1 Introduction

Linguistic theory needs to consider the fact that a language must be learnable in order to stay alive. *Learnability theory* is the part in linguistics that searches for formalized explanations for language acquisition. The study of learnability can provide insights into the mechanisms that underlie language acquisition and how languages can change. These insights have repercussions on the shaping of linguistic theory.

The general goal of this dissertation is to provide a comprehensive proposal for the learning of phonology. This is exemplified by the case of metrical phonology. I argue for a bidirectional[1] approach to phonology and its acquisition (Boersma 1997, 1998, 1999, 2000, 2001, 2005, 2006ab). The adopted grammar model uses the framework of Optimality Theory (Prince & Smolensky 1993) and formalizes two interfaces: the interface between phonology and phonetics, and the interface between phonology and semantics. For each of the three language modules there are probably multiple levels of representation; an all-embracing learnability approach would have to account for this. The present dissertation, however, restricts itself to the four representations shown in figure (1): one semantic representation (*meaning*, represented as morphemes in the lexicon), two phonological representations (the lexical *underlying form* and the abstract phonological *surface form*), and one phonetic representation (the concrete auditory *overt form*).

Throughout the book, I use the following symbols to mark the various representations, as in figure (1): meaning is represented between single quotes (‘ ’), underlying forms between pipes (| |), surface forms between slashes (/ /), and overt forms between square brackets ([ ]). For instance, meaning can be a single morpheme such as ‘day’ or a morphemic structure such as ‘day-Nom.Pl’.[2] In German, the underlying form that corresponds to this morphemic structure is |taːg+ə|, where |taːg| corresponds to the morpheme ‘day’, and |ə| corresponds to the morpheme ‘Nom.Pl’.

[1] The term ‘bidirectional’ stems from OT semanticist Blutner (2000).

[2] ‘Nom.’ stands for nominative case, ‘Pl’ for plural. This means that the term *meaning* as used here can refer to *syntactic* functions as well.

(1) The grammar model

semantic representation	\`Meaning'
phonological representations	|Underlying Form|
	/Surface Form/
phonetic representation	[Overt Form]

The corresponding surface form is /(tá$_{\mu\mu}$.gə$_{\mu}$)/, which contains hidden structures such as feet (indicated by parentheses), moras ($_{\mu}$), stress (́), and syllable boundaries (.). The corresponding *overt form* could be the detailed phonetic transcription [ˈtʰaːgə], but in this book I abstract away from segmental detail and limit myself to representing stress as in [táː.gə] (I do, however, include syllable boundaries in the overt form and therefore assume that they are known to learners from the start). Feet and moras are not contained in the overt form.

The grammar model in figure (1) is bidirectional in the sense that it can handle the two main linguistic processing tasks of *production* and *comprehension*. In production, the figure in (1) can be read top-down: a speaker is given meaning, from which she computes form: the underlying form, the corresponding surface form, and the corresponding overt form. In comprehension, figure (1) can be read bottom-up: a listener is given the overt form, from which she computes surface form, underlying form, and meaning. Within OT, this computation of representations is achieved through the constraint ranking of a language.

The child's learning task is to become a good speaker and a good listener. Her task is therefore *learning to compute* surface form, underlying form and meaning in comprehension, and underlying form, surface form and overt form in production. Within OT, the learning task is to learn an adequate constraint ranking. **The specific goal of this dissertation is to show that the child can create the whole phonology of the language, if given a sufficient number of informative pairs of overt forms and meaning in the course of acquisition. This means, in fact, that the child bootstraps into phonology by using phonetic and semantic information.** The dissertation will pass through several steps before arriving at the ultimate goal.

