Child consonant harmony and phonologization of performance errors

Tara McAllister Byun & Sharon Inkelas[1]

New York University & UC Berkeley

1.   Child-specific speech patterns: Performance or competence?

The phenomenon of child-specific phonological processes represents a longstanding challenge for efforts to arrive at a coherent model of phonological development. These processes may be robustly attested in the speech of typically developing children, yet extensive cross-linguistic investigation reveals no counterpart in adult phonological typology. Developmental consonant harmony (DCH) is one of the most frequently cited examples of this category. While adult phonologies also permit long-distance patterns of consonant agreement, the child pattern is unique in allowing assimilation with respect to major place of articulation. Examples of DCH from Pater 2002 can be seen in (1)-(2). The parameters of DCH are subject to considerable variation within and across children. However, Pater (2002:364) proposes that several implicational generalizations govern the preferred target, trigger, and direction of DCH in English, as in (3):

(1)   Regressive DCH: Velar or labial trigger, coronal or labial undergoer
    a.   [gɪːguː]   'tickle'
    b.   [gʌg]   'bug'
    c.   [pap]   'top'

(2)   Progressive CH: Velar or labial trigger, coronal or labial undergoer
    a.   [kok]   'coat'
    b.   [kʌk]   'cup'
    c.   [bɛːpˀ]   'bed'

(3)   Generalizations about DCH in English
    a.   Target/Undergoer: Non-coronal implies coronal
    b.   Trigger: Labial implies velar
    c.   Direction: Progressive implies regressive

The existence of child-specific phonological patterns is problematic for models that assume continuity between child and adult grammars (e.g. Macnamara 1982; Pinker 1984). One response is to maintain that child-specific patterns are purely the product of performance limitations of young children and are unrelated to their grammatical competence (e.g. Hale & Reiss 1998). There is certainly ample evidence that child speakers face performance pressures that are distinct from and more extensive than those experienced by adults. In addition to anatomical differences, such as the larger size and more anterior position of the child's tongue (Crelin 1987), children have a more limited ability to plan and execute precise movements or complex movement sequences (Fletcher 1992). [2] Furthermore, common child processes can be understood as a product of immature speech-motor capabilities. In the case of DCH, it is presumed that repeating the same place of articulation is a way to simplify the motor planning task.

However, there also is abundant evidence that child speech patterns are more than pure performance errors. If child errors derived directly and exclusively from phonetic pressures, we would expect to see comparable patterns across all speakers, regardless of the language of the environment. Instead, we see dissociation across languages (e.g. Li et al. 2011). Second, many accounts describe child-specific patterns as conditioned categorically by prosodically defined units (e.g. foot-initial versus foot-medial context), with no apparent influence of other factors such as voicing, vowel context, speech rate, or vocal loudness (e.g. Inkelas & Rose 2007). A final piece of evidence comes from the existence of U-shaped learning curves (e.g. Becker & Tessier 2011), in which a child is produces a sound accurately in early stages of development, then shifts to a period of systematic application of an error pattern before returning to a trajectory of increasing accuracy. This poses a challenge for the pure performance approach because the child has previously shown him/herself physically capable of approximating the adult target.

A compromise approach holds that both competence and performance play a role in child speech patterns. Here the idea is that performance pressures may become phonologized and thus take on a systematic quality. To explain the absence of any reflex of these constraints in adult typology, it is necessary to assume that the constraints arising through phonologization of children's performance limitations are somehow eliminated or inactivated by adulthood. To our knowledge, no previous model has explicitly proposed an update mechanism for the inactivation of child-specific constraints may be eliminated. We propose a mechanism (McAllister Byun, Inkelas, & Rose 2012) that explicitly models both how performance limitations are incorporated into children's grammatical computations, and how these effects are eliminated in typical maturation.

