Institute of Phonetic Sciences,
University of Amsterdam,
Proceedings 25 (2003), 101–148.

THE LEARNABILITY OF LATIN STRESS

Diana Apoussidou & Paul Boersma

Abstract

Optimality-Theoretic learning algorithms are only guaranteed to be successful if the data fed to them contain full structural descriptions of the surface forms, i.e. descriptions that include hidden structure like metrical feet. This is not realistic as a model of acquisition, because children are only exposed to *overt forms*, e.g. unstructured strings of syllables. Optimality-Theoretic learning algorithms that learn solely from overt forms turn out to sometimes succeed and sometimes fail (Tesar & Smolensky 2000). This possibility of failure is a property of both on-line learning algorithms that have been proposed for OT, namely *Error Driven Constraint Demotion* (EDCD; Tesar 1995) and the *Gradual Learning Algorithm* (GLA; Boersma 1997). The possibility of failure is not necessarily bad: one would want an algorithm to fail for languages that do not exist, and to succeed for languages that do exist. Latin exists (or existed). This paper compares the performance of the two learning algorithms for the metrical stress system of Classical Latin. It turns out that EDCD cannot learn this system from overt forms only, and that the GLA can. This suggests that the GLA may be a better model of acquisition than EDCD. The results also provide evidence in the discussion in the literature about what is the correct linguistic analysis of Latin stress: if overt forms contain main stress only, the GLA makes the child posit an analysis that makes use of uneven trochees (like the analysis by Jacobs 2000) rather than strictly bimoraic trochees (like the analysis by Mester 1994 and Hayes 1995).

To linguists, learnability theory is about creating formal models of language acquisition, i.e. it investigates what precisely is known by the beginning learner and how precisely the learner proceeds from this initial state to an adult state on the basis of language input. Some linguistic phenomena are universally observable throughout the languages of the world. For instance, children of all languages usually start their production with CV syllables, even if their language also employs VC syllables (e.g. Fikkert 1994). The generative explanation for such observations is that there exists a Universal Grammar, shared by all languages in the world, which defines the child's initial state and restricts the possible forms that her subsequent developmental grammars can take. A model of Universal Grammar that has become popular among linguists in the last ten years is Optimality Theory (henceforth OT; Prince & Smolensky 1993, McCarthy & Prince 1993, and thereafter). In this framework, universal properties of language (such as the principle that a CV syllable is more wellformed than a VC syllable) are expressed in *constraints*. In the original formulation of OT, all languages in the world share the same constraints, although the effects of all constraints might not be overtly visible in all of the languages. The difference between languages results from assigning priorities to the constraints, which is expressed in a hierarchy (*ranking*). In order to acquire his or her language, a child has to learn the correct ranking, not the constraints themselves. While the universality of constraints could be questioned in general, we will assume in this paper that at least the structural constraints that handle metrical phonology are the same in all languages. As we will see, this opens up the possibility that a descriptively simple metrical system, like that of Latin, turns out to be surprisingly complicated when described in terms of constraints proposed by linguists on the basis of cross-linguistic typology rather than in terms of constraints tailored to the specific language at hand.

To obtain information about the universal components of the grammar, one can analyse the actual language acquisition process of infants and children. Such an analysis is quite difficult in the case of phonological perception, since we cannot look inside a speaker's head to see what happens during perception, and speakers themselves, children included, have very little conscious access to the perception process, let alone the capability of reliably reporting on it. The analysis is slightly less difficult in the case of the child's language *production*, since in that situation at least part of the output of the grammar can be observed directly. But even when considering produced forms, the researcher meets with hidden structures like metrical feet, which often remain ambiguous.

Another method to identify universal aspects of the grammar is to try and simulate the acquisition process with the help of a computational learning algorithm. In that way, the universal principles derived from language acquisition data can be tested with respect to their adequacy. To make this work, a learning algorithm needs to be supplied with the universal ingredients of the grammar, which in the case of Optimality Theory means that the learning algorithm should be supplied with a universal set of constraints.

Simulating learnability has a further benefit for linguistics, namely providing evidence for or against existing analyses in the literature. By means of a learning algorithm that is based on Optimality Theory, existing OT analyses of a language can be tested with respect to their learnability. If it turns out that an analysis proposed in the literature is not learnable with a certain learning algorithm, then either this analysis or this learning algorithm should be rejected.

In this paper, we test the learnability of the metrical phonology of Latin word stress. Taking a dead language as the test subject is not as awkward as it may look. The metrical phonology of Latin has been extensively studied by linguists. Many of the principles found in Latin word stress have fed ideas about universal constraint sets for metrical phonology in general, and have been used to analyse other languages. In turn, Latin has been analysed with constraints whose cross-linguistic validity has been established in analyses of other languages. Since there exist several OT analyses of Latin word stress, we compare them with respect to their learnability. We also test different sets of data, to be able to determine the amount of information needed for a successful simulation. In addition, we run the simulation with two OT-based learning algorithms that differ with respect to their way of constraint re-ranking during the acquisition process. The paper is structured as follows. In §1 we outline the basic ideas of Optimality Theory when applied to a simple metrical example. In §2 we discuss the learnability problems associated with this example. In §3 we discuss the ingredients of metrical phonology that are needed to account for word stress. In §4 follows a description of Latin word stress, and a comparison of the various proposed analyses of the Latin stress system. In §5 we simulate the acquisition of Latin stress with a computer implementation of the two learning algorithms. In §6 we present the results, showing that in several respects the GLA performs better than EDCD. In §7, we place the findings in a larger perspective and discuss their implications for learnability theory and OT.

1 Optimality Theory

In the original version of OT (Prince & Smolensky 1993), all languages of the world share the same set of violable constraints, and the languages differ only in the *ranking* of these constraints, i.e. their relative degree of importance. The ranking of all the constraints with respect to each other then constitutes the *grammar* of a language. In this section, we will explain how the ranking of the constraints determines the actual surface forms of a language.

Consider the production part of a grammar: a speaker would like to produce an utterance. She first creates an underlying form (the *input* to the grammar) from a sequence of lexical items. On the basis of this underlying form she then chooses a surface form (the *output* of the grammar) from among a set of possible *output candidates*. The candidate that is chosen as the *optimal* one from all these candidates is the one that satisfies the constraints best. Typically, the winning candidate will satisfy most high-ranked constraints, while at the same time it could abundantly violate lower-ranked constraints. This process of *evaluation* is portrayed in *tableaus*, where the candidates are compared with respect to their fulfillment of the constraints.

Let us apply this idea to the problem of metrical structure. Phonologists generally agree that while in some languages stress could be assigned by referring to syllables only (e.g. "always stress the first syllable"), the analysis of many languages requires one to assume that syllables are grouped into hidden structures called *feet*. In every foot, one syllable is prominent, i.e., it receives *stress*. For the purposes of this section and the next, we will only consider *disyllabic* feet, i.e. feet consisting of two syllables (we will work with a larger set of foot forms when discussing Latin). There are two kinds of these (σ stands for "syllable"): the *trochee* (σ́ σ), in which the first syllable is stressed, and the *iamb* (σ σ́), in which the second syllable is stressed. Consider now the small universal constraint set in (1), where two constraints (IAMBIC and TROCHAIC) are responsible for the placement of the stressed syllable within the foot, and two constraints (ALIGNFT-R and ALIGNFT-L) are responsible for the placement of the foot within the word.

(1) *Constraints on metrical constituents*
IAMBIC: "the stressed syllable is the last syllable in its foot."
TROCHAIC: "the stressed syllable is the first syllable in its foot."
ALIGNFT-R: "align the right edge of the foot with the right edge of the word."[1]
ALIGNFT-L: "align the left edge of the foot with the left edge of the word."

These four constraints are among the many that have been proposed in the literature to account for generalizations on the phenomena of metrical phonology. The constraints IAMBIC and TROCHAIC stem from the observation that languages tend to have either iambic or trochaic feet, rather than a mix of them (McCarthy & Prince 1986, Kager 1996, Van de Vijver 1998), and ALIGNFT-R and ALIGNFT-L stem from the observation that languages tend to have feet that are either close to the end or close to the beginning of the word, or tend to assign feet iteratively starting either near the end or near the beginning of the word.

Consider now an underlying form with three syllables, |σ σ σ|.[2] If we assume that stress assignment is purely determined by the grammar (i.e. the language at hand does not have lexical stress), then we have at least the four different candidates shown in

[1] In these formulations, the *left edge* of a constituent refers to its beginning, and the *right edge* to its end.
[2] We use the following symbols to bracket the various representations: pipes for underlying forms, slashes for abstract surface structures, and square brackets for overt forms.

tableau (2). The four candidates have different main stresses, denoted here as "σ́", and different foot structures, denoted here by parentheses. Suppose now that in a specific language the highest ranked constraint is IAMBIC and the lowest ranked constraint is ALIGNFT-R. This ranking is denoted in the tableau by sorting the constraints from left to right. The asterisks (*violation marks*) in the tableau depict which candidates violate which constraints. The candidates /(σ́ σ) σ/ and /σ (σ́ σ)/ both violate the highest ranked constraint IAMBIC, since they contain a trochaic foot. These violations are marked with a "!" because they are the crucial violations that rule out these two candidates from further consideration. The choice between the remaining two candidates /(σ σ́) σ/ and /σ (σ σ́)/ cannot be made by the two highest ranked constraints IAMBIC and TROCHAIC, since these two constraints have an equal number of violations for these two candidates. The matter is decided by ALIGNFT-L, which prefers the candidate /(σ σ́) σ/, since this form has a left-aligned foot, unlike /σ (σ σ́)/. The grey cells in the tableau are those that have not contributed to the determination of the winning form. The winning candidate itself is finally denoted by a pointing finger.

(2) *An iambic left-aligning language*

Underlying: |σ σ σ|	IAMBIC	TROCHAIC	ALIGNFT-L	ALIGNFT-R
/(σ́ σ) σ/	*!			*
☞ /(σ σ́) σ/		*		*
/σ (σ́ σ)/	*!		*	
/σ (σ σ́)/		*	*!	

The language of tableau (2) can be said to be an iambic left-aligning language: the foot in the winning candidate is iambic and this foot is left-aligned within the word. The iambicity is a result of the ranking IAMBIC >> TROCHAIC, and the left alignment is the result of the ranking ALIGNFT-L >> ALIGNFT-R.

An early assumption in OT is that any ranking of the constraints should correspond to an attestable language. OT thus makes typological predictions. If we assume, for instance, a universal grammar with four constraints A, B, C, and D, these constraints can be ranked in 24 different ways, and in the extreme case this could lead to 24 different types of languages. Thus, we get a different type of language if TROCHAIC dominates IAMBIC, as in tableau (3): we get stress on the first syllable, due to a trochaic foot that is still aligned at the left edge of the word.

(3) *A trochaic left-aligning language*

Underlying: |σ σ σ|	TROCHAIC	IAMBIC	ALIGNFT-L	ALIGNFT-R
☞ /(σ́ σ) σ/		*		*
/(σ σ́) σ/	*!			*
/σ (σ́ σ)/		*	*!	
/σ (σ σ́)/	*!		*	

Let us have a look at what happens if the alignment constraints are ranked differently. If the iambic foot structure is preferred, and ALIGNFT-R is ranked over ALIGNFT-L, then stress will be on the last syllable in the output, as in (4).

(4) *An iambic right-aligning language*

Underlying: |σ σ σ|	IAMBIC	TROCHAIC	ALIGNFT-R	ALIGNFT-L
/(σ́ σ) σ/	*!		*	
/(σ σ́) σ/		*	*!	
/σ (σ́ σ)/	*!			*
☞ /σ (σ σ́)/		*		*

If, however, the trochaic foot form is preferred, and ALIGNFT-R outranks ALIGNFT-L, stress will be on the second syllable again, as in (5).

(5) *A trochaic right-aligning language*

Underlying: |σ σ σ|	TROCHAIC	IAMBIC	ALIGNFT-R	ALIGNFT-L
/(σ́ σ) σ/		*	*!	
/(σ σ́) σ/	*!		*	
☞ /σ (σ́ σ)/		*		*
/σ (σ σ́)/	*!			*

We have now exhausted the typological possibilities of these four constraints. The 24 possible rankings lead to only four different types of languages, since e.g. changing the ranking of the alignment constraints with respect to the foot form constraints does not lead to any new types of languages. Within this simple set of constraints, IAMBIC only competes with TROCHAIC, and ALIGNFT-R with ALIGNFT-L.

What the rankings predict is that in the language in (2), feet will be ranked at the left edge in every word. And since every foot in this language is iambic, stress will always be on the second syllable in a word. In language (3), stress will always be on the first syllable in a word, since here the foot form is trochaic, but feet are still ranked at the left edge of a word. Language (4) will always have final stress, since feet are iambic and are aligned at the right edge of the word. The last language in (5) will always have stress on the penultimate syllable, that is on the pre-final syllable. So whenever a linguistic principle is translated into an OT constraint, it should make predictions about languages. In other words, any ranking with this constraint should result in an attestable language. As we have seen, though, not all constraints have to be in competition with each other.

What we have also seen (and that brings us closer to the learnability problem) is that there are actually two grammars here, (2) and (5), that show the same stress pattern in trisyllabic words, namely stress on the second syllable. In other words, the two surface forms /(σ σ́) σ/ and /σ (σ́ σ)/ share the same *overt form*, namely [σ σ́ σ]. A child cannot learn the ranking of these languages from trisyllabic words alone. In this case, she will crucially depend on the presence of other, either shorter or longer forms, to figure out the exact ranking. Luckily, languages usually do employ words with more and less than three syllables. However, with the more complete constraint set that we will use later (12 constraints), it is not so obvious that informative forms will always exist, and this will be seen (in §6) to have repercussions on learnability. In the next section we have a closer look at how OT grammars can be learned.

2 Learnability

We assume (perhaps unrealistically, see §7.2) that the language-learning child already knows what the constraints for metrical phonology are. Her remaining task, then, is to rank these constraints in a way appropriate to the target language. For instance, if the language to be learned is iambic and left-aligning, she will have to establish for herself a grammar in which IAMBIC outranks TROCHAIC, and ALIGNFT-L outranks ALIGNFT-R. For the purposes of this paper, we assume that in the child's initial state, all metrical constraints are ranked at the same height, and that the child subsequently has to modify the rankings of these constraints in a way that takes her closer to the target language. There are different approaches of how this reranking of constraints into the correct hierarchy proceeds. In this article, we compare two options: Error Driven Constraint Demotion (Tesar 1995) with the Gradual Learning Algorithm (Boersma 1997). In both algorithms, learning is triggered by detecting a mismatch between the adult surface form (which is either explicitly given to the child or has to be reconstructed by the child) and the child's own surface form (i.e. the winner in her current grammar). If the two forms are different, the child takes action by changing the ranking of the constraints in her grammar. In EDCD, constraints that turn out to be violable are demoted to a lower point on the ranking scale. In the GLA, constraints can be promoted as well as demoted.

2.1 Learning from full information with Error Driven Constraint Demotion

The Error Driven Constraint Demotion algorithm (EDCD; Tesar 1995, Tesar & Smolensky 1998, 2000) is an on-line learning algorithm, i.e., it modifies the learner's grammar directly on the basis of incoming language data. The idea is that the learner considers incoming adult forms as 'correct' and her own forms (i.e. forms that her current grammar can generate) as 'incorrect'. Thus, the negative evidence needed for error-driven learning is provided internally by the learner's current grammar. The steps by which an EDCD learner acquires her language are explicitly described in (6).

(6) *Error-Driven Constraint Demotion*
   a. *Learning datum*: the learner encounters a linguistic data pair in the form of a given underlying form and a given adult surface form.
   b. *Generation*: from the given underlying form, the learner computes her own form, i.e. the surface form that is optimal in her current grammar.
   c. *Comparison*: the learner compares the adult form with her own form.
   d. *Adjustment*: if the two surface forms are different, the learner takes action by changing her grammar minimally in such a way that the adult form becomes better (more *harmonic*) than her own form: the learner makes sure that all constraints that prefer her own form to the adult form (i.e. all constraints that are violated more often in the adult form than in the learner's form) become ranked below the highest ranked constraint that prefers the adult form.[3]
   e. *Acquisition*: steps a–d are repeated for all incoming learning data. Once the learner has reached a grammar that can generate all and only adult-like forms, there will cease to be any adjustments.

[3] In the most recent version of EDCD (Tesar & Smolensky 2000), this adjustment procedure is repeated until the adult form becomes optimal in the learner's grammar ("multiple chews", see §5.3).

The following pictures illustrate how the adjustment step proceeds. Figure (7) shows a grammar with eight constraints, ranked according to C_1 >> C_2 >> C_3 >> C_4 >> C_5 >> C_6 >> C_7 >> C_8. Suppose that C_2, C_3, C_6, and C_8 prefer the learner's form (i.e., are violated more often in the adult form than in the learner's form), C_5 and C_7 prefer the adult form (i.e. are violated more often in the learner's form than in the adult form), and C_1 and C_4 have no preference (i.e., each of these constraints is violated equally in both forms). The learner now searches for the *pivotal constraint*, i.e. the highest ranked constraint that prefers the adult form. This is C_5. The learner then searches for all the even higher ranked constraints that prefer the learner's form. These are C_2 and C_3. Figure (8) now shows how these two constraints are demoted below the pivotal constraint, thus ending up in the same *stratum* (ranking layer) as C_6. Note that C_8 need not be moved, since it is already lower than C_5. Since the highest ranked constraint that makes a difference between the two forms is now C_5, the adult form has become more harmonic in the new grammar than the learner's form.

