will serve two major functions: first, if it can be shown that the two disciplines (i.e. linguistics and philosophy) work on the same notion of a theoretical construct, namely a representation, these disciplines can contribute to each other with respect to what they have to say on that theoretical construct. This is particularly important given linguistics is an empirical field and philosophy is theoretical.

Second, given this thesis aims to provide a general definition of what a mental representation must be like, this could be used as a general framework to analyze the concept of mental representation in the other fields of cognitive science.

That being said, it is particularly important that this thesis does not take any positions with regard to the existence of mental representations. That is, although the arguments and ideas that stand for the justification of representational theories of the mind will be examined in detail, the purpose here is *not* to look into phonological theory for a defense of the representationalist picture. Thus the purpose in showing that phonological theory is committed to mental representations (in the same sense that the philosophers use the term) is purely comparative: This thesis will look into a comparison between how linguists conceive of phonological representations and how philosophers of cognitive science postulate mental representations – and show that the two concepts are perfectly compatible.

Furthermore, that the notion of a phonological *representation* is used in linguistics is not something this thesis will aim to *argue* for; it is an empirical fact that an overwhelming¹ majority of phonologists make use of this concept as a theoretical tool. On the other hand, the reason why phonologists use the term ‘representation’ mainly resides in the fact that there are facts about generalities in languages and a phonologist aims to capture such generalities in a phonological representation. This thesis, aims to highlight the theoretical assumptions a phonologist is committed to when using the notion of a phonological representation,

¹ It is also true that some phonologists (see, for example: Jensen, 2001) are careful not to use the notion of ‘representation’ when talking about phonological phenomena. This is based on a concern about not being committed to any theoretical constructs that is not defined formally. For instance, when one uses the term ‘representation’, the term immediately brings about the question what is being represented in a representation – and the answer to such a question is not soundly discussed in the practice of phonology. In any case, the fact remains that an overwhelming majority of linguists use the term.

and makes its claim based on showing that these assumptions are in fact identical to the philosophers' when postulating mental representations.

However, writing on interdisciplinary grounds has its downsides. One of the major difficulties in discussing this issue will be about the audience that the thesis will aim at. It is usually the case when writing in a discipline that the audience can be assumed to know the background to the issue at hand. On the other hand, since this thesis will aim an audience of both philosophers and linguists, the amount of expository material will need to be twice as much. That is, whereas linguists will most likely find the next chapter to be somewhat obvious, philosophers will probably think the third chapter to be straightforward. However, this is a necessary dilemma given the interdisciplinary nature of cognitive science – because one of the tasks this thesis must accomplish is to bridge the gap between these two disconnected lines of studies on the notion of (mental) representations. Thus, the tone of the two following chapters will be very different and the aim will be to create a common ground in the fourth chapter to discuss the two notions of representation in comparison.

The organization of this thesis will be as follows: the second chapter will take on the task of establishing the distinction between phonetics and phonology. The first thing that is taught to linguistics students in a contemporary introduction course is that whereas *phonetics* is the study of the (physical) properties of speech sounds in isolation, *phonology* studies how clusters of speech sounds behave in linguistic environments and how speech sounds are structured (see for example: Mannell et al., 2009). I will argue in this chapter that while phonetic properties are the best way to *transcribe* sounds, they don't provide the necessary level of abstraction to explain phonological regularities in languages. I will discuss, furthermore, that the explanation of these regularities brings with it the postulation of phonological

processes – and the best way of explaining such processes has been to ascribe psychological reality to both processes and the entities that they operate on; namely to phonological representations. The details of the claim for psychological realism, however, will wait until the fourth chapter.

The third chapter will aim to provide a picture of the mainstream view on mental representation in the philosophy of cognitive science. I will discuss this view dividing it into three main components: intentional realism, the representational theory of mind, and the computational theory of mind. I will argue that these three components, when combined together, serve as the criterion for determining what it means for a theoretical entity to be a mental representation (in the sense that philosophers conceive of it). Whereas intentional realism makes an ontological claim here, the representational theory of mind outlines the formal aspects of mental representations, and the computational theory of mind describes the specifics. The discussion in this chapter will end with Jerry Fodor's "Language of Thought" hypothesis for an exemplification of a theory that combines all three of these components.

The fourth chapter will involve the main claim of this thesis: namely, that the concept of a phonological representation is a notion that participates in all three of the components I discuss in the third chapter. In order to show this, I will firstly provide a discussion from a philosophy of science perspective on the ontological commitments of contemporary phonological theory. This discussion will serve to provide a basis to compare the mental structure postulated for linguistic representation with the three components of mental representation. I will identify the phonological component of this mental structure, and argue that phonological theory

connects with the philosophy of mental representation in terms of both ontological and theoretical aspects.

CHAPTER II

PHONETIC AND PHONOLOGICAL REPRESENTATION

The purpose of the following chapter will be to show that the contemporary work in linguistics is committed to a *mental* notion of phonological representation, and to explain what this actually amounts to. In order to do that, it will be necessary to make some very general points about the aims of phonology as a discipline. I will divide this chapter into two sections. The first section, (drawing heavily on Anderson 1974, 1985), will focus on phonetics: the branch of linguistics that is concerned with the physical properties of speech sounds. I will start by explaining that a linguistic perspective on language generally conceives of it as the collection of sound and meaning pairs. The aim of the first section will be (a) to explicate phonetics as a linguistic study that represents speech events with their physical properties and (b) to discuss the structure of phonetic representation. In the second section, I will focus on phonology. In order to explain what phonology is and what it works on, this section will build on the first section and show that the endeavor of representing speech events by itself leaves some important questions (regarding regularities that are seen in speech events) unanswered. Roughly stated, these questions usually take the form of “what can a word look like in language X” or “how can a word (or a sound) change in language X in contrast to other words (or sounds)”. In order to answer questions such as these, the study of the physiological, aerodynamic and acoustic characteristics of sounds (as done in phonetics) is not enough: the linguist must study how sounds (in the mind) are organized into *systems* and utilized in languages. Doing this, however, requires the linguist to represent the sound structure of a language in a

more abstract manner. Hence, the second section will be concerned with (a) showing the need for the abstract characterization of the systems of sounds in a language and (b) defining phonology as “the study of the systems of linguistically significant sounds” (Kaye, 1989: ix).

Before we start, however, notice that it is fairly easy to see why *phonetics* needs to make use of representations: there are speech events and although they are observable through physical properties, they are not like physical objects that can be studied directly under observation; thus they must be somehow represented in order to be studied (recorded, transcribed, etc.). Furthermore, to show why it is necessary for the discipline of *phonology* to make use of representations is again equally easy: if there are regularities in speech events which, when identified, help to answer questions otherwise left unanswered, then we are better off representing those regularities somehow (albeit abstractly – and hence differently from the speech events themselves). On the other hand, it is a relatively harder task to show why a phonological representation is conceived of as a form of *mental* representation by linguists. Now in order to accomplish this harder task in this thesis, it will be necessary firstly in this chapter to explain the evolution of the concept of phonological representation throughout the historical development of the discipline. This is because, as opposed to the subject matter of, say, perception, or psychology in general, it is not the case with phonological theory that the representations it makes use of are *mental* at face value. On the contrary, the conception that a phonological representation is somehow related to what goes on in the mind is a conclusion that has evolved² with each step in the historical development of the

² Strictly speaking, the idea that phonological representations must have something to do with the way the mind works should be clear from the beginning; however, the idea that the phonological processes that are represented are *mental* is quite different. The distinction will become clear below.

discipline trying to explain what was lacking in the previous step. Therefore, given that this chapter aims to show how a phonological representation indeed has psychological reality, telling about the history of phonology as a discipline is necessary and hence is more than just adding a descriptive flavor to the exposition of the subject matter here. Because of this, this chapter will have a historical tone. After discussing the early phonetic endeavors of linguists in the first section, the second section will show how and more importantly *why* the understanding of phonological representation – that was conceived by Saussure and the structuralists as a linguist's tool to analyze languages – has later changed into the study of an actual mental phenomenon with the works of Noam Chomsky and the generativist linguists.

PHONETICS

Linguistic Function of Language: Sound and Meaning Paired

The physical event of verbal communication, considered as such, can be the subject of analyses for myriad different scientific disciplines. In practice, for instance, a neuroscientist isolates certain areas of the brain that are activated in speech production (Broca's area) and recognition (Wernicke's area) and concentrates on the chemical properties and network structures of the bundles of neurons that are involved, and their connections with other areas in the brain. A physicist is interested in the acoustic properties that are represented by spectrograms of the sound waves. A cognitive psychologist, on the other hand, aims to postulate cognitive processes and capacities that are involved in acoustic information production and recognition – in an attempt to enrich their portrayal of human cognition as a collection of various types of processual faculties. Quite a number of more approaches can be stated;

however, from a linguist's perspective, language – the structured medium that underlies such a speech event – is generally considered “as a more or less formalizable system of correlations between *sound* and *meaning*” (Anderson, 1974, p. 3, my emphasis). In the early endeavors of linguists, studies on meaning were largely absent. The problem was, although all the speakers of a language have intersubjective intuitions of what meanings are like, and although philosophers had been debating the subject for ages, even the concept of meaning itself had been resistant to attempts to give an (agreed upon) formal definition or formulation; and this was an obstacle against coming up with a systematic framework for the objective study of meaning. It has been known since early on in linguistics, on the other hand, that speech events show regular patterns that are not only observable through the physical properties of sound waves and the speaker's biology, but are at the same time more or less common across world languages.³ Now studies on semantics (the linguistic branch that studies meaning in languages) in the recent decades have been quite successful in terms of a systematic study on meaning (see for example, Saeed, 2003); but nevertheless, most of the research done in (especially early) linguistics had focused mainly on the study of the sound side of languages – that is, on the regularities of speech events that are observable in the world's languages.

Even though speech sounds are observable, however, they are not available for direct examination. That is, pretty much as the physicists employ spectrograms to examine sound waves, the linguist must make use of *phonetic transcriptions* of speech events as their primary source of data. And,

³ All spoken languages of the world, for example, have vowels and consonants.

a phonetic transcription is, *ideally*, a complete and faithful record of the physical event it describes: every possible parameter of the speech event that can be observed, extracted, and quantified would be specified in a *complete* phonetic description. (Ibid. p.4, my emphasis)

The reason why a linguist cannot simply use spectrograms for phonetic transcription lies in the fact that such a *complete* phonetic description does not seem to be possible: a speech event does not only consist in the sound waves that are exchanged between a speaker and a listener. After all, strictly speaking, the complete physical event furthermore involves chemical reactions in both the speaker's and the listener's neurology, the aerodynamics of the spatial medium throughout which the sound waves are transferred, and biological functions which are related to the event both in the speaker (vocal tract) and the listener (auditory organs) – and so on and so forth.⁴ As I have stated above, there are quite a number of approaches that a scientist can indulge in the analysis of a speech event, but no specific scientific discipline can hope to capture *everything* that pertains to it. From amongst all the physical attributes of the event, therefore, what concern the linguists are the ones that are *relevant* to the functioning of language, which, very broadly construed, is the pairing of sounds and meanings.

I believe there is one thing of importance to be careful about here. One can protest the definition of language as the coupling of sound⁵ and meaning by claiming that in the instance of a speech event, even the facial gestures – mimics of a speaker – can lead to meaning difference. After all, the way in which one utters a sentence

⁴ For a good image of these issues, see: Halle & Reiss (2008).

⁵ Apart from the objection that will be discussed here, one can also point to sign languages and indicate the inexistence of sound there. However, as it will become clear in the following pages, phonology works on abstract patterns in languages and usually these patterns are found on the sound side of languages. An exception to this would be sign languages, but one can find similar patterns in signs like those that can be found on the sound side of natural languages. It is thus possible to talk about the phonology of sound languages – but a discussion on this is out of the scope of this thesis and will be omitted here.

can lead to a difference between, say, sarcasm and sincerity depending on non-auditory aspects like such facial gestures. In a speech event – the protest could proceed – one is as competent in grasping the meaning from such gestures as one is from the uttered sounds. However, it must be emphasized that this definition for the functioning of language as the pairing of *sound* and meaning (as stated above) is from the linguistic perspective. Even though an anthropologist can be interested in the meanings of gestures in a given community, it is important for the study of linguistics that, independent from the gestural ways in which people communicate, uttered words and sentences *do* have meanings. The fact that gestures can *add to* meaning does not negate this. After all, all competent speakers of a language can make sense of an uttered sentence in the local radio, and there are no known languages (as far as I know) where verbal utterances have no meaning independently of bodily gestures. For instance, there is no language where a certain sound sequence means one thing while saying something when moving one's hand in one motion, and means something else when saying the same thing but moving the hand in the opposite direction. And although demonstrative pronouns might generally look like this, they have meaning by themselves although the content might be filled in with a hand gesture; thus everyone who speaks English knows the meaning of, say, 'this' as a pronoun that is used to identify a specific thing close at hand (or being indicated).

Now, what is probably most appealing for the study of linguistics to study speech sounds is that sounds that a given language makes use of are consistently recurrent in other languages; thus the pool of phonetic features that the linguists have so far come up with are exhaustive for the transcription of all the natural languages (even though none of them use the whole pool). And if a new sound was found in a language, linguists would add new phonetic features to capture its properties

(Ladefoged & Lindau, 1989). However, as it will become clear below where I will proceed with a discussion of those features in question, phonetic features (or even sounds) cannot be *defined* in isolation from a set of theoretical assumptions.⁶ That is, without a theory at work, one cannot even decide where a sound ends and the other begins. The discussion below will aim to clarify this issue while telling about what exactly phonetic features are.

Phonetic Structure

For most of the early work done in linguistics, the goal was to come up with an *ideal* transcription: that is, to cover as many physical properties of speech sounds as possible (Anderson, 1974, p. 4). With expanding knowledge in the sciences and technology in the nineteenth and twentieth centuries however, even though linguists became more and more precise in their observations, they also realized that the physical information they were trying to transcribe was extremely vast. Furthermore, a consensus was being reached on that most of the physical properties of speech events they were trying to capture in a phonetic transcription was not constant in all the instances of a specific event, and was most likely irrelevant for the study of languages. After all, one does not expect loudness or softness in a speaker's voice to carry any extra linguistic information (even though they might have what is sometimes called paralinguistic⁷ or emotional⁸ information). One can, for example, increase the volume if the environment is noisy and yet mean exactly the same thing. There are also constant physical properties such as gender specific properties in

⁶ This is where phonology comes into play, as I will discuss in the second section.

⁷ For the definition of "paralinguistic", see: Ladd, 1996

⁸ See, for example, Kaye, 2005 on the non-linguistic, emotional information that is also in the sound signal.

one's voice that again seem unlikely to be relevant to the properties of speech events that are responsible for coupling sound with meaning. An insight based on these intuitions was gained in the early endeavors of the nineteenth and twentieth century linguists: a phonetic transcription, accordingly – as far as it is the primary source of data for the linguist to analyze speech events – must isolate the physical features that are responsible for the coupling of sound and meaning, and leave those features out that are irrelevant to the functioning of language. That is, in a phonetic transcription, a linguist must identify the general features of speech sounds that can (a) serve to demarcate utterances from each other, (b) characterize how different languages differ from each other systematically, and finally (c) leave the remaining features out as irrelevant.

Of course, as this is a scientific approach to the subject matter, the linguist must base such an endeavor on a principled basis. What is important here is that *there is no a priori way of doing this.*⁹ That is, there is nothing necessary in the physical event itself to ground such a principled basis on; and there is no deductive way from the observed sounds to isolate what physical features are relevant to the functioning of language (as far as it is conceived of as the pairing of sound and meaning) and what features to leave out. However, it is also clear that a pre-systematic understanding of speech treats the subject matter like uttered speech sounds are collections of “smaller” sound units (e.g. toy = t+o+y) brought together. The strategy employed by linguists to represent utterances in transcription – in agreement with this pre-theoretical conception – has traditionally been to represent the sound structure regarding speech utterances as a sequence of discrete, atomic *segments*. These segments have usually been defined as sets of properties that are

⁹ Noam Chomsky also argued against discovery procedures that would assume that there is an a priori way. See for example: Chomsky, 1957, 1964.

specified with the state of the vocal organs during articulation;¹⁰ and such articulatory properties (in their widely accepted form) are based on the *place* and *manner* of articulation (See figure 2.1). The place of articulation is the place where some kind of obstruction occurs during the flow of air from the lungs; for example, when the consonant ‘d’ is uttered (say in the utterance of a word “door”), the tip of the tongue moves towards the alveolar ridge and blocks the flow of air from the lungs in the oral cavity. In this example, the tongue (since it moves to obstruct) is the *active articulator* and the alveolar ridge is the *passive articulator*. For any place of articulation, there may be several manners of articulation; in the ‘d’ example, the velum is closed – stopping the air from escaping to the nasal cavity; but the velum could be open – in which case we would get a ‘n’ and not an ‘d’.

This latter point shows one further (and very fundamental) aspect of phonetic segmentation; namely that the articulatory features are commonly specified on a *binary* scale.¹¹ That is, a segment is either alveolar (+ alveolar) or not alveolar (- alveolar) depending on the place of articulation – either voiced (+ voiced) or not (- voiced), either nasal (+ nasal) or not (- nasal), either stopped (+stop) or not (-stop) depending on the manner of articulation. Hence, the sound unit (segment) ‘d’ can be specified as the set of binary features {+ alveolar, + voiced, +stop}, whereas ‘n’ can be specified as {+ alveolar, + voiced, + nasal}.

¹⁰ This view actually suffers serious problems if it is taken to the extent that it is seen as the basis of a phonological theory, and the recent debates in phonology have mostly focused on bringing out these problems to surface. However, the discussion here is strictly on what a *phonetic* transcription is like. The problems regarding why phonetics is not enough to study sound structure and why a more abstract level of inquiry is necessary will be discussed in the next section. On the other hand, a debate on the details and problems of a *phonetically oriented phonology* (or to put it in other words, to what extent phonology should make use of phonetic data) is clearly out of the scope of this thesis.

¹¹ Notice, however, that this is not inherent in the signal, but rather is a decision based on a theoretical assumption. I address the issue in the following paragraph.




Figure 2.1: The major vocal tract articulators and places of articulation (from Mannell, 2009)

The assumption here then is that:

it is possible to represent every utterance in any language in terms of a language-neutral phonetic transcription, composed of a sequence of individual segments, for which each segment is characterized by a set of [articulatory] parameters ... The set of parameters includes all and only those factors that can be shown to be independent variables in linguistic structures, either by having distinctive function or by differing systematically from one language to another. The specification of these parameters has two aspects, which must be distinguished. Primarily, we assume that each can be specified on a binary scale as + or - (Anderson, 1974, p. 9).