2.   Does DCH reflect transient phonologization of performance pressures?

2.1   Parallels between DCH and adult speech errors

Previous work on adult CH has made note of striking parallels with patterns of assimilation in adult speech errors, e.g. *sunshine* → [ʃʌnʃaɪn] (Hansson 2001). Hansson suggested that adult CH might be a phonologized reflex of the processing or planning

---

[2]There are also perceptual differences between children and adults, but these will be largely set aside for the present.

pressures that give rise to sporadic speech errors. In this case, it seems plausible that the more significant motor limitations experienced by children could give rise to more DCH, including major place assimilation. To date, though, models of DCH have made only very limited use of this possibility. In this section, we argue that all of the implicational relations that Pater (2002) identified for DCH (see (3)) have counterparts in experimentally documented characteristics of adult speech errors.

The most striking similarity between speech errors and DCH involves the bias toward assimilation in a regressive direction. An estimated 75% of adult speech errors involve regressive assimilation (Schwartz et al. 1994). This bias finds an explanation in models where all segments/motor plans are active in a buffer at the start of the word. Plans are deactivated as they are used, so a word-final target faces less competition than an initial target (Dell, Burger, & Svec 1997). With their limited motor planning capacities, children may have particular difficulty inhibiting competing plans, which could explain why they make regressive assimilations on a broader scale than adults.

The preference for velar segments as triggers of harmony can be linked to the finding that speech sounds that have multiple phonological properties in common are more likely to interact in speech errors (e.g. Fromkin 1973; Shattuck-Hufnagel & Klatt 1979; Hansson 2001). The observation that coronal-to-velar assimilation is more persistent than coronal-to-labial or labial-to-velar assimilation in DCH can be understood as a reflection of this influence of similarity: the pressure for assimilation is greater between two lingual consonants than between targets that do not share a major articulator.

Finally, the preference for coronal segments as undergoers of harmony finds an echo in experimental work by Pouplier and colleagues. Pouplier (2008) demonstrated that a large percentage of speech errors that are perceived as categorical substitutions actually involve simultaneous production of intrusive and target gestures. Further, Pouplier & Goldstein 2005) showed that intrusive errors have asymmetrical perceptual consequences, such that coronal targets with intrusive velar gestures are perceived to have velar place, while intrusive coronal gestures during a velar target typically remain undetected. A similar predominance of labial over coronal place was reported in Byrd 1992. If DCH involves gestural coproduction, these perceptual asymmetries could explain the tendency of coronals to assimilate to velars/labials, and not vice versa.
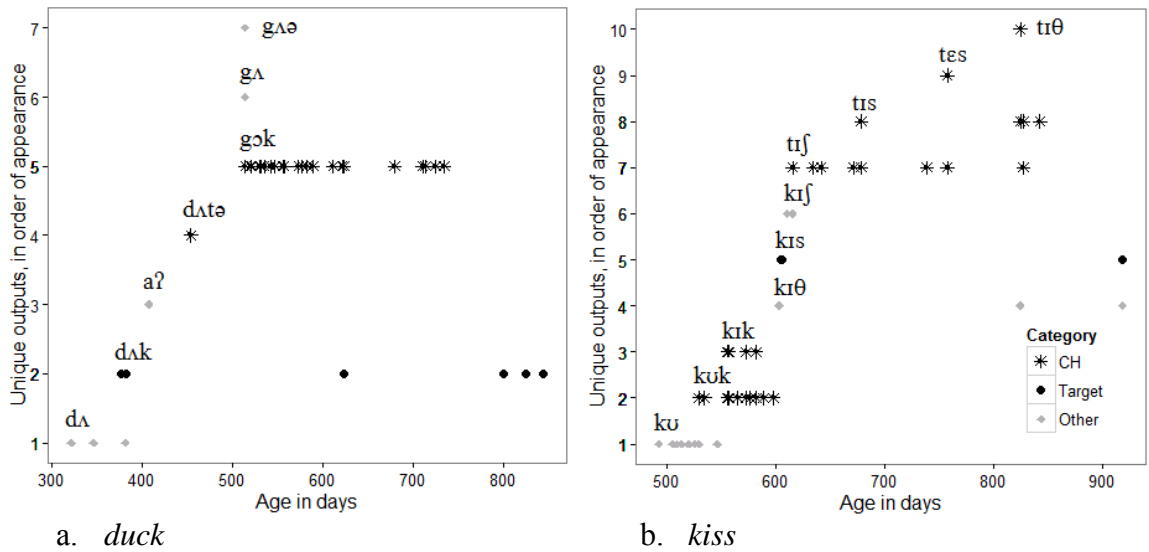
2.2   How systematic is DCH?