(7) *Before grammar adjustment*

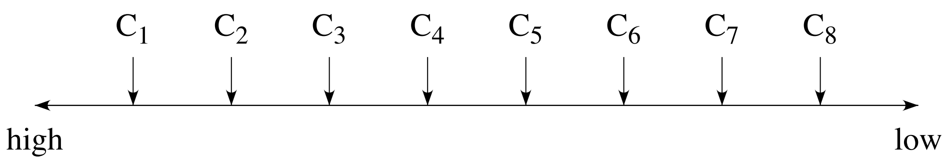

(8) *After grammar adjustment by constraint demotion*

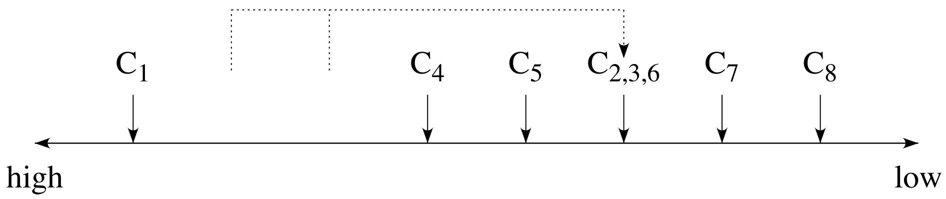

We will now illustrate how EDCD works for the example of a child trying to learn an iambic left-aligning language. Suppose that at some point during the acquisition process the child has a constraint ranking that is maximally opposite to such a language, namely the ranking in tableau (9). The child will then assign to a trisyllabic underlying form the surface structure /σ (σ́ σ)/. Now suppose (unrealistically, as we will discuss from §2.3 on) that somebody tells the child explicitly that the correct adult form is /(σ σ́) σ/ instead. This form is denoted in the tableau by a check mark (√). If the child is an EDCD learner, she can now take action by first looking up the highest-ranked constraint that prefers the adult form (this is IAMBIC) and then demoting all higher ranked constraints that prefer the child's form (in this case, only TROCHAIC) below this constraint.

(9) *A trochaic right-aligning language*

Underlying: |σ σ σ|	TROCHAIC	IAMBIC	ALIGNFT-R	ALIGNFT-L
/(σ́ σ) σ/		*	*!	
√ /(σ σ́) σ/	*!		*	
☞ /σ (σ́ σ)/		*		*
/σ (σ σ́)/	*!			*

TROCHAIC thus ends up in the same stratum as ALIGNFT-R, as we see in tableau (10). The grammar may now have improved in the sense that the adult form /(σ σ́) σ/ has become more optimal than the previous child form /σ (σ́ σ)/, but the grammar is still incorrect because the new winner is the equally incorrect form /σ (σ σ́)/.

(10) *After one EDCD step*

Underlying: |σ σ σ|	IAMBIC	TROCHAIC	ALIGNFT-R	ALIGNFT-L
/(σ́ σ) σ/	*!		*	
√ /(σ σ́) σ/		*	*!	
/σ (σ́ σ)/	*!			*
☞ /σ (σ σ́)/		*		*

But we can apply EDCD again to tableau (10). The pivotal constraint is ALIGNFT-L, and the only constraint that EDCD demotes below it is ALIGNFT-R. After this second reranking the underlying form |σ σ σ| will be produced as the surface form /(σ σ́) σ/, as in (11).

(11) *After two EDCD steps*

Underlying: |σ σ σ|	IAMBIC	TROCHAIC	ALIGNFT-L	ALIGNFT-R
/(σ́ σ) σ/	*!			*
√ ☞ /(σ σ́) σ/		*		*
/σ (σ́ σ)/	*!		*	
/σ (σ σ́)/		*	*!	

After two EDCD steps, the production of trisyllabic underlying forms is correct. It also happens to be the case that forms with other numbers of syllables are handled correctly as well. Learning has thus ended and the learner has succeeded in acquiring the target language.

However, in real life there is a considerable amount of optionality, both in what children produce and in what they encounter in the data. In this respect EDCD is not a realistic model of acquisition, since it cannot handle this phenomenon. A development towards solving this problem is the Gradual Learning Algorithm, which is described in the next section.

2.2 Learning from full information with the Gradual Learning Algorithm

The Gradual Learning Algorithm (GLA; Boersma 1997, Boersma & Hayes 2001) differs in two ways from the EDCD algorithm described in (6). First, the generation step of the GLA assumes *Stochastic OT* (Boersma 1998), in which constraints are ranked along a continuous scale and have rankings that can vary a bit between evaluations. Second, the adjustment step of the GLA is simpler than that of EDCD, since it does not depend on the relative rankings of the constraints: the GLA simply demotes *all* the constraints that prefer the child's form while *promoting* all the constraints that prefer the correct adult form. The extent to which the constraints are moved is not a big stride past a neighbouring constraint, but a small step along the continuous ranking scale.

The following pictures illustrate how the generation and adjustment steps proceed. In Stochastic OT, constraints are still ranked along a scale, but instead of marking distinct points on it, they cover *ranges* on it, in the form of a Gaussian distribution. If the current grammar of a child favours a candidate that differs from the perceived adult form, the relevant constraints start to move. The constraint that favours the child's candidate (C_1) is moved to a lower place on the ranking scale, while the constraint that favours the adult form (C_2) is moved further up the ranking scale. Because the constraints cover a range rather than a distinct point, this can lead to an overlap of these ranges, as in (13). This means that in this situation several outputs are possible. Since Stochastic OT adds a certain amount of noise to the evaluation of the candidates, the optimal output is chosen with a probability depending on the amount of overlap of the relevant constraints.

(12) *Initial ranking*

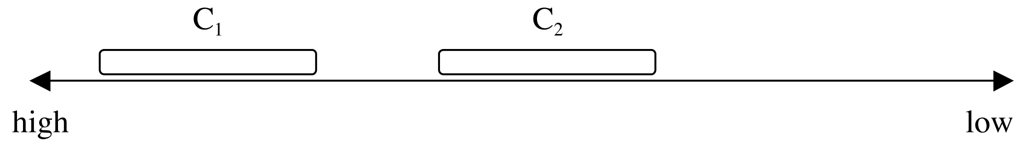

(13) *The learning process: a gradual constraint shift*

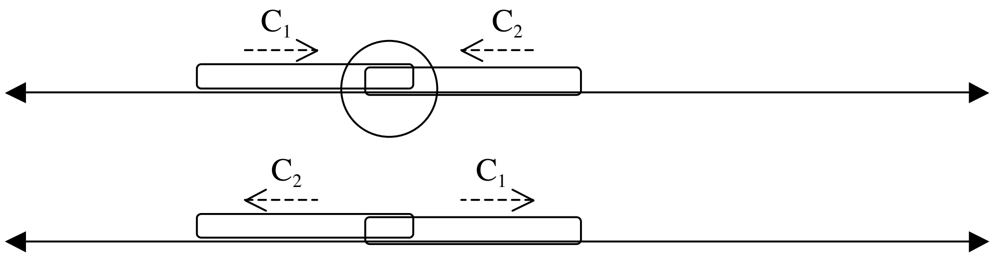

(14) *Final ranking after many learning steps*

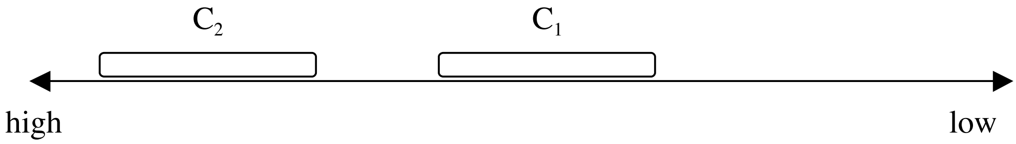

We will now illustrate how the GLA works for a child learning an iambic left-aligned language. In tableau (9), TROCHAIC and ALIGNFT-R will be demoted while IAMBIC and ALIGNFT-L will be promoted. In tableau (10), ALIGNFT-R will be demoted and ALIGNFT-L will be promoted. This does not immediately result in a complete reversal of the constraint ranking.

Both EDCD and GLA always lead to a correct adult-like grammar if fed with a sufficient number of fully specified pairs of underlying and surface forms. The power of the GLA lies in that it can handle noisy data as well as free variation in the data. It can also explain the intermediate stages in the acquisition process (Boersma & Levelt 1999, Curtin & Zuraw 2001): a child starts to produce a modified form or even the correct adult form while still using her old form at some time or other.

The GLA has been applied in the literature with considerable success when learning from fully specified metrical surface structure (Curtin & Zuraw 2001). What we expect from applying EDCD and the GLA to the Latin metrical system is to see to what extent these algorithms can cope with overt data that do *not* contain information on foot structure. The next section describes a strategy by which the learner can guess the missing information.

2.3 Robust Interpretive Parsing

The learning situations assumed in §2.1 and §2.2 were not realistic: a real child is not told the full structural description of the surface form. Consider again the iambic left-aligning language whose constraint ranking was shown in (2). A child acquiring this language hears not a fully structured input /(σ σ́)/ or /(σ σ́) σ σ/, but only the *overt* forms [σ σ́ σ] and [σ σ́ σ σ]. So the foot structure of the forms is hidden from the child. She has to establish the underlying form and the surface structure by herself. For that she needs to learn the correct mapping from the overt form [σ σ́ σ] to a surface form like /(σ σ́) σ/. Tesar (1997) and Tesar & Smolensky (2000) proposed a mechanism whereby the learner can find the surface form and the underlying form: *robust interpretive parsing* (RIP). According to this proposal, the mapping from overt to surface form (i.e. what a phonetician would call *perception*) is performed by the grammar itself; and since the grammar in turn is what the child has to acquire with EDCD or the GLA in *production*, the learner will need to move back and forth between the comprehension and production processes, as we will see.

Implemented in an OT approach, RIP looks like the following. If the child uses the same grammar for comprehension and for production (Smolensky 1996), then she uses her current grammar to deduce the hidden structure (in our case, foot structure) from the overt form she hears (Tesar 1997). So if her current grammar produces trochees (caused by the ranking TROCHAIC >> IAMBIC), then she will interpret an overt form [σ σ́ σ σ] as the surface form /σ (σ́ σ) σ/, as shown in tableau (15). Only two surface structures are compatible with the given overt form, so there are only two candidates. The second candidate satisfies the four ranked structural constraints best, so this is considered the optimal interpretation of the overt form (as a detail, note that ALIGNFT-R is interpreted as a *gradient* constraint here, i.e., it is violated twice because the foot is two syllables away from the right edge of the word).

(15) *Grammar-guided interpretation by the learner*

Overt: [σ σ́ σ σ]	TROCHAIC	IAMBIC	ALIGNFT-R	ALIGNFT-L
/(σ σ́) σ σ/	*!		**	
☞ /σ (σ́ σ) σ/		*	*	*

Having thus heard an adult form, the learner can compute two more forms. The first is rather trivial in our case: the underlying form, which must be |σ σ σ σ| (if we ignore for now the complicating possibility of lexical stress). The second form that the child can compute is her own surface form, i.e. what she herself would have produced given the underlying form |σ σ σ σ|. As shown in tableau (16), she would have produced the surface form /σ σ (σ́ σ)/, which is optimal in her current grammar.

(16) *Learning from the first interpreted input-output pair*

Underlying: |σ σ σ σ|	TROCHAIC	IAMBIC	ALIGNFT-R	ALIGNFT-L
/(σ́ σ) σ σ/		*	*!*	
/(σ σ́) σ σ/	*!		**	
√ /σ (σ́ σ) σ/		*	*!	*
/σ (σ σ́) σ/	*!		*	*
☞ /σ σ (σ́ σ)/		*		**
/σ σ (σ σ́)/	*!			**

Thus, the child assigns to the overt form that she hears a structure that minimally violates her constraint ranking, even if this structure is ungrammatical in her own production, as it is in (16).

After interpretive parsing, as in (15), and subsequent silent production, as in (16), the child can notice that the two surface forms are different. In (16), the interpreted adult form is denoted by a check mark (√) to show that the child will consider this to be the correct adult form. Realizing the mismatch between the interpreted adult form and her own corresponding production form, the child can take action by reranking one or more constraints, in order to make it more likely that her production form will match the adult form in the future. An EDCD learner would demote ALIGNFT-R below ALIGNFT-L, and thus arrive at the grammar in (17).

If the child, after learning from a quadrisyllabic form, now encounters the trisyllabic form [σ σ́ σ], her interpretation (according to her new grammar) will be /σ (σ́ σ)/, as shown in (17).

(17) *Second grammar-guided interpretation by the learner*

Overt: [σ σ́ σ]	TROCHAIC	IAMBIC	ALIGNFT-L	ALIGNFT-R
/(σ σ́) σ/	*!			*
☞ /σ (σ́ σ)/		*	*	

Once again, this is not what the learner will produce herself, as shown in (18).

(18) *Learning from the second interpreted input-output pair*

Underlying: |σ σ σ|	TROCHAIC	IAMBIC	ALIGNFT-L	ALIGNFT-R
☞ /(σ́ σ) σ/		*		*
/(σ σ́) σ/	*!			*
√ /σ (σ́ σ)/		*	*!	
/σ (σ σ́)/	*!		*	

Again the learner will compare her own produced form /(σ́ σ) σ/ with the interpreted adult form /σ (σ́ σ)/, and since the two are different she will again rerank the constraints. She will demote ALIGNFT-L below ALIGNFT-R, and this kind of reranking goes on with different input data until the child has acquired a grammar that is appropriate for the target language... But wait a minute! After the second reranking, we have arrived at the same grammar that we started with. The learner has still lacked a clue to reranking the foot form constraints. Unless the learner receives a form that will force her to demote TROCHAIC below IAMBIC, she will never arrive at the target language. In the language at hand, the disyllabic overt form [σ σ́] will do the job, since it can only be interpreted as the iambic /(σ σ́)/ (see Boersma & Levelt 2003 for the interpretation and generation tableaus), but for languages in general there is no guarantee that RIP/EDCD leads to successful acquisition. Tesar & Smolensky (2000) report that a simulation with 124 language types led to a success rate of only 60 percent, i.e., 40 percent of those 124 languages could not be learned by RIP/EDCD. The combination of RIP with GLA fares just a little bit better, but is still unable to learn 30 percent of those languages (Boersma, to appear). In RIP/GLA, the interpretation step is done within Stochastic OT, i.e. after adding a bit of evaluation noise to the constraint rankings, and this same temporary ranking is then used for the generation of the learner's own form; the adjustment step proceeds as usual, i.e. with reranking of all the contraints that prefer the adult form or the learner's form.

The failure of RIP/EDCD and RIP/GLA on some language types is not necessarily bad. A learning algorithm should work for all existing languages, and by failing on some language types it should be able to predict what kinds of languages are impossible to learn. OT learning algorithms, for instance, could predict holes in the factorial typology, i.e., they could predict what permutations of the constraints are impossible to learn; such languages would be allowed by the framework of Optimality Theory itself, yet would not exist, because there is no path by which children can acquire them. Latin is a language that has existed in reality, so its stress system must be learnable by the algorithm that real humans use to acquire their phonology. We will put both RIP/EDCD and RIP/GLA to the Latin test.

3 Metrical phonology

Let us have a look at the metrical system of a language with grammatically assigned word stress.[4] As already pointed out in §1, an important constituent for assigning stress to words is the foot. By causing a rhythmic organization of syllables, the foot underlies the metrical patterns of many languages. Feet are usually *binary*, i.e., they group syllables into pairs, resulting in a pattern of (often alternating) weak and strong syllables. In figure (19), the English word *demotion* consists of three syllables, a weak one followed by a strong one and another weak one (PrWd = *prosodic word*).

(19) *A hierarchical prosodic structure*

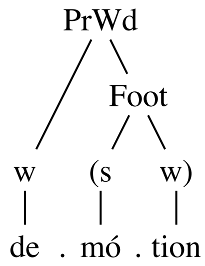

The strong syllable is prominent (stressed) in the output. According to analyses of English stress (e.g. Liberman & Prince 1977), English has trochaic, binary feet, meaning that a strong syllable forms a foot together with a following weak syllable. But how does a learner find out that the feet in her language are strong-weak sequences (trochaic) rather than weak-strong sequences (iambic)? How you group the syllables within one word will have an effect on how you will stress other words in your language. The problem is that foot structure belongs to what we have called above the fully specified surface structure and is not contained in the overt form that a learner is actually exposed to. So the learner has to find out by herself how the syllables in her language group together, e.g. as trochaic or as iambic feet.