However, as I have mentioned above, one must bear in mind that this strategy that is employed by linguists – although it constitutes a principled basis for the classification of speech sounds – is not a necessary conclusion from the observation of the physical speech event itself. First of all, like most physical phenomena, speech

is a continuous event; that is, if one looks at a spectrogram, one does not see any breaks or stops indicating units for uttered speech sounds; all one can see is a continuous wave visually representing the uttered sentence or word. Second, biologically, at any given time, the vocal organs that produce the sound waves are in continuous motion; it is not the case that while uttering one sound the organs take a specific shape, stop and break it (take a neutral position), and then, take another shape for another sound. Even though we might have intuitive grounds to differentiate sounds (like what the letters in a word represent), when uttering what might be called one sound, (say, ‘t’ as in ‘time’) the organs already start shaping into their positions for the next sound (‘i’). Yet there is a methodological need to discretize this continuous event: if one is to analyze the physical properties of a speech event, one has to conceptualize that event in terms of different moments, or motions, etc. to identify such properties. It should be kept in mind that this methodological need comes from the necessity to represent particular speech events in a principled basis; however, the problem is, there is nothing in the physical event itself to necessarily isolate and differentiate one discretized segment from the other based on their articulatory properties. Thus it should also be kept in mind that the justification for the phonetic segmentation of linguistic sounds with respect to their manner and place of articulation is a pragmatic one.¹²

¹² This means that if there is a rival theory that offers a better way of representing physical properties of speech events (i.e. without having to answer the objections about sound wave continuity or vocal organs overlapping in speech, as I discuss here), phonetics would be better off with that theory as long as it offers to do so on a principled basis. No matter how it is done, on the other hand, phonetics represents the *physical* properties of speech events. However, as I will discuss in the next section, there are regularities in speech events of different languages that cannot be captured by representing physical properties alone. A discussion of how physical properties should be represented would therefore put us in a different direction here: the aim of this chapter is to show that the contemporary work in phonology is committed to the notion of a mental representation, and for that, we need to discuss the nature of phonological representation – which is (in addition) concerned with those regularities that cannot be captured by the representation of physical properties alone.

As it will become clear below, studying the languages of the world has shown that the inventory of sounds¹³ that are used in speech events show regular and systematic behavior in certain contexts. Furthermore, one can identify certain governing rules and principles in each language that regulate such behavior. Phonetic representation, by itself, neither identifies these governing rules and principles nor explains the systematic behavior of speech sounds – it merely provides the ground for theorizing. To give an analogy, phonetic transcriptions are like musical notation; whereas such an inventory is necessary to write and reproduce a musical composition, it doesn't by itself constitute music theory.¹⁴ The following section will be concerned with the level of abstraction that is required in order to represent the systematic behavior of speech sounds – in other words, the following section will be about phonology.

PHONOLOGY

The discussion in the first section of this chapter mainly focused on the question of what would be (and is) a linguistic perspective on the physical event of verbal communication. To sum it up, linguists see this event as a tokening of language, the function of which is assumed to be nothing other than the pairing of sound and meaning. I have discussed, furthermore, that the sound side of this pairing has mostly been the center of attention. Speech sounds are conceived to be collections of units

¹³ An inventory is the set of sound units a specific language makes use of. For example, Turkish has 'ü' in its inventory but English does not.

¹⁴ This might actually provide more than an analogy: for a discussion of a *grammar* of music with the aid of generative linguistics, see: (Lerdahl & Jackendoff, 1983). A detailed discussion of this subject, however, would be irrelevant here.

(the building blocks of articulated speech in a given language) and phonetics focuses on the classification of these sound units based on their articulatory and acoustic properties. Linguists use phonetic transcriptions of the utterances as their primary source of data to represent speech events in languages. Now the whole point of such a phonetic representation, as I have discussed above, was to abstract from irrelevant information such as the different pitches in which a speaker can utter a word, loudness, gender, or environment noise that are present in the physical sound wave. For instance, there is an indefinite number of different possible utterances of the word ‘cat’ (in different environments, by different people, etc.) which would all be represented as [kæt]. Thus, phonetic transcription is supposed to be helpful in identifying the sounds that a particular language makes use of, and represent utterances on a principled basis. However, I discussed that it is not possible to ground such a principled basis on observation alone – and that theoretical assumptions are already at play where one identifies speech sounds. Irrelevant information in the sound wave of a speech event, then, is *defined* in terms of these theoretical assumptions. This shows that when in practice linguists represent utterances in a language, they determine the representation based on whatever theoretical assumptions they have *in general* about the sounds of languages. A principled inquiry on the systems of sounds in languages, however, is the job of phonology as I have mentioned in the introduction – and thus, without phonology, phonetic assumptions cannot be firmly grounded.

Now furthermore, a closer look at languages tells us that even if phonetics was a neutral and direct way of transcribing speech events – and was not loaded with theoretical assumptions already – it still could not tell us all there is to the sound side of the story: consider, for example, the French words for ‘water’ (<eau>) and ‘high’

(<haut>). Both of these words are pronounced the same, and so they share the same phonetic representation [o]. On the other hand, in French, when a word starts with a consonant, the article used is <le> (or <la>), and when a word starts with a vowel, the article used is <l'>. But when used with an article, <eau> becomes <l'eau> ([lo]) whereas <haut> becomes <le haut> ([ləo]). Interestingly, the article sounds like that of a word starting with a consonant with the latter – like in <mot> [mo] ('word') becomes <le mot> (lə mo). Given that these two words <eau> and <haut> have the same phonetic properties, and that they still behave differently when used with an article, it appears that there must be something else to them that cannot be captured in a phonetic representation.¹⁵

In the following pages, my purpose will be to discuss phonological representation. As we will see, there are certain regularities and patterns in languages that phonetic representation cannot capture but phonological representation can. I will start by describing the recurrent regularities that are seen in languages and show why they demand a more abstract level of representation. Afterwards, in order to expose the nature of this abstract level of representation, I will firstly discuss the roots of this idea in Ferdinand de Saussure and the structuralist tradition that came following his work. Here, the concept of a phoneme will become clear as the unit of phonological representation used by the structuralists to analyze the systems of languages. Finally, I will have a discussion of Noam Chomsky's work on phonology where I will discuss the limitations of structuralist phonology and show the mental aspect of phonological representations.

¹⁵ This example is taken from Markus Pöchtrager's Phonology class notes.

From Phonetics to Phonology: A Necessary Level of Abstraction

The first thing that one would notice when studying languages is that languages show patterns of regularities. Probably the most obvious of these is, for example, the ordering of different words in a sentence. Consider the following three English sentences:

- (1) Charlotte likes Max.
- (2) Max likes Charlotte.
- (3) *likes Charlotte Max.

Even if they have not necessarily heard these sentences before, all the speakers of English would agree that (1) and (2) are grammatical sentences in English although they mean different things, but (3) is not acceptable at all. This agreement is based on the fact that utterances in English consistently show a pattern of subject-verb-object (SVO). All the speakers of English have the knowledge of this pattern. Furthermore, different languages have different patterns; consider the following three Turkish sentences:

- (4) Gökçe sever çikolata. (Gökçe likes chocolate)

Gökçe likes chocolate

- (5) Gökçe çikolata sever. (Gökçe likes chocolate)

Gökçe chocolate likes

- (6) Çikolata sever Gökçe. (Gökçe likes chocolate)

Chocolate likes Gökçe

In contrast with English, all of (4), (5) and (6) are grammatical sentences in Turkish, because Turkish does not have strict word order restrictions with respect to the subject, the verb and the object. That is, as opposed to English, one might change the order of the subject, the verb and the object, and still form a grammatical sentence. This is usually done in Turkish to put emphasis on one thing or another depending on the context of the utterance; but yet, the meaning of the utterances is always the same: i.e. that Gökçe likes chocolate.¹⁶ On the other hand, the following word order restriction where a noun is followed by a postposition applies to Turkish:

(7) Senin için geldim. (I came for you)

You for I came

(8) *İçin senin geldim.

For you I came

These simple examples from English and Turkish illustrate two points:

- a) languages show patterns which determine what is possible in that language and what is not; and
- b) languages may differ from each other with respect to the patterns they show.

Now likewise, on the sound side of language, there are again patterns that can be identified in different languages. For example, speakers of English would somehow

¹⁶ It is true that Turkish usually is conceived of as having a subject-object-verb (SOV) pattern. However, it is also true that changing order of these constituents in Turkish is used in practice by native speakers for pragmatic purposes. For example, (4) might be given as an answer to the question “Who likes chocolate?” and (5) as an answer to “What does Gökçe like?”, where (6) might be the answer to “What did you say Gökçe liked?”. A discussion of this point can be found in Erguvanlı, 1984; however the further details of Turkish grammar are clearly irrelevant for this thesis.

agree, without having heard it before, that ‘bnuck’ is not a possible English word but ‘bluve’ is, although neither one exists in the vocabulary of English (Carnie, 2007, p. 4). Likewise, it is the common intuition of the native speakers of Turkish that ‘nreklar’ cannot be a Turkish word but ‘hekler’ can. These intuitions are based on the fact that English words cannot start with the sound sequence [bn] and Turkish words cannot start with [nr]. Notice furthermore that some constraints that apply to Turkish do not apply to English (e.g. ‘street’ [stri:t] is a word of English, but it cannot ever be a Turkish word). These examples show that the sound sequences that a language makes use of are not random – on the contrary, languages are *structured* in such a way that they allow certain patterns of sound sequences and they disallow certain others. Both of the points (a) and (b) above, therefore, also apply to the sound structure of languages.

However, a direct way of representing sounds, such as in a phonetic representation, by itself, certainly cannot account for the way in which sounds are structured. That is, in order to account for what a word in a language can look like, or, to account for the differences in structure between different languages, one needs to further ask the question how sounds are organized into systems and utilized in languages.

Regarding the linguistic role of sound structure, maybe the most significant of the conceptual developments in the twentieth century was made by distinguishing between phonetics and phonology. As I have mentioned in the introduction to this thesis, the first thing that is taught to linguistics students in a contemporary introduction course is that whereas the former of these is the study of the (physical) properties of speech sounds in isolation, the latter studies how clusters of speech sounds behave in linguistic environments and how speech sounds are structured (see

for example: Mannell et al., 2009). Possible questions of phonetics would be in the form of “what does the mouth do when we say ‘ü’?” or “what does the sound wave look like when we say ‘ü’?”. Questions pertaining to a phonological inquiry, on the other hand, would be in the form of “does this language *allow* ‘ü’?” or “is ‘nreklar’ a possible word in this language (and why)?”. Now I have already discussed in detail that phonetic transcription is a physical level of representation; in order to address such questions as the phonological ones, however, it is necessary to represent language in a more abstract manner. I have also discussed that early linguistic attempts to analyze linguistic sounds saw the need to isolate speech sounds from irrelevant information. However, it was with the later attempts of linguists in the nineteenth century to analyze linguistic sounds that they saw the further need to represent languages in a more abstract manner. In the following pages, my aim will be to explicate the nature of this abstract, phonological representation. For that purpose, I believe, it would be prudent to look at the development of the concept starting with its roots in nineteenth century with Saussure’s work on linguistics.

The Phoneme: Structuralism and Taxonomic Phonemics

Although his written output was quite sparse, and he did not write down most of his thoughts in publishable form, Ferdinand de Saussure is generally credited as being the founder of modern linguistics – thanks mainly to the reconstruction of his lectures in Geneva University by students and colleagues. One of the most important contributions Saussure made to the study of linguistics was the mutually exclusive distinction he postulated between *langue* and *parole*; where the former roughly translates as “language”, and the latter as “speech”. The translations do not say much, however. Saussure thought that the concept of “language” – broadly construed

in the sense that we use it every day – was not satisfactory as an object of scientific inquiry. He claimed that a linguistically significant study should refine its concepts to limit itself to more precise notions; notions such as *langue* and *parole* (Anderson, 1985, p.24). With that in mind, and having bifurcated language into two aspects, Saussure proposed that *langue* is the aspect of language that represents the speakers' knowledge of (the system of) an individual language. *Parole*, on the other hand, is nothing other than the real time application of this knowledge – that is, the way in which a particular individual puts his or her knowledge of *langue* into use in a particular occasion. Accordingly, for instance, when it is said of someone that he or she knows how to speak French, the *knowledge* in question falls into the category of *langue* – and when someone speaks in a particular occasion *in* French, that *event* falls into the category of *parole*.¹⁷

Now the primary reason for distinguishing *langue* from *parole* for Saussure was to allow the linguists to focus their attention on the former (Anderson, 1985, p.26). Like linguists before (and after) him, Saussure believed that what distinguishes a language from others is the *system* that underlies that language. The knowledge of a language, as I have discussed above, consists of the correspondences of sound and meaning – along with what utterances are possible and what not – in that language. Hence, the agreement that 'bnuck' is not a possible word in English but 'bluve' is (although neither one exists in the vocabulary of English) is granted by each individual's intersubjective knowledge of the system of English – the study of which was nothing other than the study of *langue* for Saussure. *Langue*, furthermore, is

¹⁷ As I will discuss later on, this distinction between *langue* and *parole* is practically identical to Chomsky's distinction between *competence* and *performance*, respectively – but (and that's a very important 'but' for our purpose here) whereas there is nothing in Saussure's description of the distinction to suspect an ontological claim, Chomsky's distinction is explicitly based on the claim that competence and performance are mental phenomena. I will discuss more on this later.

independent of *parole*; that is, the knowledge of a language is independent of any particular and individual utterance produced by the individual members of that community at any given occasion.¹⁸

Now Saussure proposed to study *langue* as an abstract system of arbitrary signs.¹⁹ Very briefly, each sign has two aspects: a form (*signifiant*) and content (*signifiée*). The form of a sign serves to distinguish it from the other signs in the system, whereas the content is the information that is conveyed by the sign. In any given language, it is the sound system that constitutes a structured collection of *signifiants*: signs with properties (forms) which serve the function of distinguishing one sign from the other. What gives meaning to signs (structured sounds) in a language therefore, is their relation – and mainly their *difference* – from other signs. Thus, Saussure claimed that what is particularly important about the sounds of a language does not reside in their individual properties; but on the contrary it is their negative character, i.e. their distinctive properties. Hence the study of *langue* should be the study of a network of signs – or in other words, a *semiotic system*. The individual manifestations of signs, that is, speech events themselves, tell us about neither the structure nor the organization of a semiotic system.

Hence the beginnings of the distinction between phonetics and phonology can be seen in Saussure and his distinction between *parole* and *langue*. It should be clear by now that phonetics is a subject matter that belongs to *parole*; as I have discussed above, phoneticians deal with the auditory and articulatory properties of speech

¹⁸ i.e. any tokening of language.

¹⁹ Surely, it is not possible to do justice in explaining Saussure's conception of a semiotic system in only one paragraph. However, the details of his proposals are to a large extent irrelevant to our purposes here. What should be highlighted in the paragraph that follows is that Saussure offered a structuralist analysis of languages, and that kind of an analysis conceives of a language to be a **(a)** system of **(b)** abstract and **(c)** arbitrary signs.

sounds – which can only be analyzed given a particular utterance. *Parole* furthermore covers all the physical, neurological and psychological factors that play a role in a particular speech event. These factors have their own properties which are (most of the time) independent from their use in language. What *parole* (and hence phonetics) does not cover, however, is the (semiotic) system that underlies a language. If this system can be isolated from *parole* in a systematic fashion, therefore, an important step would be taken towards understanding what sort of properties linguists should concern themselves with in order to analyze the underlying regularities that are specific to different languages. Saussure's proposal was precisely on how this isolation would be possible: through studying the *structure* of languages and explaining these linguistic structures as the systems and organizations of the networks of signs – signs that are identified via their distinctive properties. On the sound side of languages, this gives us the first definition of what *phonology* is: “the study of the systems of linguistically significant sounds” (Kaye, 1989, p. ix).

Saussure's conception of language as a network of signs was used by linguists after him as a methodological principle, and it has come to underlie the notion of ‘structural analysis’ in linguistics (as well as in many other disciplines) (Anderson, 1974, p.13). The tradition(s) of linguists that followed are commonly known as *structuralists*. Structuralist *phonologists* represented sound structures by highlighting the *distinctions* that oppose one element to the other. Now, with structuralism in phonology comes the notion of a *phonemic* analysis: the notion of a *phoneme* in such an analysis designates the minimal distinctive element of a phonological representation – the smallest unit that represents a complete sound and that leads to meaning difference. That is, a phoneme that a language is said to make

use of is a minimal *phonetic* unit in that language that is observed to lead to a meaning difference.

As I discussed above, in contemporary mainstream linguistics, a phonetic unit is identified by features on a binary scale, – but early structuralists (in particular Nikolai Trubetzkoy) thought that the smallest units of a phonological representation were phonemes; and they conceived of a phoneme as the smallest representation of a speech sound that cannot be decomposed any further. The idea of decomposing phonemes into smaller properties (i.e. binary phonetic features) was introduced later on by the linguist Roman Jakobson. For later structuralists then, a phoneme (that was conceived of as a minimal phonetic unit composed of binary features) serves the purpose of identifying distinctive sounds that lead to meaning difference in languages – and their (later structuralists’) analysis of languages was based mainly on such distinctive features of phonemes.

To clarify what a phoneme is, consider the phonetic property of aspiration: the release of a strong burst of air from the lungs that can happen after the release of obstruents. To test what aspirated sounds are like, you can put your hand in front of your mouth and say ‘spill’ [spɪl] and ‘pill’ [p^hɪl]; you will notice the flow of air in your hand with the [p^h] of the latter word. Aspiration is a perfectly valid phonetic property – the place and manner of articulation are available for description – and it can be stated in the binary scale as [+aspirated] or [-aspirated]. However, if one uses aspiration in the word ‘spill’ and (incorrectly) pronounces it as [sp^hɪl], a native English speaker would still understand this word as ‘spill’ – as opposed to, say, if one used a [k] instead of [p]; in which case the word would come out as [skɪl] and mean ‘skill’. Furthermore, if one used [m] instead of [p] in that word, it would come out as an (inexistent word) [smɪl] and a native English speaker would not understand

that word as anything – that is, as opposed to [sp^hɪl], they would *not* interpret [smɪl] as ‘spill’. It is possible to argue, therefore, that aspiration does not lead to meaning difference in English.²⁰ Consider, in contrast, that in Icelandic, aspiration is a phonetic feature that leads to meaning difference: <panna> ([p^han:a]) means ‘frying pan’ whereas <banna> ([pan:a]) means ‘to forbid’ (Gussmann, 2002, p. 5). Thus, /p^h/ and /p/ are distinct phonemes in Icelandic, but not in English.