In this section, we evaluate the extent to which actual DCH data conform to the biases identified in adult speech errors (regressive directional bias, preference for velar triggers, and preference for coronal targets). This exercise was intended to assess the viability of an account in which DCH emerges directly from limitations on speech-motor planning. For our analysis, we revisit the Trevor corpus (Compton & Streeter 1977; Pater 2002), which forms the basis for Pater's implicational generalizations about the directionality and featural preferences of DCH. Our analyses draw on the coded distillation of DCH data from Trevor's outputs from ages 0;10.11-3;1.8 (Becker & Tessier 2011; online supplement). It is worth noting that the overall rate of application of consonant harmony in this corpus is low; of 5228 relevant contexts, only 14.7% exhibit harmony.

The investigation returned mixed results. We found that the preference for a regressive direction of assimilation is upheld for coronal-velar but not for coronal-labial

pairs. In the latter case, regressive harmony (TVP→PVP) is attested in 19/461 environments (4%) and is eliminated at around 1;8. Progressive harmony (PVT→PVP) is more common, occurring in 191/1545 environments (12.4%), and is eliminated later, at around 2;0. With respect to the preferred place of the trigger and target consonant, Trevor's early outputs are consistent with prediction, with velar place typically predominating over coronal place. However, at 1;8, Trevor's outputs show a shift from velar dominance (KVT→KVK) to coronal dominance (KVT→TVT), as described by Becker & Tessier 2011. This is not predicted by an account in which the harmonized form reflects coproduction of velar and coronal gestures (e.g. Pouplier 2008).

Perhaps the most striking result of our investigation pertains to the idiosyncratic patterning of individual lexical items. To visualize changes in Trevor's output over time, we provide charts, in Figures 1-2, in which the x-axis features Trevor's age in days, while each number on the y-axis represents a unique output form attested in the corpus, in order of emergence. The phonetic transcription is superimposed above the dots representing tokens of that output form. Separate symbols mark instances of the adult target form, harmonized forms, and other deviations from the adult target. Figure 1(a-b) highlights the extensive variability that exists within Trevor's realization of a single lexical item over time. Figure 2a-b focuses on a specific phenomenon of U-shaped curves in Trevor's output. Becker & Tessier (2011) describe a U-shaped curve in Trevor's overall trajectory of acquisition of sequences of coronal and velar consonants. Figure 2 reveals that these regressions can also be observed at the level of individual lexical items. For instance, in Figure 2a, Trevor first produces the faithful form [dʌk] at 377 days, but he then enters an extended period in which the word is realized with consonant harmony; he returns to the faithful pronunciation at 800 days.

Figure 1. Variability in Trevor's realization of individual lexical items over time

a. *cat*

b. *book*

Figure 2. U-shaped trajectories in Trevor's realization of lexical items over time

a. *duck*

b. *kiss*

In sum, we see that Trevor's pattern of DCH is neither entirely random nor entirely systematic. Individual lexical items do show stable periods of harmony application, but trajectories differ across lexical items in a way that is not readily captured by simple rules or constraints. The broad generalization that best fits the observed pattern is the notion that children have a bias to continue producing their own error forms. This bias has previously been described in the context of U-shaped learning curves, lexical fossils (i.e. early-acquired or high-frequency words that continue to exhibit a pattern that has otherwise been eliminated from the active grammar), and phonological template effects (e.g. Vihman & Croft 2007). Previously, children's preference to recycle old error forms has been modeled in the framework of Error-Selective Learning (ESL; Tessier 2012; Becker & Tessier 2011), which proposes that every unique output of the child's grammar is stored in a buffer called the Cache. Even after the grammar advances to a more adult-like stage, the speaker retains the option of reusing an old form stored in the Cache.