As indicated above, feet are usually binary. However, not all languages count syllables only: some count moras as well. A mora is a smaller unit than a syllable, and determines the *weight* of the syllable: syllables with a long vowel or a diphthong contain two moras (they are *heavy*), while syllables with only a short vowel contain only one mora (they are *light*). Depending on the language, syllables that end in a

[4] Another possibility of stress assignment is lexical stress, where stress is assigned by marks for stress in the lexicon. Lexical stress would involve faithfulness constraints, which we would like to leave out of discussion here.

consonant can also count as two moras.[5] Languages in which the number of moras in a syllable influences stress (or other phonological phenomena) are called *weight*- or *quantity-sensitive*. In such languages, heavy syllables tend to be prominent in the output. In OT, this principle is captured in the constraint WEIGHT-TO-STRESS-PRINCIPLE (WSP), defined in (20).

(20) *A constraint for quantity sensitivity*
WSP: "heavy syllables are stressed."

In quantity-sensitive languages a foot ideally consists of two moras: either two light syllables or one heavy syllable. In quantity-insensitive languages, feet ideally consist of two syllables, regardless of their inner construction. This binarity is expressed in OT as a constraint FOOTBIN, defined in (21).

(21) *A constraint for foot binarity*
FOOTBIN: "feet are binary on some level of analysis (mora or syllable)."

Note that this constraint allows a foot to consist of three or four moras, as long as these moras are contained in a sequence of two syllables (heavy-light, light-heavy, or heavy-heavy).

Stress assignment is not only determined by foot-internal structure, but by the placement of feet within the word (or phrase) as well. Especially for longer words, the question is at what edge of the word the foot (or the feet) will be constructed. Some languages tend to build feet at or from the left edge of a word (word-initially), others at of from the right edge (word-finally). We have already mentioned two constraints for foot placement (ALIGNFT-R and ALIGNFT-L), and we will meet with several others when discussing Latin in the next section.

4 The metrical system of Latin

No native speakers can tell us how Latin was originally pronounced. Thus, no phonetic analysis is available. Still, its prosodic system is at least partly accessible through analyses of written text such as poems or language descriptions of contemporary witnesses. We decided to take the stress system of Latin as our test subject because it has been studied by linguists at great length. Latin is often taken as the prototypical example in general studies on metrical phonology when it comes to illustrating phenomena like weight-sensitivity and extrametricality (e.g. Allen 1973; Hayes 1985, 1987; McCarthy & Prince 1986; Prince 1990).

4.1 The facts of Latin stress

In Latin, stress is handled purely by the grammar: the foot structure of a word is predictable from the syllable structure of the word, and the mental lexicon need not contain any information about where in the word the stress is realized. Thus, a learner of Latin does not have to take into account the complexities that would arise if the language had lexically assigned stress as well. This should make it relatively easy for a learner to figure out the ranking of the relevant constraints. One would think.

Basically, Classical Latin has left-prominent feet (trochees), it is quantity-sensitive (the weight of the penultimate syllable is especially important), and the last syllable in

[5] Many more weight distinctions can be observed in the languages of the world (Hyman 1985, Gordon 2002).

a word is extrametrical (i.e., it never receives stress except if it is the only syllable of a word) (Allen 1973, and literature thereafter). Syllables ending in a short vowel are light (abbreviated here as ‘L’), while syllables with long vowels or diphthongs and syllables that end in a consonant are heavy (‘H’). In words with three or more syllables, the penultimate syllable is stressed if it is heavy. If the penultimate syllable is light, the antepenultimate syllable is stressed, regardless of its weight. In words with only one or two syllables, the leftmost syllable is stressed. Some examples are given in (22). The second column represents what we called *overt forms* in §2: phonetic representations with stress (ˈ) and vowel length (ː), enriched with some hidden phonological structure (periods indicate syllable boundaries) but without foot structure. The third column represents these overt forms without segmental information, i.e. the overt stress patterns (‘1’ for main stress).

(22) *Weight and stress in Latin*

amice ‘friend’	[a.ˈmiː.ke]	[L H1 L]
rapiditas ‘speed’	[ra.ˈpi.di.taːs]	[L L1 L H]
misericordia ‘pity’	[mi.se.ri.ˈkor.di.a]	[L L L H1 L L]
perfectus ‘perfect’	[per.ˈfek.tus]	[H H1 H]
incipio ‘I begin’	[iŋ.ˈki.pi.oː]	[H L1 L H]
domesticus ‘domestic’	[do.ˈmes.ti.kus]	[L H1 L H]
homo ‘man’	[ˈho.moː]	[L1 H]

As pointed out above, there is some discussion about the details of Classical Latin stress. The different analyses agree on the minimum size of a trochaic foot (two moras), but not on its maximum size. According to some (Mester 1994, Prince & Smolensky 1993, Hayes 1995), weight-sensitive feet are strictly bimoraic, while according to others (Hayes 1981, Jacobs 2000), trochees in Latin can be uneven, i.e. consist of up to three moras. The following section describes the analyses in general, as well as the translation of the problem into OT terms.

4.2 Linguistic analyses of Latin stress

The examples in (22) are often analysed as resulting from a combination of *extrametricality* and right-aligned feet. If we ignore for now the possibility of secondary stress, a recipe by Hayes (1995) will add foot structure to the forms in (22), ending up with the full surface structures /L (H1) L/, /L (L1 L) H/, /L L L (H1) L L/, /H (H1) H/, /H (L1 L) H/, /L (H1) L H/, and /(L1) H/. The recipe goes as follows: first make the last syllable *extrametrical*, i.e. mark it for not being able to be incorporated into a foot, then create a foot as far to the right as possible; this foot has to be *bimoraic* (i.e. it either consists of two light syllables or one heavy syllable), and if the foot is disyllabic (i.e. LL) stress falls on the first syllable.

Hayes’ (1995) bimoraic analysis is not uncontroversial. A different approach is to propose that the foot always ends just before the extrametrical syllable (Hayes 1981). The result differs from the bimoraic analysis in two forms in (22): it leads to /L L L (H1 L) L/ and /L (H1 L) H/. The foot (H1 L) has three moras. Hayes (1995) calls it an *uneven trochee*. Such an analysis satisfies the generalized principle of foot binarity (§3): feet consist either of two moras or of two syllables.

The choice between the bimoraic analysis and the uneven trochee analysis cannot be made on the basis of the overt stress patterns. Some linguists have voted for a strict bimoraic approach in Latin, on the basis of non-stress evidence like optional vowel shortening (Mester 1994; Hayes 1987, 1995; Kager 1993; McCarthy & Prince 1986, 1990). According to these authors, an unfooted non-final syllable, like *ti* in /do.(ˈmes).ti.kus/, is better than an uneven trochee, as in /do.(ˈmes.ti).kus/ (Hayes 1995: 91). We will not model these shortening phenomena in the simulated language data fed to the child, and leave it to the simulated child to construct either a bimoraic or an uneven trochee analysis (though see §7.1 for a brief discussion of iambic shortening).

Disyllabic words cause several complications. The underlying sequence |LL| is pronounced as the overt form [L1 L]. But what about its hidden surface structure? Is it footed as /(L1) L/, violating bimoraicity and foot binarity, or as /(L1 L)/, violating extrametricality? Likewise, |LH| is pronounced [L1 H]. Is it footed as /(L1) H/, violating foot binarity and bimoraicity, or as /(L1 H)/, violating bimoraicity and extrametricality? These sound like questions about the ranking of constraints, so it is natural to express all these conflicting principles in constraints. The ones often seen in the literature are those in (21) and (23) (a specific constraint for foot bimoraicity will be introduced later).

(23) *A constraint for extrametricality*
NONFINAL: "the last syllable is not contained in a foot."

There have been several proposals for the foot structure of underlying |LL| words in Latin. Prince (1980) and McCarthy & Prince (1986) argue that the structure must be /(L1 L)/. The argument runs as follows. Latin has a so-called *minimal word* requirement: there are no monosyllabic words in Latin that consist of only a light syllable. This observation can be explained by a combination of two requirements: every word must contain at least one foot, and Latin satisfies a ban on degenerate feet, i.e., the (L1) foot is prohibited completely from Latin surface structure. Apart from ruling out monomoraic words, these requirements also demand that words consisting of two light syllables must incorporate the final syllable into the foot: /(L1 L)/. Expressed in a constraint ranking, this would mean that FOOTBIN would have to outrank NONFINAL (Prince & Smolensky 1993). Such a ranking also predicts that |LH| is footed as /(L1 H)/, because a form with a degenerate foot /(L1) H/ would violate high-ranking FOOTBIN. Prince & Smolensky (1993: 63) abandon feet with the form (H1 L) because this form is "marked or even absent in trochaic systems" (they refer to Hayes 1987, Prince 1990, and Mester 1994); they formulate this as the constraint *(HL) or RHYTHMICHARMONY.

The foot in Latin thus ideally consists of two moras. The analyses also agree that feet containing four moras, like (H1 H), are forbidden in Latin. Jacobs (2000) accepts the uneven trochee (H1 L), but abandons (L1 H) feet.

In the following section we list the constraints most commonly used in the OT literature and the different ways in which they have been interpreted.

4.3 Latin Stress in OT

For our simulations we use the same underlying forms, candidate generator, and set of constraints as Tesar & Smolensky (2000) did in their simulations of 124 types of languages with metrical stress. To accommodate the Latin analyses by Jacobs (2000) and Prince & Smolensky (1993), we also investigate some constraint sets that are slight modifications of the Tesar & Smolensky set. In total, we consider six different constraint sets, but stay with the same underlying forms and generator. In the following we discuss these ingredients in detail.

Underlying forms. With Tesar & Smolensky (2000), we consider underlying forms that consist of two to seven syllables. For the forms with two to five syllables, all possible sequences of heavy and light syllables are taken into account. Thus, the underlying disyllables are |L L|, |L H|, |H L|, and |H H|. Likewise, there are eight trisyllabic underlying forms: |L L L|, |L L H|, |L H L|, |L H H|, |H L L|, |H L H|, |H H L|, and |H H H|. In the same vein, there are 16 forms with four syllables, and 32 with five. For the forms with six or seven syllables, we ignore those with heavy syllables (for computational reasons that can be deduced from column 6 in table (41)), thus leaving only |L L L L L L| and |L L L L L L L|. In total, therefore, there are 62 different underlying forms, the same ones that Tesar & Smolensky used. Unlike Tesar & Smolensky, who taught the learners all 62 possible overt forms in their simulations, we teach the learners only the 28 forms (i.e. 28 underlying-surface pairs or 28 overt forms) that have maximally four syllables; if a learner then arrives at a grammar appropriate for these 28 forms, we can have a look at how she generalizes this grammar to the 34 forms that consist of five syllables or more.

The generator. In production, each of the 62 underlying forms comes with a tableau. For each of the 62 tableaus, we construct a candidate set in the same way as Tesar & Smolensky (2000) did. The candidate set consists of all surface forms that meet the following criteria: the sequence of syllables is identical to the sequence of syllables in the underlying form with respect to number, weight, and order; every foot contains exactly one primary-stressed or secondary-stressed syllable; every word contains exactly one foot that contains a primary-stressed syllable; every primary-stressed or secondary-stressed syllable is contained in a foot; no foot contains more than two syllables. For each of the four disyllabic underlying forms, there are six candidates for the surface form. For instance, an underlying |H L| has the following candidates: /(H1) L/, /(H1 L)/, /(H1) (L2)/, /H (L1)/, /(H L1)/, /(H2) (L1)/. Each of the 8 underlying trisyllables has 24 candidates, e.g., |H L L| has: /(H1) L L/, /(H1 L) L/, /(H1) L (L2)/, /(H1) (L L2)/, /(H1 L) (L2)/, /(H1) (L2) L/, /(H1) (L2 L)/, /(H1) (L2) (L2)/, /H (L1) L/, /H (L1 L)/, /(H L1) L/, /H (L1) (L2)/, /(H L1) (L2)/, /(H2) (L1) L/, /(H2) (L1 L)/, /(H2) (L1) (L2)/, /H L (L1)/, /H (L L1)/, /H (L2) (L1)/, /(H L2) (L1)/, /(H2) L (L1)/, /(H2) (L L1)/, /(H2 L) (L1)/, /(H2) (L2) (L1)/. The 16 forms with four syllables have 88 candidates each; the 32 forms with five syllables have 300 candidates each. The single form with six syllables has 984 candidates, the form with seven syllables has 3136.[6]

Constraints. The basic constraint set we use is the one adopted by Tesar & Smolensky (2000). This constraint set takes into account the restrictions enforced by the generator: since the generator does not generate candidates with trisyllabic feet, we need no constraints against trisyllabic feet;[7] and since the generator does not generate candidates without main stress, we need no constraints to enforce that every word should contain stress (such as the constraint LEX~PR by Prince & Smolensky 1993). Tesar & Smolensky's 12 constraints are listed in (24).

[6] The monosyllabic underlying form |H| has only a single output candidate: /(H1)/. This makes it impossible for the learner to learn anything from such a form. This is the reason why we do not consider monosyllabic forms in any tableaus or simulations.
[7] This means that the only reason for including FOOTBIN is that it militates against monomoraic feet.

(24) *The Tesar & Smolensky constraint set*

ALL-FEET-LEFT (AFL): "align each foot with the word, left edge."
ALL-FEET-RIGHT (AFR): "align each foot with the word, right edge."
MAIN-LEFT: "align the head-foot with the word, left edge."
MAIN-RIGHT: "align the head-foot with the word, right edge."
WORD-FOOT-LEFT (WFL): "align the word with some foot, left edge."
WORD-FOOT-RIGHT (WFR): "align the word with some foot, right edge."
NONFINAL: "do not foot the final syllable of the word."
PARSE: "each syllable must be footed."
FOOTNONFINAL: "each head syllable must not be final in its foot."
IAMBIC: "align each foot with its head syllable, right edge."
WEIGHT-TO-STRESS-PRINCIPLE (WSP): "each heavy syllable must be stressed."
FOOTBIN: "each foot must be either bimoraic or bisyllabic."

We now discuss the precise meaning of each constraint.

The alignment constraints AFL and AFR make sure that a foot is aligned with one of the edges of a word. They are identical to the constraints ALIGNFT-L and ALIGNFT-R that we used in the simple example of §1 and §2. Their violation is *gradient*: AFL is assigned one violation mark for every syllable between the left edge of the word and the left edge of every foot. In the candidate /L (L2 L) (L1 L)/, for example, AFL is violated four times: once for the first foot, three times for the second foot. Both AFL and AFR tend to minimize the number of feet in a word, since non-existing feet cannot cause violations of these constraints.

The constraints MAIN-L and MAIN-R do the same as AFL and AFR, but only for the foot that contains the main stress. Thus, the candidate /L (L2 L) (L1 L)/ violates MAIN-L three times, and MAIN-R not at all.

The two WORDFOOT alignment constraints favour candidates where at least one foot is aligned with the word edge. These constraints are not gradient, but *binary*: they are assigned a single violation mark if there is an unfooted syllable at the edge of the word. Thus, the candidate /L L (L1) (L2 L)/ violates WFL (once), but not WFR.

The constraint NONFINAL expresses extrametricality: it is violated if the last syllable is *parsed* (included) in a foot. This constraint thus prefers /(L1) L/ to /(L1 L)/. Note that WFR and NONFINAL have *complementary violations* on the word level: a word that violates WFR does not violate NONFINAL, and a word that violates NONFINAL does not violate WFR.

The constraint PARSE favours candidates in which all syllables are parsed into feet. It is assigned one violation mark for each unfooted syllable. Thus, the candidate /L (L1 L) L L/ violates PARSE three times. This constraint works so as to maximize the number of feet (or the size of a foot): if it outranks both AFL and AFR, the language tends to have secondary stress, even in light syllables. If AFL or AFR outranks PARSE, words tend to contain a single foot with the main stress (remember that candidates without any foot at all are not generated).

The constraint FOOTNONFINAL favours candidates with trochaic feet like (L1 L), (L2 L), (L1 H), and so on. However, *degenerate feet* consisting of only one syllable, like (L1) and (H2), violate this constraint. The constraint IAMBIC favours candidates with iambic feet like (L L1), and this constraint is *not* violated in degenerate feet like (L1). This asymmetry in the formulation of the two foot form constraints (other than in the TROCHAIC-IAMBIC pair of §1 and §2) leads to complementary violations on the foot level: FOOTNONFINAL is assigned one violation mark for each occurrence of (L1) and (L L1), and IAMBIC for each occurrence of (L1 L).

The WEIGHT-TO-STRESS-PRINCIPLE favours candidates that have stress on a heavy syllable. Every heavy syllable that is not stressed in a form causes a violation of this constraint. Thus, /(L2 H) H (H1) L/ violates WSP twice (once for the unfooted H, once for the H in the first foot's weak position), whereas /(L H2) (H2) (H1) L/ does not violate WSP. Like PARSE, this constraint tends to maximize the number of feet (though less strongly), and it prefers sequences of heavy monosyllabic feet like (H1)(H2) to superheavy feet like (H1 H). This constraint can also override the foot form constraints: even in a basically trochaic language (TROCHAIC >> IAMBIC), a high-ranked WSP can force the occurrence of the iambic foot (L H1).

FOOTBIN is the constraint for foot size: feet should be binary either regarding syllables or moras. In the candidate set under discussion here, this constraint is only assigned a violation mark for each monosyllabic light foot, i.e. (L1) and (L2), whereas feet like (L1 H) and (H H2) do not violate this constraint (remember that candidates with more than two syllables are not generated).