Now the notion of a phoneme highlights the abstract nature of phonological representation in Structuralist linguistics: besides having a direct, physical, transcriptive level of representation for [sp^hɪl] and [spɪl], we also have a phonological level of representing both of these utterances as the instances of the same abstract form /spɪl/. The phonetic units [p] and [p^h] are called the *allophones* of the same phoneme /p/ in English; whereas /p^h/ and /p/ are distinct phonemes in Icelandic. While phonetic realizations belong to *parole*; phonemes belong to *langue*. This example demonstrates that a pattern that differentiates English from Icelandic, which cannot be captured in a phonetic representation, can be represented with the help of a more abstract, phonological representation – and this is at the heart of the Saussurean claim for the linguist to study *langue* as opposed to *parole*.

It should now be clear why we need an abstract level of phonological representation: we need to capture some linguistic properties in the sound structure

²⁰ This way of arguing about the phonemic inventory of languages – that is, *defining* phonemes with respect to their distinctive phonetic features that lead to meaning difference in languages – is (as I discussed above) a method of analysis peculiar to later structuralists. The method of analysis exemplified in this paragraph is generally known as the “minimal pairs test”. As I will discuss later on, however, this kind of distinctive analysis has serious flaws and was shown to be problematic by Noam Chomsky and the contemporary generative phonology he influenced. Furthermore, the claim here might not be true: one can point to the difference between the English words ‘bin’ [bɪn] and ‘pin’ [pɪn] – and claim that the only difference between them is the property of aspiration and not of voicing. In any case, it should be emphasized here that the present argument is neither defending distinctive analysis in phonology, nor claiming the importance of minimal pairs in phonological analysis. The only purpose of the example here, as it will become clear in the next paragraph, is to show the abstract nature of phonemes *as they were conceived of by the later structuralists*.

of languages that cannot be captured within phonetic representation. The next question I will address here is related to the *nature* of this abstract representation. Notice that in the structuralist conception of the system of languages as networks of signs, the sign's existence is purely formal. The system of signs that constitutes *langue*, therefore, is a purely formal pattern of relationships among linguistic forms (Anderson, 1985, p.28). Notice furthermore, that in such a conception, there is no explanation of *why* these relationships are the way there are. The task of the linguist in the view of the structuralists then is to identify the relationships among linguistically significant forms and come up with a taxonomic picture for each language – i.e. to make taxonomic analyses of the structures of languages. Because of the way in which they represented sound structure, the structuralist linguists' work is usually referred to as *taxonomic phonemics*. Therefore, a phonological representation is conceived of as the linguist's tool for analyzing languages in the Saussurean tradition: there is no claim for the reality of phonological representations, and hence no claim for truth.²¹

Up to now, I have explained what it is that is being represented in phonological representation, why such representation is necessary, and what kind of a function representations are supposed to serve in linguistics. But now we need to look at another issue, namely that of what kind of entities phonological representations are – abstractions for linguist's arsenal of tools, or representations with a mental nature? It is clear that the structuralist conception of phonological representation doesn't ascribe to it any sort of psychological reality; in fact what is

²¹ I mention this in passing here, but I will discuss the notion of truth and ontological commitment of phonological theory in detail in the fourth chapter – as it will turn out, the commitments of *contemporary* phonological theory with respect to these aspects weigh heavily over how we should understand a phonological representation; and it is precisely because of these commitments that a phonological representation is a mental representation *exactly* in the sense that philosophers of cognitive science conceive of it. But it is too early at this point to argue for that.

missing in the discussion above, is the structuralists' ascription of *any* kind of reality to phonological representations – the issue doesn't seem to be relevant from a structuralist perspective. In the following, I will contrast this with Noam Chomsky's position – which was basically grounded on the rejection of structuralist taxonomies. Thus the following pages will aim to show the psychological aspect of phonological representations.

Noam Chomsky: From Abstract Representation to Mental Rules

The discussion above suggests that the Saussurean take on the issue of phonological representation serves to distinguish an abstract level of representation from the speech events – a level of representation in which questions such as “what would a possible word (in a given language) look like”, and “is this sound a minimally distinctive sound in this language” can be addressed. However, a significant linguistic theory must also address questions such as “how do word forms relate to each other?” and “how can a word change in a given language?” – namely, the questions regarding *phonological derivation*. Consider, for instance the Turkish stems ‘ev’ (a house) and ‘at’ (a horse). The plural for ‘ev’ is ‘ev-ler’ whereas the plural for ‘at’ is ‘at-lar’. Now the plural forms in Turkish always end in either ‘-ler’ or ‘-lar’, and which form is to be used is consistently recurrent (*‘ev-lar’ would not be a possible plural in Turkish). There is, therefore, a pattern to which plural form will follow a word stem in Turkish – to such an extent that every native speaker of Turkish would agree that the plural form for the made up (but phonologically possible) Turkish word ‘hek’ would be ‘hek-ler’ and never ‘*hek-lar’.²² These questions (as it will become clear below) ultimately pertain to derivations of

²² I owe this example to Markus Pöchttrager.

phonological representations in accordance with general rules – I will henceforth classify questions regarding such patterns in the sound structure of languages under the title ‘phonological derivation’. Now as we have seen, the structuralist account of linguistics can address the question whether ‘hek’ is a possible word of Turkish or not. However, their taxonomic endeavor does not address questions such as why do we have an alternation between ‘-ler’ and ‘-lar’ based on the nature of the previous vowels – and more importantly *why* there is an agreement among native speakers that ‘hek’ would change into ‘hek-ler’, that is, if ‘hek’ existed in Turkish.

In the 1960s phonological theory was expanded to address the questions regarding phonological derivation following the lead of Noam Chomsky.²³ This resulted in the widespread acceptance of *generative* phonology and effectively in the end of structuralist phonology. In the following pages, I will explicate Chomsky’s generativist approach addressing it from three fronts: I will firstly discuss the linguistic need for this approach, secondly address the generativist meta-theoretical concerns on what a linguistic theory should be like – with a criticism of structuralist phonology. Finally I will show the ontological outcome of the generativist program in linguistics – which, as it will turn out – is the most important one for our purposes here; for with the acceptance of the generativist position, also comes the ontological commitment to the *psychological* reality of phonological representations.

The starting point of the generativist program in linguistics can be identified with the following question: How is it possible that

²³This of course was embedded in a larger scale paradigm change, as it will become apparent later on in this thesis.

a mature speaker can produce a new sentence of his [or her] language on the appropriate occasion, and other speakers can understand it immediately, though it is equally new to them (Chomsky, 1964, p: 7).

For instance, for most speakers of English it would be the first time to read the sentence: The green shirt that is in the green box which is in the green house that is in Green Street that is in a green country of a green world has a red taint on it. Yet every speaker of English would understand it, and accept it to be well-formed, although it might be somewhat of little literary value. Sentences can potentially be of infinite length (in all languages) when they are conjoined together as in the above example – it is therefore impossible for a speaker of a language to have learned all the possible sentences from experience. Likewise, speakers of Turkish agree on how the plural form of the inexistent word ‘hek’ would sound like (namely, as [hæklær]) – and the speakers of English agree on the plural form of ‘bluve’ as ‘bluves’ (i.e. [blu:vz] and not *[bluvs]) without any learning (regarding these words)²⁴ from experience involved. These remarks suggest that there is a *creative* aspect of language (or in Saussurean terms, of *langue*) that the structuralist linguists have largely overlooked.

Now Chomsky’s argument on the subject matter is that, on the basis of limited experience, every speaker of a language has developed a form of *competence* in that language – which can best be represented as a *system of rules* (Ibid, pp. 8-9). Chomsky famously calls this system the *grammar* of one’s language:

To each phonetically possible utterance [...], the grammar assigns a certain *structural description* that specifies the linguistic elements of which it is constituted and their structural relations (or, in the case of

²⁴ Of course experience is involved in learning that making the plural consists of adding –s. What is meant here is that we know how to apply such a rule even with a word that we have no previous experience of.

ambiguity, several such descriptions). For some utterances, the structural description will indicate, in particular, that they are perfectly well formed sentences. This set we can call the *language generated by the grammar* (Ibid, p: 9, original emphasis).

Notice furthermore, that *competence* is independent of *performance*. Grammar for Chomsky, as it can be seen from the above quote, is a generative device that produces the set of well-formed sentences in a language, but not necessarily in action. One can (for any cognitive reason) utter a badly formed sentence, but immediately recognize that it is badly formed. This is possible only because the structural descriptions that are assigned by grammar also indicate the manner of deviation from perfect well-formedness. For example, if one utters a spoonerism like “is it kisstomary to cuss the bride?”, it will immediately be corrected by both the speaker and the listener to “customary to kiss the bride”. The reason for this is that both the speaker and the listener share the competence of English, and whereas performance can be in error, competence never can. In a sense, there is a striking resemblance between the generativist competence/performance bifurcation and the structuralist *langue/parole* one (respectively). Both traditions agree on the claim that, when dealing with the sound side of a language, the linguist must study the abstract level of phonological representation (as opposed to the phonetic), namely competence or *langue*. However, there is also a striking difference: whereas for the structuralists, the phonological representation consists of a set of possible words (as in a lexicon); for the generativists the phonological representation consists of a set of generative rules *plus* the medium (i.e. representation) on which these rules operate.²⁵

²⁵ This, of course, is the crux of this thesis if the representation on which rules operate can be shown to be a mental representation – but I will reserve the details until the fourth chapter, and until after I discuss the conception of mental representation in the next chapter. Presently, what is highlighted here is the contrast of generative phonology with structuralist phonology in terms of their conception of phonological representations.

There is an extremely important point here: whereas the structuralists had the luxury of conceiving of their lexicon-like phonological representations as *arbitrary*, generativists claim that there is no arbitrariness involved with the rule based derivations of phonological representations: The set of generative rules of the generativists consistently assign the same set of structural description to each phonetic utterance and such consistency *cannot be arbitrary*. Now, given that the set of generative rules (i.e. grammar) is the same for every individual that share competence in a language, it furthermore follows that these rules belong to *a capacity that is shared by all the individuals*. This means that where the arbitrary phonological representations for the structuralists were mere academic conventions for the linguists to use, they have to be *actually existing (mental) processes* for generative phonology.

Perhaps an analogy could clarify this last point: Think of a computer program which takes an input (say, certain numbers) and returns an output (say, some other numbers). The structuralist conception of phonological representation here is analogical to a taxonomic analysis of the relations between the output numbers. This endeavor serves the function of speculating on what the possible number sequences can be like in the future outputs. But as the analogy goes, we have seen that there are certain regularities to the output number sequences (e.g. phonological derivation). We should be convinced now that the output sequences are not random or arbitrary, but they are generated through an algorithm. What we must concern ourselves with should no longer be the taxonomy of outputs. What we must now be concerned with is the *algorithm* that generates the output (i.e. the rules of grammar).

On the other hand, from a meta-theoretical front, Chomsky furthermore argued for different levels of success for a linguistic theory (Ibid, pp. 28-55).

Accordingly, there are three levels of adequacy that a linguistic theory might achieve. The lowest level can be achieved through *observational adequacy*, where the theory solely accounts for the observed data accurately. On a higher level of success, a theory can achieve *descriptive adequacy* if it supersedes observational adequacy by further specifying the observed data by highlighting generalities that express underlying regularities in languages. An example of this would be the plural case regularity in Turkish that I have discussed above. A theory that achieves descriptive adequacy furthermore gives “a correct account of the linguistic intuition of the native speaker” (Ibid, p: 28), such as in the case of the constructed Turkish word ‘hekler’ above. However, as Chomsky argues, the highest level of adequacy a theory might hope to achieve is one of *explanatory adequacy*, where a theory, in addition to having observational and descriptive power, furthermore can provide a general basis showing *where* its descriptive account is coming from and why it is better than other adequate descriptive alternatives.²⁶

Now, Chomsky rejects structuralist phonology on the grounds that it has largely been concerned with observational adequacy. This rejection seems to be in place: after all, as I have discussed above, structuralist phonology was concerned with what a word can look like in a language, but not with the patterns that are seen behind phonological alternations - or why such patterns exist. That is, structuralist theories were *representational* in the sense that their sound structure analyses were aimed for taxonomic representations of phonological phenomena. Generativist theories, on the other hand, aim to provide sets of rules for generalizing the patterns behind how words change (e.g. -ler, -lar) as well as providing rules for what words

²⁶ Chomsky’s concern for explanatory adequacy is usually discussed in connection to language acquisition – that it explains how children can acquire what they acquire. I will not go into the details of these discussions; but this is one of the *achievements* of generative linguistics as an explanatorily adequate theory in the way that is conveyed here.

can look like (i.e. how grammar can generate them)²⁷. In this sense, generativist theories are certainly *rule-based*, and that is a big difference from the representational theories of the structuralists. However, if this was the end of the story for the generativists, their account would be stuck at descriptive adequacy by Chomsky's own categorization and would not achieve explanatory power. What brings in this last component is Chomsky's insistence on a linguistic theory for being "concerned with the internal structure of the device" (Ibid, p: 29) – that is, with the explanation of how the grammar works. For Chomsky, this device in question that is responsible for generating phonetic representations is without question *the mind*. In Saussurean terminology then, what gives Chomsky's model explanatory power is the claim that the explanation of *langue* is an explanation *because of the fact that* it is expository of (at least part of) how the mind works. In Chomsky's own words:

In any event, there is surely no reason why the linguist must necessarily limit himself to 'the study of phenomena and their correlations', avoiding any attempt to gain insight into such data by means of an explanatory theory of language, a theory which is, of course, 'mentalistic', in that it deals with the character of mental processes rather than with their physical basis (Ibid, p: 99).

Now this is even a bigger difference; for whereas the structuralist linguists with their account of phonological representations made no ontological commitments and no claims for *truth*, generative phonology is certainly committed to an ontological picture of psychological realism.

There are, of course, many other problems that Chomsky saw with the structuralist notions. However, for our purposes, what we are concerned here are these two major differences between structuralist and generativist linguistics –

²⁷ Plus, of course, they were also concerned with a universal theory of grammar.

namely, the concern for truth in linguistic theories and the claim for psychological realism that it gives rise to. That is, given the influence generative linguistics has on contemporary linguistic theory, contemporary phonology is certainly committed to this ontological picture as well – and the implication of this ontological picture, furthermore, indicates that a phonological representation is a representation of *something* that goes on in the mind. Notice, however, that there is nothing about mental *representations* in the quotation above: All it shows is that for Chomsky and the generativists, an explanatory theory of linguistics is necessarily *mentalist* and that phonology as such is committed to the existence of mental *processes*. Thus although the discussion above shows that phonology is committed to the existence of mental entities, at this point we cannot fathom whether a phonological representation is indeed a mental representation or not – in the sense that the philosophers use the term. But in order to decide this, we now need to look at what philosophers of cognitive science understand when they use the term “mental representation”.

The following chapter will therefore take on the job of defining and explaining in detail what a mental representation is as it is conceived of in the philosophy of cognitive science. Once I identify what a mental representation is, I will turn back in the fourth chapter to where we leave the debate on phonology here, and show in detail what the philosophical implications of the claim for truth and psychological realism are for contemporary phonology. As it will turn out, it is the precise notion of mental representations that will be defined in the next chapter that one can see at work in the actual practice of phonology.

CHAPTER III

MENTAL REPRESENTATION

Without any doubt, one of the main theoretical constructs of contemporary philosophical work on cognition has been the concept of mental representation. The idea that we mentally represent the outer world is at least as old as the early modern philosophers – if not the ancient Greeks. However, the idea that the mind works through *processes that operate on* mental representations of the outer world is an idea that is entirely novel: it is an idea that should be credited to the cognitive revolution that took place in the mid-twentieth century²⁸ – and it is this latter idea that has been the center of attention since. Scientific psychology that is based on information processing with a commitment to realism about mental representations is pretty much definitive of this cognitive revolution. Furthermore, as I will discuss below, such cognitivism in psychology by and large can be seen as a reaction to the *behaviorists* of the early twentieth century – who in turn had banished all forms of mentalism from their psychology. For now, it suffices to say that the central claim to behaviorism is that, since the inner states of an organism are not subject to (direct) observation, psychologists should proceed with the investigation of stimulus-response relationships of organisms – where the internal machinery of the organism (that mediates such stimulus-response) are to be left unexplored. Now, in the days of the emergence of primitive computers, cognitivists (both philosophers and

²⁸ It can also be argued that, given a certain reading, a similar idea can be found in Kant's *Critique of Pure Reason*, where he discusses the schematism of the pure forms of intuition. In any case, however, Kant's purposes were far different than the cognitivists; where the former was exploring the limits of experience through introspectionist means, the latter is interested in a naturalistic theory in order to explain how the mind works. I am therefore implying the naturalistic treatment on this conception of mental representation where I ascribe the credit to the cognitivists.

psychologists) thought of the idea of computation to be the key to overcoming the behaviorist pessimism: they wanted to study the mind, and they were not satisfied with the stimulus-response psychology that conceived of the mind as a *black box* – and that hence abolished the study of mental phenomena. Computers are basically symbol manipulators – they operate on a physical medium and yet they do abstract operations. Computational symbol manipulation, furthermore, is based on purely formal, syntactic (non-semantic) procedures and yet it gives rise to semantic output. Computers are therefore a great analogy if what you want is a working analogy of something that is physical, observable, and that has internal states that are private.

Now, if one pushes this analogy further to explain how the mind works, the concept of symbol manipulation essentially brings with it the question what the supposed manipulation is manipulation *of* – and this is precisely where the notion of mental representation comes in. Starting with cognitivists, therefore, this notion has more or less been the center of attention in the literature concerning the philosophy of psychology and cognitive science. Despite this, on the other hand, there is still nothing near to a consensus on what the exact nature or even the content of a mental representation is. In this chapter, my purpose will be to introduce the debates surrounding the issue of mental representation in order to present the issue as clearly as possible. I will firstly begin with explaining the representational theory of mind. Jerry Fodor is a central figure in this discussion, for two main reasons: first, he consistently wrote from a philosophy of cognitive science perspective and as such his work is expository of the background and nature of the issue – second, the current formulation of both the representational and the computational theories of mind owe much to his work. Hence, after I abstractly outline the borders of what a representational theory of mind is like, I will expand this discussion by introducing

Fodor's arguments for intentional realism: this position regarding the ontological status of mental states can be seen as a direct call for a representational theory of mind. After presenting this position, I will discuss behaviorism and show that the main reason in calling for a representationalist approach in psychology was to overcome the behaviorist skepticism about studying the mind (and mental phenomena in general). Overall, my aim with the discussion in the first section of this chapter is to contextualize and situate the very idea of mental representation.

In the second section, I will begin the discussion by eliminating the possible alternatives to intentional realism that are in the literature – this is essential, because these positions (neo-behaviorism, connectionism and instrumentalism) all deny intentional realism one way or the other; and all the classical forms of the representational theory of mind (and realism about mental representations) depend on intentional realism. After clarifying what I will *not* be talking about, I will firstly present the computational theories of mind. Such theories, in essence, assume the representational theory of mind and provide an account of *reasoning* (or more generally, *intelligence*). The central claim here is that intentional *processes* are computational (i.e. symbolic) operations that are defined over the syntactic properties of mental representations. The explanation of this claim will be of central importance for this thesis, because this claim stands for a definition of what a theory committed to mental representations would be like – and the next chapter will show that phonological theory is indeed one such theory. Hence I will explain that a computational theory of mind conceives of mental representations as entities with *syntactic organization*. Accordingly, furthermore, it is this syntactic constituency that is responsible for the semantics of mental representation. Secondly, I will exemplify the computational theory of mind with Fodor's "language of thought" hypothesis.