The first step is to model the acquisition of the mapping between surface form and underlying form, the two middle forms in figure (1), when both underlying forms and surface forms are given. This has already been done in former OT approaches on learnability: if the learner is given a sufficient number of informative pairs of underlying and surface forms, the learner can learn how to compute a surface form for any given underlying form (Tesar & Smolensky 1993, 1996, 1998, 2000; Tesar 1995; Boersma & Hayes 2001; Apoussidou and Boersma 2003[3]).[4] Likewise, it has been shown that if the learner is given a sufficient number of informative pairs of overt and surface forms (the two bottom representations in figure (1)), the learner can learn to compute a surface form for any given overt form (Boersma 1997, Escudero & Boersma 2001, 2004). There are various proposals of how this learning can proceed: Tesar & Smolensky (1993) and Tesar (1995) formalized this learning as *Constraint Demotion* (henceforth CD), while Boersma (1997) and Boersma & Hayes (2001) formalized this learning as the *Gradual Learning Algorithm* (henceforth GLA). I explain the mapping between surface form and underlying form and the formalized learning approaches in chapter 3, and apply them in chapter 4 on Latin stress.

The second step is to model the acquisition of the mapping between overt form, surface form and underlying form, the three bottom representations in figure (1). It has already been shown that if a learner is given a sufficient number of informative pairs of overt and underlying forms, the learner can learn how to compute surface and overt forms for any given underlying form (in a CD manner shown by Tesar 1997, 1998ab, 1999; and Tesar & Smolensky 1996, 1998, 2000; in a GLA manner shown by Boersma 1998, 2003; Apoussidou & Boersma 2003, 2004ab[5]; and Apoussidou 2006a[6]). In this learning situation, surface forms are hidden; they have to be created by the learner. This requires a translation mechanism for the comprehension process. In the OT formalizations the child is enabled

---

[3] This paper has been incorporated into this dissertation in chapter 4.
[4] The reader familiar with the literature might object that in the listed learnability approaches, the learners were not fed underlying forms; however, the underlying forms were known to the learner. The learner had no other option than mapping a surface form like /(σ σ́) σ/ onto an underlying form |σ σ σ|, or a surface form like /σ (σ́ σ) σ/ onto |σ σ σ σ|. This boils down to saying that the learner was provided with both surface form and underlying form in the learning process.
[5] These papers have been incorporated into this dissertation in chapter 4.
[6] This paper has been incorporated into this dissertation in chapter 5.

to interpret the speech signal by using her current grammar to assign hidden structure, introduced as *Robust Interpretive Parsing* by Tesar & Smolensky (henceforth RIP; 1996, 1998, 2000). While the interpretation is only a first guess and might be wrong, the mechanism itself is robust, in that it never fails to give an interpretation. This mechanism enables the child to detect what might be wrong with her grammar and needs to be changed by comparing her interpretation of a heard form to what she would produce herself. In OT, the grammar is the language-specific ranking of constraints. Children's grammars differ from adult ones, and although children have no idea about what the adult constraint ranking looks like, they have access to the constraint violations of the forms they process, and thereby the means to change their grammar in response to the language data they encounter. I pick up on the ideas of RIP, CD and the GLA, and explain them in chapter 3. I apply them in chapter 4 on Latin stress and in chapter 5 on Pintupi stress.

The third step is to model the acquisition of the mapping between surface form, underlying form and meaning, the three top representations in figure (1). I propose an *on-line* learning approach, where the learner does not have to wait and gather data before she can make a learning move, but where lexicon and grammar are learned concurrently. In fact, underlying forms are learned through the grammar. I show how a learner can learn to compute underlying forms and surface forms if provided with a sufficient and informative number of pairs of surface forms and meaning by applying the mechanism of RIP (see also Apoussidou 2006b[7]). The situation of learning underlying form and grammar, given surface form, has indeed been tackled by e.g. Tesar et al. (2003) with a rather complicated learning algorithm that switches back and forth between the modification of the underlying forms in the lexicon and modification of the grammar, and by Tesar (2004, 2006). The relation between meaning and underlying form was not handled by the grammar. Crucial for these approaches is the ability of a learner to compare paradigms, implying *off-line* learning: the learner gathers data and stores them for later processing. I argue that this is not a realistic learning situation and therefore propose the on-line alternative, outlined in chapter 3 and applied in combination with the GLA in chapter 6.

As a fourth and final step I model the acquisition of the mapping between meaning, underlying form, surface form and overt form that are shown in figure (1). I demonstrate how a learner, if given a sufficient

[7] This paper has been incorporated into this dissertation in chapter 6.

number of informative pairs of overt forms and meaning, can learn to compute underlying forms, surface forms, and overt forms for any given meaning. This approach to the learning of underlying forms relaxes the demarcation between grammar and lexicon: the connection between underlying forms and surface forms, and between underlying forms and meaning, is handled by the grammar. Again I make use of RIP; this time to model the two levels of hidden structure, surface and underlying form. Step 4, the goal of this dissertation, is achieved in chapter 6 on Modern Greek.