Apart from the set of 12 constraints that Tesar & Smolensky used (the *T&S set*), we investigated five slightly different constraint sets. These sets involve the three constraints listed in (25).

(25) *Additional or alternative constraints for the Jacobs set and the Prince & Smolensky set*

TROCHAIC: "align each foot with its head syllable, left edge."

HEADNONFINAL: "the head foot is not aligned with the right edge of the word, and the head syllable (i.e. the stressed syllable in the head foot) is not the last syllable in the word."

FOOTBIMORAIC: "each foot must be bimoraic."

In order to replicate the idea behind the analysis of Jacobs (2000), we define the *Jacobs constraint set*. This set of 12 constraints (Jacobs himself used only six) is similar to the T&S set, but FOOTNONFINAL is replaced with the constraint TROCHAIC that we met with earlier in §1 and §2. The formulation now mirrors that of IAMBIC. The difference with the T&S set is that degenerate feet like (L2) violate FOOTNONFINAL, but not TROCHAIC. In the Jacobs constraint set,[8] the constraints TROCHAIC and IAMBIC conspire to minimize the number of syllables in a foot, since monosyllabic feet violate neither.

Our third constraint set is based on the analysis by Prince & Smolensky (1993). Like the Jacobs set, this *P&S constraint set* contains TROCHAIC rather than FOOTNONFINAL. Moreover, the constraint NONFINAL is replaced with HEADNONFINAL, which demands that neither the head syllable of a foot nor the head foot itself are in a final position. Both of these conditions can assign a violation mark: HEADNONFINAL is violated once in /(H2) (L1 L)/, twice in /(H2) (L L1)/, and not at all in /(H1) (L L2)/. Prince & Smolensky called this constraint NONFINAL, but we changed the name in order to make the meaning behind every constraint name unambiguous.

Since FOOTBIN does not distinguish between disyllabic and bimoraic feet, we felt it might be worthwhile to investigate the workings of an explicit bimoraic analysis. To this end, we considered three more constraint sets, all consisting of 13 constraints. These sets were constructed by adding to the T&S, Jacobs, and P&S sets (all of which contain 12 constraints) a straightforward constraint FOOTBIMORAIC. Other than FOOTBIN, this constraint is assigned a violation for every monomoraic foot such as (L1) or (L2), every trimoraic foot such as (H1 L) or (L2 H), and every quadrimoraic foot such as (H H1).

[8] Jacobs did not actually include IAMBIC in his set. This will turn out to be crucial in §4.4. He also did not include most alignment constraints.

4.4 Assessment of Jacobs' OT analysis of Latin stress

Jacobs (2000) prefers the constraint NONFINAL to Prince & Smolensky's HEADNONFINAL because the formulation of NONFINAL is much simpler and because HEADNONFINAL seems to predict unattested ('quarternary') stress patterns. We can translate Jacobs' analysis to the set of constraints and candidates discussed in §4.3, by first noting that one of the constraints employed by Jacobs (LEX~PR: "every word must contain a foot with the main stress", from Prince & Smolensky 1993) is now part of the candidate generator, so that this constraint can be left out of consideration. For right alignment, Jacobs uses a constraint that we will call LAST-FOOT-RIGHT "align the last foot with the word, right edge" (LFR). This constraint is assigned one violation mark for every syllable that follows the last foot; it can thus be seen as a gradient version of WFR, and for words with a single foot it has the same number of violations as AFR and MAIN-R. Jacobs' article happens to contain a ranking that handles all the forms that he considers.[9] This ranking is TROCHAIC >> NONFINAL >> FOOTBIN >> LFR >> WSP >> PARSE. Since Jacobs does not supply most of the relevant tableaus, we will show here how this ranking correctly predicts the forms /(L1) L/ (*fala* /(ˈfa).la/ 'siege tower'), /(L1) H/ (*fames* /(ˈfa).meːs/ 'hunger'), /(H1) L/ (*fama* /(ˈfaː).ma/ 'rumor'), /(H1) H/ (*fagus* /(ˈfaː).gus/ 'beech'), /(L1 L) L/ (*fabula* /(ˈfa.bu).la/ 'little bean'), /(L1 L) H/ (*fragilis* /(ˈfra.gi).lis/ 'fragile'), /L (H1) L/ (/a.(ˈmiː).ke/), /L (H1) H/ (*facultas* /fa.(ˈkul).taːs/ 'opportunity'), /(H1 L) L/ (*fabula* /(ˈfaː.bu).la/ 'story') and /(H1 L) H/ (*flammulae* /(ˈflam.mu).lai/ 'little flames'), several of which contain the uneven trochee (H1 L).

Tableaus (26) to (29) for the disyllabic forms illustrate the ranking of NONFINAL above the four constraints FOOTBIN, LFR, WSP, and PARSE. If NONFINAL were ranked below any of these constraints, at least one of these tableaus would have had a different winner. These four tableaus show no evidence for the ranking of TROCHAIC, nor for the relative rankings of FOOTBIN, LFR, WSP, and PARSE with respect to each other.

(26) *Extrametricality beats last-foot-right and PARSE*

|H L|	TROCHAIC	NONFINAL	FOOTBIN	LFR	WSP	PARSE
☞ /(H1) L/				*		*
/(H1 L)/		*!				
/H (L1)/		*!	*		*	*
/(H L1)/	*!	*			*	

[9] For Classical Latin, Jacobs actually proposes the ranking { FOOTBIN, TROCHAIC } >> NONFINAL >> LFR >> WSP >> PARSE (p. 345). However, this ranking must be incorrect, because it would give final stress in /L (H1)/ in our tableau (28), unless PARSE outranks WSP. This latter ranking confusingly occurs in the tableau on Jacobs' page 342, so that /(L1 H)/ becomes the winning candidate; but the form /L (H1) L/ requires WSP >> PARSE in order to beat /(L1 H) L/, as Jacobs notes himself and the reader can see in tableau (31). At the very end of his article, in a seemingly unrelated discussion on handling some facts about shortening processes in non-classical varieties of Latin, Jacobs introduces the ranking TROCHAIC >> NONFINAL >> FOOTBIN >> LFR >> WSP >> PARSE in order to account for the failure of syncope to apply to disyllabic forms, apparently without noting that this must be the only correct ranking for Classical Latin as well.

(27) *Extrametricality also beats foot binarity*

|L L|	TROCHAIC	NONFINAL	FOOTBIN	LFR	WSP	PARSE
☞ /(L1) L/			*	*		*
/(L1 L)/		*!				
/L (L1)/		*!	*			*
/(L L1)/	*!	*				

(28) *Extrametricality also beats the weight-to-stress principle*

|L H|	TROCHAIC	NONFINAL	FOOTBIN	LFR	WSP	PARSE
☞ /(L1) H/			*	*	*	*
/(L1 H)/		*!			*	
/L (H1)/		*!				*
/(L H1)/	*!	*				

(29) *Extrametricality beats LFR, WSP, and PARSE*

|H H|	TROCHAIC	NONFINAL	FOOTBIN	LFR	WSP	PARSE
☞ /(H1) H/				*	*	*
/(H1 H)/		*!			*	
/H (H1)/		*!			*	*
/(H H1)/	*!	*			*	
/(H1) (H2)/		*!				

The tableaus for the trisyllabic forms show more detailed evidence for rankings. Tableau (30) shows evidence for the existence of TROCHAIC, which prefers /(L1 L) L/ to /(L L1) L/. But there is no evidence for the *ranking* of TROCHAIC; it could just as well be ranked at the bottom, as far as the pair |L L L| – /(L1 L) L/ is concerned. This freedom of ranking of TROCHAIC is caused, of course, by the absence of the counteracting constraint IAMBIC from the tableau. If we had included IAMBIC, the choice of /(L1 L) L/ instead of /(L L1) L/ would have been direct evidence for the ranking TROCHAIC >> IAMBIC.

(30) *Evidence for trochaicity*

|L L L|	TROCHAIC	NONFINAL	FOOTBIN	LFR	WSP	PARSE
/(L1) L L/			*!	**		**
☞ /(L1 L) L/				*		*
/L (L1) L/			*!	*		**
/L (L1 L)/		*!				*
/(L L1) L/	*!			*		*

We do not show tableaus for trisyllabic forms that end in a heavy syllable, because the high ranking of NONFINAL ensures that L-final and H-final words are always handled in the same way. The next form to consider, then, is |L H L|. Tableau (31) shows direct evidence that PARSE is dominated by WSP as well as by TROCHAIC. If the ranking of WSP and PARSE had been reversed (with high-ranked TROCHAIC), the candidate /(L1 H) L/ would have won. It is apparently worse to have an unstressed H than to have an unfooted L. If TROCHAIC had been ranked below WSP and PARSE, the iambic candidate /(L H1) L/ would have won.

(31) *Weight-to-stress and trochaicity beat PARSE*

|L H L|	TROCHAIC	NONFINAL	FOOTBIN	LFR	WSP	PARSE
/(L1) H L/			*!	**	*	**
/(L1 H) L/				*	*!	*
☞ /L (H1) L/				*		**
/L (H1 L)/		*!				*
/(L H1) L/	*!			*		*

The next underlying form to consider is |H L L|. Tableau (32) shows that the winner yields evidence for the existence of the constraints LFR or PARSE. Without these constraints, the candidate /(H1) L L/ would have been equally harmonic as /(H1 L) L/. Jacobs' constraint set thus favours the uneven trochee analysis /(H1 L) L/ over the bimoraic analysis /(H1) L L/, irrespectively of the ranking of the constraints, since the violations of /(H1) L L/ form a superset of those of the violations of /(H1 L) L/. In order to turn the bimoraic analysis /(H1) L L/ into a winner, we would need the help of an extra constraint that is ranked above LFR and PARSE, perhaps FOOTBIMORAIC.

(32) *Evidence for LFR or PARSE*

|H L L|	TROCHAIC	NONFINAL	FOOTBIN	LFR	WSP	PARSE
/(H1) L L/				**!		**
☞ /(H1 L) L/				*		*
/H (L1) L/			*!	*	*	**
/H (L1 L)/		*!			*	*
/(H L1) L/	*!			*	*	*

We have now discussed all the forms that Jacobs considers. Conspicuously absent from his article, though, is the underlying form |H H L|. Tableau (33) shows that this form is problematic.

(33) *A stress clash or a superheavy foot?*

|H H L|	TROCHAIC	NONFINAL	FOOTBIN	LFR	WSP	PARSE
/(H1) H L/				**!	*	**
/(H1 H) L/				*	*!	*
/H (H1) L/				*	*!	**
/H (H1 L)/		*!			*	*
/(H H1) L/	*!			*	*	*
☞ /(H1) (H2) L/				*		*
☞ /(H2) (H1) L/				*		*

In tableau (33), two forms with two feet are optimal. To make the form with penultimate main stress win, we would have to include the constraints for the placement of main stress, and rank them in the order MAIN-RIGHT >> MAIN-LEFT, perhaps at the bottom of the hierarchy. The results for the |H H L| forms generalize to quadrisyllabic and longer forms. Because of the ranking WSP >> PARSE, these forms will tend to have secondary-stressed feet around every heavy syllable, and because of the presence of PARSE, light syllables will tend to be footed as well if this does not create iambs. Examples are /(L2 L) (H1) L/ (*manifesta* [ˌma.ni.ˈfes.ta] 'caught in the act'), /(H2 L) (H1) H/ (*militaris* [ˌmiː.li.ˈtaː.ris] 'military'), /L (H2) (L1 L) L/ (*amicitia* [a.ˌmiː.ˈki.ti.a] 'friendship', /(H2) (H2) (H2) (H1) H/ (*definitivus* [ˌdeː.ˌfiː.ˌniː.ˈtiː.vus] 'definitive'), and /(H2 L) L (H1) L/ (*deliciosa* [ˌdeː.li.ki.ˈoː.sa] 'spoiled'). The choice between /(L2 L) L (H1 L) L/ ([ˌmi.se.ri.ˈkor.di.a]) and /L (L2 L) (H1 L) L/ ([mi.ˌse.ri.ˈkor.di.a]) would probably have to be made by constraints such as ALL-FEET-LEFT and ALL-FEET-RIGHT.

But it is a question whether we should allow secondary-stressed forms at all, especially those with *stress clashes* (consecutive stressed syllables) like those in (33). Jacobs obviously considered it outside the scope of his paper to discuss the complexities of secondary stress. So let us assume that the correct form in (33) should be /H (H1) L/, with a single foot. We can get rid of the two bipedal candidates in (33) by replacing LFR with AFR. This would not change anything in tableaus (26) to (32), but the two last candidates in (33) would get three violations of AFR. If we make sure that AFR outranks WSP, the last two candidates in (33) perish. However, the winner will now be the form /(H1 H) L/, with a superheavy foot. This form is observationally incorrect, with its antepenultimate main stress (we say [au.ˈdiː.re] 'to hear', not [ˈau.diː.re]); the correct form is /H (H1) L/. However, we can see from (33) that /H (H1) L/ has superset violations when compared with /(H H1) L/. In order to make /H (H1) L/ more harmonic than /(H H1) L/, then, we would have to use an extra constraint. An obvious choice is IAMBIC, and it should be ranked above PARSE, as (34) shows.

(34) *Jacobs' hierarchy patched up*

|H H L|	TROCHAIC	NONFINAL	FOOTBIN	AFR	WSP	IAMBIC	PARSE
/(H1) H L/				**!	*		**
/(H1 H) L/				*	*	*!	*
☞ /H (H1) L/				*	*		**
/H (H1 L)/		*!			*		*
/(H H1) L/	*!			*	*		*
/(H1) (H2) L/				**!*			*
/(H2) (H1) L/				**!*			*

Tableau (34) gives us a ranking that will work for all Latin forms. It correctly generalizes to words of more than three syllables, causing all of them to end in /... (H1) X/ or /... (X1 L) X/. The remaining question is where IAMBIC has to be inserted into the hierarchy. According to (34), it has to outrank PARSE. Obviously, it has to be ranked below TROCHAIC, otherwise /(L L1) L/ would be better than /(L1 L) L/; that would be observationally incorrect, since *iacere* 'to throw' is pronounced [ˈja.ke.re], not [ja.ˈke.re]. Given the current set of seven constraints, and the low ranking of PARSE, IAMBIC has to be ranked below FOOTBIN, because /(L1 L) L/ has to be better than /L (L1) L/. Finally, AFR has to outrank both WSP and IAMBIC in order to make /H (L1 L) X/ (e.g. *nobilitas* [noː.ˈbi.li.taːs] 'fame') better than /(H1) L L X/. The complete set of crucial rankings is shown in (35). The rankings not marked by lines in this figure are not fixed. Thus, TROCHAIC could be ranked anywhere between the very top and a position below WSP, as long as it outranks IAMBIC; FOOTBIN could be ranked above AFR or below WSP, as long as it is ranked below NONFINAL and above IAMBIC; and so on.

(35) *The crucial rankings for the uneven trochee analysis without secondary stress*

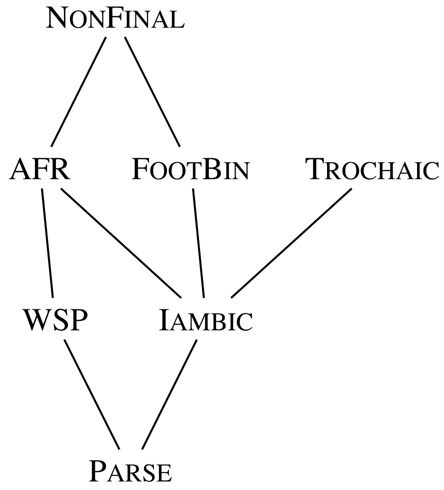

Note the *conspiracy* of the constraints TROCHAIC and IAMBIC. Together they have a preference for monosyllabic feet, since such feet violate neither of these constraints. In our ranking, this monosyllabic bias is just enough to outrule /(H1 H) L/, because TROCHAIC and IAMBIC are both ranked above PARSE. The bias is not enough to outrule /(L1 L) L/ and /H (L1 L) L/, because IAMBIC is still ranked below FOOTBIN and AFR. This combination of requirements on the ranking of IAMBIC brings about a relatively *deep* grammar: the tree in (35) shows that we need four levels of constraints to describe the Latin stress rule with its relatively simple formulation of “stress the penultimate if it’s heavy, else the antepenultimate”.

4.5 Assessment of the Tesar & Smolensky constraint set for Latin stress

The effects of the bias of TROCHAIC and IAMBIC for short feet is not found for Tesar & Smolensky’s combination of FOOTNONFINAL and IAMBIC. This is because these two constraints have complementary violations on the foot level: monosyllabic and iambic feet, i.e. (X1) and (X X1), violate FOOTNONFINAL, while trochaic feet, i.e. (X1 X), violate IAMBIC. Thus, the sum of the number of violations of FOOTNONFINAL and IAMBIC is equal to the number of feet in the word. This means that these two constraints together still have the side effect of minimizing the number of feet in a word, but they are not capable of forcing a specific foot form in the way TROCHAIC and IAMBIC could. The reduced power of FOOTNONFINAL as compared to TROCHAIC turns out to make it impossible for Tesar & Smolensky’s set of 12 constraints to handle the facts of Classical Latin stress in the uneven trochee analysis without secondary stress. This is proved in §6.1.