The Language of thought hypothesis is consistent with the computational theories of mind; but the central claim of this hypothesis will not be important for our discussion: Fodor's main claim here is on the nature of mental *states* rather than on mental *representations* – that all (or most) thought is linguistic. I will argue, however, that the arguments that the language of thought hypothesis provides to establish this claim tell us interesting facts about mental representations. The discussion here will involve the characteristics of mental representations that reflect what Fodor called the *productivity* and *systematicity* of language. Overall, the second section will aim to present the specifics of the nature and content of mental representations as the philosophers have conceived of them.

REPRESENTATIONAL THEORY OF MIND

The concept of mental representation (in theories of the mind) constitutes the grounds for one of the most widely debated discussions in the contemporary philosophy of mind. There is, naturally, no single (specific) theory to call *the* representational theory of mind. On the other hand, the theories that commit to mental representation, of course, share something – something that allows for the categorization of each such theory as *a* representational theory of mind. In their most general and abstract form, representational theories of the mind share two aspects – one involving ontological, and the other involving theoretical commitments.

Ontologically, a representational theory is built on the assumptions that (a) we inhabit a world occupied by objects and their properties, and that (b) we internally represent those objects and their properties. Furthermore, internal representation of the outer world is (c) mediated through the existence of mental representations as

intermediaries. Finally, most of the representational theories of mind are more or less consistent with naturalism: the claim that only natural laws (such as those endorsed in physics, chemistry, biology, etc.) operate on this world and all scientific endeavor must be consistent with an acknowledgement of such laws. Hence commonly, (d) the accounts for the existence of mental representations are hoped to be (ultimately) provided in terms of the features of the brain. On the other hand, representational theories of the mind commonly share the theoretical assumption that (a) mental states are about things in the world – that they are *intentional*. (b) Each mental state, construed as such, is defined as a specific *relation* to a mental representation, and (c) the intentionality of mental states is explained in terms of the semantic properties of mental representations²⁹ (Pitt, 2008). Henceforth, I will use the abbreviation “RTM” to refer to a theory that shares in all these ontological and theoretical commitments.

Now RTM was first and foremost endorsed by the cognitive psychologists, with research mostly conducted on the nature of mental processes that operate on mental representations – but the notion of mental representation was operationally defined by them, and it was largely articulated (and defended) by the philosopher Jerry Fodor. In the following pages, I will discuss Fodor and his arguments on mental state psychology in order to clarify and contextualize the abstract definition I have provided for RTM above. Fodor calls the position he is defending *intentional realism*, and throughout the defense of this position, he explicates the nature of mental *states* – as well as what it means for a mental state to be *a specific relation* to a mental representation.

²⁹ What these theoretical claims amount to will become clear in the discussion that follows.

Intentional Realism and Common Sense (“Folk”) Psychology

In one of his prominent works, *Psychosemantics*, Jerry Fodor proposed a theory of psychology that he calls commonsense psychology. This version of psychological theory, which is sometimes also referred to as *folk psychology*, embraces and assents to the explanatory value of everyday concepts such as “beliefs”, “desires”, “hopes” and “fears”, etc. among its theoretical inventory. Fodor claims that:

We have no reason to doubt – indeed we have substantial reason to believe – that it is possible to have a scientific psychology that vindicates commonsense belief/desire explanation (Fodor, 1987, p. 16).

According to Fodor’s own version of scientific psychology,³⁰ mental states (assuming there are such things) such as beliefs and desires can be understood as what is often called *propositional attitudes*. They are *propositional* in that they have a linguistic form; for instance my belief in the roundness of the world is expressed in the form “I believe that the world is round” – and for Fodor this linguistic form is the only expressible form for my mental state in believing that the world is round. My desire to eat a hamburger similarly amounts to the sentence “I desire to eat a hamburger”. What makes mental states *attitudes*, on the other hand, is the more interesting claim that such states have (intrinsic) intentionality: the claim that they are essentially *about*, or directed *towards* something. For instance my belief that is expressed by the proposition “the world is round” is *about* the world; my desire to eat the hamburger on the table is directed *towards* (eating) the hamburger on the

³⁰ I will show later on that, strictly speaking, Fodor’s own version of scientific psychology turns out to be a form of the representational theory of mind (i.e. the language of thought hypothesis), but this should become clear only after I discuss the debates on mental states and propositional attitudes in what follows.

table.³¹ Furthermore, Fodor argues that this latter property of the propositional attitudes – namely intentionality – seems to be a unique property: trees and rocks are not about something, and neither are numbers or people (Ibid, p. 11), *in the same way that* beliefs and desires are.

But in what way are mental states *about* things? Now Fodor discusses the intentionality of mental states by highlighting two aspects of propositional attitudes: firstly, they (propositional attitudes – and hence mental states) are *semantically evaluable*, and second, they have *causal powers*. First, consider what the *content* of a mental state is: suppose I believe that Immanuel Kant is a dead philosopher. Now the content of my belief is (obviously) that Immanuel Kant is a dead philosopher – but notice two facts: (a) that my belief is about the world only because its content is about a dead philosopher who lived (and died) in the world, and (b) that this content can be completely expressed *only*³² by the proposition “Immanuel Kant is a dead philosopher”. Suppose now, that I desire for Immanuel Kant to be a dead philosopher. The content of my desire is the same as that of my belief – namely that Immanuel Kant is a dead philosopher. Notice, furthermore, that the two facts that held for the content of my belief also hold for the content of my desire: namely that (a) my desire is about the world only because its content is about a (desired to be) dead philosopher (who, in compliance with my desire) lived and died in the world,

³¹ If intentionality by itself doesn’t clarify why philosophers call mental states “*attitudes*”, think of it this way: I have a settled way of thinking or feeling *about* something in both these examples, which in turn is reflected in my behavior (I can try to travel around the world, eat that hamburger, etc.) – hence the preference to call mental states “attitudes”. In any case, it is the *intentionality* of the mental states that is important here.

³² One might argue that this is not true, that, for instance, a picture is worth a thousand words (or a hundred propositions, for that matter). However, pictorial (and otherwise perceptual) representations are always open to interpretation, and are not the kind of things one can judge to be true or false, but beliefs are. The objection can of course, press by claiming that propositions are similarly open to paraphrasing and that we should not be hasty in ascribing the same truth value in all possible paraphrases, and so on and so forth. In any case, it should be remembered here that I am presenting Fodor’s (and a vast majority of philosophers’) conception of mental states as propositional attitudes and not providing a critique of it.

and that (b) the content of this desire can be completely expressed *only* in terms of the proposition “Immanuel Kant is a dead philosopher”. Finally, notice that the contents of both my mental states (in believing and desiring that Kant is dead) *are the same*, and they are both exhaustively expressed both by the same proposition (“Kant is dead”). That they (contents) are *propositions*, should not be surprising, given that “propositions exist to be what beliefs and desires are attitudes *toward*” (Ibid, p.11); that they are *the same* propositions, could be surprising – but it shows that these two mental states of mine are *in two different functional relations to the same propositional mental content*. That is, believing that Kant is dead, and desiring that Kant is dead are two functionally different relations to the same propositional content that “Kant is dead” – the former can be true or false, whereas the latter can be fulfilled or frustrated. On the other hand, (in this case) what makes my belief indeed *true* is the fact that Immanuel Kant is indeed a dead philosopher: in other words, what makes a belief true or false is the evaluation of its propositional content with respect to the actual world. Similarly, what makes my desire *fulfilled* is the fact that Immanuel Kant is indeed a dead philosopher: again, what makes a desire fulfilled or frustrated, therefore, is the evaluation of its propositional content with respect to the actual state of the world. Now since *all* mental states can be evaluated in terms of their propositional contents one way or the other; and since propositions are *always* meaningful (they represent something in the world) – Fodor claims that one consistent aspect of mental states is their *semantic evaluability*:

If you know what the content of a belief [or any other kind of mental state/ABE] is, then you know what it is about the world that determines the semantic evaluation of that belief [or any other mental state, respectively/ABE] (Ibid, p.11 – original emphasis).

In short, the constraint of semantic evaluability consists in the fact that mental states (i.e. the propositional attitudes that one has about the world) can be *evaluated* with respect to their specific (functional) relation to the non-psychological world: beliefs can be true or false; desires can be fulfilled or frustrated, etc.

On the other hand, mental states are *causally involved*: they can *somehow*³³ cause behavior (hence we can say that attitudes are reflected in behavior), they can be formed (caused) by environmental events and they can evoke other mental states (as it happens in a chain of thought). However, this is not the end of the story. Semantic evaluability shows that mental states have meaningful content – but it is also the case that meaningful contents hold consistent relations among each other. That is, any coherent theory of semantics must acknowledge the fact that the propositions “Kant is an alive and healthy philosopher” and “Kant lived and died in Prussia” cannot be true (or false) at the same time. Now it is apparently true also that mental states have causal relations among each other that *reflect* such consistent relations among meaningful contents. To clarify this point, suppose that I did not know that Kant is dead, and I for some reason – say, out of ignorance – held the belief that Kant is an alive and healthy philosopher. And suppose, also, that I read in a book (that I trust the author completely) that Kant never left Prussia – he lived and died there. Now the most likely thing to happen is that I will abandon my belief that Kant is alive, and form a new belief that Kant is dead³⁴ – but it is unlikely that I will believe that Kant is alive and *also* believe that Kant lived and died in Prussia. What

³³ At this stage, no discussion on a theory of causality is necessary – as it only suffices, for the present purpose, that mental states somehow cause behavior, and are somehow caused by external stimuli.

³⁴ I believe, however, that it is perfectly plausible to object to this: After all, people can be absent minded or irrational at times which will not result in the scenario I have just conveyed – which assumes people to be rational agents. I am not sure if that assumption holds (at least at all times), but then again I don’t need to be, for I am presenting Fodor’s position and he apparently makes this assumption. What is important here is that people indeed have a capacity to be rational – which is the scenario described here.

happens in this scenario shows that the causal relations between mental states are *symmetric* with the consistent relations that hold for their meaningful content (i.e. propositions) – that is, causal relations among mental states seem to somehow parallel the logical consistency of semantics. I will return to this point later on, where I will discuss the language of thought hypothesis; but for now, the upshot of the story pertaining to the semantic evaluability and causal involvement of mental states according to Fodor is that trees, rocks and numbers can never satisfy these two constraints (not even people considered as entities) and thus they are never *about* things in the same *intentional* way as mental states are.

Now Fodor's main argument in *Psychosemantics* seems to be the claim that, since our everyday concepts such as believing, desiring, fearing etc. have explanatory value in that they are tokens of (semantically evaluable and causally involved) propositional attitudes, it follows that the implicit generalizations of commonsense belief/desire psychology are largely true of our mental states (Ibid, p.10) – thus we better treat every day folk psychology as a serious scientific psychology. As he colorfully puts it:

[T]he predictive adequacy of commonsense psychology is beyond rational dispute; nor is there any reason to suppose that it's obtained by cheating. If you want to know where my physical body will be next Thursday, mechanics – our best science of middle-sized objects after all and reputed to be pretty good in its field – *is no use to you at all*. [And importantly neither is neuroscience, biology, nor physics.] Far the best way to find out (usually, in practice, the only way to find out) is: *ask me!* (Ibid, p. 6 – original emphasis)

Fodor's proposal of commonsense psychology hence boils down to his defense of the position he calls *intentional realism*: beliefs, desires and such other mental states are real; they are causally involved (have causal powers), determinate and contentful

states that are subject to semantic evaluation. Folk psychology is hence in fact a scientific theory of the mind that gives us accurate predictions by using the language of mental states (i.e. propositional attitudes), and social life would be impossible without it.

The question here is why Fodor is insisting on such a seemingly controversial claim. After all, it is not so commonsensical to claim that *everyone* is actually a psychologist; nor is it commonsensical to think that psychologists from Freud to cognitivists tried to come up with theories of psychology in vain – where literally everyone already holds and practices a working one. What then, should we make of the claim that folk psychology is in fact a working scientific theory? At the heart of Fodor's argument for common sense psychology as a scientific theory is the idea that mental states are *representational* in nature: mental states are various functional relations to the mental representations of the world. As I have discussed, this idea comes from the semantic evaluability constraint on the intentionality of mental states. For instance, my belief that there is a hamburger on the table is semantically evaluable *because* it has a truth-function relation to the mental representation of the real object in the world (namely the hamburger); and my desire to eat that hamburger has a different kind of functional relation to the same mental representation – this should be clear enough given the discussion above. Furthermore, Fodor explains the causality constraint on the intentionality of mental states by claiming that mental *processes* are causal sequences of tokenings of mental representations (ibid, p. 17) – as in the event of my belief in the existence of a hamburger on the table causing a desire to eat it. Now in the order Fodor chooses to present his argument in *Psychosemantics*, his indulgence into RTM comes second to his defense of intentional realism:

Because believing and desiring are representational states, realism about belief/desire explanations leads one, by a short route, to worry about representation (Ibid, p. x).

Hence, it appears that *alongside* his defense of intentional realism, Fodor argues that the reasons for believing why RTM may be true reside in (a) the fact that one version or the other of RTM underlies practically all the scientific research being done in contemporary cognitive science,³⁵ and that “our best science is ipso facto our best estimate of what there is and what it’s made of.” (Ibid, p. 17) – and more convincingly that (b) there is no alternative to RTM in explaining the symmetry between causal relations among mental states (given intentional realism) and the semantic relations that hold among their propositional objects (Ibid, p. 18). In other words, and as I have discussed above, the objects of mental states (i.e. propositions) have meaningful content and hold consistent relations between such content; and mental states have causal relations among each other that reflect these consistent relations between the content of their objects – and RTM is (Fodor goes on claiming) *the only* theory we have that explains this reflection (i.e. symmetry). Hence, in the way Fodor presents his story, it appears the argument is that RTM is the best candidate to establish folk psychology as a scientific theory (given one intends to be a realist about the theoretical entities a scientific theory postulates), because *there is no other alternative*.

However, I believe this order of presentation is misleading: It still is not clear why Fodor is committed to everyday folk psychology as a scientific theory or why he is strongly arguing for realism about mental states and representations. It even seems circular: why folk psychology? Because RTM works for it. Why RTM? Because it is

³⁵ This claim is after all what we are looking for showing in this thesis, whether we can say this as well for the actual practice of the linguistic branch of *phonology* – as it is part of the cognitive scientific endeavor.

the only alternative that can explain folk psychology. The question we must address is this: there is no alternative (to RTM) *for what purpose?* And then the answer is clear: there is no alternative if what you want is *a working theory of how the mind works*. In what follows, in order to make sense of the intentional realist position, I will discuss what Fodor is responding to when he demands realism about the theoretical entities a scientific psychology must make use of. The position Fodor is reacting against is nothing other than behaviorism (and sometimes a reductive equal of it – eliminativism). As it will become clear below, Fodor's demand for representationalism has the main purpose of establishing a theory of how the mind works in the face of a tradition of psychologists that denied such a purpose – and it is the only alternative in doing so.

The Motivation for RTM: The Rise and Fall of Behaviorism

In the seventeenth and early eighteenth centuries, it was characteristic of philosophers to speculate on psychology without indulging into much empirical work. The then dominant views on the mind (and ontology in general) were roughly divided into the rationalist (Descartes, Leibniz, Spinoza, etc.) and the empiricist schools (Locke, Berkeley, Hume, etc.) – both of which endorsed an introspectionist approach in their philosophy (Smith & Thomasson 2005). Late eighteenth and nineteenth century philosophy can by and large be seen as a critique of these early modern philosophical schools; but this critique was a critique of their epistemology and ontology – not their introspective methodology. However, with the advancements in the natural sciences in the late nineteenth and early twentieth centuries, a growing tendency among some philosophers³⁶ surfaced as a widespread

³⁶ One good example here is the neo-Kantian philosophers of the early twentieth century and especially those in the Vienna Circle. One of these neo-Kantian schools of philosophy, The Marburg

belief that in general philosophy should orient itself with respect to the natural sciences rather than striving to forge a metaphysical system of its own – that is, independent from the results of the positive sciences (Kim, 2003). This belief was (setting socio-political factors aside)³⁷ largely based on the dynamicity of the principles underlying the study of natural sciences: scientific explanations are based on hypothetical statements that are verifiable or falsifiable through methodical observation of the phenomenon in question, and such explanations are subject to change accordingly – as opposed to the fixed axioms analogous to those of introspective and speculative metaphysics.

Hence the twentieth century witnessed a widespread skepticism among philosophers towards introspective, a priori theorizing – and the general abandonment of metaphysics in favor of naturalistic explanations of what is observable. The mind, however, is *prima facie not* an entity that is observable through naturalistic means: the supposedly mental phenomena, such as emotions, thoughts, sensations, beliefs and desires, etc. are at best private – observable only to oneself and only through introspection, that is, if observable at all. What is publicly observable, on the other hand, is an organism's behavior that is (at best) expressive of mental phenomena. All these considerations gave rise to two skeptical ideas that dominated the early twentieth century thinking: firstly, the idea that the existence of

School, can be seen as the historical ancestor of analytic philosophy and in particular the field of the philosophy of science – and was founded by Hermann Cohen and continued afterwards by Paul Natorp and Ernst Cassirer until the mid-twentieth century. For a detailed discussion, see: Kim 2003, Friedman 2000.

³⁷ Social and political aspects of science were largely ignored by the early philosophers of science, but this is changing recently. With the involvement of sociologists (such as Robert Merton, see for example: Merton, 1973) and philosophers that are sensitive to social issues such as feminism (e.g. Noretta Koertge, Koertge, 2003), philosophers of science are becoming more and more reflexive about the social and political influences on science and its values. For a detailed and comprehensive discussion of this subject, see for instance: Longino, 2006. Even if such socio-political factors were at play during the call for the naturalization of philosophy, however, such factors do not weigh over the arguments and assumptions I am discussing here.

mental phenomena is doomed to be a discussion topic of metaphysics and not of naturalistic science, and secondly, the idea that if we are to do psychology at all, we need to focus on behavior rather than on mental phenomena – the ontological status of which was seen as questionable.