While Becker & Tessier suggest that retrieving a cached form might require less effort or processing cost than generating the correct form through the grammar, in general the ESL model does not elaborate on the question of *why* children would prefer to recycle their own error forms. Elsewhere it has been argued that error forms may have a more stable motor plan than faithful forms, since the old form has been practiced many times (e.g. Ota & Green 2013). In the following section, we propose a grammatical constraint that favors continued production of a candidate with a stable motor-acoustic mapping, even if this comes at the expense of perfect faithfulness.

3.   The A-Map model

The core insight of our model is that children's grammatical computations are influenced not only by the desire to be *accurate*, i.e. to produce an output that is acoustically similar to the adult target, but also to be *precise*, i.e. to select a target that can be realized reliably across multiple attempts. Adopting an exemplar-based model of phonology (e.g. Johnson

1997; Pierrehumbert 2001), we assume that all phonetic forms experienced in the act of producing and perceiving speech are stored as detailed traces in a high-dimensional map of the phonetic properties of speech. In the case of the speaker's own outputs, we assume that stored exemplars contain detail about both the motor plan executed and its acoustic consequences, with links between them. The nature of this motor-acoustic mapping will vary across targets. A simple, reliably executed motor plan will map to a narrowly defined region of acoustic space (high precision), whereas a complex, unreliable motor plan will be associated with extensive scatter in acoustic space (low precision).

We illustrate these concepts in Figure 3 with an example from DCH. We assume, uncontroversially, that a motor plan for a syllable containing two different consonant places (e.g. /dʌk/) is more difficult to execute than a plan in which the same place of articulation is repeated twice (e.g. /gʌk/). Therefore, the motor-acoustic traces associated with the target form /dʌk/ will be divided across numerous errors in addition to some correct outputs. Due to the anticipatory bias of speech errors, errors reflecting regressive assimilation are expected to outnumber progressive errors ([gʌk] > [dʌt]). By contrast, few performance errors arise in connection with the motorically simpler sequence /gʌk/, and its record in exemplar space reflects a more reliable, precise motor-acoustic mapping.

Figure 3. Differing degrees of precision in motor-acoustic mapping
    a.   A motorically complex target yields a diffuse mapping; precision is low.
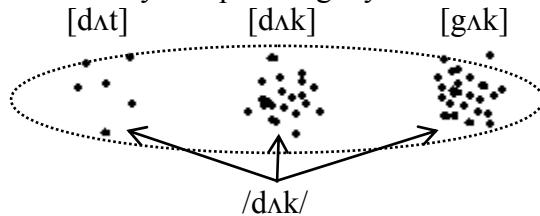
    b.   A motorically simple target yields a high-precision mapping.
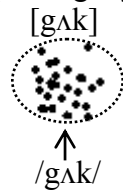
We propose that information about the accuracy and precision of motor-acoustic mappings is stored in an Articulatory-map, or A-map—analogous to Steriade's 2001) P(erceptual)-map—that can then be referenced by grammatical constraints. An A-map entry is a vector with three components: $<MP_{mean}, A_{mean}, A_{SD}>$. MP_{mean} represents a stored motor plan, as averaged over cloud of previous executions of closely related motor plans. A_{mean} represents the center, in multidimensional acoustic space, of the cloud of acoustic outcomes associated with past executions of motor plan MP. Finally, A_{SD} is the standard deviation of the entire distribution of acoustic outcomes associated with past executions of MP; it serves as an index of the precision of the motor-acoustic mapping.

We implement the A-map in the framework of Harmonic Grammar (Legendre, Miyata & Smolensky 1990), in which weighted constraints take on scalar-valued violations. The pressure to select a form whose acoustic output will resemble the adult target is implemented by the constraint ACCURATE. The magnitude of a candidate's ACCURATE violation is dictated by the distance between *MP* and the center of *T,* the cloud of acoustic traces representing (the child's perceptions of) adult productions of the target. Meanwhile, the pressure to choose a form with a stable motor-acoustic mapping is expressed by a constraint called PRECISE, whose violation magnitude is dictated by the value of A_{SD}. A broader, more scattered cloud (large A_{SD}) incurs a greater PRECISE violation than a compact cloud. We assume that PRECISE and ACCURATE coexist with conventional markedness and faithfulness constraints, which take on increasing importance as lexical representations become increasingly segmentalized over time (e.g. Munson et al. 2005; Werker & Curtin 2005; Curtin et al. 2011). Because children who exhibit DCH tend to be very young, we assume that their lexical representations have the coarse-grained, holistic quality that is typical of early stages of acquisition. We thus reason that PRECISE and ACCURATE play a substantial role in grammatical computations of children in this stage of development.