4.6 Assessment of Prince & Smolensky's OT analysis of Latin stress

In the analysis of Prince & Smolensky (1993), the following constraint ranking was proposed: { FootBin, Trochaic } >> HeadNonfinal >> { WSP, AFR } >> Parse. For a form like [ˈa.mo:] this leads to the incorporation of both syllables into a foot, as shown in (36). It is crucial in this analysis that the candidate /L (H1)/ violates HeadNonfinal twice since both the head foot and the head syllable are word-final.

(36) *Exhaustive parsing into a 'wretched' trochee (Prince & Smolensky's term)*

|L H|	FootBin	Trochaic	HeadNonfinal	WSP	AFR	Parse
/(L1) H/	*!			*	*	*
☞ /(L1 H)/			*	*		
/L (H1)/			**!			*
/(L H1)/		*!	**			

Since HeadNonfinal has nothing against secondary stress in final syllables, such analyses must pop up, as shown in tableau (37), unless prevented by constraints against stress clash.

(37) *Stress clash in a disyllabic word*

|H H|	FootBin	Trochaic	HeadNonfinal	WSP	AFR	Parse
/(H1) H/				*!	*	*
/(H1 H)/			*!	*		
/H (H1)/			*!*	*		*
/(H H1)/		*!	**	*		
☞ /(H1) (H2)/					*	

Analogously, this analysis predicts forms like /(L1 L) (H2)/.

Even though Prince & Smolensky's account was meant to be the translation of Mester's (1994) bimoraic approach into OT, the resulting foot structures do include trimoraic feet, unless ruled out by straightforward constraints like *(HL). This is shown in tableau (38).

(38) *An uneven trochee again*

|H L L|	FootBin	Trochaic	HeadNonfinal	WSP	AFR	Parse
/(H1) L L/					**!	**
☞ /(H1 L) L/					*	*
/H (L1) L/	*!			*	*	**
/H (L1 L)/			*!	*		*
/(H L1) L/		*!		*	*	*
/(H2) (L1 L)/			*!		**	
/(H1) (L2 L)/					**!	

But like Jacobs' analysis, Prince & Smolensky's analysis predicts a wrong result for underlying |H H L|, as shown in tableau (39).

(39) *Main stress on the antepenultimate despite a heavy penult*

|H H L|	FOOTBIN	TROCHAIC	HEADNONFINAL	WSP	AFR	PARSE
/(H1) H L/				*	**!	**
(☞) /(H1 H) L/				*	*	*
/H (H1) L/				*	*	**
/H (H1 L)/			*!	*		*
/(H H1) L/		*!		*	*	*
/(H1) (H2) L/					**!*	*
/(H2) (H1) L/					**!*	*
(☞) /(H1) (H2 L)/					**	
/(H2) (H1 L)/			*!		**	

We can see that if WSP >> AFR, candidate /(H1 H) L/ wins, as in Jacobs' analysis. If AFR >> WSP (or if WSP and AFR are crucially tied so that their violation marks add up and the buck is passed to PARSE), candidate /(H1) (H2 L)/ wins. As with Jacobs' analysis, we can save Prince & Smolensky's analysis by assuming AFR >> WSP and by inserting IAMBIC into the hierarchy, as (40) shows.

(40) *Prince & Smolensky's hierarchy patched up*

|H H L|	FOOTBIN	TROCHAIC	HEADNONFINAL	AFR	WSP	IAMBIC	PARSE
/(H1) H L/				**!	*		**
/(H1 H) L/				*	*	*!	*
☞ /H (H1) L/				*	*		**
/H (H1 L)/			*!		*		*
/(H H1) L/		*!		*	*		*
/(H1) (H2) L/				**!*			*
/(H2) (H1) L/				**!*			*
/(H1) (H2 L)/				**!		*	
/(H2) (H1 L)/			*!	**!		*	

Although the hierarchy in (40) positions main stress correctly in all forms, it leads to several secondary stresses after the main stress, and no secondary stresses before it. This starkly diverges from the common standpoint that if Latin has secondary stresses at all, these occur before the main stress rather than after it. A ranking of AFL above WSP is likely to be able to get rid of these final monosyllabic feet. We will see whether the simulated learners will be able to come up with such rankings.

5 The simulations

The six constraint sets introduced in §4.3 were tested by computer simulations on three training sets of underlying-surface pairs and two training sets of overt forms, for 10 virtual EDCD learners and 10 virtual GLA learners. This adds up to a total of 6x5x20 = 600 simulated acquisition processes. The following sections describe the constraint sets, the training sets, the details of the algorithms, and the details of the acquisition processes.

5.1 Constraint sets

The six constraint sets were described in §4.3. Table (41) summarizes them. In order that the reader can perform a simple though perhaps tedious check on the correctness of our evaluator, we include in the last column a count of the total number of constraint violations in the 15344 candidates in the 62 tableaus.

(41) *Statistics on the six constraint sets*

Constraint set	Constraints	Trochaicity constraint	Extrametricality constraint	Bimoraicity constraint	Violations
T&S	12	FOOTNONFINAL	NONFINAL	(none)	370404
Jacobs	12	TROCHAIC	NONFINAL	(none)	340028
P&S	12	TROCHAIC	HEADNONFINAL	(none)	335932
T&S + FTBIMOR	13	FOOTNONFINAL	NONFINAL	FTBIMOR	398062
Jacobs + FTBIMOR	13	TROCHAIC	NONFINAL	FTBIMOR	367686
P&S + FTBIMOR	13	TROCHAIC	HEADNONFINAL	FTBIMOR	363590

5.2 Training sets

As mentioned before, every training set contains 28 different forms: no words with five or more syllables are fed to the listener during acquisition. Each of the first three training sets consists of 28 pairs of given underlying forms together with the fully specified surface forms. The complete list is in table (42); where two or more analyses predict the same form, we saved some ink. The ‘uneven trochee’ set is meant to replicate Jacobs’ uneven trochee analysis. The ‘at most bimoraic’ and ‘at least bimoraic’ sets are meant to give an analysis that is even more bimoraic than Prince & Smolensky’s (which still includes HL feet if no special constraints are added). At this point we decided not to include any explicit analyses with secondary stress.

The three analyses in table (42) all share the same overt forms, which can be seen in the first column of table (43). For our simulations with overt forms, we use these.

An important question when dealing with stress systems is whether the language employs secondary stress. For Latin, this question is not trivial (§7.1), and *if* Latin has secondary stress, there are many different possibilities to place it (it could be quantity-sensitive or quantity-insensitive, stress clash could be permitted or not, and so on), so many different sets of overt forms with secondary stress are thinkable. We decided to include one secondary-stressed overt data set, which is shown in the second column of (43). This set has weight-sensitive secondary stress before the main stress: every H is footed,[10] as is every remaining LL; the ambiguity that this will lead to in cases like |L L L H L| will have to be solved by the learner on the basis of her own generalization from the shorter forms in the training set to forms longer than four syllables.

We thus have five training sets, although one could think of several more, both for the underlying-surface pairs and for the overt forms. All thinkable training sets, however, must be identical with respect to where the main stress falls: on the penultimate syllable if this is heavy, and on the antepenultimate otherwise.

[10] Unlike Jacobs (2003), who assumes *CLASH to prevent two heavy syllables next to each other from being both stressed, we put stress on every heavy syllable to the left of the main stressed syllable.

(42) *Three training sets with fully structured surface forms*

<table>
<tr><th rowspan="2">underlying forms</th><th colspan="3">surface forms</th></tr>
<tr><th>uneven trochee</th><th>at most bimoraic</th><th>at least bimoraic</th></tr>
<tr><td>|L L|</td><td colspan="2">/(L1) L/</td><td>/(L1 L)/</td></tr>
<tr><td>|L H|</td><td colspan="2">/(L1) H/</td><td>/(L1 H)/</td></tr>
<tr><td>|H L|</td><td colspan="3">/(H1) L/</td></tr>
<tr><td>|H H|</td><td colspan="3">/(H1) H/</td></tr>
<tr><td>|L L L|</td><td colspan="3">/(L1 L) L/</td></tr>
<tr><td>|L L H|</td><td colspan="3">/(L1 L) H/</td></tr>
<tr><td>|L H L|</td><td colspan="3">/L (H1) L/</td></tr>
<tr><td>|L H H|</td><td colspan="3">/L (H1) H/</td></tr>
<tr><td>|H L L|</td><td>/(H1 L) L/</td><td colspan="2">/(H1) L L/</td></tr>
<tr><td>|H L H|</td><td>/(H1 L) H/</td><td colspan="2">/(H1) L H/</td></tr>
<tr><td>|H H L|</td><td colspan="3">/H (H1) L/</td></tr>
<tr><td>|H H H|</td><td colspan="3">/H (H1) H/</td></tr>
<tr><td>|L L L L|</td><td colspan="3">/L (L1 L) L/</td></tr>
<tr><td>|L L L H|</td><td colspan="3">/L (L1 L) H/</td></tr>
<tr><td>|L L H L|</td><td colspan="3">/L L (H1) L/</td></tr>
<tr><td>|L L H H|</td><td colspan="3">/L L (H1) H/</td></tr>
<tr><td>|L H L L|</td><td>/L (H1 L) L/</td><td colspan="2">/L (H1) L L/</td></tr>
<tr><td>|L H L H|</td><td>/L (H1 L) H/</td><td colspan="2">/L (H1) L H/</td></tr>
<tr><td>|L H H L|</td><td colspan="3">/L H (H1) L/</td></tr>
<tr><td>|L H H H|</td><td colspan="3">/L H (H1) H/</td></tr>
<tr><td>|H L L L|</td><td colspan="3">/H (L1 L) L/</td></tr>
<tr><td>|H L L H|</td><td colspan="3">/H (L1 L) H/</td></tr>
<tr><td>|H L H L|</td><td colspan="3">/H L (H1) L/</td></tr>
<tr><td>|H L H H|</td><td colspan="3">/H L (H1) L/</td></tr>
<tr><td>|H H L L|</td><td>/H (H1 L) L/</td><td colspan="2">/H (H1) L L/</td></tr>
<tr><td>|H H L H|</td><td>/H (H1 L) H/</td><td colspan="2">/H (H1) L H/</td></tr>
<tr><td>|H H H L|</td><td colspan="3">/H H (H1) L/</td></tr>
<tr><td>|H H H H|</td><td colspan="3">/H H (H1) H/</td></tr>
</table>

5.3 The learning algorithms

For the two learning algorithms we used the implementation in the Praat program (www.praat.org). The evaluation model for EDCD was OT with crucial ties, i.e. the violations of constraints that are ranked equally high are added to each other as if these constraints formed a single constraint; in Praat, this can be simulated by setting the evaluation noise to zero. As in Tesar & Smolensky (2000), we allowed the algorithm to chew five times on every piece of language data, with backtracking if the quintuple chews did not succeed in making the (alleged) correct adult form optimal in the learner's grammar. A slight difference with Tesar & Smolensky's evaluation model was that when two forms were equally harmonic, we chose a winner randomly

(43) *Two training sets with overt forms*

overt forms	
main stress only	secondary stress
[L1 L]	
[L1 H]	
[H1 L]	
[H1 H]	
[L1 L L]	
[L1 L H]	
[L H1 L]	
[L H1 H]	
[H1 L L]	
[H1 L H]	
[H H1 L]	[H2 H1 L]
[H H1 H]	[H2 H1 H]
[L L1 L L]	
[L L1 L H]	
[L L H1 L]	[L2 L H1 L]
[L L H1 H]	[L2 L H1 H]
[L H1 L L]	
[L H1 L H]	
[L H H1 L]	[L H2 H1 L]
[L H H1 H]	[L H2 H1 H]
[H L1 L L]	[H2 L1 L L]
[H L1 L H]	[H2 L1 L H]
[H L H1 L]	[H2 L H1 L]
[H L H1 H]	[H2 L H1 H]
[H H1 L L]	[H2 H1 L L]
[H H1 L H]	[H2 H1 L H]
[H H H1 L]	[H2 H2 H1 L]
[H H H1 H]	[H2 H2 H1 H]

from among them, whereas Tesar & Smolensky somewhat less realistically chose the form that occurred first in the tableau (Bruce Tesar, p.c.).

The evaluation model for the GLA was Stochastic OT with an evaluation noise of 2.0. This noise leads to slightly different rankings of the constraints at each evaluation. Within an evaluation of an overt form, however, the ranking stayed constant: the same ranking values drawn from the Gaussian distributions were used first to interpret the overt form into a surface form and an underlying form, then to produce the learner's surface form from the interpreted underlying form.

5.4 The acquisition processes

For each of the 600 virtual learners, the initial state was that all 12 or 13 constraints were ranked at a height of 100. After this, language data were drawn randomly with equal probability from the 28 underlying-surface pairs or from the 28 overt forms. All learners therefore heard the forms in different orders and with (very slightly) different frequencies. When a pair or a form caused a mismatch between the learner's own produced surface form and (her guess of) the correct adult form, the EDCD learner had an adjustment model that would demote the ranking of one constraint by a distance of 1 (e.g. to 99 when a constraint is demoted for the first time), and the GLA learner had an adjustment model that would raise the rankings of some constraints by 0.1 and lower the rankings of some others by 0.1; in the case of the GLA learner, this *plasticity* of 0.1 was further randomized by a relative plasticity standard deviation of 0.1.

An EDCD learner was allowed to listen to maximally 1000 underlying-surface pairs or 1000 overt forms. After every 100 pairs or forms, however, we checked whether the learner had already arrived at a grammar in which all 28 pairs or forms were *singly grammatical*. An underlying-surface pair is singly grammatical if the surface form is the only optimal candidate for the underlying form, i.e. if it is optimal in its tableau and no other candidate in the same tableau is equally harmonic. An overt form is singly grammatical if for all the tableaus in which it occurs (in the current case this is always a single tableau), this overt form is shared by all optimal candidates. For instance, the overt form [H H1 L L] can be considered singly grammatical if all the optimal candidates in the tableau for |H H L L| are among the set consisting of /(H H1) L L/, /H (H1) L L/, and /H (H1 L) L/ (in practice, of course, there will usually be a single optimal candidate). If all 28 pairs or forms are singly grammatical, we can be sure that the learner will not be capable of any more learning with these forms. When this occurred, we considered learning successful and stopped the simulation (i.e., we stopped feeding any more forms to the learner). An EDCD learner usually either successfully acquired the language within the first or second round of 100 pairs or forms, or she did not acquire the language even after 1000 pairs or forms; in the latter case we can be certain that the learner will never succeed, as we will exemplify in the discussions on tables (48) and (55).

GLA learners (who take much smaller steps than EDCD learners) were allowed to listen to maximally 40,000 underlying-surface pairs or 40,000 overt forms. After every 1000 pairs or forms, we checked whether the learner had arrived at a grammar in which all of the pairs or forms were singly grammatical, and if so, we stopped the simulation. When deciding whether a pair or form was singly grammatical, we set the evaluation noise to zero and computed the optimal candidates in the 28 relevant tableaus, then proceeding as above. Although the learner would still be likely in this situation to make several mistakes if the evaluation noise had the usual value of 2.0, we still decided that learning had succeeded, because we could be sure that the learner's constraints were already ranked in the correct order and that future learning would reduce the error rate but not change the crucial rankings.

6 Results

Table (44) shows the results for the 600 learners. In each cell, the result is indicated as x/y, where x is the number of EDCD learners that succeeded and y is the number of GLA learners that succeeded. When none of the 10 learners succeeded, this is indicated by "–"; when all 10 learners succeeded, this is indicated by "√".

(44) *Simulation results for 600 learners, in the form "EDCD/GLA"*

	Learning from pairs of underlying and surface forms			Learning from overt forms	
Constraint set	uneven trochees	at most bimoraic	at least bimoraic	main stress only	secondary stress
T&S	–/–	–/–	–/–	–/–	–/√
Jacobs	√/√	–/–	–/–	–/√	–/√
P&S	√/√	–/–	–/–	1/–	–/–
T&S + FOOTBIMORAIC	√/√	√/√	–/–	–/–	–/√
Jacobs + FOOTBIMORAIC	√/√	√/√	–/–	–/7	–/√
P&S + FOOTBIMORAIC	√/√	√/√	√/√	9/–	–/–

6.1 Simulations with fully given underlying-surface pairs

Table (44) shows that EDCD and GLA were equally successful in learning from pairs of underlying and surface forms: every cell in the first three columns either contains "√/√" (all 10 EDCD learners and all 10 GLA learners succeeded) or "–/–" (all 20 learners failed). This does not surprise us. EDCD is a generally applicable OT learning algorithm that when supplied with fully specified underlying-surface pairs, is guaranteed to find a ranking that can generate those forms, if there *is* such a ranking. Thus, from the first "–" in every cell with "–/–" we can derive that there is no ranking at all that can generate the 28 underlying-surface pairs at hand with the constraint set at hand. This necessarily means that the GLA will not be able to find an appropriate ranking either (as confirmed by the second "–" in all these cells). From the first "√" in every cell with "√/√" we can derive that there *is* a ranking, and the second "√" in these cells tells us that the GLA has also been able to find it. Since there are no cells with "√/–" in the first three columns, we can conclude that in all the cases in which EDCD works, the GLA works as well. This confirms our suspicion that the GLA, like EDCD, is a generally applicable OT learning algorithm that works in all cases in which a correct ranking exists, if fully specified underlying-surface pairs are given.