These skeptical ideas were further intensified by the success of the work being done by the psychologists of the time, such as the reflex systems research that was being executed by Ivan Pavlov. This led to a movement in psychology that is commonly known as *behaviorism*. In 1913, John Watson, a distinguished psychologist who was as well reflexive to the methodology of psychology as was influenced by Pavlov's work, published an article titled "Psychology as the Behaviorist Views it" – an article that is close to a manifesto of the behaviorist school in psychology with its opening words:

Psychology as the behaviorist views it is a purely objective experimental branch of natural science. Its theoretical goal is the prediction and control of behavior. Introspection forms no essential part of its methods, nor is the scientific value of its data dependent upon the readiness with which they lend themselves to interpretation in terms of consciousness. The behaviorist, in his efforts to get a unitary scheme of animal response, recognizes no dividing line between man and brute. The behavior of man, with all of its refinement and complexity, forms only a part of the behaviorist's total scheme of investigation (Watson 1913, p. 158).

Behaviorism took its toll on the study of mind in philosophy: philosophers of the behaviorist tradition conceived of the mind as an entity analogous to a black box, where the inside of the box represents an organism's internal states. The idea here is that since you can say nothing of the box, you have to make do with what you put in the box, and what you get from it; that is, behavior, the only object of psychological study, can only be studied through the responses organisms display given certain

stimuli – and the internal mechanisms responsible for binding stimulus to response are hopelessly out of reach for one to talk about.

What is basically behind this story is that behaviorists had the supposition that the logically necessary and sufficient conditions for behavior can be given *only* in terms of stimulus-response variables – and thus they claimed that there is *no way* to analyze an organism's internal states (i.e. mental states – that are involved in the *causation and explanation* of behavior for Fodor.) On the other hand, behind Fodor's argumentation, as I have discussed above, is the idea that the representational theory of mind is *the only way* to analyze an organism's internal states: there is no alternative. Put this way, it should be clear that behind Fodor's intentional realism is the reaction that, in the face of behaviorist pessimism about studying how the mind works, the only resolution is to study the mind through the conceptual tool that is mental representation. It should also be noteworthy that when Fodor appeals to the fact that one version or the other of RTM underlies practically all the scientific research being done in contemporary cognitive science, as I have discussed above, he is echoing the early twentieth century tendency in believing that philosophy should orient itself with respect to the natural sciences. And behind his interpretation of the success of cognitive psychology is the statement that, "if you have a trouble with the thing, get it fixed" (Fodor 1978, p. xi) – where the "thing" in question is nothing other than mental representation. That is, rather than pointing out to the (ontological and epistemological) difficulties in understanding mental representation and claiming we should stick with the stimulus-response psychology (as the behaviorists do) – one should instead try to reflect on these difficulties and come up with a coherent representational picture of the mind (as the cognitive psychologists have done so by operationally defining the notion of mental representation). For Fodor, what remains

for the philosopher to do is to articulate the content and nature of this operationally defined concept of mental representation – and as I will discuss further below, this is precisely what he aims to do.

Hence, *prima facie* Fodor appears to be mainly defending folk psychology as a perfectly valid and coherent theory of the mind (via assuming intentional realism). However, given the consideration that he is responding to a behaviorist attitude about studying the mind, it seems in fact to be the case that the strength of folk psychology stands in Fodor's argumentation as a *vindication* of intentional realism – and if intentional realism is true, we can study the mind through studying intentional states. If this is the case the strength of folk psychology must be shown; indeed vindicated itself. But this is precisely what RTM stands for:

RTM vindicates commonsense psychology for what RTM identifies as the core cases; in those cases, what common sense takes to be tokenings of propositional attitudes are indeed tokenings of a relation between an organism and a mental representation (Fodor 1987, p. 21).

RTM in turn is a working theory that we can evaluate from the success of cognitive psychology. What remains, then, is to leave the behaviorist worries aside and study how the mind works through articulating specifics for RTM – and the nature and content of mental representations.

THE NATURE AND CONTENT OF MENTAL REPRESENTATIONS

Let me capture the story so far. Fodor defends *intentional realism* – a position that claims:

- (1) *mental states are real*: our everyday “folk” concepts, such as beliefs, desires, hopes and fears refer to actual states of the mind (those that can be explained in terms of the brain).
- (2) *mental states are intentional states*: they are essentially about, and directed towards the (non-psychological) world.
- (3) *the intentional object of a mental state is a proposition*: mental states are propositional attitudes; that is, mental states are (intentional) attitudes that are *toward* propositions. A mental state of desiring X is an attitude that is directed towards making true that X – where X is a proposition, such as, “I am eating a hamburger”.
- (4) *mental states are semantically evaluable*: to say of a belief true or false is to evaluate that belief in terms of its relation to the world. To say of a desire fulfilled or frustrated is, again, to evaluate it in terms of its relation to the world. Mental states are semantically evaluable as such because of the fact that (c) their intentional objects (i.e. propositions) are semantically evaluable (Fodor 1987, p. 11). Hence the belief that X is Y – is true or false with respect to the evaluation of the proposition “X is Y” in terms of the actual world. Similarly, to desire X to be Y – is fulfilled or frustrated with respect to the evaluation of the proposition

“X is Y” (such that, the desire for X to be Y is to desire for the proposition “X is Y” to be true.)

(5) *mental states have causal powers*: they can cause behavior as well as they can cause other mental states, and they can be caused by environmental stimuli – for instance, the event of my seeing a bear in front of me can cause my mental state in believing that there is a bear in front of me, which, in turn, can cause my mental state of fearing to be attacked by that bear.

What RTM adds to this is the claim that such intentional mental states are various relations to mental representations – *which, in turn are the intentional objects of mental states*. Mental representations as such mediate cognition between the world and mental states. Now it is possible to get off the bus at this stage by denying intentional realism (even before it comes to RTM): firstly, it is important to see that behaviorism is an epistemological *attitude*. As I have discussed above, the behaviorist tendency to conceive of the mind as a black box comes not from ontological arguments, but on the contrary from skepticism towards those kinds of arguments. Hence, one can stick with the behaviorist principles and deny the proposal for replacing the black box with an “*intention box*” (cf. Fodor, 1987, p. 137) – after all, the proposal for the intentional realist position is not grounded on evidence that is based on observation.³⁸ Gilbert Ryle can be seen as the paradigm example of a philosopher that shares this behaviorist attitude – where he argues that a modern version of the Cartesian (substance) dualism lies implicitly behind all mental

³⁸ Sure, you can observe people reporting states such as beliefs and desires; or maybe even observe them being in such *behavioral* states, but that doesn’t mean you observe people’s *mental states*. Even if Fodor is right in thinking that our everyday, folk psychological concepts – that we use to understand and predict each other’s behavior – entail the reality of mental states, this entailment comes from ontological assumptions and commitments and not from observation.

(and any kind of inner) state talk (Ryle, 1949). Ryle thinks that these implicit Cartesian assumptions lead to a category mistake when philosophers use a mentalistic language – as in Fodor’s argumentation for realism about mental states. Secondly, one can refute the claim for realism about mental states (directly (a) above) by refuting the coherency and consistency of the folk psychological concepts that mental states are vindicated by. This kind of *elimination* of intentional mental states from one’s theory of psychology has generally been defended by Paul and Patricia Churchland within the position known as “eliminative materialism”.³⁹ The Churchlands roughly claim that our everyday folk-psychological concepts (such as believing, desiring, etc.) are poorly defined – and since there can be no mental state talk without reference to such (ill-defined) folk concepts, no coherent neural basis could ever be found for mental states. Accordingly, psychological concepts in a well-defined theory of the mind must be judged by how well they reduce to the biological level (Lycan & Pappas, 1972) – and we must not bother trying to find neural correlates of ill-defined concepts.⁴⁰ Third, one can be an *instrumentalist* in believing that while mental state ascriptions and explanations are not *literally* true, they can be quite useful *theoretical* (fictive) constructs in interpreting behavior. Such a conception therefore denies realism about mental states (one cannot find their correlates in the brain), but does not refute the study of them (one can study the mind

³⁹ See, for example: Churchland, 1981.

⁴⁰ From such a perspective, the Churchlands have been championing the cause of *connectionism* in cognitive science – where the mind is studied through *neural networks*. Such networks are conceptualized structures of nodes and their connections. These structures can be studied in terms of how well they reduce to the biological level (the brain) that they aim to replicate. They can also be implemented using computers and hence virtually simulate the brain. In such a conception of cognitive science, however, there is no apparent reason to claim for realism about mental representations. Fodor and Pylyshyn have provided a detailed discussion and critique of connectionism, see: Fodor & Pylyshyn, 1988. However, because connectionism in cognitive science is not necessarily committed to the notion of mental representation, it would be irrelevant to discuss further in this thesis.

qua representational). Daniel Dennett can be seen arguing for this point – where he claims that intentional concepts (beliefs, desires, etc.) help with behavior interpretation and explanation by providing a unique explanatory *stance*; which is however distinct from the physical stance (Dennett, 1987). If one adopts the intentional stance, however, one neither finds it plausible to discuss the nature or content of mental states, nor can expect to find neural correlates of them – for they are accordingly only instrumental and conceptual tools in explaining behavior and have no essential reality to explore. Notice, furthermore, that all these alternative suggestions to intentional realism are not necessarily committed to the concept of a mental representation – this is because, as I have discussed above, realism about mental representations makes sense only when one is committed to realism about intentional states – and all these positions deny intentional realism one way or the other. However, since my purpose in this thesis is dealing with the issue of whether what philosophers say about mental representations fits in with the practice in phonology or not – explaining these positions in further detail would be irrelevant; I will hence discuss them no further. What still remains on the other hand, is to discuss the specifics of mental representation – and theories that fill in those specifics.

The Computational Theories of Mind

Since the demise of behaviorism in the mid twentieth century, a general tendency has risen to conceive of the human brain as a biological computer. In the philosophy of mind, this tendency has resulted in the widespread acceptance of the computational theories of mind – which, accepting RTM in the abstract form I discussed above, expand on it by providing an account of *reasoning* – that is, an account of *intelligence*. A computational theory of mind (henceforth CTM) conceives of a

mental representation as a *symbolic* entity – and accordingly a symbolic mental representation has both *syntactic* and *semantic* properties. As I have discussed above, RTM postulates mental representations to be the intentional objects of mental states – and they have semantic properties because they are propositional. Now, the main thesis of CTM is that the processes of reasoning are responsive only to the syntactic properties of mental representations – that:

CTM: Cognitive processes are *computational* operations defined over the *syntactic properties* of *symbolic* mental representations (that have *intentional content*).

I have already explained above what it means for a mental representation to have *intentional content*: it means accepting intentional realism. On the other hand, writing an extensive definition for the notion of computation would be a master's thesis of its own – maybe more. Given the shortness of space I have that I can reserve for this purpose here, it will have to suffice that computation – in the technical sense that it matters here – is *formal manipulation of symbols*.⁴¹ Now consider what a symbol is: something that stands for, *represents* something else.

Given that mental representations are indeed *representations* of the world, it should not be confusing why they are symbolic – but CTM goes further in claiming that

⁴¹ The notion of “formal symbol manipulation” (and hence, computation) ultimately traces back to the late nineteenth and early twentieth century crisis in mathematics. As I have hinted above (while I was discussing the rise and fall of behaviorism) those years saw a widespread call for naturalization, and skepticism towards introspectionism. Part of this was related to the fact that mathematics was undergoing a reformation for *formalization*. Euclidean geometry had the postulation that parallel lines never cross; yet the plausibility of this postulate lied not in the axioms of the Euclidean system but in its being common sensical. The crisis and the project of formalization began when mathematicians (such as Lobachevsky) postulated consistent geometries where parallel lines crossed (in sharp contrast to the Euclidian intuition) – the bare possibility of this shows that intuition cannot be trusted and all mathematics has to be based on formal, axiomatic systems. Such formal systems demand mathematical functions (and postulations) to be *symbolic*: where the rules of derivation are explicitly expressed in terms of the *syntactic properties* of the symbols that are employed. Some mathematical functions, however, cannot be formalized as such – and the ones that *can be* have been famously called “computable” by Alan Turing (Turing 1936).

such symbolic mental representations are “syntactically organized”⁴² (Fodor, 2001, p. 4 n.5) with respect to formal (well defined) rules. That is (according to CTM) a mental representation is an organized medium of brought together smaller units – and this organization is structured (i.e. brought together) with respect to the formal rules of computation. *Reasoning* according to CTM, then, is the execution of computational operations on this medium.

This is precisely where the analogy between brain and computer originates. A (Von Neumann architecture) computer works pretty much like a formal system of symbolic manipulation. Hardwired in the CPU of a computer, there are electric circuits that are set up such that the patterns of electrical activity in the computer can be interpreted as *instructions* (e.g. add two numbers) – that is, the circuitry of a computer is such that, the causal (electric) events on the CPU match the meanings of instructions. Now when you run a computer program, what happens is that the CPU stores the program into the computer memory, and assigns the program a process id⁴³ – which helps the CPU to treat (and track) that program as a *symbolic entity*. Throughout the runtime of the program the CPU executes instructions in accordance with the algorithm of the program – which, in short, amounts to various symbol manipulations given that the computer treats the program as a symbolic entity. It should be clear that the program that is stored into the computer memory as such has both syntactic properties and semantic properties: it is stored and executed via its

⁴² From ancient Greek, *syn*: “together” and *tassein*: “arrange” - syntax by definition already means “arranged (organized) together”. What the idea of computation adds is that the organization is with respect to formal rules.

⁴³ Modern computers are famous for being *multitasking* – a feature that allows the computer to run more than one tasks at the same time. However, conceptually, what the computer does is *not* running different programs at the same time; on the contrary a CPU runs bits of all the programs (with different process IDs) one by one in rotation and this works seamlessly only because modern computers have exceptional raw power, that is, they run too fast to notice this. Hence, multitasking has nothing philosophical to do with the idea of computation itself.

syntactic properties, but, in the end, it amounts to something semantic – such as the addition of two numbers. It is because computer programs are *such* symbolic entities that CTM comes forward with the analogy of mind as computer – where, given the naturalistic commitments, the computational medium for the mind is the brain.

Now Jerry Fodor has been one of the main proponents of CTM:

[T]he Computational Theory of Mind ... is, in my view, by far the best theory of *cognition* that we've got; indeed, the only one we've got that's worth the bother of a serious discussion. There are facts about the *mind* that it accounts for and that we would be utterly at a loss to explain without it; and its central idea – that intentional processes are syntactic operations defined on mental representations – is strikingly elegant. There is, in short, every reason to suppose that the Computational Theory is *part of the truth* about cognition (Fodor, 2001, p. 1 – my emphasis).

However, as it should be clear from the quotation above, Fodor thinks CTM has its limitations.⁴⁴ First and foremost, CTM is a theory of *reasoning* and it tells *part of the truth* about the mind because it postulates only on *one* cognitive aspect of the mind (i.e. reasoning) – hence only part of the story.⁴⁵ That said, CTM provides a striking insight about the nature of mental representations: namely that they are symbolic entities with both semantic and syntactic properties. Now mental representations

⁴⁴ Fodor's *The Mind Doesn't Work That Way* is indeed all about these limitations. He argues that CTM, for example, has nothing to do with *consciousness*, (Ibid, p.1, n.1) and although it is the best theory we have on how the mind works, it couldn't possibly comprise more than a fragment of a full and satisfactory cognitive science. Practitioners of AI and sometimes cognitive psychologist (especially nowadays) have a tendency to miss this point.

⁴⁵ Antonio Damasio, for example, argues that this rationalist conception of cognition (as formal symbol manipulation) does not work in most cases of intense emotion or feeling – even in some cases of decision making (Damasio, 1994). CTM provides an account for *the cases it works*, however; and hence it provides an account of one of the human cognitive *capacities* – namely *reasoning*. This, however, has also been recently challenged: “We must not assume cognition to be purely internal, symbolic, computational, and disembodied, but seek out the gross and detailed ways that language and thought are inextricably shaped by *embodied action*.” (Gibbs, 2006, p. 9). Once again, however, our aim here is not to discuss whether CTM is true or not – but rather, to show what it tells about mental representations. And given that there are no (necessary) representations in the “Embodiment” thesis on cognition, it would not be relevant to discuss here.

have their semantic properties through their propositional nature (propositions have meanings) – but what exactly are syntactic properties?

Well, to begin with: Syntactic properties are peculiar. On the one hand, they're among the 'local' properties of representations, which is to say that they are constituted entirely by what parts a representation has and how these parts are arranged. You don't as it were, have to look 'outside' a sentence to see what its syntactic structure is, any more than you have to look outside a word to see how it is spelled. But though it's true that the syntax of a representation is a local property in that sense, it's also true that the syntax of a representation determines certain of its relations to other representations. Syntax, as it were, faces inward and outward at the same time (Fodor, 2001, p. 20).

There are two important issues here that tell us about the nature of mental representations: first; that (a) its syntactic nature renders a representation as a constituency (a collection) of organized parts – and second, (b) that representations are related to each other through their syntactic organization. Both of these bring emphasis on the importance of *organization*: “the identity of a mental representation cannot survive alteration of its syntax” (Ibid, p. 24). For instance, a mental representation of the sentence “John loves Mary” cannot be reorganized into “Mary loves John” (or “*loves John Mary”) without it being another representation (or losing its semantics). This furthermore shows that semantic properties of a mental representation *supervene upon* syntactic properties, as is evident with the sentences “Mary loves John” and “John loves Mary” where different syntactic constituents result in different meanings. The fact that syntactic properties are responsible for the relations among mental representations further adds to this; the mental representation of the sentence “John loves Mary” is related to the question form “who loves Mary?” – and although it seems that the relation is semantic in this example, the syntactic

constituency is the reason why “Who loves John” is not related to “John loves Mary”.⁴⁶

Furthermore, since mental representations are related to each other through their syntactic properties, it is only possible to get from one representation to the other (as it happens in a chain of thought – e.g. “who loves Mary?”, “John loves Mary.”) through syntactic properties. Hence, ultimately, CTM is a theory of reasoning – but what it brings to the table in terms of mental representation is that, the syntactic properties of mental representations are the most important ones: it is because of their syntactic nature (constituent structure) that representations are responsible for reasoning – it is because of their syntax that mental representations constitute domains over which mental *processes* are defined.⁴⁷ As John Haugeland has once commented; “if you take care of the syntax, *the semantics will take care of itself!*” (Haugeland, 1981 – original emphasis) – because semantics supervenes upon syntax. Now, given the widespread acceptance of CTM in the philosophy of cognitive science, and the lack of alternatives that build on RTM, the details of syntactic organization become critical in telling the story about how mental representations are conceived by philosophers. I will aim to shed some light on this below where I will discuss one of Fodor’s arguments for his language of thought hypothesis.

⁴⁶ And of course, the syntactic properties for sentences also apply to mental representations, if RTM is true.

⁴⁷ “If you think of a mental process – extensionally, as it were – as a sequence of mental states each specified with reference to its intentional content [i.e. a mental representation], then mental representations provide a mechanism for the construction of these sequences; they allow you to get, in a mechanical way [i.e. syntactic/computational], from one such state to the next *by performing operations on the representations.*” (Fodor, 1987, p. 145 – original emphasis). Accordingly, if one does not acknowledge mental representations as domains for defining mental processes, one is left (as far as our current understanding goes) with no other choice but to define mental processes as “unknown neurological mechanisms” (Ibid, p. 147).

