

How to Choose Successful Losers in Error-Driven Phonotactic Learning

Giorgio Magri

SFL UMR 7023 (CNRS and University of Paris 8)
59/61 rue Pouchet, 75017 Paris
France
magrigrg@gmail.com

René Kager

UiL-OTS (Utrecht University)
Trans 10, 3512 JK Utrecht
The Netherlands
RWJKager@uu.nl

Abstract

An error-driven phonotactic learner is trained on a stream of licit phonological forms. Each piece of training data counts as a *winner* in terms of Optimality Theory. In order to test its current grammar, the learner needs to compare the current winner with a properly chosen *loser*. This paper advocates a new subroutine for the choice of the loser, based on the idea of minimizing the “distance” from the given winner.

1 Error-driven phonotactic learning and the problem of the choice of the loser

Phonotactics is knowledge of the distinction between licit and illicit phonological forms (Chomsky and Halle, 1965). We adopt the model of phonotactics developed within *Optimality Theory* (Prince and Smolensky, 2004), briefly reviewed here. A *candidate set* is a collection of pairs (x, y) ; the first element x is called the *underlying* form; the second element y is called the *candidate surface* form. A *constraint* assigns to each candidate pair (x, y) a non-negative number of *violations*, which measures how the mapping of the underlying form x to the surface form y deviates from the ideal relative to the specific perspective valued by that constraint. Constraints come in two types. *Faithfulness constraints* punish a candidate pair (x, y) based on the discrepancy between the underlying form x and the surface form y . For instance, the faithfulness constraint IDENT[VOICE] is violated by a candidate pair of segments differing in voicing. *Markedness constraints* punish a candidate pair (x, y) based on the

ill-formedness of the surface form y . For instance, the markedness constraint NOVOICEOBS is violated by a candidate pair of segments whose surface segment is a voiced obstruent.

A *constraint ranking* \gg is a linear order over the constraint set. The *grammar* G_{\gg} corresponding to the ranking \gg takes an underlying form x and returns a candidate (x, y) which \gg -beats any other candidate (x, z) whose underlying form is x and whose surface form z is different from y , in the sense of condition (1). The candidate (x, y) is then called the *winner* while (x, z) is called a *loser*. Losers are stricken out as a mnemonic.

- (1) There exists a constraint which is *winner-prefering* (i.e., assigns *less* violations to the winner candidate (x, y) than to the loser candidate (x, z)) which is \gg -ranked above every constraint which is *loser-prefering* (i.e., assigns *more* violations to the winner candidate (x, y) than to the loser candidate (x, z)).

A surface form y is *phonotactically licit* according to the grammar G_{\gg} provided there exists some underlying form x which is mapped to y by that grammar.

An *error-driven* phonotactic learner is trained on a sequence of phonological forms all phonotactically licit according to a target grammar and it tries to infer that grammar as follows. It maintains a current hypothesis of the target grammar, which is initialized to a most restrictive grammar, namely one which deems illicit as many forms as possible. The current grammar is then slightly updated in the direction of a looser phonotactics, whenever it incorrectly predicts the current piece of data to be il-

licit. This learning scheme is formalized in OT as the *error-driven ranking algorithm* (EDRA) outlined in the following pseudo-code and detailed below (Tesar and Smolensky, 1998; Boersma, 1998).

-
- 1: **Initialize** the current ranking vector θ
 - 2: **repeat**
 - 3: get a licit *winner* surface form y
 - 4: choose a corresponding *underlying* form x
 - 5: choose a *loser* form z to compare to y
 - 6: **If** the winner (x, y) does not beat the loser (x, z) according to the ranking vector θ :
 - 7: update the current ranking vector θ
 - 8: **until** no more errors are made at line 6
-

The EDRA knows the underlying constraint set C_1, \dots, C_n . The current constraint ranking is represented by assigning to each constraint C_k a numerical *ranking value* θ_k , with the understanding that high ranking values correspond to high ranked constraints. The ranking values are collected into a *ranking vector* $\theta = (\theta_1 \dots, \theta_n)$. At line 1, the ranking values of the faithfulness and the markedness constraints are initialized to 0 and to a large positive constant $\theta > 0$, respectively. Thus, the markedness constraints start out above the faithfulness ones, yielding a grammar which is phonotactically maximally restrictive (Smolensky, 1996).

At line 3, the EDRA is fed a piece of training data, consisting of a surface form y licit according to the target constraint ranking. No assumptions are made on the sequence of training data (e.g., no assumptions are made on the frequency with which various licit forms are fed to the learner). At line 4, the EDRA needs to reconstruct an underlying form x corresponding to the current winner surface form y . A common choice is to set x identical to y , under the assumption that the underlying OT typology is *idempotent*, namely it maps every phonotactically licit form faithfully into itself (Magri, 2015b). The proper definition of the subroutine for the choice of the loser form z at line 5 is the topic of this paper and will thus be discussed in detail below.

At line 6, the EDRA checks whether the current ranking vector θ satisfies condition (2), where W and L are the sets of winner- and loser-preferring constraints relative to the intended winner and loser

candidates (x, y) and (x, z) .

$$(2) \quad \max_{h \in W} \theta_h > \max_{k \in L} \theta_k$$

If condition (2) holds, any ranking \gg which respects the current ranking vector θ (in the sense that $C_h \gg C_k$ whenever $\theta_h > \theta_k$) satisfies the OT condition (1), namely succeeds at making the intended winner (x, y) beat the intended loser (x, z) (Boersma, 2009). In this case, the learner has nothing to learn from the comparison between the current winner and loser forms.

Failure of condition (2) instead suggests that the current ranking values of the loser-preferring (winner-preferring) constraints are too large (too small, respectively) and thus need to be updated at line 7. We assume the re-ranking rule (3) (Tesar and Smolensky, 1998; Boersma, 1998; Magri, 2012).

- (3) a. Increase the ranking values of the w winner-preferring constraints by $\frac{1}{w+1}$;
- b. decrease the ranking values of the *undominated* loser-preferring constraints by 1.

Each winner-preferring constraint is promoted by $\frac{1}{w+1}$, where w is the total number of winner-preferring constraints. The loser-preferring constraints are demoted by 1. Only those loser-preferring constraints that really need to be demoted are indeed demoted, namely those which are not currently ranked underneath a winner-preferring constraint and are therefore called *undominated*.

The only implementation detail which has been left open in this outline of the EDRA model is the proper definition of subroutine for the choice of the loser at line 5. This is the topic of this paper.

2 Two test cases to evaluate subroutines for the choice of the current loser

The EDRA model is guaranteed to *converge* (under the assumption that the target grammar is idempotent): after a finite (small) number of iterations, it is not possible to sample from the set of target licit forms any surface form y which would force the learner to make an update in the if-loop at lines 6-8. Convergence holds irrespectively of the subroutine for the choice of the current loser used at line 5. Suppose now that this subroutine satisfies the basic condition (4). This condition says that the EDRA never wastes data (Tesar and Smolensky, 1998): if

there is an opportunity to learn something from the current winner (i.e., if there exists at least a loser which is able to trigger an update), the EDRA will not “waste” that opportunity (i.e., the chosen loser indeed triggers an update).

- (4) The subroutine for the choice of the loser returns a loser which triggers an update at line 7, whenever such a loser exists.