With the OT learnability approaches, and especially when multiple representations are considered, many of the learning problems stated in generative linguistics can be solved. One of the main learning problems is the 'poverty of the stimulus' (Chomsky 1986:7): the fact that during the learning period a child never gets to hear every possible sentence that the adult language can construct (the child is exposed to an 'impoverished input'), yet is able to learn the language. It has therefore been argued that a child does not learn language by heart in memorizing every sentence she ever gets to hear, but that she can learn the language by inferring hidden structure from the speech samples that she encounters. By abstraction, the learner can adapt her grammar and create new, meaningful utterances. This abstraction can be handled in phonology by the mentioned mechanism Robust Interpretive Parsing. Another problem in the literature on learnability is the problem that children can learn only from positive evidence in the data. A child cannot deduce from the fact that she does *not* hear a certain structure that this structure is *not* permitted. She can only deduce from a certain structure she *does* hear that it *is* permitted. OT deals with this problem by providing the learner with implicit negative evidence: any form that is regarded as optimal by the child's grammar constitutes the positive evidence, and all other candidates that are discarded as not optimal constitute the negative evidence (Tesar & Smolensky 2000:33). The mechanisms of interpretation and constraint reranking are outlined in chapter 3, and applied in chapters 4, 5 and 6.

I exemplify the proposed approach to learning with word stress of different languages. The dissertation at hand models three different kinds of stress systems, outlined in chapter 2. I distinguish between *grammatically assigned stress* (modelled in chapters 4 and 5) and *lexically assigned stress* (modelled in chapter 6). Within grammatical stress systems I distinguish between *weight-sensitive* stress and *weight-insensitive* stress. The first language is Classical Latin in chapter 4, a language with weight-sensitive stress. Although a dead language, the prosodic system of Latin is well-studied in phonology. The study of Latin provided insights into cross-linguistic principles of phonology, and the study of the learnability of Latin stress can provide insights into cross-linguistic principles of learnability. Stress in Latin is largely determined by *heavy* syllables, and therefore provides an example for a weight-sensitive language. With the modelling of Latin stress it will be shown that under the learning model argued for, learners are not only capable of inferring hidden structures from what they hear, but can create structure on their own, if permitted (e.g. create secondary stress from data with only primary stress). They come up with slight variations in their grammar. This chapter furthermore provides a comparison between different analyses that have been proposed for Latin stress, and their learnability. It turns out that some of the analyses are better learnable than others, and that some are not learnable at all. The chapter also provides a comparison between the two reranking strategies CD and GLA, and a comparison of different constraint sets that the learners are equipped with. Three representational levels are involved: one phonetic level and the two phonological levels of surface form and underlying form (the three bottom representations of figure (1)).

The second language is Pintupi, a Pama-Nyungan language spoken in Western Australia, modelled in chapter 5. Pintupi has weight-insensitive stress. The modelling of Pintupi stress makes clear that the overt forms of the learners are uniform, yet the grammars of the different learners can vary. Communication among each other is still guaranteed. Again, CD and the GLA are tested with respect to their performance in learning, as well as different constraint sets. As in the Latin simulations, three levels of representation are modelled: phonetic form, surface form, and underlying form.

The third language that is modelled is Modern Greek in chapter 6, where stress is largely determined by the lexicon. Lexically assigned stress is interesting to model, because it is not predictable form the grammar and makes it necessary to account for a learning of underlying forms. In this chapter, all four levels of representation that are shown in figure (1) are modelled.

I conclude in chapter 7 that within the proposed learnability approach:

1. learners are able to create structure (shown in chapter 4),
2. the grammars of speakers can differ while they are still speaking the same language: they can have different surface and underlying structures, but as long as their overt forms are the same, communication is guaranteed (shown in chapters 5 and 6), and
3. parts of the lexicon can be learned by acquiring the phonological grammar (shown in chapter 6).