4. Modeling child speech in the A-Map framework

4.1 An A-Map model of DCH

The tableau in (4) depicts the ACCURATE and PRECISE violations incurred by three competing candidates for the adult target [dʌk]. As described above, for a young speaker like Trevor, past attempts at the disharmonic target /dʌk/ (candidate (a)) will have resulted in a number of speech errors, producing scatter in motor-acoustic exemplar space. Thus, A_{SD} is relatively large for candidate (a), with a correspondingly large violation of PRECISE (here arbitrarily given the value 1). However, target /dʌk/ incurs no ACCURATE violation because the predicted acoustic outcome, A_{mean}, is identical to the adult target. By contrast, the harmonized target /gʌk/ (candidate (b)) diverges perceptually from the adult target, violating ACCURATE, but incurs only a trivially small violation of PRECISE because attempts at /gʌk/ are reliably realized as [gʌk]. To model the grammar of a young child, we assume that PRECISE carries a higher weight than ACCURATE. Thus candidate (b) is more harmonic than faithful candidate (a).

Another competitor is candidate (c), /dʌt/, which features progressive coronal harmony. The preference for regressive over progressive harmony (which is present in adult as well as DCH; Hansson 2001) is challenging to model in a constraint-based grammar that takes perceptual cue strength into account. This is because both the progressively and regressively harmonized forms will satisfy all articulatory or featural well-formedness constraints, and the progressive form has the added advantage of preserving the features that occur in the perceptually privileged word- or syllable-initial position. (This difference in perceptual salience is reflected in (4) with a slightly lower ACCURATE violation for candidate (c) relative to (b).) In the context of DCH, the preference for regressive over progressive harmony can be explained if we focus on the A-map at the level of the individual lexical item. At a global level, it is true that a child speaker like Trevor should have no difficulty executing the motor plan /dʌt/ and

producing the output [dʌt]. However, because of the regressive directional bias of speech errors, he only rarely produces [dʌt] in connection with the lexical target /dʌk/. The more densely populated exemplar cloud for [gʌk] will have a smaller standard deviation than the more sparsely populated cloud for [dʌt]. As a result, the mapping /gʌk/ → [gʌk] is favored by PRECISE over the reliable mapping from /dʌt/ to the more sparsely populated cloud [dʌt], in (b), and over the unreliable mapping /gʌk/→[gʌk] in candidate (a):

(4)   Comparison of candidates for target /dʌk/

	Adult target: [dʌk]	PRECISE	ACCURATE	*H*
		w = 2	*w* = 1	
a.	</dʌk/, [dʌk], 2>	-1		-2
☞ b.	</gʌk/, [gʌk], .25>	.25	-1	-1.5
c.	</dʌt/, [dʌt], .5>	-.5	-.75	-1.75

Up to this point, we have defined PRECISE only in relation to individual lexical items. Does this mean that child patterns arising from PRECISE will remain lexically specific and not generalize across words with a similar shape? This would be problematic, since child speech patterns can change abruptly across the lexicon, and children may immediately apply a pattern such as DCH to a newly presented nonword (McAllister 2009). We propose that PRECISE can encompass multiple sub-constraints that apply over different-sized chunks of speech (at the word level, at the level of syllable- or foot-sized chunks of speech, and at a sub-syllabic or segmental level). These levels are not simultaneously present from the earliest stage, but emerge gradually in tandem with more refined levels of representation over the course of exposure to linguistic inputs. As PRECISE violations come to be calculated over smaller chunks, its effects begin to appear more systematic and rule-like. When changes in the A-map make it possible to produce a new segment or diphone, associated substitutions may thus be eliminated in an across-the-board fashion.