This condition (4) ensures that, if a surface form y is licit according to the target grammar the EDRA has been trained on, then it is also licit according to any ranking \gg which respects the final ranking vector θ^{fin} entertained by the EDRA at convergence (in the sense that $C_h \gg C_k$ whenever $\theta_h^{\text{fin}} > \theta_k^{\text{fin}}$). In other words, the EDRA succeeds at half of the learning problem: it has learned to recognize licit forms as such. The ranking learned by the EDRA could nonetheless deem licit too many forms. In other words, it could describe a phonotactics which, although *consistent* with the target one, is not sufficiently *restrictive*. Are there guarantees that the EDRA also learns to recognize illicit forms as such?

Consider a phonotactic pattern which has the following property: there exists a subset of the markedness constraints which punish exactly all and only the illicit forms. This phonotactic pattern can thus be analyzed in terms of a constraint ranking such as (5): the designated subset of markedness constraints hold sway at the top while the remaining markedness constraints are silent at the bottom. The relative ranking of the faithfulness constraints sandwiched in between is irrelevant. We therefore refer to these phonotactic patterns as *\mathcal{F} -irrelevant*.

- (5)
$$\begin{array}{c} \text{a subset of } \mathcal{M} \text{ constraints} \\ | \\ \text{all } \mathcal{F} \text{ constraints} \\ | \\ \text{the remaining } \mathcal{M} \text{ constraints} \end{array}$$

For instance, suppose that the constraint set contains a markedness constraint against voiced velar obstruents and a markedness constraint against dorsal fricatives. The velar inventory $[\text{g k } \text{ɣ } \text{x}]$, which only admits the voiceless velar stop (illicit segments are stricken out), follows by just letting those two markedness constraints hold sway at the top.

The EDRA model described in section 1 has been shown to be restrictive when the target phonotactics

is \mathcal{F} -irrelevant (Magri, 2013a; Magri, 2014a; Magri, 2015c; Magri, 2015a). In other words, it succeeds at learning the target phonotactics. This success holds irrespectively of the details of the phonological analysis (e.g., the content of the markedness and faithfulness constraints or any other properties of the target ranking, besides it being \mathcal{F} -irrelevant). It also holds irrespectively of how the current loser is chosen at line 5 of the pseudo-code—as long as condition (4) is respected. In order to tackle the issue of the proper definition of the subroutine for the choice of the loser at line 5, we thus need to look at the behavior of the EDRA model on phonotactic patterns which are *not* \mathcal{F} -irrelevant. Let’s briefly recall two two examples of such phonotactic patterns which are not \mathcal{F} -irrelevant (Magri and Kager, 2015).

Voicing is especially effortful at the velar place: due to the small oral volume behind the velar constriction, the supra-glottal pressure quickly equalizes the sub-glottal pressure, hindering vocal cords vibration (Ohala, 1983). Many attested velar inventories comply with phonetic markedness, namely have voiceless stops or fricatives without the voiced ones. Yet, UPSID (Maddieson, 1984) documents two inventories $[\text{g k } \text{ɣ } \text{x}]$ and $[\text{g k } \text{ɣ } \text{x}]$ which are phonetically *counterintuitive*, as they admit the voiced fricative at the exclusion of the voiceless one. If we could posit a markedness constraint which punishes $[\text{x}]$ at the exclusion of $[\text{ɣ}]$, these inventories could be generated by letting that markedness constraint hold sway at the top of the ranking (5). Yet, such a markedness constraint would be incompatible with the *grounding hypothesis* (Hayes and Steriade, 2004): $[\text{x}]$ is not any worse than $[\text{ɣ}]$ from any phonetic perspective. Fortunately, the desired inventory can be generated in compliance with the grounding hypothesis whenever $/\text{ɣ}/$ is harder to neutralize than $/\text{x}/$, so that the former surfaces at the exclusion of the latter. This neutralization pattern requires some faithfulness constraints (which preserve $/\text{ɣ}/$ from neutralizing) to be ranked above some other faithfulness constraints (those violated by the neutralization of $/\text{x}/$). In conclusion, the grounding hypothesis forces us to posit a crucial relative ranking among the faithfulness constraints. The learnability guarantees recalled above for \mathcal{F} -irrelevant target rankings (5) thus do not apply in these cases.

Let's look closer at the inventory [g k ʎ x], which lacks the voiced stop and the voiceless fricative. We analyze this inventory as follows: /x/ can be neutralized to [k] preserving voicing, while /ʎ/ cannot be neutralized preserving voicing, because [g] is independently ruled out by a dedicated constraint. This intuition can be cashed out as follows. We assume that only velar obstruents are candidates of the velar obstruents, as stated in (6a). The ranking (6b) then yields the target inventory [g k ʎ x].

- (6) a. $Gen(/g k ʎ x/) = [g k ʎ x]$
 b. $\begin{array}{c} \text{IDENT[VOI]} \quad \text{NOVOISTOP} \\ | \quad \diagdown \quad / \\ \text{NODORFRIC} \\ | \quad \diagup \quad | \\ \text{NOVOIFRIC} \quad \text{IDENT[CONT]} \end{array}$

The inventory [g k ʎ x] only lacks the voiceless fricative and thus differs from the inventory considered above only because the voiced velar stop [g] is now licit. We analyze this inventory as follows: /x/ can be neutralized to [h] preserving voicing, while /ʎ/ cannot be neutralized while preserving voicing, because [fi] is independently ruled out by a dedicated constraint. This intuition can be cashed out as follows. We assume place impermeability apart from the velar/glottal border: only the velar and glottal obstruents are candidates of the velar and glottal obstruents, as stated in (7a). The ranking (7b) then generates the target inventory [g k ʎ x].

- (7) a. $Gen(/g k ʎ x ? fi h/) = [g k ʎ x ? fi h]$
 b. $\begin{array}{c} M = *[fi] \\ | \\ \text{IDENT[VOICE]} \quad \text{IDENT[CONT]} \\ | \quad \diagdown \quad / \\ \text{NOVOISTOP} \quad \text{NOVOIFRIC} \\ | \quad \diagup \quad | \\ \text{NODORFRIC} \\ | \quad \diagdown \quad / \\ \text{IDENT[DOR]} \quad *[h] \end{array}$

From the perspective of the phonological analysis, the assumption (6a) of velar place impermeability and the assumption (7a) of place impermeability apart from the velar/glottal border are not restrictive: they can be reinterpreted as the assumption that the faithfulness constraints for place features are high ranked. Yet, this interpretation introduces additional relative ranking conditions among the faithfulness constraints which would need to be carefully con-

sidered in the learnability analyses. To start from the simplest case, the learnability analyses developed in the rest of the paper explicitly adopt the restrictive assumptions (6a) and (7a) on candidacy.

The only detail in the description of the EDRA model which has been left open in section 1 concerns the proper definition of subroutine for the choice of the loser at line 5. The rest of this paper tackles this issue using the test cases of the two velar inventories just discussed.

3 Motivating a new subroutine for the choice of the current loser

Suppose that the EDRA model is trained on the velar inventory [g k ʎ x] and thus needs to learn the corresponding constraint ranking (6b). At line 3, the model is effectively always fed the surface form [ʎ], since [ʎ] and [k] are the only two licit forms and the latter can be ignored, because it is unmarked relative to the constraints assumed (it violates none) and thus never prompts an update. At line 4, the model assumes the faithful underlying form /ʎ/. At line 5, the model has to choose the current loser candidate among [g], [k], or [x]. According to the *classical subroutine* for the choice of the loser, the learner chooses as the current loser the candidate predicted to win by the current ranking values—more precisely, by an arbitrary ranking consistent with the current ranking values (Tesar and Smolensky, 1998; Magri, 2013b). This classical subroutine for the choice of the loser satisfies condition (4), namely it never wastes data: if there exists at least a loser which is able to trigger an update, the chosen loser can be shown to indeed trigger an update.