CTM Exemplified: The Language of Thought Hypothesis

A decade before publishing *Psychosemantics*, Fodor argued in *The Language of Thought* for a hypothesis that he called “the Language of Thought hypothesis” (LOT, from now on). LOT is a claim that is grounded on the intentional realist position – and it assumes CTM (and of course, therefore RTM). Its essential claim is that there is a language of thought – a *mentalese* – that is innate, unconscious, and equal in expressive power to any natural language – albeit it is not identical to any of the natural languages. After its first articulation in *The Language of Thought*, Fodor developed LOT through *Psychosemantics*, various articles, and finally recently *LOT 2* – throughout which he has been consistent with his arguments.

Firstly, LOT borrows the idea regarding the intentionality of mental states from intentional realism – and claims that there *are*⁴⁸ semantically evaluable mental states with intentional content. Second, from RTM, it borrows the idea that mental states are functional relations to mental representations – and claims that **(a)** there *are* mental representations that are the intentional objects of mental states and **(b)** that those mental representations are the kind of entities that explain the nature of thought because of this. It finally borrows the idea that mental representations are symbolic entities from CTM – and claims that mental representations have constituent structure with respect to their syntax. Now, what it goes further and adds to these – and this is therefore the central claim of LOT – is the claim that *mental representations are symbolic entities that have not just any constituent structure but a linguistic one*. And since what you say on mental representations explains the nature of thought (by RTM), it follows that thought has a linguistic structure as well – that there is a language of thought; a *mentalese*. However, here, we are concerned

⁴⁸ Ontological commitment emphasized.

more with the *arguments* LOT provides to establish this claim and less with what it actually *claims* – LOT is mainly a claim on the nature of thought (i.e. on mental *states*), but it establishes this claim through arguing on the nature of mental representations. In what follows, I will highlight the relevant aspects of this argument and show what it implies about the nature of mental representations.

Now the classical argument for LOT assumes that there are two important features of thought: *productivity* and *systematicity*. But anything that explains the nature of thought must then be productive and systematic as well – so mental representations are productive and systematic. Furthermore, we can see these features in *language*; and not only that – language is also the *only* phenomenon we know of that has both a constituent structure (like mental representations, given CTM) *and* is productive and systematic. But then, the best way to explain the productivity and systematicity of mental representations should be to claim that they are linguistic entities – and since they explain thought, this carries over to thought as well: we have a language of thought. Notice that such an argument is an inference to the best explanation that aims to show the productivity and systematicity of thought through the productivity and systematicity of language:

Aunty, reading over my shoulder, remarks that this has the form of affirmation of the consequent. So be it; one man's affirmation of the consequent is another man's inference to the best explanation (Fodor, 1987, p. 149).

I discussed in the previous chapter that there is no limit to the number of sentences in a natural language. For instance, in all the languages, one can conjoin two different sentences, change nouns or verbs, iterate adverbs (such as “very ...” and “very, very ...”), or even use numbers to modify nouns (“1 cat”, “3 cats”, etc.) – in short, it is

possible to play with sentential constituents, and the possibilities are literally endless. Different constituents can come together to *produce* new sentences. This aspect of language is, for Fodor, the *productivity* of language. But it would, of course, be impossible for languages to embody this feature if sentences did not have *constituent structure*. On the other hand, closely related to the productivity of language, the feature of systematicity can best be understood if one reflects on the linguistic ability to understand and produce sentences: “the ability to produce/understand some of the sentences is *intrinsically* connected to the ability to produce/understand many of the others” (Fodor, 1987, p. 149). Fodor explains this by comparing how one learns their native language to *trying*⁴⁹ to learn a language from a phrase book. It is not possible to find anyone, for instance, who is a native speaker of English (or who can be judged to know how to speak English by native speakers) that can entertain the sentence “John loves Mary” – but yet cannot entertain “Mary loves John”. The story with the person who consults a phrase book, on the other hand, is entirely different: if you didn’t know any English, you could learn from a phrase book that the sentence “John loves Mary” means John loves Mary (in your own language) but wouldn’t have the slightest idea how to say “Mary loves John”. Now if our linguistic capacity had a nature similar to that of the phrase book story, it would be possible to have native English speakers around that knew how to say “John loves Mary”, but not “Mary loves John”. But it is not – it is not, because productivity and systematicity are features of languages and languages have constituent structure with a certain kind of syntactic organization. Thus productivity and systematicity run together on the basis of syntactic organization; and “if you postulate mechanisms adequate to account for

⁴⁹ Emphasis is on *not knowing* that language.

the one, then – assuming you’re prepared to idealize⁵⁰ – you get the other automatically” (Ibid, p. 150).

Now as I have said, Fodor claims *thought* has systematicity and productivity, too: it is *ideally* possible to think ad infinitum. That is, one can think *productively* as well as one can express these thoughts in sentences. But thought is *systematic* as well: if I can think of a cat on the mat, I can think of a *mat on the cat*, too. Thus anything that should explain thought should also be productive and systematic as well – and since we are committed to RTM, and our explanation of thought is through mental representations, it follows that mental representations should be systematic and productive as well. LOT, therefore, hypothesizes that mental representations also have productivity and systematicity – and by assuming CTM, it answers how this is possible by claiming that mental representations have constituent structure. This is one reason why CTM is so popular: the symbolic conception of representations can cover these essential properties of productivity and systematicity by assigning the constituent structure to (computational) syntactic properties. The mechanisms responsible for productivity and systematicity of both thought and language capacity here in question, then, are mental representations for LOT. The alternative to a “phrase book” account of language acquisition is to claim that sentences of natural languages have *combinatorial semantics* – that learning a language is learning a “perfectly general procedure for determining the meaning of a sentence from a specification of its syntactic structure together with the meanings of

⁵⁰ Fodor is referring here to a potential counter argument for the productivity of thought – namely that it requires idealization to claim for the infinity of the possible entertainment of *thoughts* – given, for example, the facts of mortality being what they are, people can only *ideally* entertain infinite thoughts. This counter argument, however, does not work for systematicity, so it doesn’t pose a serious threat for LOT.

its lexical elements”⁵¹ (Ibid). Furthermore, given RTM, this carries over to mental representations and shows that they have a combinatorial semantics as well. But this should not be surprising, for as I have discussed above, the symbolic conception of mental representations provided by CTM emphasizes the syntactic organization of mental representations – and the property of systematicity adds to this only marginally by exemplifying the combinatorial semantics behind such syntactic organization.

Remember the historical context I provided above in the first section: That the behaviorists diverged from the mainline thought of psychological theorizing (dominant since the early modern philosophers) with regard to their introspectionist methods and their mentalism. Now Fodor claimed recently in *LOT 2* that whereas the behaviorist psychology denied the mentalistic aspect of their precedents, they kept their *associationism*. Accordingly, psychologists and philosophers (of particularly the empiricist tradition) before behaviorism held that the sequential occurrence of perceptions caused their respective ideas to be associated in the mind – where association was explained as a relation between ideas. For example, from seeing white swans all the time, one relates (associates) the ideas “white” and “swan” – and hence is able to entertain the thought (and thus the sentence) “swans are white”. Now even though behaviorists spurned all forms of mentalism (and hence any talk on “ideas” or “thoughts”), they apparently kept an implicit associationism – which can be seen in how they related stimulus to response: “responses become conditioned to stimuli as a function of their temporal contiguity to reinforcers” (Fodor, 2008, p. 103). That is, for the behaviorists, even though the study of the black box was a taboo – they assumed that the sequential occurrence (or proximity) of stimulus and

⁵¹ This is of course in complete agreement with generative linguistics.

response caused their associations (somehow) in the black box. Behaviorists therefore answered questions like “how do people entertain a sentence like ‘swans are white’” in terms of an associationism between repeatedly seeing white swans and the linguistic utterance (behavior) “swans are white”. Notice, however, that such an associationist conception of linguistic capacity cannot explain the *productive* and *systematic* entertainment (utterance and understanding) of language because it fails to explain how people entertain sentences they have never heard before. Recall from the previous chapter that this is precisely the starting point of Chomsky’s generativist program in linguistics. The Fodor of *LOT 2* sees this as the paramount achievement of Chomsky and his ascription of productivity and systematicity of language to mental processes:

[W]hat now seems most important about Chomsky’s critique [of behaviorism] is the dilemma that he posed for associationists; namely, that you can have *associative* mental processes or you can have *productive* mental processes, but you can’t have both (Ibid, p. 103).

I have so far discussed various elements in the classical view on mental representation. While intentional realism defines the ontological commitments of this view, RTM outlines what representational theories should be like, and CTM shows the particular approach endorsed by the mainstream philosophy of mental representation. LOT, as we can see from the discussion above, exemplifies what a theory that is based on CTM can be like. The precise concern of LOT – that *thought* is linguistic – does not concern us to a great extent here. On the other hand, the arguments for LOT also show what mental representations should be like if they are to account for language *capacity* – and acquisition – as well; namely, that mental representations should have constituent structure that is suitable for the combinatorial

semantics of languages. It is, furthermore, not a coincidence that Fodor ends up praising Chomsky for his generative paradigm: as I will argue in the next chapter, the mental structure that Chomsky postulates for language capacity (and acquisition) is *precisely* what Fodor has in mind when he claims that mental representations have linguistic constituent structure. Now, we turn to the generative account of language (which is pretty much definitive of contemporary linguistics) and see what this linguistic mental structure looks like – so we can identify the phonological component in it, and show that a phonological representation is indeed a mental representation in the sense that I have discussed in this chapter.

CHAPTER IV

THE PHILOSOPHICAL NATURE OF PHONOLOGICAL REPRESENTATION

The crux of the present chapter is this: *what kind of entities are phonological representations?* On the one hand, there is a scientific need to systematically represent a linguistic phenomenon – that of a phonological nature; and on the other hand a phonological representation is not just an abstraction among the linguists' arsenal of tools – it is conceived to be psychologically real. In this chapter my aim will be to clarify what exactly this latter conception amounts to; and the conclusion will be that a phonological representation is a mental representation *precisely* in the sense that the classical view in the philosophy of cognitive science conceives of it. In order to show this, I will bring the two previous chapters together – comparing the philosopher's conception of mental representations with what I have discussed about contemporary phonology.

In the first section, I will start by exploring the claim for psychological realism in linguistics. The argument for realism of linguistic representations, as I will discuss, comes in the form of an inference to the best explanation; but such an argument for realism is not the last word in the philosophy of science. Hence, I will contrast this with an alternative view from the philosophy of science – namely, van Fraassen's constructive empiricism. The aim of this discussion is to contextualize the claim for psychological realism in a language that will make a comparison with the philosophical view on mental representations possible. I will then clarify two possible conceptions of the notion of a psychologically real representation; one being a representation *of* the mind, and the other a representation *in* the mind. I will show

the latter to be the conception philosophers endorse when they think about mental representations.

In the second section, I will provide a detailed analysis of the (linguistic) mental structure that is postulated by Chomsky – and the remainder of the chapter will aim to compare this mental structure and its phonological component with the elements of the philosophy of mental representation that I have discussed in the previous chapter. These elements are intentional realism and the representational and computational theories of mind – and I will argue that a phonological representation that is situated as such in the mental structure that Chomsky postulates, must indeed be a representation in the mind – in the same sense that philosophers conceive of it.

PSYCHOLOGICAL REALISM

Psychological Realism after Chomsky

As I had discussed previously, one of the most important points regarding the wide acceptance of the generativist program in linguistics is about the ontological outcome it portrays: the idea here is that phonological (and in general linguistic) representations have *psychological reality*. But what does this mean, and where exactly does such a commitment come from? First and foremost, remember that Chomsky had criticized the structuralist linguists before him mainly with the claim that the latter had largely overlooked the *creative* aspect of languages. A consideration of how it is possible, for example, that someone can utter a sentence (in their own language) that they never heard before – or how it is possible, that people understand sentences that they are hearing for the first time – gives rise to the possibility that there are *rules* that govern the formation of sentences that people hear

and utter. This opens up a wider possibility that other aspects of languages (like phonology) might as well be rule governed. Asserting that such rules are based on generative principles is no necessary conclusion, but it is one possible way to go (and there doesn't seem to be an alternative) – but the taxonomic endeavor of the structuralists that didn't postulate rules cannot, for sure, give any insight into these questions. Furthermore, from a phonological point of view, there are regularities in languages regarding (as I have discussed before; for instance) how word forms relate to each other or how words can change in certain contexts. But if one's approach to the phonology of languages (like that of the structuralists) is based on observing them and formally representing, say, with respect to the distinctive inventories they make use of (e.g. phonemes) – then (as I discussed) this understanding renders phonological phenomena as *arbitrary*. Arbitrary; because even if you can identify *what* words in a given language can look like (e.g. what phonemes are used by that language), you have nothing to say about *why* the sounds *behave* the way they do in a given language. For instance, showing that the plural form 'cat' in English is [kæt] + [s] and the plural form of 'dog' is [dɒg] + [z]; and claiming that /s/ and /z/ are the phonetic realizations of the same phoneme /z/ does not tell you *why* the respective plural forms *have to be* used with these words. That is, a taxonomic approach cannot give you an *explanation* of phonological regularities – and again, a linguist needs to postulate rules (e.g. voicing assimilation) and check them with data from languages if they are to capture such regularities. Furthermore, a linguist should look for universal patterns among such regularities – something that structuralism explicitly rejected. Thus by overlooking this creative, rule-based, universal aspect of languages, the structuralists – according to Chomsky – had limited themselves to a

taxonomic linguistics; and by doing so, they left some very important questions unanswered.

Now, the structuralist approach had no claim for *truth*: first of all, since a taxonomic approach doesn't need one, the structuralists didn't (necessarily) have any ontological commitments – and if there is no claim for realism about anything, then there is no claim for truth in that (ontological) sense. Second, they used linguistic representations mainly to be able to represent their observations – and did not try to postulate rules that explain what governs the regularities that can be seen in these observations. Again, if there are no principles postulated, then there is no claim about the truth of principles, either. Now Chomsky, on the other hand, is after truth and in both senses: there is both an ontological commitment and rules that can be judged as accurate or not in his representational system of (generative) rules – the system that he calls *grammar*. But still, where does the ontological commitment come from, and what exactly is this a commitment *to*? Well, first of all, the rules postulated in his system must operate *on something*. The way Chomsky presents his system, a phonological (or, in general a linguistic) representation amounts to a medium (that *something* – e.g. “a string of formatives of specified syntactic structure” (Chomsky, 1964, p. 9)) on which a set of generative rules – as processes – must operate. The main task of the linguist here is to find out the rules specific to the language where the phonological phenomenon occurs – but in terms of *the object* of their study, the linguist is committed to a realism about (both the rules and) what those rules are operating on; namely the medium of representation.

Second, one way to make sense of the distinction between *competence* and *performance* – or *langue* and *parole* – is to assert that the system of generative rules

embodies the speaker's competence in a given language.⁵² Thus, (for Chomsky) *competence* in a language is nothing other than *internalizing* the grammar (the set of generative rules) of that language (Ibid, p. 10) – and thus conceived, one can have a story to tell (as opposed to the structuralists) about why a speaker can err in *performance but not in competence*: the facts of humanity being as they are, no performance is immune to error – but grammar, while ascribing structural descriptions to utterances, also can tell the derivations from the perfect forms; so a listener (who is also competent) can correct a badly formed utterance upon hearing. Now – and here is the leap – (a) given that the set of generative rules is the same for every individual that shares competence in a language, and that (b) there doesn't seem to be any other way to account for these rules – it must be the case that these rules belong to a *mental capacity* that is shared by all the individuals.

These two claims, taken together, shed light on the question why there is an ontological commitment to psychological realism in linguistics starting with Chomsky. In short, without the postulation of rules, a linguistic endeavor is limited to observational power and is bound to conceive of its object of study as an arbitrary phenomenon. But with the postulation of rules also comes a commitment to the *existence* of a medium (i.e. representation) that the rules operate on. Moreover, given there is a difference between *competence* and *performance* – and that the linguist wants to study the former – it seems that the only way to explain this distinction is to

⁵²This should be clear given the discussion on the system of generative rules (grammar) from the second chapter: for Chomsky, competence in a language is explained by the speaker's knowledge of the grammar of that language – where the grammar's main function is to assign structural descriptions to the (seemingly) arbitrary utterances. "To each phonetically possible utterance [...], the grammar assigns a certain *structural description* that specifies the linguistic elements of which it is constituted and their structural relations [...]. For some utterances, the structural description will indicate, in particular, that they are perfectly well formed sentences. This set we can call the *language generated by the grammar*." (Chomsky, 1964, p:9)

think of this representational medium as a mental phenomenon. Hence Chomsky's conclusion:

In any event, there is surely no reason why the linguist must necessarily limit himself to 'the study of phenomena and their correlations', avoiding any attempt to gain insight into such data by means of an explanatory theory of language, a theory which is, of course, 'mentalistic', in that it deals with the character of mental processes rather than with their physical basis (Chomsky, 1964, p:99).

Explanatory Adequacy vs. Empirical Adequacy

Notice, however, that the ontological commitment in psychological realism comes *after* the claim for the postulation of rules. This can be seen explicitly in Chomsky's meta-theoretical concerns: when he distinguishes between three levels of success for linguistic theories, he claims that the structuralists were interested in a methodical but mainly *observational* study – a level of success that is achieved by presenting the observed data correctly. But a higher level of success is accordingly possible when the linguistic theory “gives a correct account of the linguistic intuition of the native speaker” – that is, by specifying “the observed data (in particular) in terms of generalizations that express underlying regularities in the language” (Ibid, p. 28). This level of success (that is, the *descriptive* level) is where a linguistic theory postulates rules and principles to account for the accurately presented observations. But of course, it is never the case that such rules and principles can be *a priori* true: there is never a necessary conclusion from data to generalizations that confirms that the postulated principles are expressive of the observation. Thus, (Chomsky claims) there has to be a yet third level of success where the linguistic theory comes forward and provides a *general basis* for selecting which of the postulated principles (that

are, of course, all relevant and consistent with the observed data). If a theory is successful in providing such a general basis, then, it is achieving *explanatory* adequacy – in that it explains, not merely observes or accounts, for the linguistic intuition of the native speaker.

It should be clear at this point that the “general basis” in question for Chomsky and his generativist scheme is the postulation of some kind of a mental structure. Thus, the ontological commitment to this mental structure (and therefore psychological realism) comes from a meta-theoretical concern for explanatory adequacy – a concern related to the question why the linguist⁵³ must select generative rules over and above others that are equal in *descriptive* power. In other words, we (as linguists) start by accepting the linguistic need to postulate rules and generalizations if we want our linguistic theory to *account for* the observed data – but we also *don’t* want to say that this account is *arbitrary*. But the best explanation we have that can show us that our account is not arbitrary seems to be an explanation where *the mind* is responsible for these rules. Thus, we commit ourselves to realism about a mental structure in order to explain (rather than merely observe or describe) the linguistic phenomena in addition to the rules and principles we postulate.