4.2   Capturing U-shaped curves with the A-map

As discussed above, U-shaped learning trajectories represent a particular puzzle in developmental phonology. We saw that U-shaped curves were present in Trevor's patterns of DCH, both at a global level and in the trajectories followed by individual lexical items. In our model, these developmental reversals can be understood as a consequence of changes in the properties of the A-map in the earliest stages of learning. It takes time for traces representing the child's own productions of various target sounds to build up into well-defined regions of clustering and separation. Until well-formed distributions coalesce out of the noise, all candidates will have similarly high A_{SD} values. With PRECISE thus failing to distinguish among candidates, the decision will fall to ACCURATE, and the faithful candidate will be targeted for production. This does not mean that the candidate will always be realized accurately; at this point in the child's development, there is still a very high probability that some performance error will occur. In (5), we depict a comparison of candidates at a developmentally very early time point, where sufficient observations have not been collected for a well-specified A-map.

(5)   Comparison of candidates for target /dʌk/

	Adult target: [dʌk]	PRECISE	ACCURATE	*H*
		w = 2	*w* = 1	
☞ a.	</dʌk/, [dʌk], 2>	-2		-4
b.	</gʌk/, [gʌk], .25>	-2	-1	-5
c.	</dʌt/, [dʌt], .5>	-2	-.75	-4.75

4.3   Elimination of child-specific patterns in the A-map model

A crucial argument in favor of the A-map model is its ability to model both the origin and obsolescence of child-specific speech processes. Previous models have proposed that child-specific markedness constraints are constructed in response to articulatory or perceptual pressures (e.g. Pater 1997; Becker & Tessier 2011). On such approaches it is unclear how child-specific constraints are inactivated prior to adulthood. In the A-map model, there are no true child-specific constraints; PRECISE and ACCURATE remain in the grammar through adulthood. Child-specific patterns disappear as the influence of PRECISE is attenuated due to maturation and experience-driven changes in the A-map.

The A-map account actually suggests two paths, which are by no means mutually exclusive, to the elimination of a child phonological process. First, we anticipate that some changes over the course of phonological acquisition will occur due to increases in the weight of ACCURATE relative to PRECISE. The Gradual Learning Algorithm for Harmonic Grammar (Boersma & Pater 2008) predicts that the weight of markedness constraints like PRECISE will decrease incrementally in each cycle of evaluation in which the stored form favored by PRECISE differs from the adult target form.

Supplementing this gradual learning of constraint weights is a second type of learning, in which changes in the A-map decrease the magnitude of the PRECISE violation incurred by a given target. It is therefore crucial to understand what conditions can bring about significant developmental changes in the A-map. Clearly, maturational changes in the anatomy and motor control of the speech structures influence the reliability with which a certain speech target can be realized. However, the A-map will change only if the child has opportunities to observe that a particular target can now be executed with greater reliability. Such opportunities will not arise if the child's grammar, influenced by PRECISE, uniformly selects a non-faithful (but reliable) candidate for a given target. As the weights of PRECISE and ACCURATE get closer together or cross over time, the child is more likely to attempt the fully faithful target. If the faithful target is selected and maturation has occurred, the greater reliability of the mapping from the motor plan to acoustic space will be encoded as a lower A-map score, translating to a smaller-magnitude PRECISE violation. This will increase the frequency with which the faithful target is selected, creating more opportunities for modification of the A-map. This "virtuous cycle" will arise only if the child's motor and/or anatomical limitations have in fact been lifted. If continuing performance limitations keep the error rate high, the A-map will remain unchanged, and PRECISE will continue to favor the substitution of a more stable target. Thus, the A-map provides a mechanism to explain and model differences in the trajectory of elimination of phonological patterns across children.