Unfortunately, the classical subroutine for the choice of the loser leads to trouble when the EDRA model is trained on the inventory [g k ʎ x]. Here is why. The comparison between the winner mapping (/ʎ/, [ʎ]) and the three loser mappings (/ʎ/, [g]), (/ʎ/, [k]), and (/ʎ/, [x]) sorts the constraints into winner-/loser-preferring as represented in (8) with Elementary Ranking Conditions (Prince, 2002).

(8)

| | IDENT[VOI] ID[CON] | NODORFRIC NOVOISTOP NOVOIFRI | |
|-------------------------|-----------------------|------------------------------------|-----|
| (/ʎ/, [ʎ]) ~ (/ʎ/, [g]) | W | L | W L |
| (/ʎ/, [ʎ]) ~ (/ʎ/, [k]) | W | L | L L |
| (/ʎ/, [ʎ]) ~ (/ʎ/, [x]) | W | | L L |

The loser [g] will (almost) never be chosen, because the constraint NOVOICEDSTOP is winner-prefering in the corresponding first ERC in (8), starts ranked at the top, and is never demoted (because never loser-prefering). The choice of the current loser thus effectively boils down to [k] and [x]. The markedness constraints start out above the faithfulness constraints. That ranking configuration is preserved for a large initial portion of the run. Throughout that portion, the choice of the current loser is thus completely determined by the markedness constraints. Since [k] is unmarked, the classical subroutine for the choice of the loser always chooses [k]. Unfortunately, the corresponding second ERC in (8) promotes the two faithfulness constraints IDENT[VOICE] and IDENT[CONT] on a par. If the EDRA stubbornly always chooses [k] as the current loser, it will always promote the two faithfulness constraints on a par until they reach the top of the ranking, thus failing at learning the target ranking (6b) which instead requires IDENT[VOICE] to be ranked above IDENT[CONT].

We thus replace the classical subroutine with the *new subroutine* described below in pseudo-code. Here, we consider an arbitrary underlying form x (while the EDRA model always chooses x equal to the winner form y). For a related proposal, see (Riggle, 2004). Three remarks are in order. First,

Require: a current winner (x, y) candidate:

- 1: construct the ERC matrix corresponding to the comparisons $(x, y) \sim (x, z)$ for all possible loser candidates z for the underlying form x
 - 2: split any ERC with multiple L's into multiple ERCs with a single L;
 - 3: determine the smallest number \hat{n} such that there exists an ERC with \hat{n} winner-preferrers which is inconsistent with the current ranking values;
 - 4: pick at random among the inconsistent ERCs with \hat{n} winner-preferrers.
-

the new subroutine satisfies condition (4): if there exists a loser which is able to prompt an update, the new subroutine will return one such loser, as it only searches among losers whose corresponding ERC is inconsistent with the current ranking values. Second, the new subroutine chooses among such losers the one(s) which minimizes the “differ-

ence” between the winner and the loser, as measured in terms of the number of winner-prefering constraints which distinguish between them. Third, the new subroutine is computationally less expensive than the classical one, because it circumvents the computation of the predicted optimal candidate.

When the new subroutine for the choice of the loser is deployed on the surface form $y = [\gamma]$ corresponding to the ERC block (8), it prevents the EDRA from choosing the loser $\{g\}$, because the corresponding first ERC is already consistent with the current ranking values (through the high ranked winner-prefering markedness constraint NOVOISTOP). And it also prevents it from choosing the loser $\{k\}$, because the corresponding second ERC has “too many” w’s. The EDRA is thus biased towards choosing the loser $\{x\}$. The corresponding third ERC promotes IDENT[VOICE] but not IDENT[CONT], leading to the target ranking (6b). We thus obtain the following

Theorem 1 *When trained on an arbitrary sequence of data sampled from the velar inventory $\{g\ k\ \gamma\ x\}$, the EDRA model with the new subroutine for the choice of the loser succeeds at learning the target ranking (6b).*

4 How to analyze the new subroutine

The preceding section has motivated a new subroutine for the choice of the current loser. This section highlights some formal properties of this new subroutine which turn out useful in the analysis of EDRA’s restrictiveness. For concreteness, we focus on the velar inventory $\{g\ k\ \gamma\ x\}$, with the analysis (7) recalled in section 2. This analysis requires the voiceless glottal fricative [h] to be licit and the voiced glottal fricative [ɦ] to be instead illicit. To start from the simplest case, we minimize the number of licit forms that the model is trained on, by assuming that the glottal stop [ʔ] is illicit as well (say, because of a dedicated high ranked markedness constraint $*[\text{ʔ}]$). Thus, suppose that the EDRA model is trained on the velar/glottal inventory $\{g\ k\ \gamma\ x\ \text{ʔ}\ \text{ɦ}\ h\}$. At line 3, the model is effectively always fed one of the surface forms [g], [γ], and [h], since the only other licit form [k] is unmarked and can thus be ignored. At line 4, the model assumes the corresponding faithful underlying forms. At line 5, the model

chooses a current loser form. The loser candidates [ʔ] and [ɦ] can be ignored, because the corresponding constraints *[ʔ] and *[ɦ] are never violated and thus stay ranked at the top—just like NOVOISTOP in (8) above. The ERCs corresponding to the other losers are listed in (9).

(9)

| | ID[DOR] | IDENT[VOI] | ID[CON] | NODORFRIC | NOVOISTOP | NOVOIFRI | *[ɦ] |
|-------------------------|---------|------------|---------|-----------|-----------|----------|------|
| (/g/, [g]) ~ (/g/, [k]) | W | | | L | | | |
| (/g/, [g]) ~ (/g/, [t]) | | W | | W | L | W | |
| (/g/, [g]) ~ (/g/, [x]) | | W | W | W | L | | |
| (/g/, [g]) ~ (/g/, [h]) | W | W | W | | L | | W |
| (/y/, [y]) ~ (/y/, [t]) | | W | | L | W | L | |
| (/y/, [y]) ~ (/y/, [k]) | | W | W | L | | L | |
| (/y/, [y]) ~ (/y/, [x]) | | W | | | | L | |
| (/y/, [y]) ~ (/y/, [h]) | W | W | | L | | L | W |
| (/h/, [h]) ~ (/h/, [t]) | W | W | W | | W | | L |
| (/h/, [h]) ~ (/h/, [k]) | W | | W | | | | L |
| (/h/, [h]) ~ (/h/, [x]) | W | W | | W | | W | L |
| (/h/, [h]) ~ (/h/, [ʔ]) | W | | | W | | | L |

Because of the new subroutine for the choice of the loser, the original ERC matrix (9) can effectively be simplified block-by-block as in (10). To illustrate, consider for instance the top block, corresponding to the surface form [g]. The bottom two ERCs of the block corresponding to the losers [x] and [h] can be ignored: since these losers are too different from the intended winner [g], their corresponding ERCs are entailed by the other ERCs in the block and will therefore never be selected by the new subroutine for the choice of the current loser.


(10)

| | ID[DOR] | IDENT[VOI] | ID[CON] | NODORFRIC | NOVOISTOP | NOVOIFRI | *[ɦ] |
|-------|---------|------------|---------|-----------|-----------|----------|------|
| ERC 1 | W | | | L | | | |
| ERC 2 | | W | | W | L | W | |
| ERC 3 | | W | | L | W | | |
| ERC 4 | | W | | | W | L | |
| ERC 5 | | W | W | L | | | |
| ERC 6 | | W | | | | L | |
| ERC 7 | W | W | | L | | | W |
| ERC 8 | W | | W | | | | L |
| ERC 9 | W | | | W | | | L |

This illustrates an important formal property of the new subroutine: ERCs entailed by other ERCs in the same block (namely ERCs which correspond to too dissimilar losers) can be ignored. That's instead not

always the case with the classical subroutine for the choice of the loser, as shown above with the choice of the loser [k] in the case of the ERC matrix (8).