(45) *Idealized results for Jacobs EDCD learning of the uneven-trochee analysis*

Constraints	ranking
NONFINAL	100.000
TROCHAIC	100.000
AFR	99.000
FOOTBIN	99.000
MAIN-R	99.000
WFR	99.000
IAMBIC	98.000
WSP	98.000
AFL	97.000
MAIN-L	97.000
PARSE	97.000
WFL	97.000

As expected from the ranking we found in §4.4, the uneven trochee analysis was learnable with the Jacobs constraint set. Two of the EDCD learners arrived at the ranking in (45). This ranking is exactly what can be predicted from the crucial rankings in (35). EDCD is an algorithm that is claimed to rank every constraint maximally high. In (35), NONFINAL and TROCHAIC are undominated, so their ranking stays at the original 100 in our simulations. In (35), the constraint AFR and FOOTBIN are outranked only be undominated constraints, so in our simulations they end up at a height of 99. In (35), each of the constraints WSP and IAMBIC is dominated by a constraint from the second level, so they end up at 98. The deepest constraint in (35) is PARSE; it must end up at 97. The remaining five constraints end up as high as they can without altering any of the optimal candidates: MAIN-R can end up at 99 because it is assigned the same number of violations in all winning candidates as AFR. WFR has to go below NONFINAL, with which it has complementary violations.

But the constraints are not always ranked maximally high. One learner ends up with a ranking similar to (45), but with WFR ranked at 98; and another learner ends up with both WFR and MAIN-R ranked at 98. While this makes no difference in the output forms, the maximally-high-ranking claim of EDCD is violated here, probably because of the existence of solutions with *crucial ties*, for which we will now see some more dramatic examples.

Four EDCD learners ended up with what is probably the minimum number of strata: they collapsed the AFR – FOOTBIN – MAIN-R – WFR stratum with the IAMBIC – WSP stratum (at 99), and had the four bottom constraints (AFL – MAIN-L – PARSE – WFL) end up at 98. At first sight this violates the crucial rankings established in §4.4. But in tableau (34) we can see that an equal ranking of AFR and WSP is possible if we allow crucial ties: the three violations of AFR in /(H2) (H1) L/ outnumber the single violations of WSP and AFR in /H (H1) L/! What's more, even FOOTBIN can be ranked equally high as IAMBIC, at least if we allow crucial ties, as tableau (46) shows.

(46) *A crucial tie between FOOTBIN and IAMBIC at work*

|L L L|	NONFINAL	TROCHAIC	AFR	FOOTBIN	IAMBIC	WSP	PARSE
/(L1) L L/			**(!)	*(!)			**
☞ /(L1 L) L/			*		*		*
/L (L1) L/			*	*			**!
/L (L1 L)/	*!				*		*
/(L L1) L/		*!	*				*

We think that the concept of the crucial tie, although inherited from the early days of OT, is not worth pursuing. After all, how can we say that two violations of the doubly gradient constraint AFR (which counts feet as well as distance) are worse than a single violation of the singly gradient constraint WSP (which counts syllables)? Under a more realistic interpretation of tied constraints, namely that by Anttila (1997), an equal ranking of AFR and WSP in (34) would mean that both /H (H1) L/ (/au.(ˈdiː).re/) and /(H2) (H1) L/ (/(ˌau).(ˈdiː).re/) would win in 50% of the cases, and an equal ranking of IAMBIC and FOOTBIN in (46) would mean that both /(L1 L) L/ (/(ˈja.ke).re/) and /L (L1) L/ (the overtly incorrect /ja.(ˈke).re/) would win in 50% of the cases. This optionality could be introduced in our simulations by taking a tiny evaluation noise, say 0.000001, for the EDCD simulations performed with Praat. All 10 EDCD learners would end up in the ranking in (45).

It remains to be said what the remaining two EDCD learners did. Like the four crucial tie reliers just discussed, they had three strata, but one of them had WSP in the bottom stratum (at 98), and the other had MAIN-R, WFR, and WSP in that stratum. Apparently, both of these had managed to learn the language, but again by relying on the crucial tie principle.

The ranking in (45) and those discussed after (45) produce all 28 forms in the first column of (42). The rankings also correctly generalize to the 34 longer forms that the learner has never heard: they predict for instance /H L (H1 L) H/ (*indigentia* /in.di.(ˈgen.ti).a/ 'want') and /L L L L (L1 L) L/ (a form that we have not heard ourselves).

The next step is to see how the GLA learners have performed. They cannot be bothered by crucial ties, because with a non-zero evaluation noise the probability that two constraints are ranked equally high at evaluation time is zero. If two constraints are ranked at nearly the same height, the distribution of outputs of the grammar will be very similar to the Anttila interpretation of a pair of tied constraints. All GLA learners end up with the ranking in table (47), although the precise ranking values differ among the learners, and half of the learners have a reversed ranking for the two bottom-ranked constraints WFL and PARSE.

(47) *A typical result for Jacobs GLA learning of the uneven-trochee analysis*

Constraints	ranking
NONFINAL	110.027
TROCHAIC	105.725
AFR	105.057
FOOTBIN	104.664
IAMBIC	100.539
WSP	99.984
MAIN-R	99.826
AFL	97.967
MAIN-L	94.105
WFR	89.973
WFL	89.702
PARSE	88.618

We see that the constraints divide up in strata. All the crucial rankings of figure (35) can be found as large ranking distances in table (47). WSP and IAMBIC have stayed where they began, around 100. The three constraints that crucially outrank these two in (35) have been pushed up to about 105. The single constraint that crucially outranks two of the constraints around 105 has been pushed up to a height of 110. The constraint crucially dominated by WSP and IAMBIC (i.e. PARSE) has fallen a double distance, to the region near 90. This deep falling of weak constraints is typical of what the GLA does in general; in this case it is not a result of a domination by MAIN-L or so.

So far, all EDCD and GLA learners have performed equally well, although the EDCD learners have practiced fancy behaviour by ingeniously inventing analyses with crucial ties whereas the GLA learners have boringly but reliably mimicked the expected ranking of (35).

For the T&S constraint set, table (44) shows that there exists no ranking that produces the forms associated with the uneven trochee analysis. According to table (41), this can only be due to a difference between TROCHAIC and FOOTNONFINAL.

Indeed we saw in §4.5 that the combination of FOOTNONFINAL and IAMBIC is not capable of performing the conspiracy that led the combination of TROCHAIC and IAMBIC to force a winner with a monosyllabic foot in tableau (34). If we replace TROCHAIC with FOOTNONFINAL in tableau (34) or graph (35), we see that /(H1 H) L/ becomes the winner, because /H (H1) L/ now violates FOOTNONFINAL, which is higher ranked than IAMBIC, which remains the only constraint in (34) and (35) that favours /H (H1) L/ over /(H1 H) L/. But it is still instructive to see how EDCD and GLA learners perform with the T&S set. Table (48) shows where one EDCD learner was after our simulations had to stop, i.e. after 1000 language data.

(48) *A result for T&S EDCD learning of the uneven-trochee analysis*

Constraints	ranking
NONFINAL	100.000
AFR	99.000
FOOTBIN	99.000
MAIN-R	99.000
WFR	99.000
WSP	99.000
WFL	–102.000
IAMBIC	–109.000
AFL	–110.000
FOOTNONFINAL	–110.000
MAIN-L	–110.000
PARSE	–110.000

This learner has not ended up in a stable grammar. If we go on feeding her language data, the six constraints at the bottom will continue tumbling down the hierarchy. All ten EDCD learners have these six constraints ranked in different orders, but all in the vicinity of –110, which will be around –320 after 2000 language data. At this snapshot in time, the learner of table (48) has iambic forms like /(L L1) L/ (/(ja.ˈke).re/). When being told that the form should be /(L1 L) L/ (/(ˈja.ke).re/), she will demote IAMBIC to –111. Unfortunately, this will in turn lead her to generate a trochaic /(H1 H) L/ (/(ˈau.diː).re/). When being told that this should have been /H (H1) L/ (/(au.ˈdiː).re/), she will demote WFL, AFL, FOOTNONFINAL, MAIN-L, and PARSE to –112, because all of these constraints prefer /(H1 H) L/ to /H (H1) L/ (and are higher ranked than IAMBIC, the highest constraint that prefers /H (H1) L/). This will go on forever. To measure how well these learners behave as speakers of Latin, we computed their *error rates* in the following way. We randomly drew 1000 underlying-surface pairs, chosen with equal probability from the 28 underlying-surface pairs that we had used for training (therefore, each pair was chosen approximately 36 times on average), and computed the learner's surface form for the given underlying form. We then compared each learner's form with the given adult surface form, and considered the learner correct if the surface forms were identical. If a learner had e.g. 600 forms correct, her error rate was 40%. Eight of the EDCD learners turned out to have error rates of approximately 65%, the remaining two had error rates of about 44%.

The GLA learners also fail with the T&S set, but in a different way from the EDCD learners. The GLA learners all end up in a stable grammar. Table (49) shows the result for one learner.

(49) *GLA learner with T&S constraints, fed with the uneven-trochee analysis*

Constraints	ranking
NONFINAL	156.752
AFR	150.041
WSP	144.622
IAMBIC	139.944
FOOTNONFINAL	139.618
FOOTBIN	95.926
MAIN-R	70.001
WFR	43.248
WFL	–401.795
AFL	–1078.115
MAIN-L	–1095.503
PARSE	–1204.395

At the top of the hierarchy, we see a division into strata as before, with a distance of about 6 between constraints that the algorithm thinks are crucially ranked with respect to each other. IAMBIC and FOOTNONFINAL are ranked very close together. Half of the 10 GLA learners have the same ranking as in (49), half have IAMBIC and FOOTNONFINAL reversed. The distance between these two constraints is always small, so that if the learner has evaluation noise during her productions, she will have the ranking IAMBIC >> FOOTNONFINAL approximately half of the time, and FOOTNONFINAL >> IAMBIC the other half of the time. The error rates computed with an evaluation noise of 1.0 (smaller than the noise during training) are between 48% and 58%; the typical errors are that the learners show variation between /(L1 L) L/ and /(L L1) L/ and between /(H1 H) L/ and /(H H1) L/.

Table (44) shows that none of the three constraint sets without FOOTBIMORAIC is capable of learning a truly bimoraic analysis, like the 'at most bimoraic' and 'at least bimoraic' analyses of table (42). This is not so surprising. We have already seen in §4.6 that without constraints that favour strictly bimoraic feet, like *(HL) or FOOT-BIMORAIC, one cannot expect the grammars to be able to learn bimoraic data. Still, the simulations with the P&S constraint set were successful in learning the uneven trochee analysis. Table (50) shows the resulting grammar for an EDCD learner.

(50) *The generic result for P&S EDCD learning of the uneven-trochee analysis*

Constraints	ranking
HEADNONFINAL	100.000
TROCHAIC	100.000
AFR	99.000
FOOTBIN	99.000
MAIN-R	99.000
IAMBIC	98.000
AFL	97.000
MAIN-L	97.000
WFL	97.000
PARSE	96.000
WFR	96.000
WSP	96.000

This complicated grammar (5 strata deep) was reached by six of the ten EDCD learners. This grammar works for both versions of non-stochastic OT: that with crucial ties and that with variationist (Anttila) ties. The remaining EDCD learners had a grammar with a depth of 4, in which IAMBIC was one stratum higher, at 99 (now that the third stratum had been vacated, the ranking of the six constraints dominated by IAMBIC was of course 1 higher as well); this grammar relied on a crucial tie between IAMBIC and FOOTBIN, as above in the case of the T&S constraint set. Table (51) shows the result for a GLA learner.

(51) *A result for P&S GLA learning of the uneven-trochee analysis*

Constraints	ranking
HEADNONFINAL	110.744
TROCHAIC	106.141
FOOTBIN	105.867
AFR	105.058
IAMBIC	103.180
MAIN-R	102.464
AFL	100.627
WSP	97.805
MAIN-L	88.622
WFL	85.479
PARSE	78.739
WFR	78.037

In (51), no clear layering has yet been established. This could be due to the fact that learning was designed to stop when the error rate was 0% if the evaluation noise was set to zero. With an evaluation of 2.0, i.e. the same as during learning, the error rate for the learner in (51) is still 30%. This means that in the grammar state of (51), the learner will detect mismatches for 30% of the incoming data and therefore take another learning step in 30% of the cases. These learning steps will continue to increase the separation between the constraints in (51). In order to see whether the crucial rankings had been established in (51), we computed the error rate for an evaluation noise of 1.0. It was 7%; this number tells us something about how the learner will behave after making twice as many learning steps as she has made before reaching the state in (51).

We can note that in all cases the resulting ranking for the P&S constraint set is rather different from Prince & Smolensky's (1993) proposal, which we discussed in §4.6. As predicted, we find that a high ranking of AFR outrules secondary stress before the main stress, and that AFL >> WSP outrules secondary stress *after* the main stress.

From table (44) we see that including the FOOTBIMORAIC constraint improves learnability from underlying-surface pairs. It is not surprising that if the Jacobs set and the P&S set succeeded in learning the uneven trochee analysis, this analysis is still learnable with these constraint sets if we add a constraint. But the addition of FOOTBIMORAIC seems to be just enough for the T&S set to achieve successful acquisition. The cause of this is that FOOTBIMORAIC is capable of ruling out HH feet but not LL feet: FOOTNONFINAL can now outrank IAMBIC in order to produce /(L1 L) L/ rather than /(L L1) L/, without fear of producing /(H1 H) L/, beause this form is ruled out by FOOTBIMORAIC. Otherwise, uneven trochees remain, as in /(H1 L) L/.

In general, the uneven trochee analysis seems to require fewer constraints (twelve) than the bimoraic analyses. The Jacobs and P&S constraint sets seem to be more successful than the T&S constraint set. We can compare the complexities of the resulting grammar between these two sets: the Jacobs set leads to a grammar with four crucial strata, the P&S set to a grammar with five crucial strata. It is likely that grammars with fewer strata are easier to learn. But the differences between the constraint sets are small, especially regarding the success of the P&S+FOOTBIMORAIC set with the 'at least bimoraic' analysis. In the next section, we will see whether there are any differences between the constraint sets when learning from overt forms only. The current section has at least shown that there were *some* combinations of constraint sets and analyses that were capable of learning the Latin stress system, so that we can now turn with confidence to the more realistic simulations, those for learning from overt forms, where hidden structures like foot boundaries are not explicitly provided to the learner but where she will have to construct them by herself.

6.2 Simulations with primary-stressed-only overt forms

Table (44) shows that the T&S constraint set is not capable of learning a ranking for primary-stress-only overt data. This is not surprising, since the three primary-stress-only analyses (i.e. sets of given underlying-surface pairs) are not learnable with the T&S set either. Of course the learners could have invented a fourth analysis, perhaps one that includes /(L1 L)/ and /(L1) H/ or so, but they did not, so it is possible that there exists no analysis at all for primary-stress-only Latin with the T&S set.

The simulations with the Jacobs constraint set are more interesting: EDCD fails with this constraint set, the GLA succeeds. The first question now is: what analysis did the GLA learners come up with? The answer is that all learners came up with the same analysis, namely the uneven trochee analysis, i.e. for each of the 28 underlying forms in (42) they would produce the corresponding surface form in the 'uneven trochee' column (we computed these surface forms by running the 28 underlying forms through the learner's final grammar with an evaluation noise of zero). These learners ended up with the ranking in (52), sometimes with a different permutation of the very closely ranked constraints FOOTBIN, WSP, and TROCHAIC, or of IAMBIC, AFL, and MAIN-R.

(52) *A typical result for Jacobs GLA learning from overt forms: creation of the uneven-trochee analysis*

Constraints	ranking
NONFINAL	114.290
AFR	108.639
FOOTBIN	104.784
WSP	104.476
TROCHAIC	104.470
IAMBIC	101.302
AFL	100.739
MAIN-R	99.521
MAIN-L	95.039
WFR	85.710
PARSE	82.381
WFL	82.127

We see that the stratification is very different from that in (47). Still, all of the crucial rankings in (35) are satisfied. It will come to no surprise that these learners also correctly generalize the uneven-trochee analysis to forms of more than four syllables.

The learners were rather slow in constructing the uneven trochee analysis by themselves. Whereas in the case of the underlying-surface pairs of table (47) all GLA learners had succeeded after the first 1000 data, the learners of the overt forms needed 3 to 35 rounds of 1000 data to arrive at an appropriate ranking. But they all succeeded.

Table (44) shows that for the three sets of 12 constraints, 29 out of 30 EDCD learners of primary-stressed overt forms fail. There is only one EDCD learner who happens to acquire an appropriate 12-constraint grammar; this learner uses the P&S set and invents an analysis that we have not considered, combining the two 'at least bimoraic' forms /(L1 L)/ and /(L1 H)/ of table (42) with the uneven trochees /... (H1 L) X/; but this cannot be considered a success for EDCD, since if only 10% of the children had been capable of learning Latin, this language would have perished much faster than it did.