Now, one way of reading this claim is as an “inference to the best explanation” argument for (psychological) realism. That is, the claim here is that (as

⁵³ It may be necessary to emphasize that the concern for explanatory adequacy here is *indeed* a meta-theoretical concern. Chomsky’s account for explanatory adequacy is usually related to studies about children’s language acquisition – and the folk psychological background assumptions (see previous chapter) in the related literature usually give way to the description of a child as a “theorist”. Thus, if conceived in that sense, this line could be misunderstood here – as if the concern for explanatory adequacy is related to the question why *the child* must select generative rules over and above others that are equal in *descriptive* power. On the contrary, Chomsky makes it very clear that when he talks about the levels of success for grammatical descriptions, he indeed talks about their associated *linguistic* theories – and *theories*, in the sense that what the linguists are postulating when they are doing science. See, for instance: “[explanatory adequacy is] achieved when the associated *linguistic theory* provides a general basis for ...” (Chomsky, 1969, p: 9) or “a *linguistic theory* that aims for explanatory adequacy is concerned with...” (Ibid, p: 10) – and he goes on criticizing what particular schools of linguistics have established with observational and descriptive adequacy, such as the post-Bloomfieldian American linguistics (structuralists) and Sapir’s traditional grammar, respectively.

far as we know) *only* if⁵⁴ you are postulating some kind of mental structure will your linguistic theory be explanatorily adequate – or in other words, committing to psychological realism seems to provide the best explanation we have for the rule based nature of linguistic phenomena. But now, we are in the domain of the philosophy of science – for inference to the best explanation, as an argumentative style, has been under extensive discussion⁵⁵ and criticism in contemporary philosophy of science. As an argumentative style, an inference to the best explanation is arriving at an explanatory hypothesis from an observation where the hypothesis is (mostly) sufficient, but not logically necessary for the observed phenomenon. The obvious criticism, of course, resides in its not being logically necessary: this kind of reasoning has striking resemblance to the logical fallacy that is known as “affirming the consequent”⁵⁶. But on the other hand, proponents usually claim that in practice this is how actual science works (Lipton, 2004) – and always has been. Now the critical question here, therefore, is whether Chomsky is justified in inferring realism for mental structure given that generative rules express language and define the *competence/performance* distinction: psychological realism seems indeed to be sufficient for explaining those, but Chomsky has no argument that it is necessary.

Perhaps the most prominent contemporary critic of scientific realism has been the philosopher of science Bas van Fraassen. In his now classic work *The Scientific*

⁵⁴ “only if” here does not imply any form of *logical* necessity – it’s just that we (yet) have no other way of having explanatory adequacy for linguistic theories.

⁵⁵ For an extensive presentation and defense of inference to the best explanation (IBE), see: Lipton, 2004. Lipton defends IBE as both an accurate description of the actual processes of science, and as an epistemic rationalization of the scientific inferential processes – that is, he provides a discussion of why scientists are rightfully engaged in the inferential processes while doing science.

⁵⁶ *Affirming the consequent*: from $P \rightarrow Q$, where we know Q , we infer P . But we are in a fallacy because we don’t know; it could also be the case that $\sim P \ \& \ R \rightarrow Q$ and both R and $\sim P$ are true.

Image (1980), van Fraassen developed an empiricist view of science that reinstated the discussions on scientific realism and anti-realism – a view he calls *constructive empiricism*. Constructive empiricism as a view is intended to be both *descriptive* of the actual scientific endeavor, and *normative*, in that it asserts what science should be concerned about and how it should be done. Now van Fraassen’s claim is that science aims (and should aim) to give us theories that are *empirically adequate*; and a theory is empirically adequate if (and only if) what it says about the *observable* phenomena is true. As his slogan goes, it is empirically adequate if it “saves the phenomena” (van Fraassen, 1980, p. 12). Notice, however, that this view essentially differs from instrumentalism: whereas an instrumentalist thinks that truth doesn’t (and shouldn’t) play a role in the scientific agenda, the constructive empiricist thinks that theories should be true about the observable aspects of the world, but not (necessarily) about the unobservables. Now as far as the structuralists’ conception of linguistics went, it seems to have been a form of instrumentalism; at least as far as truth is concerned, for the structuralists did not have any claim for truth – neither in the ontological sense, nor in the logical sense. On the other hand, in the face of Chomsky’s claim for explanatory adequacy, someone with the constructive empiricist tendency would accept that the postulation of a mental structure would be the best explanation for the generative rules that describe linguistic phenomena – but deny that its being the best explanation actually provides any reason to believe in the theoretical entities (i.e. the mental structure) postulated by that explanation. That is, a constructive empiricist would acknowledge the need of postulated principles and their truth (diverging themselves from the structuralists) – but wouldn’t make the leap from best explanation to the commitment to psychological realism.

What the difference between Chomsky's claim for realism and van Fraassen's anti-realism boils down to, then, is a *choice* between *explanatory adequacy* and *empirical adequacy* – between going for the best explanation and believing in its consequences, or sticking with the observable at hand and being safe in terms of what we are committed to. In other words, Chomsky's claim for explanatory adequacy is not really the last word that can be said in terms of a criterion for accepting or denying the postulated linguistic generalizations; and accepting his claim depends on our choice for being a realist or an anti-realist with respect to the unobservable aspects of the world. The choice here, then, is between the idea that (a) we should acknowledge the *existence* of the entities that our best explanatory theories postulate, (even, as in the case of psychological realism, if those entities are unobservable) and decide the truth of our theories by judging them with the postulated entities⁵⁷ – and the idea that (b) we should only be concerned with the observable aspects of the world; that we should accept theories or reject them on the basis of our (empirical) observations and limit the notion of *truth* to the observable.

On the other hand, as I have mentioned beforehand, the debate between scientific realism and anti-realism in the philosophy of science has two aspects; one descriptive, and one normative – whereas in the latter sense this debate aims to provide a guideline for how science *must be* done, in the former sense this aims to be descriptive of how science *is being* done. We are however, concerned with the descriptive aspect here, for what we are concerned with is how linguists are doing their study – specifically when they do phonology. Now, it goes without saying that

⁵⁷ For example: why are the generative principles (theory) true rather than any other sorts of principles? Because the mind (postulated entity) works that way. This is precisely what it means to say, at the same time, both **(a)** that an explanatory theory “provides a *general basis* for selecting a grammar that achieves the second [descriptive/ABE] level of success over other grammars consistent with the relevant observed data that do not achieve this level of success” (Chomsky, 1964, p. 28) – and **(b)** that such an explanatory theory is “‘mentalistic’ in that it deals with the character of mental processes” (Ibid, p. 99).

Chomsky (and his claim for explanatory adequacy) has been extremely influential for contemporary linguistics, and this means that the choice between realism and anti-realism is pretty clear when it comes to understanding the linguists' practice: contemporary linguistics values explanation over empirical adequacy and is therefore committed to the existence of a mental structure that is responsible for the truth of the rules and generalizations that are postulated by linguistic theories – even though the existence of such a mental structure is beyond observation.

Representation of the Mind or in the Mind

So what gives Chomsky's generative model explanatory power is the claim that linguistic explanations are explanations *because of the fact that* they are expository of what goes on in the mind: the generative principles are *true* because the mind works that way. But what does this tell us about the nature of phonological representations? That is, so far we have seen that phonological representation as such is a representation of a *mental nature*, but what kind of a mental representation exactly is it? To clarify this, consider the following distinction: that of being a representation *of* the mind and being a representation *in* the mind. Now, if a phonological representation is a representation of *what goes on in the mind*, one way of understanding this is to think that phonology is *describing* real phenomena "in the head" (e.g. the brain, the propositional attitudes, the network of nodes, the cognitive mechanism, etc. – whatever one's philosophical outlook requires the mental structure to be). As far as this understanding goes, it amounts to the claim that a phonological representation is a representation *of* the mind – or strictly speaking, a certain (phonological) aspect of the mind. If this is the end to the claim, Chomsky (while showing the shortcomings of structuralism) is making a very clear and persuasive

argument for it – he is showing that the best explanation we have for linguistic phenomena postulates a mental structure that is responsible for our linguistic capacity; and *phonological* representation as such must be part of that mental structure. Very persuasive – as far as we are willing to be scientific realists. On the other hand, as it should be apparent from the discussion in the previous chapter, this is *not exactly* what philosophers mean when they use the term “mental representation”. That is, philosophers with a commitment to realism about mental representations are not only claiming that certain phenomena are representations of how the mind does what it does; they are also *making a claim about how the mind works* – namely, by *using* internal representations of the world. A mental representation in this latter sense is a representation *in* the mind; a mental entity that the mind makes use of in order to work the way it does. Now one can be committed to psychological realism about phonological representations in the former sense (i.e. representation *of*) without being committed to the latter claim (i.e. that there are representations *in* the mind). For instance, it can be argued that our best understanding of doing science is to represent phenomena in order to study them, and OK, what Chomsky shows us is that what we are representing when we are studying phonological phenomena is a representation of something that happens in the mind – but that doesn’t necessarily tell us that what happens in the mind *itself* has a representational nature.

Now, the question that has been the central concern of this thesis can finally be postulated as follows: A linguistic representation after Chomsky is believed to be a representation *of* the mind, but is it compatible with the further claim (asserted by the philosophers of cognitive science) that there are representations *in* the mind? – And, if so, what do the phonological phenomena tell us about this picture? Is a

phonological representation, a representation *in the mind*? In chapter two I discussed the linguistic approaches for representing phonetic and phonological phenomena, and in chapter three I explored the mainstream view in the philosophy of mental representation. In the remainder of this thesis, my aim will be to bring these two expository chapters together to answer this question. I will start by exploring Chomsky's description of the mental structure that is responsible for linguistic capacity in detail – and by identifying the phonological component in such a description. Afterwards, I will compare this mental structure with the elements of mental representation from the previous chapter.

PHONOLOGICAL REPRESENTATION *IN THE MIND*

Phonological Component within Grammar as a whole

To start with, let us clarify the mental structure that Chomsky postulates⁵⁸ in detail:

The generative grammar of a language should, ideally, contain a central *syntactic component* and two *interpretive components*, a *phonological component* and a *semantic component*. The syntactic component generates strings of minimal syntactically functioning elements (following Bolinger, 1948, let us call them *formatives*) and specifies the categories, functions and structural interrelations of the formatives and systems of formatives. The phonological component converts a string of formatives of specified syntactic structure into phonetic representation. The semantic component, correspondingly, assigns a semantic interpretation to an abstract structure generated by the syntactic component. Thus each of the two interpretive components maps a syntactically generated structure onto a 'concrete' interpretation, in one case phonetic and in the other, semantic. The grammar as a whole can thus be regarded as, ultimately, a device for

⁵⁸ Chomsky's following quotation (that we are repeating here) is about the "generative grammar of a language"; but it should be clear at this point that given the claim for psychological realism, talking about the grammar of a language doesn't differ in any way from talking about the mental structure related to it.

pairing phonetically represented signals with semantic interpretations, this pairing being mediated through a system of abstract structures generated by the syntactic component. Thus the syntactic component must provide for each sentence (actually, for each interpretation of each sentence) a semantically interpretable *deep structure* and a phonetically interpretable *surface structure*, and, in the event that these are distinct, a statement of the relation between these structures (Chomsky, 1964, pp. 9–10).

There are a number of points in this quote that will be important for our purpose here. The following is a list of these points:⁵⁹

- (1) Grammar is a *device* that pairs “phonetically represented signals” (a representation of the physical form of sounds) with “semantic interpretations” (meaning).
- (2) Grammar has a syntactic component, which is responsible for
 - (a) the generation of “strings of minimal syntactically functioning elements” (a constituency of elements)
 - (b) the specification of “the categories, functions and structural interrelations” of those elements (organization of that constituency)
 - (c) feeding the *phonological* and *semantic components* these elements (a constituency of organized elements) as input
- (3) Grammar has a *phonological* component that functions as a sub-module of the grammatical system – it takes input from the syntactic component and outputs it in the form of phonetic representation.

⁵⁹ The notes in parentheses are for future reference, they will be important when I discuss the compatibility of the mental grammar postulated as such with the representational and computational theories of mind.

- (4) Grammar has a *semantic* component that (somehow) assigns a semantic interpretation to the organized elements that are supplied by the syntactic component.
- (5) Grammar as a device thus conceived, *mediates* the pairing of sound and meaning through its syntactic component.
- (6) And given psychological realism, grammar as such is the mental structure that is responsible for our linguistic capacity.

First and foremost, what this picture shows is that the phonological component is intimately tied to the syntactic component of the mental structure that is being postulated – and hence cannot work in isolation from it. That is, anything that is being heard by a competent speaker of a language is linguistically intelligible to her *only if* it can be identified as one of the “minimal syntactically functioning elements” of that language. For instance, for a word to be understood it must *first be heard as a word* – and although it is in the first place the phonological component that is responsible for this, by (2b) the category and function of that word is determined by the syntactic component. Thus the whole system functions as a whole when the competent speaker cognizes (even) a word (let alone a sentence) as a full-fledged, functioning element of language. If this is too abstract, recall the example from the French words “water” <eau> and “high” <haut> in the second chapter: both of these words are pronounced the same – that is, they are both phonetically realized with the same representation [o]. In French, the article used for words that start with a vowel is <l’>. Now, when it is used with an article, <eau> becomes <l’eau> and it is phonetically realized as [lo] - but <haut> becomes <le haut> and it is realized as [ləo]. In the context of the second chapter, this example showed that phonetics by

itself cannot capture all there is to the sound side of languages: given that these two words <eau> and <haut> have the same phonetic properties, and that they still behave differently when used with an article, it follows that there must be an abstract level of phonological representation behind the “phonetically interpretable *surface structure*” – as Chomsky calls it in the quotation above. However, here, this example has a further implication: it shows that (a) the two words <eau> and <haut> have two different phonological representations and (b) while syntax puts them after an article, and then phonology takes its course, syntax also – in a sense⁶⁰ – tells the phonological component which of the two phonological representations it is dealing with.

Furthermore, notice that one other aspect of the grammatical picture above is the claim that semantics always goes hand in hand with a phonological representation:⁶¹ (1), (2c), (4) and (5) above all show that the phonology of a word is essentially bound to its meaning, and it is again the syntactic component of this mental structure that is responsible for this binding (5). Thus, given the quote above, where the syntactic component “provides for each sentence a semantically interpretable *deep structure*⁶² and a phonetically interpretable *surface structure*” – when these are distinct, it also provides “a statement of the relation between these structures”. What this amounts to can be seen explicitly in our example from French:

⁶⁰ This, of course, doesn’t mean that syntax *determines* which of the phonological representations is being used – syntax just supplies the phonological component with the necessary information that indicates what article is being used; but it is the phonological component that *interprets* the representation and *produces* a phonetic signal.

⁶¹ This, of course, does not mean that “meaning only changes if and only if the phonological representation changes”; after all, a word can mean several things, but the phonological representation for each of the different meanings can be the same (e.g. “bow”). On the contrary, what is meant here is that the phonology and meaning(s) of a word is consistently in correlation.

⁶² It should be noted that since Chomsky has described this structure, a lot has been debated on the nature of the semantic component. This, therefore, is an outdated model in terms of the semantic component in contemporary linguistics – but since we are concerned with the phonological component here, nothing in what follows hangs on this.

in the case that <eau> and <haut> are used in a sentence and with an article, the syntactic component provides the phonological component the necessary input enabling the latter to produce the phonetic interpretation [lo] or [ləo] respectively. But while doing that, it (the syntactic component) also provides which semantic interpretation will match with this phonetic output – hence providing a statement of the relation between the phonetic (*surface*) and the semantic (*deep*) structures.⁶³

These two points show that the phonological component of this mental structure is embodied within the grammatical system as a whole – and it follows that a phonological representation should be understood in terms of the whole grammatical system. That is, whether it is a word or a sentence that is being represented, the whole system functions together, and thus, even though the phonological properties of a representation can be conceptually separated from its syntactic properties, it should be kept in mind that there can't be any phonology of a language without its syntax – and likewise, there can't be a syntax to a language without phonology⁶⁴ (and hence, by (2c), semantics).

Grammar and Mental Representation

We can now answer the question at hand. What we are now exploring is a phonological representation embodied within the grammatical system – and the question is whether it fits in with the claim that the mind makes use of representations in order to work the way it does. Now, first of all, what we are taking as compatible with the latter claim – that there are representations *in* the mind – is to

⁶³ Cf. n61.

⁶⁴ This indeed means that sign languages also have a phonological component – as it should be clear by now that what the phonological component works on is an abstract level of representation, and in the case of sign languages, the output can be argued to be hand gestures. Such an argument, however, is out of the scope of this thesis and a further discussion will be omitted here.

be compatible with the elements of the classical view in the philosophy of mental representation that I discussed in the previous chapter. These elements are intentional realism (IR), the representational theory of mind (RTM), the computational theory of mind (CTM). What remains, therefore, is to discuss the compatibility of the phonological component of grammar as a mental device with these elements – and to determine whether we can call a phonological representation a representation *in* the mind, given it is embodied in a device as such.

As I discussed in the preceding chapter, the definition for intentional realism is as follows: beliefs, desires and such other mental states are *real*; they are causally involved (have *causal powers*), determinate and contentful states that are subject to semantic evaluation. Folk psychology is a working psychological theory that gives us accurate predictions – because people have mental states and they express them via propositions. Now commitment to intentional realism as such has two major aspects: first, a (scientific) realism about mental entities, and second, a claim about those entities (i.e. intentionality – which, according to the classical view, can be exhaustively explained by the causality and semantic evaluability of mental states). Put that way, it should be clear that intentional realism indeed connects with generativist linguistics directly in terms of the ontological commitment they share – namely, scientific (psychological) realism for mental entities. In other words, both the intentional realist and the linguist are committed to a scientific conception where they infer the existence of the entities that are postulated by their best explanatory theories – and in both their cases, those theories postulate mental entities. On the other hand, since there is nothing in the grammatical mental structure postulated by Chomsky that is contradictory with the claim for intentionality; and since intentional realism is first and foremost an ontological claim – it follows that intentional realism

and generativist linguistics are perfectly compatible with each other. With regard to *phonological representation*, however, the main question is of course whether a phonological representation is compatible with *RTM* or not – for after all, the theory that outlines what a representation *in* the mind amounts to is *RTM*. Intentional realism is basically the ontology that *RTM* is based on, and the common commitment of linguists and philosophers regarding psychological realism here shows that the generativist mental structure – and of course the phonological component embodied in such a structure – fits the bill for that ontology.