Let S be the first time when the current ranking vector entertained by the EDRA becomes consistent with either ERC 1 or ERC 6 in (10)—convergence ensures that such a time S exists. Because of the new subroutine for the choice of the current loser, ERC 2 cannot trigger any update before time S , because ERC 2 belongs to the same block as ERC 1, because ERC 2 has more w's than ERC 1, and because the current ranking vector is never consistent with ERC 1 before time S . Analogously, ERCs 3, 4, 5, and 7 cannot trigger any update before time S . In other words, the run up to time S is determined by ERCs 1, 6, 8 and 9 alone. Consider next ERC 7. Since it has more w's than the other ERCs 3-6 in the same block, it cannot trigger any update until the current ranking values have become consistent with the other ERCs 3-6. In other words, if ERC 7 triggers updates at all in the run considered, it will start triggering updates only late into the run. Thus, let the time T be defined as follows. If ERC 7 triggers at least an update in the run considered, T is the smallest time such that ERC 7 triggers an update between times T and $T + 1$; if ERC 7 triggers no updates in the run considered, T is the final time of the run. Of course, $S \leq T$ (before time S , the current ranking values are still inconsistent with ERC 6 and the EDRA is thus forbidden to consider ERC7, which has more w's). In conclusion, a generic run of the EDRA model on the current test case can be split into three stages, as in (11).



This reasoning illustrates another formal property of the new subroutine for the choice of the loser: that it makes the various ERCs enter the scene in stages, ordered by their complexity, namely by the number of winner-preferring constraints they re-rank. This means in turn that the analysis of a generic run can be split into different stages, with an increasing number of ERCs active at each stage.

This turns out to be very useful for the analysis of EDRA's restrictiveness. In fact, establishing restrictiveness requires a characterization of the final rank-

ing vector entertained by the EDRA at convergence. To obtain that characterization, we start from the initial stage and work towards the end. For each stage, we characterize the ranking vectors the learner can end up with at the end of that stage. Obviously, we have to do that for each ranking vector the learner could end up with at the end of the preceding stage. This logics is illustrated in (12). Suppose that the analysis of the first stage (in between the beginning of the run and time S) concludes that the EDRA can end up with one of two ranking vectors θ_1^S and θ_2^S at the time S when that stage ends. The analysis at the second stage (in between times S and T) will then have to be repeated twice, for each of the two ranking vectors viable at time S . And so on.



These considerations suggest that we aim for particularly tight analyses of the ranking vectors entertained at the end of the initial stages, in order to avoid a combinatorial explosion of the analyses required at later stages. Of course, tight analyses are readily possible when only a few training ERCs can trigger updates and thus mold the current ranking vector. As we increase the number of training ERCs which trigger updates, the analysis becomes more involved, and the characterization of the stage-final ranking vectors becomes looser. As illustrated in (11), the new subroutine for the choice of the loser thus comes very handy for the analysis of restrictiveness, as it ensures that the EDRA is trained on few ERCs at the beginning of the run, with additional ERCs entering the scene only at later stages.

The final appendix makes these considerations concrete through a detailed analysis of the behavior of the EDRA model with the new subroutine for the choice of the loser trained on the ERC matrix (10) corresponding to the inventory [g k ʏ x ʔ ɦ h]. The

resulting analysis establishes the following result.

Theorem 2 *When trained on an arbitrary sequence of data sampled from the glottal/velar inventory [g k ʏ x ʔ ɦ h], the EDRA model with the new subroutine for the choice of the loser succeeds at learning the ranking (7b).*

5 Conclusion

This paper has motivated a new subroutine for the choice of the current loser in phonotactic error-driven learning. Informally, the new subroutine chooses a loser which is as similar as possible to the intended winner, while being able to trigger an update. Similarity is measured in terms of the number of winner-preferring constraints in the corresponding ERC. Crucially, this new subroutine allows the various training ERCs to become active in stages, ordered by their complexity, measured in terms of the number of winner-preferring constraints. This allows for careful restrictiveness guarantees, such as the one provided by theorem 2. The proof of the theorem illustrates a number of techniques for the restrictiveness analysis of the EDRA model with the new subroutine for the choice of the loser.

A Proof of theorem 2

A.1 Analysis at an arbitrary time in the run

The markedness constraint *[h] starts high and the faithfulness constraint IDENT[DOR] starts low. Lemma 1 says that *[h] can never drop by more than 5/4 underneath IDENT[DOR]. This follows from the fact that only ERCs 7, 8, and 9 in (10) re-rank these two constraints. And that ERC 7 promotes them in tandem, and thus does not contribute to their separation. While ERCs 8 and 9 cannot demote *[h] a long way underneath IDENT[DOR], as IDENT[DOR] is winner-preferring in both ERCs 8 and 9.

Lemma 1 *The ranking values of the markedness constraint *[h] and the faithfulness constraint IDENT[DOR] satisfy the following inequality:*

$$(13) \quad \theta_{*[h]}^t \geq \theta_{\text{ID[DOR]}}^t - \frac{5}{4}$$

at any time t in any run.

Proof. The proof is by induction on time t . The inequality (13) trivially holds at the initial time $t = 0$, because of the choice of the initial ranking values

$\theta_{*[h]}^{t=0} = \theta > 0$ and $\theta_{\text{ID}[\text{DOR}]}^{t=0} = 0$. Assume that the inequality holds at time t and let me show that it then holds at time $t + 1$ as well. If the update between times t and $t + 1$ has been triggered by the ERCs 1 through 6, then the inequality holds at time $t + 1$ because it held at time t and the two constraints $*[h]$ and $\text{IDENT}[\text{DOR}]$ have not been re-ranked between times t and $t + 1$. If the update between times t and $t + 1$ has been triggered by the ERC 7, then the inequality holds at time $t + 1$ because it held at time t and both constraints $*[h]$ and $\text{IDENT}[\text{DOR}]$ have been promoted by the same amount in between times t and $t + 1$. Finally, if the update between times t and $t + 1$ has been triggered by the ERCs 8 or 9, then the inequality holds at time $t + 1$ because of the following chain of inequalities.

$$(14) \quad \theta_{*[h]}^{t+1} \stackrel{(a)}{=} \theta_{*[h]}^t - 1 \stackrel{(b)}{\geq} \theta_{\text{ID}[\text{DOR}]}^t - 1 \\ \stackrel{(c)}{=} \theta_{\text{ID}[\text{DOR}]}^{t+1} - \frac{1}{4} - 1$$

At step (14a), I have used the fact that the update by ERCs 8 or 9 in between times t and $t+1$ has demoted the constraint $*[h]$ by 1, according to the re-ranking rule (3b). At step (14b), I have used the fact that, in order for ERCs 8 or 9 to have been able to trigger an update in between times t and $t + 1$, the current ranking value $\theta_{*[h]}^t$ of the loser-preferring constraint $*[h]$ must have been larger than or equal to the ranking value $\theta_{\text{ID}[\text{DOR}]}^t$ of the winner-preferring constraint $\text{IDENT}[\text{DOR}]$. Finally at step (14c), I have used the fact that the update by ERCs 8 or 9 in between times t and $t + 1$ has promoted the constraint $\text{IDENT}[\text{DOR}]$ by $1/4$, as that ERC has $w = 3$ winner-preferring constraints and the re-ranking rule (3a) set the promotion amount equal to $\frac{1}{w+1}$. ■

If ERCs 8 and 9 were to trigger lots of updates, $\text{IDENT}[\text{DOR}]$ would be promoted a lot and $*[h]$ would be demoted a lot. In the end, $*[h]$ would thus find itself underneath $\text{IDENT}[\text{DOR}]$ separated by a large distance. But lemma 1 says that is impossible. Hence, ERCs 8 and 9 can never trigger too many updates, as stated by lemma 2.