When the constraint sets are enriched with FOOTBIMORAIC, the performance of EDCD improves. With the P&S+FOOTBIMORAIC set, nine learners managed to construct a functioning analysis. Seven of these came up with the uneven trochee analysis with the 'at least bimoraic' form /(L1 X)/ mentioned before. The rankings of these learners slightly varied, as before. Table (53) shows an example.

(53) *A typical result for P&S+FOOTBIMORAIC EDCD learning from overt forms: creation of the at-least-bimoraic uneven-trochee analysis, with empty strata*

Constraints	ranking
FOOTBIN	100.000
HEADNONFINAL	100.000
AFR	99.000
FOOTBIMORAIC	99.000
MAIN-R	99.000
AFL	98.000
MAIN-L	98.000
WFL	98.000
PARSE	97.000
WFR	97.000
TROCHAIC	94.000
WSP	94.000
IAMBIC	93.000

A conspicuous property of seven of the resulting rankings is that they contained *empty strata*. In table (53), which is an average case, strata 95 and 96 are empty. Such empty strata can never occur when EDCD learns from fully specified underlying-surface pairs, but they can when EDCD learns from overt forms only. Another conspicuous property of the seven rankings is that none of them is correct under the variational interpretation of tied constraints. To see whether there exists such a ranking at all, we would have to run a simulation in which the P&S+FOOTBIMORAIC constraint set learns an explicitly given at-least-bimoraic uneven trochee analysis. If so, and if we want to see whether EDCD can also learn it from overt forms, the simulations that led to table (53) will have to be rerun with a tiny evaluation noise.

Two of the EDCD learners constructed an at-least-bimoraic analysis. Both relied on crucial ties. Table (54) shows one of the rankings.

(54) *Another result for P&S+FOOTBIMORAIC EDCD learning from overt forms: creation of the at-least-bimoraic analysis*

Constraints	ranking
FOOTBIMORAIC	100.000
FOOTBIN	100.000
HEADNONFINAL	100.000
AFR	99.000
MAIN-R	99.000
AFL	98.000
MAIN-L	98.000
WFL	98.000
PARSE	97.000
TROCHAIC	97.000
WFR	97.000
WSP	97.000
IAMBIC	96.000

The ranking of FOOTBIMORAIC above AFR and MAIN-R causes the preference for /(H1) L L/ over /(H1 L) L/. We can compare this to ranking (53), where the crucial tie between these three constraints favours the uneven trochee /(H1 L) L/ over the bimoraic /(H1) L L/: FOOTBIMORAIC casts a single vote in favour of /(H1) L L/ whereas AFR and MAIN-R gang up with two votes in favour of /(H1 L) L/.

The tenth P&S+FOOTBIMORAIC EDCD learner did not succeed in learning Latin. Her ranking after 1000 data is given in (55).

(55) *The single failure for P&S+FOOTBIMORAIC EDCD learning from overt forms*

Constraints	ranking
FOOTBIN	100.000
HEADNONFINAL	99.000
AFR	98.000
MAIN-R	98.000
AFL	97.000
FOOTBIMORAIC	97.000
MAIN-L	97.000
WFL	97.000
PARSE	96.000
WFR	96.000
WSP	96.000
IAMBIC	–104.000
TROCHAIC	–105.000

This learner has experienced IAMBIC and TROCHAIC tumbling down the hierarchy, alternatingly making the by now usual mistakes of /(H1 H) L/ and /(L L1) L/. To see whether this learner would learn the language later, we taught her 10,000 extra overt forms. This had no other effect than demoting IAMBIC and TROCHAIC down to –2238 and –2239. We conclude that this learner, in contrast with the tenth GLA learner of the Jacobs set discussed above, has really got trapped in a sequence of grammars that she can never get out of (a 'non-globally-optimal limit cycle'). This may mean that the '9' in table (44) indicates that primary-stressed-only Latin is not learnable by the whole generation of learners if they entertain the P&S+FOOTBIMORAIC constraint set. Whether this situation means that this combination of constraint set, training set, and learning algorithm can be ruled out as a proposal for how Latin children acquired their language, or whether it is just a predictor of sound change, depends on the exact fraction of failures. Our best guess at this point is of course 10%, but this number could be estimated more accurately after a future simulation of, say, 1000 learners. Our 10 GLA learners, by the way, were consistent in not learning with the P&S+FOOTBIMORAIC constraint set at all.

The remaining interesting figure for the main-stress-only forms in table (44) is the '7' for the GLA learners with the Jacobs+FOOTBIMORAIC constraint set. Apparently, adding the FOOTBIMORAIC constraint to the set made the language *less* learnable from overt forms. The seven successful learners ended up with rankings that follow the stratification in (56), with varying rankings within the three strata that contain more than one constraint.

(56) *A typical success for Jacobs+FOOTBIMORAIC GLA learning from overt forms*

Constraints	ranking
NONFINAL	120.563
AFR	113.294
WSP	106.877
FOOTBIMORAIC	105.181
AFL	103.154
TROCHAIC	103.016
FOOTBIN	102.790
MAIN-R	102.439
IAMBIC	100.746
MAIN-L	97.615
PARSE	81.554
WFR	79.437
WFL	73.763

The remaining three learners were not lucky. Even after 50,000 data, they stuck with grammars like (57). The relative rankings within the third and fourth strata can vary, but TROCHAIC and IAMBIC are always ranked very closely. Grammar (57) is of the type that we have seen several times before: since TROCHAIC and IAMBIC are very closely ranked, these learners end up producing one of the two mistakes /(H1 H) L/ or /(L L1) L/. The cause of the problem is that these learners have moved AFL too high up, and not managed to raise FOOTBIMORAIC above it. If FOOTBIMORAIC is ranked higher than AFL, it is capable of outruling /(H1 H) L/, so that IAMBIC is freed from the task of outruling /(H1 H) L/; this allows IAMBIC to fall to a stratum below TROCHAIC, so that the learner also stops producing /(L L1) L/ errors. Apparently, adding a constraint does not necessarily improve learnability from overt forms.

(57) *A typical failure for Jacobs+FOOTBIMORAIC GLA learning from overt forms*

Constraints	ranking
NONFINAL	124.255
AFR	115.923
AFL	107.371
WSP	107.226
MAIN-R	103.226
FOOTBIN	101.215
MAIN-L	100.711
TROCHAIC	98.868
IAMBIC	98.633
FOOTBIMORAIC	95.829
PARSE	83.892
WFR	75.745
WFL	71.173

Conclusions from the primary-stressed-only overt forms. First it has to be said that learning Latin from overt data turns out to be possible. However, it also brings about some instances of the expected failures of RIP/EDCD and RIP/GLA (§2.3), since the overt forms very often contain ambiguous structures. In fact, the only combination of constraint set and algorithm that was capable of learning from main-stress-only overt forms for all 10 learners was the Jacobs set with the GLA. A combination that got close to this performance was the P&S+FOOTBIMORAIC set with EDCD, where nine out of ten learners detected a correct ranking. In order to reliably prove that the former combination is better than the latter, we would have to show that it is nearly 100% correct, for instance by teaching 1000 learners with the Jacobs – GLA combination and computing the percentage correct. That will take us two weeks of computer time.

Since both RIP/EDCD and RIP/GLA make use of the same interpreting mechanism (Robust Interpretive Parsing), any crucial differences in performance between the two have to be attributed to the different kinds of reranking strategy (demotion-only vs. demotion-and-promotion, and one-shot learning vs. graduality).

6.3 Simulations with overt forms that contain secondary stress

Table (44) shows that EDCD is not capable of learning from overt forms with secondary stress listed in the last column of table (43), with any constraint set. By contrast, the GLA is successful with the T&S and Jacobs constraint sets, regardless of whether FOOTBIMORAIC is included or not. This looks better than the performance with the primary-stress-only forms, which could mean that additional information such as secondary stress does support learning.

Apart from the striking difference between the learning algorithms, the most conspicuous result in table (44) is that the T&S constraint set is successful for the first time. The 10 GLA learners created grammars very similar to the one in table (58).

(58) *The result for T&S GLA learning from secondary-stressed overt forms*

Constraints	ranking
NONFINAL	108.705
WSP	104.865
MAIN-R	102.437
FOOTBIN	101.430
WFL	100.773
FOOTNONFINAL	99.888
AFL	99.852
PARSE	99.273
IAMBIC	97.686
MAIN-L	95.353
AFR	91.682
WFR	91.295

For the 28 overt forms in table (43), this learner constructs an analysis with rather exhaustively parsed syllables and both iambic and uneven trochaic feet, with a preference for trochees: /(L1) X/, /(H1) X/, /(L1 L) X/, /(L H1) X/, /(H1 L) X/, /(H2) (H1) X/, /L (L1 L) X/, /(L2 L) (H1) X/, /L (H1 L) X/, /(L H2) (H1) X/, /(H2) (L1 L) X/, /(H2 L) (H1) X/, /(H2) (H1 L) X/, /(H2) (H2) (H1) X/. The learner generalizes this exhaustivity to the 32 forms with five syllables: /(L2 L) (L1 L) X/, /(L2 L) (L H1) X/, /(L2 L) (H1 L) X/, /(L2 L) (H2) (H1) X/, /(L H2) (L1 L) X/, /(L H2) (L H1) X/, /(L H2) (H1 L) X/, /(L H2) (H2) (H1) X/, /(H2 L) (L1 L) X/, /(H2 L) (H1 L) X/, /(H2 L) (H2) (H1) X/, /(H2) (H2) (L1 L) X/, /(H2) (H2 L) (H1) X/, /(H2) (H2) (H1 L) X/, /(H2) (H2) (H2) (H1) X/, and, importantly, /(H2 L) (L H1) X/ (not /(H2) (L2 L) (H1) X/). Note that in these cases, the learner has created her own patterns of overt forms, e.g. [L2 L L H1 X], which were not in the training set. This means that the learner will produce reasonably good pronunciations for five-syllable forms, even if she has never heard them before; for instance, if the learner is familiar with the nominative /ra.(ˈpi.di).taːs/ 'speed', she will come up with the form /(ˌra.pi).(di.ˈtaː).te/ for the ablative singular even if she has never heard that form. For the longest forms consisting of light syllables only, the analyses have a single left-aligned foot: /(L2 L) L (L1 L) L/ and /(L2 L) L L (L1 L) L/. The exhaustivity noted above thus reduces (only in the case of light syllables) to a right-aligned main foot and a left-aligned secondary foot, which is caused by a high ranking of WFL, and a ranking of AFL above PARSE. Three other learners have exactly the same language, and three learners have a slightly different ranking that leads to the exact same forms as above except that the form with seven syllables scans as /(L2 L) (L2 L) (L1 L) L/. This even more exhaustive parsing of syllables is caused by the ranking PARSE >> AFL. Actually, the speaker in table (58), with her close ranking of AFL and PARSE, can be expected to waver between the two forms with seven syllables. To us this variation (both between speakers and within speakers) seems to be similar to what real speakers of English, German or Dutch do with longer words (it could even depend on speaking rate, i.e., you could rank PARSE a bit lower when speaking fast). The remaining three GLA learners have /L (H2 L) (H1) X/ instead of /(L H2) (L H1) X/ (with variation in the seven-syllable form) caused by a ranking of FOOTNONFINAL over WFL and PARSE; it is hard to hear the difference between those two forms, so large-scale interspeaker variation for such hidden structures within the speech community should come to no surprise.

The results with the Jacobs constraint set are quite different. Table (59) shows the final ranking of one learner.

(59) *One result for Jacobs GLA learning from secondary-stressed overt forms*

Constraints	ranking
NONFINAL	108.761
FOOTBIN	104.053
WSP	103.321
MAIN-R	102.504
TROCHAIC	102.322
PARSE	99.965
AFL	99.818
WFL	97.991
IAMBIC	97.735
MAIN-L	95.671
AFR	94.134
WFR	91.239

This learner avoids iambic forms: she has /L (H1) X/, /L (H2) (H1) X/, /(L2 L) L (H1) X/, /L (H2 L) (H1) X/, and, this time, no other choice than the exhaustive form /(H2) (L2 L) (H1) X/. Since PARSE >> AFL, the seven-syllable form is /(L2 L) (L2 L) (L1 L) L/. All nine other Jacobs GLA learners have AFL >> PARSE, and therefore the forms /(H2 L) L (H1) H/ and /(L2 L) L L (L1 L) L/.

Adding FOOTBIMORAIC to the Jacobs constraint set can give the ranking in (60).

(60) *Jacobs+FtBim GLA learning from secondary-stressed overt forms*

Constraints	ranking
NONFINAL	108.926
WSP	103.858
MAIN-R	103.641
FOOTBIN	103.238
FOOTBIMORAIC	102.983
PARSE	101.836
AFL	100.866
TROCHAIC	100.785
WFL	97.937
MAIN-L	97.039
IAMBIC	96.771
AFR	94.656
WFR	91.074

The learner in (60) has come up with an analysis that has uneven trochees for main stress (caused by MAIN-R >> FOOTBIMORAIC), but avoids uneven trochees for secondary stress (caused by the ranking FOOTBIMORAIC above PARSE and AFR): both phenomena can be seen in /(H2) L (H1 L) L/. This learner also has /(H2) (L2 L) (H1) X/ and /(L2 L) (L2 L) (L1 L) L/. Eight other GLA learners arrive in the same language as the learner in (60), except that three of them have /(H2) L L (H1) X/ and /(L2 L) L L (L1 L) L/, caused, as usual, by a reverse ranking of PARSE and AFL. Finally, the remaining learner happened to come up with a real bimoraic analysis, avoiding all uneven trochees, e.g. /(H2) L (H1) L L/. Her ranking is in table (61).

(61) *Jacobs+FTBIM GLA learning with secondary-stressed overt forms: creation of the at-most-bimoraic analysis*

Constraints	ranking
NONFINAL	108.711
WSP	103.607
FOOTBIMORAIC	102.917
MAIN-R	102.891
FOOTBIN	102.450
TROCHAIC	101.201
PARSE	100.534
AFL	100.430
WFL	97.738
IAMBIC	96.920
MAIN-L	96.906
AFR	94.262
WFR	91.289

The 10 GLA learners with the T&S constraint set and FOOTBIMORAIC behaved similarly: eight created the bimoraic analysis with the exhaustive forms /(H2) (L2 L) (H1) X/ and /(L2 L) (L2 L) (L1 L) L/, one a bimoraic analysis with medially unfooted light syllables, i.e. /(H2) L L (H1) X/ and /(L2 L) L L (L1 L) L/, and one allowed uneven trochees in main feet only.

Conclusions from the secondary-stressed overt forms. Again, learning Latin from overt data turns out to be possible, at least with the RIP/GLA algorithm. Whether this means that RIP/EDCD should be ruled out as a candidate for describing Latin with secondary stress remains to be seen, since different secondary stress patterns than the one investigated here are thinkable.

The learners came up with ten different analyses for the overt data with secondary-stressed forms, with a total of 109 different surface forms for the 62 underlying forms. For the forms with at most four syllables, the overt forms associated with these ten analyses were of course identical. Differences between the analyses showed up only in a couple of overt forms with five syllables (namely [H2 L L H1 X] versus [H2 L2 L H1 X]) and in a form with seven syllables (namely [L2 L L L L1 L L] versus [L2 L L2 L L1 L L]), and this difference occurred with all five pairs of analyses that we have seen (i.e. ‘iambic&trochaic’, ‘less iambic’, ‘trochaic’, ‘bimoraic in secondary feet’, ‘bimoraic everywhere’), only depending on the relative ranking of AFL and PARSE, which were always closely ranked. Attested Latin forms with more than four syllables, if weight-sensitively secondary-stressed as here, would therefore give us no information about whether Latin learners used the T&S set or the Jacobs set, with or without FOOTBIMORAIC, and which of the five analysis types they created. Whether other patterns of secondary stress would give us such information remains to be investigated. See §7.1 for an interesting next step.

7 Discussion

In this rather long paper on the learnability of Latin stress, we investigated the performance of two learning algorithms, six constraint sets, three analyses, and two kinds of overt forms. In §7.1 we report some results on all these issues and indicate how several more constraint sets, analyses, and kinds of overt forms should be investigated in the future. In §7.2 we finally express our remaining general uneasiness about the Optimality-Theoretic approach to metrical phonology.

7.1 What has been achieved and what has not

Correct analyses. Since previous OT accounts of Latin stress in the literature turned out to be capable of accounting for only part of the forms in which we were interested, our investigation had to start with discovering a couple of analyses that were capable of handling the positioning of Latin main stress correctly. One of those analyses (§4.4) derived from Jacobs' (2000) original ranking, which was first corrected (footnote 9), then augmented with IAMBIC in order to handle |HHL|. The other analysis (§4.6) derived from Prince & Smolensky's (1993) ranking, which was augmented with the ranking AFR >> WSP and with the constraint IAMBIC, again in order to handle |HHL|.