But *RTM* itself, as I have discussed, is an abstract characterization of the representationalist picture; it outlines the main idea by asserting that there *are* representations *in* the mind that mediate cognition between the world and the cognitive system – but it does not postulate anything regarding the nature and content of mental representations. That is, *RTM* provides a theoretical framework for theories to fill in the specifics for the nature and content of mental representation – but it is the *computational* approach to *RTM* (i.e. *CTM*) that fills in those specifics in the classical view on mental representation. Now this means two things: first, whatever is compatible with *CTM* should be compatible with *RTM*, and second, if we are interested in the pertinence of a particular kind of representation (in our case phonological representation) with the representationalist picture in the philosophy of cognitive science, we should look at *CTM* for a comparison.

There are two major claims that *CTM* brings to the table: first, a mental representation is a symbolic entity that has syntactic and semantic properties, and second, mental representations are syntactically organized – and it is the organization of syntactic constituents that define the content of mental representations. Now a quick look at the list above shows us that the grammatical system as a whole is

indeed on the same boat with CTM: from (2) and (4), it is self-evident that grammar as a device has both syntactic and semantic components – and of course, given (6), i.e. psychological realism, what this device represents also has syntactic and semantic properties. Furthermore, (2b) and (2c) tell us that syntactic properties of such a representation are organized, and it is this syntactic organization that defines semantic content. So far so good, but notice that the phonological component is lacking from this picture. What remains therefore, is to indulge into the details of a phonological representation and to show that it is indeed a mental representation that fits in with CTM.

Phonological Representation as Representation *in* the Mind

The definition I provided for CTM in the previous chapter stated that *cognitive processes are computational operations defined over the syntactic properties of symbolic mental representations*. Thus, if we are looking for the answer to the question whether a phonological representation is a mental representation in the sense that CTM conceives of it – what we are ultimately looking for is whether we can say in the case of phonology that (a) *phonological* processes are *computational* operations (b) defined over the *syntactic properties* of (c) *symbolic* phonological representations.

Now first and foremost, let me put the terminology straight. Remember that *computation*, in the sense that is relevant here, is *formal manipulation of symbols*. A symbol is something that stands for, represents something else. A *symbolic representation* is an organized medium of brought together smaller units (i.e. symbols) – and this organization is structured (i.e. brought together) with respect to formal rules. Thus, to make the case for phonology, we first need to identify two

things: (a) what are the building blocks (i.e. symbols) in question for a phonological representation? – and (b) what do these symbols represent?

Now as I discussed in the second section, early structuralists (in particular Trubetzkoy) thought that the atomic building blocks of a phonological representation were *phonemes*; where a phoneme was conceived of as a complete, simple (i.e. cannot be decomposed), smallest unit that was the representation of a speech sound. Later linguists, however, starting with Roman Jakobson, thought that phonemes were further composed of distinctive binary⁶⁵ features.⁶⁶ Jakobson held that it is features, not phonemes that should be the fundamental units of linguistic analysis. However, this view has undergone some significant changes since Jakobson. As I have discussed, the kind of phonological questions from the structuralists to generativists can be seen changing from those of *distinctive analysis* (e.g. what is a possible word or sound in this language?) to those regarding explanations of what I have called *phonological derivation* in the second chapter (i.e. alternations – e.g. how do word forms change in certain contexts?). Now for Chomsky, there are two ideas related to phonemes: the first is the conception of phonemes as the individual building blocks – which he calls *segments* – of a phonological representation, and the second is the idea that phonemes are *defined* by their contrast to and distinctiveness from other phonemes. Distinctive features come into play in the latter idea, and whereas Chomsky adheres with the former idea, he explicitly rejects the latter (which, as I discussed, he attributes to structuralist taxonomics) on the grounds that the criteria for establishing it don't work⁶⁷ (Chomsky, 1964, pp. 83-84). However, in

⁶⁵ Remember the discussion on phonetic features.

⁶⁶ For an extensive discussion of this, see: Anderson, 1985, pp. 117-139.

⁶⁷ Chomsky's criticism of phonemes as contrastive elements is aimed to be a criticism of the structuralists' taxonomic analysis. I will not go into the details of this criticism here since what we are

contemporary phonology, Chomsky's view has undergone some significant improvements as well. Chomsky's conception of a phonological representation was a *linear* string of segments: a string of successive phonemes with binary features. Thus, for instance, the English word "cat" is composed of three phonological segments: /k/ + /æ/ + /t/. Now in such a conception, the segments of a phonological representation are not related to each other in any way. However, think, for example, of the case with Turkish plural suffixes '-ler' and '-lar': The plural for 'ev' (a house) is 'ev-ler', and 'ev-lar' is never a possible plural.⁶⁸ The plural for 'at' (a horse) is 'at-lar', and 'at-ler' is never a possible plural. What this shows is that the middle segments of Turkish plural suffixes somehow *change with respect to a sound* (or, strictly speaking, a segment) in the word stem that they are being added to. But if one's conception of a phonological representation is a collection of segments that are not related to each other (as is in Chomsky's linear account), then one cannot give a phonological account of what is going on in this example. In 1976, John Goldsmith wrote a PhD dissertation titled "Autosegmental Phonology" that has been extremely influential for contemporary phonology – and in that, he aimed to improve generative phonology. Autosegmental phonology addresses the relations between phonological segments by introducing *tiers* – overlaying levels of a phonological representation that determine the constituency of the segments of that phonological representation. Tiers in autosegmental phonology serve the function of explaining – in terms of formal rules – why the segments of a phonological representation are brought together the way they are. Such rules are specific to languages, but universal among

concerned with is the former conception of phonemes, namely that they are constituents of phonological representations.

⁶⁸ Although it can never be a plural, however, a word such as 'evlar' could of course be a borrowed stem from another language – but this doesn't change the argument that follows.

the phonology of each language. Thus, in the case of the Turkish ‘ev-ler’ and ‘at-lar’, the reason for the plural suffixes changing with respect to the features of the sound in the stem word is explained by an additional tier of their representations which – in this case – is responsible for the regulation of Turkish vowel harmony.⁶⁹

A phonological representation, therefore, is a medium of brought together units (i.e. segments), and these units are indeed *symbols* – in that they *represent* the speech sounds utilized in languages.⁷⁰ Furthermore, given what autosegmental phonology brings to the table, this medium of brought together units is *organized* – and this organization is with respect to *formal rules*. It follows, that it is safe to call a phonological representation a *symbolic representation* – and in the same sense that CTM conceptualizes it. Now this answers part of the question at hand – and we still have to show that (a) phonological processes are computational operations that are (b) defined over the syntactic properties of (symbolic) phonological representations.

First, let’s consider what phonological processes are like: they are the application of rules to phonological representations that produce *phonetic* strings (recall the function of the phonological component in Chomsky’s mental structure: (3)). Here is an instance of a typical phonological process known as the “voicing assimilation” in English: the regular English plural suffix with /z/ is changed to [s] after a voiceless sound:

⁶⁹ The details of autosegmental phonology and the exact nature of the tiers it introduces to a phonological representation are vague here. However, for our purposes, the details are not crucial – what matters here is that in the contemporary conception of a phonological representation, there are symbolic units (i.e. segments) that are a collection of *organized* units, and this organization is provided by the additional layers that are tiers. The addition of tiers to a phonological representation is extremely helpful in explaining, for instance, the phonology of tone languages such as those of African languages. For a detailed discussion, see: Goldsmith, 1976, Kaye, 1989, pp. 99-112.

⁷⁰ It should be noted that saying “speech sounds” is not identical to saying “sound signal” in terms of what is being represented by the constituent symbols of a phonological representation. As I discussed in the second chapter, the very idea of segmenting speech sounds is loaded with background assumptions regarding, for instance, where a sound begins and the other ends. When one says “speech sounds” are represented, therefore, one already is talking about a level of abstraction involved with what is being represented.

<cat> + plu. → <cats>

/kæt/ + /z/ → [kæts]

What exactly happens here is that the [+voiced] feature of /z/ is changed into [-voiced]. But given we have decided that phonological representations are symbolic representations, one other way of putting this is that the phonological symbol /z/ is being manipulated with respect to a well-defined rule (i.e. voicing assimilation). But if that is the case, then phonological processes can be conceived of as *formal manipulations of symbols* – and therefore are computational operations.

But are they defined over the *syntactic properties* of phonological representations? What the above example shows is that phonological processes are defined over the binary features of phonological segments. That is, in the above example /z/ (+voiced, +alveolar, +anterior, +coronal) changes into [s] (-voiced, +alveolar, +anterior, +coronal) where one of the binary features is manipulated. Are the binary features among syntactic properties then? As I have discussed in the previous chapter, one way of understanding syntactic properties is understanding them as non-semantic properties – and binary features are nothing near semantic. But furthermore, consider what syntactic⁷¹ properties are according to Fodor – to re-quote:

Syntactic properties ... [o]n the one hand, [are] among the ‘local’ properties of representations, which is to say that they are constituted entirely by what parts a representation has and how these parts are arranged ... it’s also true that the syntax of a representation determines certain of its relations to other representations (Fodor, 2001, p. 20).

⁷¹ Of course “syntactic” here is used in a literal sense and should not be confused with the syntactic component of the grammatical mental structure Chomsky postulates.

Well, binary features are certainly among the ‘local’ properties of phonological representations – they are entirely constituted by what parts (segments) a representation has and how these parts are organized. But more importantly, it is also true that the features define how a representation differs from another. Binary features, therefore, are among the syntactic properties of phonological representations – and given phonological processes computationally operate on these features we now have our answer: *phonological* processes are *computational* operations (b) defined over the *syntactic properties* of (c) *symbolic* phonological representations – and such conceived, a phonological representation conforms to CTM. By what has been said so far, this shows that a phonological representation is indeed a mental representation *in the same sense that the philosophers of cognitive science conceive of it*.

CHAPTER V

CONCLUSION

The main point of the second chapter was to present the reader the two kinds of representations that are related to the sound side of languages. I have discussed first that the need for phonetic representation comes from a need to represent speech events in languages. In their early attempts at this, linguists tried representing as much as they could capture in a transcription. But growing technology afterwards showed that the signal was too rich to represent all – even if such a representation was possible, there are irrelevant aspects of speech events that the linguist does not want to represent. Thus, phonetic representation is the necessary level of abstraction from irrelevant noises. The reason why I introduced phonetic representations was, furthermore, to contrast them to phonological representations in the hopes to define the latter. A phonological representation, accordingly, is an attempt to capture the regularities in languages that cannot be captured in a phonetic representation. Phonetic representation is a tool the linguists make use of when they transcribe surface sound sequences in languages. Furthermore, phonetic features are represented with respect to the physical features of speech production. However, the important thing to keep in mind about a phonetic representation is – although the general tendency among linguists is to think of it as a language neutral transcription – we have also seen that the choice of how to represent with respect to physical properties of speech sounds is not inherent in the sound signal. That is, phonetic features are not a priori conclusions from the observation of speech events. Phonetic segmentation therefore, is not that neutral as it pretends to be since it involves

decisions and those might be language-specific. A phonological representation, on the other hand, is a representation of the underlying abstract level of what mentally goes on when a speech event occurs. The task of the phonologists is to identify patterns in languages and postulate rules and generalizations that describe these patterns.

In the third chapter, I presented the mainstream view on mental representations in the philosophy of cognitive science – and argued that it is composed of three fundamental components. Whereas intentional realism defines the ontological commitments behind the idea of mental representation, RTM and CTM define what theories committed to the notion of mental representations can be like. These three components, taken together can be used as a criterion for deciding what would be a theory that refers to mental representations and what not. In this discussion, I intentionally left out the views in the philosophy of cognitive science that oppose the Fodorian picture. In such views, (e.g. connectionism, embodied cognition, etc.) there is no apparent reason to believe in the existence of mental representations, and what we aimed for was comparing phonological theory with the philosophical positions that have *arguments* for the existence of mental representations.

Given what we have discussed in these two chapters, two striking commonalities surface that I believe intensifies the coherency of what I argued for in this thesis. The first of these is that both the generativist linguistics and cognitivism in philosophy emerge as reactions against behaviorism: Chomsky argues against the behaviorists on the basis that they cannot explain language acquisition – and Fodor argues against them on the basis that their pessimism regarding studying the mind is not grounded on firm grounds. Second, as I have discussed, both generative

linguistics and cognitivism establish their ontological claims through a form of inference to the best explanation: the former infers the existence of a linguistic mental structure as an inference to the best explanation, and the latter infers the existence of intentional mental states likewise. I believe this is not a coincidence: it is because they ultimately postulate the same kind of mental representations that they follow the same argumentative path.

In the fourth chapter, I presented the main claim of this thesis; that what the mainstream philosophy of cognitive science (i.e. cognitivism) conceived of as a mental representation can be reflected in the actual practice of phonology. In order to argue for this, I firstly provided a discussion from the philosophy of science perspective – highlighting the ontological commitments of the contemporary phonological theory. The purpose of such a discussion was to serve as the ground to compare the mental structure postulated for linguistic representation with the three components of mental representation – on primarily an ontological basis. After identifying the phonological component of this mental structure, I argued that phonological theory connects with the philosophical claims for mental representation with respect to their shared claim for psychological realism. Ultimately, I argued that a phonological representation can be defined in terms of CTM, which I believe, with the addition of the shared ontological position, establishes the main claim of this thesis.

The conclusion reached here (if true) has some interesting implications. First of all, it implies that the rules and generalizations that the phonologists are after tell us actually about how the mind works. Thus when a consistent formal pattern is found in phonological phenomena, our conclusion here indicates that there is a good chance that that pattern can be found in other cognitive phenomena too. This, I

believe can pave the way for future interdisciplinary studies on cognition.

Furthermore, it implies that philosophers should not contradict phonological findings in their postulations of mental representations – which means that it would be prudent for the philosophers that are interested in mental representations to follow the phonological literature.

REFERENCES

- Anderson, S. R. (1985). *Phonology in the Twentieth Century: Theories of Rules and Theories of Representations*. Chicago: University of Chicago Press.
- Anderson, S. R. (1974). *The Organization of Phonology*. London: Academic Press, inc.
- Bolinger, D. L. (1948). On Defining the Morpheme. *Word*, 4, 18-23.
- Carnie, A. (2007). *Syntax: A Generative Introduction*. Oxford: Blackwell Publishing.
- Churchland, P. M. (1981). Eliminative Materialism and the Propositional Attitudes. *Journal of Philosophy*, LXXVIII(2), 67-90.
- Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.
- Chomsky, N. (1964). *Current Issues in Linguistic Theory*. The Hague: Mouton.
- Damasio, A. (1994). *Descartes' Error*. London: Vintage.
- Dennett, D. (1987). *The Intentional Stance*. Cambridge, MA: MIT Press.
- Erguvanlı, E. E. (1984). *The Function of Word Order in Turkish Grammar*. Berkeley: University of California Press.
- Fodor, J. A. (1987). *Psychosemantics*. Cambridge, Massachusetts: The MIT Press
- Fodor, J. A. (2001). *The Mind Doesn't Work That Way*. Cambridge: MIT Press.
- Fodor, J. A. (2008). *LOT 2: The Language of Thought Revisited*. Oxford: Clarendon Press.
- Fodor, J. A. & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: a critical analysis. *Cognition*, 28(1-2), 3-71.
- Friedman, M. (2000). *Parting of the Ways*. Illinois: Open Court Publishing Co.
- Gibbs, R. (2006). *Embodiment and Cognitive Science*. Cambridge: Cambridge University Press
- Goldsmith, J. (1976). *Autosegmental Phonology*. (Doctoral Dissertation, MIT, 1976).

- Gussmann, E. (2002). *Phonology: Analysis and Theory*. Cambridge: Cambridge University Press.
- Halle, M. & Reiss, C. (2008). *The Phonological Enterprise*. Oxford: Oxford University Press.
- Haugeland, J. (1981). Semantic engines: an introduction to mind design. In J. Haugeland (ed). *Mind Design*. Cambridge, MA: MIT Press.
- Jensen, S. A. G. (2001). *A Computational Approach to the Phonology of Connected Speech*. (Doctoral Dissertation, University of London, 2001).
- Kaye, J. (1989). *Phonology: A Cognitive View*. New Jersey: Lawrence Erlbaum Associates, Inc..
- Kaye, J. (2005). GP, I'll have to put your flat feet on the ground. In Broekhuis, H., Corver, N., Huybregts, R., Kleinhenz, U. & Koster, J. (eds.) *Organizing Grammar. Studies in Honor of Henk van Riemsdijk*. Berlin: Mouton de Gruyter.
- Kim, A, 2003. Paul Natorp. In *The Stanford Encyclopedia of Philosophy* (Fall Edition), E. N. Zalta (ed.). Retrieved February 29, 2011, from <http://plato.stanford.edu/archives/fall2003/entries/natorp/>.
- Koertge, N. (2003). Feminist values and the value of science. In C. Pinnick, N. Koertge and R. Almeder (eds.), *Scrutinizing Feminist Epistemology*. New Brunswick: Rutgers University Press.
- Ladd, R. D. (1996). *Intonational Phonology*. Cambridge: Cambridge University Press.
- Ladefoged, P & Lindau, M. (1989). Modeling articulatory-acoustics relations: a comment on Stevens' 'On the quantal nature of speech'. *Journal of Phonetics*, 17, 99-106.
- Lerdahl, F. and Jackendoff, R. (1983). *A Generative Theory of Tonal Music*. Cambridge: The MIT Press.
- Lipton, P. (2004). *Inference to the Best Explanation, 2nd Edition*. London: Routledge.
- Longino, H. (2006). The Social Dimensions of Scientific Knowledge. In *The Stanford Encyclopaedia. of Philosophy*, E.N. Zalta, (ed.). Retrieved October 21, 2010, from <http://plato.stanford.edu/entries/scientific-knowledge-social/>
- Lycan, W. G. & Pappas, G. (1972). What is eliminative materialism? *Australasian Journal of Philosophy*, 50, 149-59.

- Mannell, R., Cox, F. and Harrington, J. (2009). An Introduction to Phonetics and Phonology, *Macquarie University*. Retrieved May 2, 2010, from <http://clas.mq.edu.au/phonetics/index.html>
- Merton, R. (1973). *The Sociology of Science*. Chicago: Chicago University Press.
- Pitt, D. (2008). Mental Representation. In *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, E. N. Zalta (ed.). Retrieved May 5, 2011, from <http://plato.stanford.edu/archives/fall2008/entries/mental-representation/>
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson.
- Saeed, J. I. (2003). *Semantics*. Oxford: Blackwell Publishing.
- Smith, D. W. & Thomasson A. L. (2005). *Phenomenology and Philosophy of Mind*. New York: Oxford University Press.
- Turing, A. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Maths. Soc., ser. 2*, 42, 230–265.
- Watson, J. B. (1913). Psychology as the behaviorist views it. *Psychological Review*, 20, 158-177.