Lemma 2 *The numbers α_8^t and α_9^t of updates triggered by ERCs 8 and 9 up to an arbitrary time t in an arbitrary run can be bound as follows:*

$$(15) \quad \alpha_8^t + \alpha_9^t \leq \frac{4}{5}\theta + 1$$

where θ is the initial ranking value of the markedness constraints.

Proof. The ranking values $\theta_{*[h]}^t$ and $\theta_{\text{ID}[\text{DOR}]}^t$ of the constraints $*[h]$ and $\text{IDENT}[\text{DOR}]$ at an arbitrary time t can be expressed as follows in terms of the numbers of updates $\alpha_7^t, \alpha_8^t, \alpha_9^t$ triggered by the ERCs 7, 8, and 9 up to time t .

$$(16) \quad \text{a.} \quad \theta_{*[h]}^t = \theta + \frac{1}{5}\alpha_7^t - \alpha_8^t - \alpha_9^t \\ \text{b.} \quad \theta_{\text{ID}[\text{DOR}]}^t = \frac{1}{5}\alpha_7^t + \frac{1}{4}\alpha_8^t + \frac{1}{4}\alpha_9^t$$

The inequality (15) follows by plugging the expressions (16a) and (16b) into (13). ■

A.2 Analysis up to time T

Recall from subsection 4 that time T is the smallest time such that ERC 7 triggers an update between times T and $T + 1$ (or the time when the run ends, in case ERC 7 triggers no updates). Constraint $\text{IDENT}[\text{DOR}]$ is only promoted by ERCs 8 and 9 up to time T (ERC 7 triggers no updates before time T). Since these two ERCs cannot trigger too many updates by lemma 2, $\text{IDENT}[\text{DOR}]$ cannot raise too high up to time T , as stated by the following lemma.

Lemma 3 *The ranking value of the faithfulness constraint $\text{IDENT}[\text{DOR}]$ satisfies*

$$(17) \quad \theta_{\text{ID}[\text{DOR}]}^t \leq \frac{1}{5}\theta + \frac{1}{4}$$

at an arbitrary time $t \leq T$.

Proof. The faithfulness constraint $\text{IDENT}[\text{DOR}]$ is only promoted by ERCs 8 and 9 up to time T (ERC 7 triggers no updates before time T). The ranking value of $\text{IDENT}[\text{DOR}]$ at an arbitrary time $t \leq T$ can then be expressed as follows in terms of the numbers α_8^t, α_9^t of updates triggered by ERCs 8 and 9 up to time t .

$$(18) \quad \theta_{\text{ID}[\text{DOR}]}^t = \frac{1}{4}(\alpha_8^t + \alpha_9^t)$$

Plugging (15) into (18) yields (17). ■

The following lemma says that the markedness constraint NODORFRIC cannot have dropped too much before time T . This follows from the fact that only ERCs 2 and 5 demote NODORFRIC up to time T (ERC 7 has not triggered any update yet). In order for NODORFRIC to have been demoted a long

way, these two ERCs must have triggered many updates. Yet, the faithfulness constraint IDENT[CON] is winner-preferring in both ERCs and is thus promoted by each update they trigger. These ERCs thus cannot trigger too many updates, because they cannot demote NODORFRIC a long way underneath IDENT[CON].

Lemma 4 *The ranking value of the markedness constraint NODORFRIC satisfies*

$$(19) \quad \theta_{\text{NODORFRIC}}^t > \frac{1}{5}\theta + \frac{1}{4}$$

at an arbitrary time $t \leq T$.

Proof. Suppose by contradiction that the claim is false. This means that there exists some time $t < T$ such that the markedness constraint NODORFRIC is demoted in between times $t - 1$ and t and its ranking value at time t is smaller than or equal to the forbidden threshold $\theta/5 + 1/4$. Since constraints are demoted by 1, its ranking value at the time $t - 1$ preceding the update must have been already smaller than or equal to $\theta/5 + 1/4 + 1$, as stated in (20a). Only ERCs 3 and 5 can have triggered this demotion (ERC 7 triggers no updates before time T). Crucially, the constraint IDENT[CONT] is winner-preferring relative to both ERCs 3 and 5. In order for either ERC 3 or 5 to have been able to demote NODORFRIC in between times $t - 1$ and t , the ranking value of the winner-preferring constraint IDENT[CONT] at time $t - 1$ must thus have been smaller than or equal to the ranking value of the loser-preferring constraint NODORFRIC at time $t - 1$, as stated in (20b). The rest of the proof derives a contradiction from these two inequalities (20).

$$(20) \quad \text{a.} \quad \theta_{\text{NODORFRIC}}^{t-1} \leq \frac{1}{5}\theta + \frac{1}{4} + 1$$

$$\text{b.} \quad \theta_{\text{NODORFRIC}}^{t-1} \geq \theta_{\text{ID[CON]}}^{t-1}$$

The ranking value of the markedness constraint NODORFRIC at time $t - 1$ can be lower bounded as $\theta_{\text{NODORFRIC}}^{t-1} \geq \theta - \alpha_3^{t-1} - \alpha_5^{t-1}$, by only considering the contribution of the ERCs 3 and 5 which demote it, while ignoring the contribution of the ERCs 2 and 9 which promote it. Plugging this bound into (20a) yields the following bound on the number of updates triggered by ERCs 3 and 5 up to time $t - 1$.

$$(21) \quad \alpha_3^{t-1} + \alpha_5^{t-1} \geq \frac{4}{5}\theta - \frac{5}{4}$$

The ranking value of the faithfulness constraint IDENT[CON] at time $t - 1$ can be lower bounded as $\theta_{\text{ID[CON]}}^{t-1} \geq \frac{1}{3}\alpha_3^{t-1} + \frac{1}{3}\alpha_5^{t-1}$, by only considering the contribution of ERCs 3 and 5. Using the bound (21) on the number of updates triggered by ERCs 3 and 5, we obtain the following lower bound on the ranking value of IDENT[CON].

$$(22) \quad \theta_{\text{ID[CON]}}^{t-1} \geq \frac{4}{15}\theta - \frac{5}{12}$$

The inequalities (20a), (20b), and (22) are contradictory (provided θ is large), because they require the ranking value $\theta_{\text{NODORFRIC}}^{t-1}$ to be smaller than $\frac{1}{5}\theta + \frac{5}{4} \simeq \frac{3}{15}\theta$ but larger than $\frac{4}{15}\theta - \frac{5}{12} \simeq \frac{4}{15}\theta$. ■

A.3 Analysis at time S

Recall that time S is the first time when the current ranking vector becomes consistent with either ERC 1 or ERC 6. Because of the new subroutine for the choice of the loser, the run up to time S is determined by ERCs 1, 6, 8, and 9, as noted above. Since the faithfulness constraint IDENT[VOICE] is the only winner-preferring constraint in both ERCs 1 and 6, it must raise a long way in order for the current ranking vector to become consistent with either ERC 1 or ERC 6 at time S , as stated by lemma 5.