Learning algorithms. EDCD and the GLA performed strikingly differently in the simulations with overt forms. As summarized in table (44), the GLA succeeded with five combinations of constraint sets and kinds of overt forms. All of these five successes were among the eight combinations with the four constraint sets currently considered plausible (the Tesar & Smolensky set and the Jacobs set, with or without FOOTBIMORAIC). None of the 80 EDCD learners managed to acquire Latin with any of these four sets. EDCD performed a bit better on the two constraint sets that contain the currently deprecated and perhaps implausible constraint HEADNONFINAL, although each of these four groups of 10 EDCD learners contained at least one learner who did not acquire Latin (see the rows 'P&S' and 'P&S+FOOTBIMORAIC' in table (44), in combination with the columns 'main stress only' and 'secondary stress').

Constraint sets. The simulations seem to reveal that some of the proposed constraint sets are more adequate than others. For instance, TROCHAIC seems to be a better formalization for a trochaic foot pattern than FOOTNONFINAL, which seems to be too restrictive. Also, NONFINAL seems to be a more effective formalization of extra-metricality than HEADNONFINAL. But no constraint set can be ruled out completely yet. As usual in OT, the legitimacy of a constraint set ultimately has to be proven in combination with systems of other languages than the specific language under study.

Analyses. The uneven trochee analysis was better learnable than either of the two bimoraic analyses.

However, from the simulations with overt forms a fourth analysis transpired that we had not considered before: an analysis with uneven trochees, as in Jacobs (2000), but with at-least-bimoraic feet, so that the light-initial disyllables become /(L1 L)/ and /(L1 H)/.[11] This fourth analysis may well lie at the basis of the process of *iambic shortening* in Pre-Classical Latin (underlying |L H|, e.g. the concatenation of the verb

[11] These two forms actually occurred in Jacobs' original analysis for Classical Latin, but as shown in §4.4, these forms require (with the Jacobs constraint set) a ranking of FOOTBIN >> NONFINAL and of PARSE >> WSP, the latter of which fails to handle |LHL| correctly.

stem |am-| 'love' with the first singular ending |-oː|, becomes /(L1 L)/, e.g. /(ˈa.mo)/ 'I love'), which many authors discuss (Prince & Smolensky 1993, Mester 1994, Jacobs 2000). Future research will have to take this analysis into account.

Overt forms. Learning from forms with a certain type of secondary stress turned out to be easier than learning from forms with primary stress only. There is disagreement in the literature about whether Latin had secondary stresses, and therefore feet, before the primary-stressed foot, and if it had, where these secondary stresses were: they could have been weight-sensitive (Allen 1973) or not (Jacobs 1989). It is unlikely that every H syllable was stressed in Latin as it was in our series of forms, the last column in table (43), given that languages tend to avoid stress clash. This series already led to ten different analyses, and other secondary stress patterns will lead to many more. A possible solution to this problem is to *let the learner decide:* give her only overt forms with primary-stressed syllables, but allow her to invent a full foot structure with secondary-stressed syllables. This could be implemented by saying that the overt form that corresponds to the surface form /(H2 L) L (H1) L/ is not [H2 L L H1 L], but just [H L L H1 L]. A learner who is allowed to fill in secondary stresses, i.e. to turn [H L L H1 L] into /(H2 L) L (H1) L/, would perhaps also in turn ignore secondary stresses if she hears them.

Frequency. In our simulations, we fed the learner every type of underlying form equally often. But it is likely that heavy syllables were more frequent in Latin than light syllables. Since the typical mistakes of our learners were superheavy trochees in /(H1 H) X/ and iambs in /(L L1) X/, and the former type must have been more frequent, the GLA would cause the learner to end up with rankings that are less likely to make mistakes for |H H X| than for |L L X|. In most cases, the mistakes just mentioned were caused by a close ranking of TROCHAIC (or FOOTNONFINAL) and IAMBIC without a compensatory ranking somewhere else in the hierarchy. We could have fed the learner with more |H H X| than |L L X| forms, probably ending up with a ranking of TROCHAIC slightly above IAMBIC. It is not unlikely that such a ranking would have helped the learner to avoid non-global optima, but we will leave this for later investigation.

Sound change. We have seen cases in which a small percentage of the learners did not succeed in acquiring Latin, while the great majority of learners did succeed. Such cases can be predictors of acquisition-induced sound change. It is possible, for instance, that not all constraints are innate, but that they are instead constructed by the learner. In that case, some learners may well entertain constraint sets that we have shown to lead to unlearnability. The typical mistakes were trochaicity in /(H1 H) X/ and iambicity in /(L L1) X/. The disadvantage of taking a dead language to study acquisition can thus turn into an advantage, since we know a lot about what happened in the daughter languages. With some luck, later investigations may also be able to model the historical change from initial stress in Pre-Classical Latin to weight-sensitive right-aligned stress in Classical Latin.

More realistic models of metrical acquisition. We have simplifyingly been assuming that the learner's productions contained the same number of syllables as their underlying forms and the adult forms. However, it is likely that Latin childen were similar to Dutch children (Fikkert 1994) and English children (Gnanadesikan 1995) in that they started out by truncating longer words, e.g. by turning trisyllabic words into disyllabic words consisting of a single foot, and that segmental structure interfered. Such a situation would have strong implications for all of the steps in our modelling of acquisition. For instance, this could mean that learners start out by acquiring everything there is to know about short words, before they go on to consider longer words. For a non-OT metrical acquisition model that takes into account selective attention to specific structures, see Dresher (1999).

7.2 Conspiracies between constraints

As we have seen in §4.4, TROCHAIC and IAMBIC conspire to minimize foot size. This is problematic because OT was invented in order to get rid of conspiracies between rules. The alternative constraint pair is FOOTNONFINAL and IAMBIC. These two have fewer side effects since they have complementary violations on the foot level. This means that doing OT with the pair FOOTNONFINAL-IAMBIC is close to having a *parameter* "foot direction" in the grammar that is set to one of the values *nonfinal* or *iambic* (not entirely, because these constraints still conspire to minimize the number of feet).

The situation is even worse with the constraints for weight sensitivity. In a weight-sensitive language (WSP high), a bimoraic analyst would feel that feet like (H1 L) ought to be ruled out by their trimoraicity. One would like to interpret FOOTBIN as "a foot has two moras" in such languages. In a quantity-insensitive language (WSP low), (H1 L) should be fine, and one would like to interpret FOOTBIN as "a foot has two syllables". But in OT, the interpretation of FOOTBIN cannot depend on the ranking of WSP, because the syllable vs. mora interpretation of FOOTBIN is a binary decision while the ranking of WSP is not binary (it can not only be ranked at the top or at the bottom, but also somewhere in the middle). Splitting up FOOTBIN into constraints like FOOTBIMORAIC and FOOTDISYLLABIC does not really solve the problem, since one would still like FOOTDISYLLABIC to be ranked low when WSP is ranked high (a quantity-sensitive language, hence no preference for HL or HH feet), and one would like FOOTBIMORAIC to be ranked low when WSP is ranked low (a quantity-insensitive language, hence no preference for monosyllabic H feet). We think that the real problem lies in the representations. What we called H and L in this paper are really abbreviations for $(\mu\mu)_\sigma$ and $(\mu)_\sigma$ (μ for mora). What we would really want to be able to state is that in a quantity-insensitive language syllables are not represented as H or L but simply as σ, a syllable without internal weight structure. The constraint FOOTBIN would then simply prefer $(\sigma\sigma)_F$ feet. Thus, the syllable pronounced [paː] is to be represented as $(\mu\mu)_\sigma$ in a weight-sensitive language, as σ in a weight-insensitive language. But OT does not easily allow representational solutions; to implement a representational solution one would probably have to model *perception* explicitly: in a weight-insensitive language, the overt form [paː] is simply perceived (interpreted) as /σ/, and the mora is just not represented at any level of analysis; in a weight-insensitive language, [paː] is perceived as $/(\mu\mu)_\sigma/$ (or perhaps as /μμ/ if syllables play no role at all in the language at hand). The current paper actually does model perception: it is nothing else than the overt-to-surface mapping that we have been calling 'interpretation' and 'robust interpretive parsing' (Tesar & Smolensky's terms). However, the example of the overt-to-surface mapping discussed in this paper is markedness-only, since all 12 or 13 constraints evaluate the output of interpretation (i.e. the result of perception) only, while none of these constraints take into account the relation between the overt form and the surface form. The cause of this restriction is the source of the example, namely Tesar & Smolensky's insistence on the idea that interpretation uses the same constraints as production. In general, however, interpretation (perception) must use more than just structural constraints: there must also be faithfulness-like constraints that are specific to the overt-to-surface mapping, simply because this part of the mapping is language-specific. It is true that one could think that a structural constraint like *μ "the surface form does not contain moras" could help in mapping [paː] to /σ/ rather than to $/(\mu\mu)_\sigma/$. But this does not work, for two reasons: (1) if *μ were an OT constraint it would be violable, and if high ranked it would lead to a production of /σ/ in some cases, $/(\mu\,\mu)_\sigma/$ in others;[12] (2) faithfulness-like constraints are still needed, as we will see. We propose that the mora is a language-specific construct, just like any other phonological category. Briefly speaking: English does not have clicks, Pintupi does not have moras. In perception, Pintupi listeners map the syllable-like overt forms [pa], [pap], and [paː] to the mora-less surface forms $/(\text{pa})_\sigma/$, $/(\text{pap})_\sigma/$, and $/(\text{paː})_\sigma/$. Latin listeners map the same three forms to $/((\text{pa})_\mu)_\sigma/$, $/((\text{pa})_\mu(\text{p})_\mu)_\sigma/$, and $/((\text{pa})_\mu(\text{a})_\mu)_\sigma/$. We now see that this mapping has to involve overt-surface relationships: long vowels have to map to two moras, short vowels to one mora in Latin. The mapping could be described by the various kinds of categorization constraints proposed by Boersma (1998) and Escudero & Boersma (2003), and the optional development of the mora as a language-specific structure could be handled in OT in a way analogous to the category creation model of Boersma, Escudero & Hayes (2003). We will not fall short of future research possibilities.

Acknowledgments

This work was funded by the NWO Vernieuwingsimpuls grant 016.024.018. We would like to thank Louis Pols for his comments.

References

Allen, W. Sydney (1973). *Vox Latina. The pronunciation of Classical Latin.* Second edition. Cambridge University Press.

Anttila, Arto (1997). *Variation in Finnish phonology and morphology.* Doctoral thesis, Stanford University.

Boersma, Paul (1997). How we learn variation, optionality, and probability. *IFA Proceedings* 21, 43–58.

Boersma, Paul (1998). *Functional phonology: formalizing the interactions between articulatory and perceptual drives.* Doctoral thesis, University of Amsterdam.

Boersma, Paul (to appear). Review of Tesar & Smolensky (2000), *Learnability in Optimality Theory.* To appear in *Phonology.*

Boersma, Paul, Paola Escudero & Rachel Hayes (2003). Learning abstract phonological from auditory phonetic categories: an integrated model for the acquisition of language-specific sound categories. *Proceedings of the 15th International Congress of Phonetic Sciences*, Barcelona, pp. 1013–1016.

Boersma, Paul & Bruce Hayes (2001). Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32, 45–86.

Boersma, Paul & Claartje Levelt (1999). Gradual constraint-ranking learning algorithm predicts acquisition order. In Eve Clark (ed.), *Proceedings of Child Language Research Forum* 30, Stanford, California, pp. 229–237.

Boersma, Paul & Claartje Levelt (2003). Optimality Theory and phonological acquisition. *Annual Review of Language Acquisition* 3, 1–50.

Curtin, Suzanne & Kie Zuraw (2001). Explaining constraint demotion in a developing system. In Barbora Skarabela, Sarah Fish & Anna H.-J. Do (eds.), *Proceedings of the 26th annual Boston*

[12] This might actually not be so bad. It would allow us to express extrametricality in a more principled way, namely with the constraint $*\mu)_{\text{PrWd}}$ "a word does not end in a mora". The Latin overt form [a.ˈmiː.kus], for instance, could then be analysed up to the syllable level as $/((\text{a})_\mu)_\sigma\ ((\text{mi})_\mu(\text{i})_\mu)_\sigma\ (\text{kus})_\sigma/$. The current paper cannot begin to address the issue of how foot structure could be built on such representations.

University Conference on Language Development, pp. 118–129. Somerville, Mass.: Cascadilla Press.
Dresher, B. Elan (1999). Charting the learning path: cues to parameter setting. *Linguistic Inquiry* 30, 27–67.
Escudero, Paola & Paul Boersma (2003). Modelling the perceptual development of phonological contrasts with Optimality Theory and the Gradual Learning Algorithm. In Sudha Arunachalam, Elsi Kaiser & Alexander Williams (eds.), *Proceedings of the 25th Annual Penn Linguistics Colloquium* (*Penn Working Papers in Linguistics* 8.1), pp. 71–85.
Fikkert, Paula (1994). *On the acquisition of prosodic structure*. Doctoral thesis, Leiden.
Gnanadesikan, Amalia (1995). Markedness and faithfulness constraints in child phonology. Ms., University of Massachusetts, Amherst.
Gordon, Matthew (2002). A phonetically-driven account of syllable weight. *Language* 78, 51–80.
Hayes, Bruce (1981). *A metrical theory of stress rules*. Doctoral thesis (1980), MIT.
Hayes, Bruce (1985). Iambic and trochaic rhythm in stress rules. In M. Niepokuj et al. (eds.), *Berkeley Linguistics Society* 13, 429–446.
Hayes, Bruce (1987). A revised parametric metrical theory. *Northeastern Linguistic Society* 17, 274–289.
Hayes, Bruce (1995). *Metrical stress theory*. University of Chicago Press.
Hyman, Larry (1985). *A theory of phonological weight*. Dordrecht: Foris.
Jacobs, Haike (1989). *Nonlinear studies in the historical phonology of French*. Doctoral thesis, Katholieke Universiteit Nijmegen.
Jacobs, Haike (2000). The revenge of the uneven trochee: Latin main stress, metrical constituency, stress-related phenomena and OT. In Aditi Lahiri (ed.), *Analogy, levelling, markedness*, pp. 333–352. Berlin: Mouton de Gruyter.
Jacobs, Haike (2003). The emergence of quantity-sensitivity in Latin. In D. Eric Holt (ed.), *Optimality Theory and language change*, pp. 229–247. Dordrecht: Kluwer.
Kager, René (1993). The moraic iamb. In L. Dobrin, L. Nichols & R. Rodriguez (eds.), *CLS 27: Papers from the 27th Regional Meeting of the Chicago Linguistic Society. Volume One: The General Session*, pp. 291–305. Chicago: Chicago Linguistic Society.
Kager, René (1996). Stem disyllabicity in Guugu Yimidhirr. In Marina Nespor & Norval Smith (eds.), *Dam phonology. HIL Phonology Papers II*, pp. 59–101. The Hague: Holland Academic Graphics.
Liberman, Mark & Alan Prince (1977). On stress and linguistic rhythm. *Linguistic Inquiry* 8, 249–336.
McCarthy, John & Alan Prince (1986). Prosodic morphology. Unpublished ms., University of Massachusetts at Amherst and Brandeis University.
McCarthy, John & Alan Prince (1990). Foot and word in prosodic morphology: The Arabic broken plural. *Natural Language and Linguistic Theory* 8, pp. 209-283.
McCarthy, John & Alan Prince (1993). Generalized alignment. In Geert Booij & Jaap van Marle (eds.) *Yearbook of Morphology 1993*, pp. 79–153. Dordrecht: Kluwer.
Mester, R. Armin (1994). The quantitative trochee in Latin. *Natural Language and Linguistic Theory* 12, 1–61.
Prince, Alan (1980). A metrical theory of Estonian quantity. *Linguistic Inquiry* 11, 511–562.
Prince, Alan (1990). Quantitative consequences of rhythmic organization. In K. Deaton, M. Noske & M. Ziolkowski (eds.), *CLS 26-II: Papers from the Parasession on the Syllable in Phonetics and Phonology*, pp. 355–398. Chicago: Chicago Linguistic Society.
Prince, Alan & Paul Smolensky (1993). *Optimality Theory: constraint interaction in generative grammar*. Ms., Rutgers University, New Jersey.
Smolensky, Paul (1996). On the comprehension/production dilemma in child language. *Linguistic Inquiry* 27, 720–731.
Tesar, Bruce (1995). *Computational Optimality Theory*. Doctoral thesis, University of Colorado.
Tesar, Bruce (1997). An iterative strategy for learning metrical stress in Optimality Theory. In Elizabeth Hughes, Mary Hughes & Annabell Greenhill (eds.), *Proceedings of the 21st annual Boston University Conference on Language Development*, pp. 615–626. Somerville, Mass.: Cascadilla Press.
Tesar, Bruce & Paul Smolensky (1998). Learnability in Optimality Theory. *Linguistic Inquiry* 29, 229–268.
Tesar, Bruce & Paul Smolensky (2000). *Learnability in Optimality Theory*. MIT Press.
Van de Vijver, Ruben (1998). *The iambic issue: iambs as a result of constraint interaction*. Doctoral dissertation, Free University of Amsterdam. HIL Dissertations 37.