Lemma 5 *The ranking value of the faithfulness constraint IDENT[VOICE] satisfies the following inequality at time S :*

$$(23) \quad \theta_{\text{ID[VOI]}}^S \geq \frac{1}{3}\theta$$

Proof. For concreteness, suppose it is ERC 1 which becomes consistent with the current ranking vector at time S (the reasoning is identical if it is ERC 6 instead). This means that the ranking value of the winner-preferring constraint IDENT[VOICE] is larger than the ranking value of the loser-preferring constraint NOVOICEDSTOP at time S , as stated by the following inequality.

$$(24) \quad \theta_{\text{ID[VOI]}}^S \geq \theta_{\text{NOVOICEDSTOP}}^S$$

By definition of time S , only ERCs 1 and 6 have promoted IDENT[VOICE] up to time S and only ERC 1 has demoted NOVOICEDSTOP. Their ranking values can thus be expressed as follows.

$$(25) \quad \text{a.} \quad \theta_{\text{ID[VOI]}}^S = \frac{1}{2}\alpha_1^S + \frac{1}{2}\alpha_6^S$$

$$\text{b.} \quad \theta_{\text{NOVOICEDSTOP}}^S = \theta - \alpha_1^S$$

Plugging (25a) and (25b) into (24) yields the following bound on the number of updates triggered by ERC 1 up to time S .

$$(26) \quad \alpha_1^S \geq \frac{2}{3}\theta - \frac{1}{3}\alpha_6^S.$$

The following chain of inequalities then yields the desired bound on the ranking value of IDENT[VOICE] at time S .

$$(27) \quad \begin{aligned} \theta_{\text{Id[voi]}^S} &\stackrel{(a)}{=} \frac{1}{2}\alpha_1^S + \frac{1}{2}\alpha_6^S \\ &\stackrel{(b)}{\geq} \frac{1}{3}\theta - \frac{1}{6}\alpha_6^S + \frac{1}{2}\alpha_6^S \stackrel{(c)}{\geq} \frac{1}{3}\theta \end{aligned}$$

At step (27a), I have used the expression (25a) of the ranking value of IDENT[VOICE]. At step (27b), I have used the bound (26) on α_1^S . At step (27c), I have lower bounded by getting rid of the contribution of α_6^S , which is crucially multiplied by a positive coefficient. ■

A.4 Analysis after time S

The faithfulness constraint IDENT[DOR] is only promoted by ERCs 7, 8, and 9. The latter two ERCs 8 and 9 promote IDENT[DOR] and not IDENT[VOICE]. Yet, they can only trigger few updates by lemma 2, and thus cannot give a substantial advantage to the former constraint over the latter. Furthermore, ERC 7 promotes both IDENT[DOR] and IDENT[VOI], and thus does not give the former any advantage over the latter. The following lemma thus concludes that IDENT[DOR] will never be able to surpass IDENT[VOICE], which already sits high at time S by lemma 5.

Lemma 6 *The ranking values of the faithfulness constraints IDENT[VOICE] and IDENT[DOR] satisfy the following inequality at any time $t \geq S$:*

$$(28) \quad \theta_{\text{Id[voi]}^t} \geq \theta_{\text{Id[DOR]}^t} + 2$$

Proof. Suppose by contradiction that (28) fails at some time $t \geq S$, as stated in (29).

$$(29) \quad \theta_{\text{Id[DOR]}^t} > \theta_{\text{Id[voi]}^t} - 2$$

From now on, let $\alpha_i^{S,t}$ denote the number of updates triggered by the i th ERC in between times S and t . Thus, $\alpha_i^t = \alpha_i^S + \alpha_i^{S,t}$. The ranking value of the faithfulness constraint IDENT[DOR] at time t can be expressed as in (30). At step (30a), I have used the fact that this constraint is promoted only by ERCs

7, 8, and 9. At step (30b), I have used the fact that ERC 7 triggers no updates before time T and thus also no updates before time S (because $S \leq T$), so that $\alpha_7^S = 0$ and thus $\alpha_7^t = \alpha_7^{S,t}$.

$$(30) \quad \begin{aligned} \theta_{\text{Id[DOR]}^t} &\stackrel{(a)}{=} \frac{1}{4}\alpha_8^t + \frac{1}{4}\alpha_9^t + \frac{1}{5}\alpha_7^t \\ &\stackrel{(b)}{=} \frac{1}{4}\alpha_8^t + \frac{1}{4}\alpha_9^t + \frac{1}{5}\alpha_7^{S,t} \end{aligned}$$

The ranking of the faithfulness constraint IDENT[VOICE] at time t can be expressed as in (31). At step (31a), I have expressed the ranking value at time $t \geq S$ as the ranking value at time S plus the increment in the ranking value due to the promotions between times S and t . At step (31b), I have lower bounded the ranking value of IDENT[VOICE] at time S using (23).

$$(31) \quad \begin{aligned} \theta_{\text{Id[voi]}^t} &= \\ &\stackrel{(a)}{=} \theta_{\text{Id[voi]}^S} + \frac{1}{2}\alpha_1^{S,t} + \frac{1}{3}\alpha_5^{S,t} + \frac{1}{2}\alpha_6^{S,t} + \frac{1}{5}\alpha_7^{S,t} \\ &\geq \theta_{\text{Id[voi]}^S} + \frac{1}{5}\alpha_7^{S,t} \\ &\stackrel{(b)}{\geq} \frac{1}{3}\theta + \frac{1}{5}\alpha_7^{S,t} \end{aligned}$$

Plugging (30) and (31) into (29) yields $\alpha_8^t + \alpha_9^t > \frac{4}{3}\theta - 2$, which contradicts (15). ■

A.5 An auxiliary result

The next step in the analysis (namely, the proof of lemma 7 below) rests on theorem 3 (Magri, 2014b).

Theorem 3 *Consider an arbitrary run of the EDRA with the re-ranking rule (3). Assume that each training ERC has a unique L. Focus on a specific training ERC, say the \bar{i} th one. Let C_ℓ be its unique loser-preferring constraint and let C_h be one of its (possibly many) winner-preferring constraints, as in (32).*

$$(32) \quad \bar{i}\text{th ERC} = \begin{bmatrix} \dots & C_h & \dots & C_\ell & \dots \\ \dots & \text{W} & \dots & \text{L} & \dots \end{bmatrix}$$

Define the coefficient δ_i as follows:

$$(33) \quad \delta_i = \begin{cases} \frac{1}{w_i+1} & \text{if } i\text{th ERC} = [\dots e/\text{W}/\text{L}\dots \text{W} \dots] \\ \frac{w_i+2}{w_i+1} & \text{if } i\text{th ERC} = [\dots \text{L} \dots \text{W} \dots] \\ 1 & \text{if } i\text{th ERC} = [\dots \text{L} \dots \text{L}/e/\text{W}\dots] \\ 0 & \text{otherwise} \end{cases}$$

The number of updates $\alpha_{\bar{i}}$ triggered by the \bar{i} th input ERC is either null or else bounded as follows

$$(34) \quad \alpha_{\bar{i}} \leq \frac{w_{\bar{i}} + 1}{w_{\bar{i}} + 2} \left(\underbrace{\theta_{\ell}^{\text{init}} - \theta_h^{\text{init}}}_{(a)} + \underbrace{1}_{(b)} + \underbrace{\sum_i \delta_i \cdot \alpha_i}_{(c)} \right)$$

where $\theta_{\ell}^{\text{init}}$ and θ_h^{init} are the initial raking values of the two constraints C_{ℓ} and C_h ; α_i is the number of updates triggered by the i th training ERC; w_i is the number of its winner-preferring constraints; the sum in (34c) runs over all training ERCs.

Here is the intuitive idea. Suppose that the initial ranking value $\theta_{\ell}^{\text{init}}$ of the loser-preferrer C_{ℓ} is larger than the initial ranking value θ_h^{init} of the winner-preferrer C_h . A certain number of updates by the \bar{i} th ERC are thus justified just in order to compensate for this bad choice of the initial ranking values, as quantified by the term (34a). At that point, the two constraints could in principle have exactly the same ranking values. An additional update is thus justified in order to bring the winner-preferring constraint C_h above the loser-preferrer C_{ℓ} , yielding the term (34b). Further updates by this \bar{i} th ERC are only justified if this ranking configuration $C_h \gg C_{\ell}$ is disrupted by updates triggered by some other training ERCs, as quantified by the term (34c). This term sums the number of updates α_i triggered by the generic i th training ERC multiplied by the ‘‘amount of disruption’’ δ_i caused by that ERC to the ranking configuration $C_h \gg C_{\ell}$. For instance, suppose that the i th training ERC looks like the top ERC listed in (33). The amount δ_i of disruption caused by that ERC is $\delta_i = \frac{1}{w_i+1}$, because that ERC disrupts the ranking configuration $C_h \gg C_{\ell}$ by promoting C_{ℓ} by $\frac{1}{w_i+1}$.

A.6 Analysis after time T

Lemma 7 says that IDENT[CONT] is always ranked above IDENT[DOR] after time T , with a sufficient distance in between the two faithfulness constraints (at least 2). The proof of this lemma is more involved than the proof of the preceding lemmas. The difficulty is due to the fact that ERC 2 promotes IDENT[CONT] at the exclusion of IDENT[DOR] while ERC 7 promotes IDENT[DOR] at the exclusion of IDENT[CONT]. In order to compare the ranking values of these two faithfulness constraints, we thus need some connection between the numbers

of updates α_2^t and α_7^t triggered by the two ERCs 2 and 7. What allows this connection to be established is the fact that the constraint NODORFRIC is winner-preferring in ERC 2 but loser-preferring in ERC 7. Since ERC 2 thus promotes the constraint NODORFRIC which ERC 7 tries to demote, updates by ERC 2 ‘‘buy’’ extra updates by ERC 7. If ERC 2 happens to trigger few updates (and thus to contribute little to the height of IDENT[CONT]), then it will buy only few updates by ERC 7 (which will therefore contribute little to the height of IDENT[DOR]). Theorem 3 is used to formalize this intuition, yielding the link between α_2^t and α_7^t in (38).

Lemma 7 *Suppose that in the run considered, ERC 7 does trigger at least an update. The ranking values of the faithfulness constraints IDENT[CON] and IDENT[DOR] satisfy the following inequality:*

$$(35) \quad \theta_{\text{ID[CON]}}^t \geq \theta_{\text{ID[DOR]}}^t + 2$$

at any time time $t \geq T$.

Proof. Since ERC 7 triggers an update between times T and $T + 1$, the loser-preferring constraint NODORFRIC cannot be underneath the winner-preferring constraint IDENT[VOICE] at time T , as stated in (36a). Furthermore, the current ranking vector at time T must be consistent with ERC 5 (otherwise, the algorithm would have chosen ERC 5 instead of ERC 7, as the former has less w’s). This means that the loser-preferring constraint NODORFRIC is already underneath IDENT[CON] at time T , as stated in (36b).

$$(36) \quad \begin{aligned} \text{a.} \quad & \theta_{\text{NODORFRIC}}^T \geq \theta_{\text{ID[VOI]}}^T \\ \text{b.} \quad & \theta_{\text{ID[CON]}}^T > \theta_{\text{NODORFRIC}}^T \end{aligned}$$

The chain of inequalities in (37) thus holds. In step (37a), I have used (36). In step (37b), I have used the fact that $T \geq S$ and that the ranking values of the faithfulness constraints can only grow with time (because they are never demoted). In step (37c), I have used the inequality (23).

$$(37) \quad \theta_{\text{ID[CON]}}^T \stackrel{(a)}{>} \theta_{\text{ID[VOI]}}^T \stackrel{(b)}{\geq} \theta_{\text{ID[VOI]}}^S \stackrel{(c)}{\geq} \frac{1}{3} \theta$$

Since ERC 7 triggers no updates before time T , the number α_7^t of updates it has triggered up to time t is equal to the number $\alpha_7^{T,t}$ of updates it has triggered between times T and t , as stated in (38a). Ap-

plying theorem 3 to ERC 7 pivoting on its winner-preferring constraint IDENT[DOR] and considering time T as the initial time yields the inequality (38b). In step (38c), I have used (36b) together with the fact that $\theta_{\text{ID[DOR]}}^T \geq 0$.

$$(38) \quad \alpha_7^t \stackrel{(a)}{=} \alpha_7^{T,t} \leq \stackrel{(b)}{\leq} \frac{5}{6} \left(\theta_{\text{NoDORFRIC}}^T - \theta_{\text{ID[DOR]}}^T \right) + 1 + \frac{5}{24} \left(\alpha_2^{T,t} + \alpha_9^{T,t} \right) \stackrel{(c)}{\leq} \frac{5}{6} \theta_{\text{ID[CON]}}^T + 1 + \frac{5}{24} \left(\alpha_2^{T,t} + \alpha_9^{T,t} \right)$$

The inequality (39) completes the proof.

$$(39) \quad \theta_{\text{ID[DOR]}}^t = \stackrel{(a)}{=} \frac{1}{4} (\alpha_8^t + \alpha_9^t) + \frac{1}{5} \alpha_7^t \leq \stackrel{(b)}{\leq} \frac{1}{4} (\alpha_8^t + \alpha_9^t) + \frac{1}{6} \theta_{\text{ID[CON]}}^T + \frac{1}{5} + \frac{1}{24} \alpha_2^{T,t} + \frac{1}{24} \alpha_9^{T,t} \stackrel{(c)}{\leq} \left(\frac{1}{4} + \frac{1}{24} \right) (\alpha_8^t + \alpha_9^t) + \frac{1}{6} \theta_{\text{ID[CON]}}^T + \frac{1}{5} + \frac{1}{24} \alpha_2^{T,t} \stackrel{(d)}{\leq} \frac{7}{24} \left(\frac{4}{5} \theta + 1 \right) + \frac{1}{6} \theta_{\text{ID[CON]}}^T + \frac{1}{5} + \frac{1}{24} \alpha_2^{T,t} \stackrel{(e)}{\leq} \frac{7}{24} \left(\frac{4}{5} 3\theta_{\text{ID[CON]}}^T + 1 \right) + \frac{1}{6} \theta_{\text{ID[CON]}}^T + \frac{1}{5} + \frac{1}{24} \alpha_2^{T,t} = \frac{7}{10} \theta_{\text{ID[CON]}}^T + \frac{1}{6} \theta_{\text{ID[CON]}}^T + \left(\frac{1}{5} + \frac{7}{24} \right) + \frac{1}{24} \alpha_2^{T,t} = \theta_{\text{ID[CON]}}^T - \frac{8}{60} \theta_{\text{ID[CON]}}^T + \left(\frac{1}{5} + \frac{7}{24} \right) + \frac{1}{24} \alpha_2^{T,t} \stackrel{(f)}{\leq} \theta_{\text{ID[CON]}}^T - \frac{8}{60} 3\theta + \left(\frac{1}{5} + \frac{7}{24} \right) + \frac{1}{24} \alpha_2^{T,t} \stackrel{(g)}{\leq} \theta_{\text{ID[CON]}}^T - 2 + \frac{1}{24} \alpha_2^{T,t} \stackrel{(h)}{\leq} \theta_{\text{ID[CON]}}^t - 2$$

The ranking value of the faithfulness constraint IDENT[DOR] can be expressed as in (39a) in terms of the contribution of the three ERCs 7, 8, and 9 which promote it. At step (39b), I have upper bounded the number α_7^t of updates triggered by ERC 7 through (38). At step (c), I have upper bounded the sum of the two terms marked with a wiggly line with $\left(\frac{1}{4} + \frac{1}{24}\right) (\alpha_8^t + \alpha_9^t)$. At step (39d), I have upper bounded $\alpha_8^t + \alpha_9^t$ using (15). At steps (39e) and (39f), I have used the inequality $\theta < 3\theta_{\text{ID[CON]}}^T$ obtained in (37). At step (39g), I have used the fact that the quantity marked by the wiggly line is smaller than -2 , provided the initial ranking value θ is large enough. Finally at step (39h), I have used the fact that $\theta_{\text{ID[CON]}}^t \geq \theta_{\text{ID[CON]}}^T + \frac{1}{4} \alpha_2^{T,t}$, because ERC 2 promotes IDENT[CONT] by $1/4$. ■

A.7 Putting the pieces together

It is easy to see that, if some rankings values are consistent with the ERC matrix (9b) and furthermore the ranking values of the two constraints sc NoDorFric and IDENT[DOR] satisfy the strict inequality (40), then each of the rankings consistent with those ranking values neutralizes [x] and thus generate the target inventory [g k y x].

$$(40) \quad \theta_{\text{NoDORFRIC}} > \theta_{\text{ID[DOR]}}$$

In order to guarantee that the EDRA model succeeds at learning this inventory and thus complete the proof of the theorem, it thus suffices to prove that its final ranking values satisfy this inequality (40), as convergence ensures consistency with the training ERC matrix (9b). To this end, we repeat below some of the inequalities which have been obtained with the preceding lemmas.

$$(17) \quad \theta_{\text{ID[DOR]}}^t \leq \frac{1}{5} \theta + \frac{1}{4} \quad \text{for any } t \leq T \quad (\text{lemma 3})$$

$$(19) \quad \theta_{\text{NoDORFRIC}}^t > \frac{1}{5} \theta + \frac{1}{4} \quad \text{for any } t \leq T \quad (\text{lemma 4})$$

$$(28) \quad \theta_{\text{ID[VOI]}}^t \geq \theta_{\text{ID[DOR]}}^t + 2 \quad \text{for any } t \geq S \quad (\text{lemma 6})$$

$$(35) \quad \theta_{\text{ID[CON]}}^t \geq \theta_{\text{ID[DOR]}}^t + 2 \quad \text{for any } t \geq T \quad (\text{lemma 7})$$

The two inequalities (17) and (19) guarantee that the markedness constraint NODORFRIC is indeed ranked above IDENT[DOR] up to time T . In other words, that the inequality (40) holds up to time T . If ERC 7 never triggers any update in the run considered, time T is the final time of the run, and the proof is completed. Thus, suppose that T is not the end of the run. The inequalities (28) and (35) ensure that the two faithfulness constraints IDENT[VOICE] and IDENT[CONT] are always above IDENT[DOR] from time T on, with enough space in between (namely at least 2). Since NODORFRIC is ranked above IDENT[DOR] at time T , it can never demoted below it, because any ERC where it is loser-preferring counts at least one of the two constraints IDENT[VOICE] or IDENT[CONT] as winner-preferring.

References

- Paul Boersma. 1998. *Functional Phonology*. Ph.D. thesis, University of Amsterdam, The Netherlands. The Hague: Holland Academic Graphics.
- Paul Boersma. 2009. Some correct error-driven versions of the constraint demotion algorithm. *Linguistic Inquiry*, 40:667–686.
- Noam Chomsky and Morris Halle. 1965. Some controversial questions in phonological theory. *Journal of Linguistics*, 1:97–138.
- Bruce Hayes and Donca Steriade. 2004. Introduction: the phonetic bases of phonological markedness. In Bruce Hayes, Robert Kirchner, and Donca Steriade, editors, *Phonetically based phonology*, pages 1–33. Cambridge University Press.
- Ian Maddieson. 1984. *Patterns of Sounds*. Cambridge University Press.
- Giorgio Magri and René Kager. 2015. How to account for phonetically counterintuitive segment inventories using only phonetically grounded markedness constraints. In Thuy Bui and Deniz Ozyildiz, editors, *Proceedings of NELS 45*, Amherst, MA. GLSA Publications.
- Giorgio Magri. 2012. Convergence of error-driven ranking algorithms. *Phonology*, 29(2):213–269.
- Giorgio Magri. 2013a. The error-driven ranking model of the early stage of the acquisition of phonotactics: an initial result on restrictiveness. In Hsin-Lun Huang, Ethan Poole, and Amanda Rysling, editors, *Proceedings of NELS 43: the 43rd annual meeting of the North East Linguistic Society*, pages 277–290.
- Giorgio Magri. 2013b. A note on the GLA’s choice of the current loser from the perspective of factorizability. *Journal of Logic, Language, and Information*, 22:231–247.
- Giorgio Magri. 2014a. The EDRA model of the acquisition of phonotactics: the problem of \mathcal{F} -controlling. In Özlem Çetinoglu, Jeffrey Heinz, Andreas Maletti, and Jason Riggle, editors, *Proceedings of MORPHFSM 2014*. Association for Computational Linguistics.
- Giorgio Magri. 2014b. Tools for the robust analysis of error-driven ranking algorithms and their implications for modeling the child’s acquisition of phonotactics. *Journal of Logic and Computation*, 24.1:135–186.
- Giorgio Magri. 2015a. How to control the height of the faithfulness constraints. Ms.
- Giorgio Magri. 2015b. Idempotency in constraint-based phonology and the formal underpinning of correspondence theory. Submitted manuscript.
- Giorgio Magri. 2015c. Restrictiveness and the relative ranking of the markedness constraints. Ms.
- John J. Ohala. 1983. The origin of sound patterns in vocal tract constraints. In P. F. MacNeilage, editor, *The production of speech*, pages 189–216. Springer Verlag, New York.
- Alan Prince and Paul Smolensky. 2004. *Optimality Theory: Constraint Interaction in generative grammar*. Blackwell, Oxford. As Technical Report CU-CS-696-93, Department of Computer Science, University of Colorado at Boulder, and Technical Report TR-2, Rutgers Center for Cognitive Science, Rutgers University, New Brunswick, NJ, April 1993. Also available as ROA 537 version.
- Alan Prince. 2002. Entailed ranking arguments. Ms., Rutgers University, New Brunswick, NJ. Rutgers Optimality Archive, ROA 500. Available at <http://www.roa.rutgers.edu>.
- Jason Riggle. 2004. Contenders and learning. In *Proceedings of the 23rd annual meeting of the West Coast Conference on Formal Linguistics*, pages 101–114.
- Paul Smolensky. 1996. The initial state and Richness of the Base in Optimality Theory. John Hopkins Technical Report.
- Bruce Tesar and Paul Smolensky. 1998. Learnability in Optimality Theory. *Linguistic Inquiry*, 29:229–268.