Since PRECISE is not eliminated from the grammar, we can ask whether it has any reflexes in adult speech patterns. For most adults, it is trivially easy to execute virtually any sound (combination) allowed by the native language, and speech errors occur with sufficiently low frequency that there are no meaningful differences in the reliability of the motor-acoustic mapping across targets. With motor pressures thus leveled, violations of PRECISE converge on similar values, and conventional markedness and faithfulness constraints emerge as the primary driving forces of grammar. Nonetheless, some of the pressures that lead to systematic errors in children remain present at a low level in adults, where they drive gradient phonetic tendencies or influence the rate of occurrence of sporadic speech errors. The A-map model also predicts that effects of PRECISE could reemerge in adult speakers after a stroke or brain injury negatively impacts the reliability of the motor-acoustic mapping for speech. This is consistent with reports (e.g. Buchwald 2009) that patterns of error in adults with acquired speech impairment have a systematic character amenable to analysis with a constraint-based grammar. In fact, adults with aphasia have been reported to exhibit consonant harmony processes (Kohn, Melvold, & Smith 1995), although CH for major place of articulation has not been reported.

5.   Conclusions

This paper set out to answer several intersecting questions. One is why children exhibit speech patterns not seen in adults. Our answer to this question is that children are subject to child-specific performance limitations. These limitations are incorporated into the computations of the grammar by means of the A-map, which includes an index of the reliability of the mapping from motor plans to acoustic space. A second question is why children persist in 'incorrect' patterns even after demonstrating the ability to pronounce words in a more adult-like fashion. Why is there a lag between the waning of performance limitations and the cessation of the associated phonological patterns? The answer to this question lies in the influence of the A-Map and the PRECISE constraint, through which performance limitations are incorporated into the grammar. As a distillation of traces of past productions, the A-Map is intrinsically conservative; old errors may continue to have an influence even after the child has outgrown the relevant performance limitation. A third question is what causes child-specific speech patterns to disappear. This, too, can be traced to the A-Map and PRECISE. As motor maturation takes place, the rate of occurrence of motor deviations and speech errors declines, even for complex speech targets. As old error traces decay and are replaced with traces of correct productions, A-Map values become more homogeneous across targets, and the influence of PRECISE is attenuated. In the past, some discussions of child-specific phonology have questioned whether children's deviant productions are the output of their grammar (e.g. Pater 2002), or whether they reflect performance limitations (e.g. Hale & Reiss 1998). In the A-map model, this is not an either-or question; the two factors are inseparable.

References

Becker, M., and A.-M. Tessier. 2011. Trajectories of faithfulness in child-specific phonology. *Phonology* 28:163-196.

Boersma, P., and J. Pater. 2008. Convergence properties of a gradual learning algorithm for Harmonic Grammar. Ms, University of Amsterdam and University of Massachusetts Amherst [ROA-970].

Buchwald, A. 2009. Minimizing and optimizing structure in phonology: evidence from aphasia. *Lingua* 119:1380–1395.

Byrd, D. 1992. Perception of assimilation in consonant clusters: a gestural model. *Phonetica* 49:1–24.

Compton, A. and M. Streeter. 1977. Child phonology: data collection and preliminary analyses. *Papers and Reports on Child Language Development 7*. Stanford University.

Crelin, E. S. 1987. *The human vocal tract: Anatomy, function, development, and evolution*. New York: Vantage Press.

Curtin, S., K. Byers-Heinlein and J. Werker. 2011. Bilingual beginnings as a lens for theory development: PRIMIR in focus. *Journal of Phonetics* 39:492–504.

Dell, G., L. Burger and W. Svec. 1997. Language production and serial order: a functional analysis and a model. *Psychological Review* 104:123–147.

Fletcher, S. 1992. *Articulation: A physiological approach*. San Diego: Singular.

Fromkin, V. 1973. *Speech errors as linguistic evidence*. The Hague: Mouton.

Hale, M. and C. Reiss. 1998. Formal and empirical arguments concerning phonological acquisition. *Linguistic Inquiry* 29:656–83.

Hansson, G. 2001. Theoretical and typological issues in consonant harmony. Doctoral dissertation, UC Berkeley.

Inkelas, S. and Y. Rose. 2007. Positional neutralization: a case study from child language. *Language* 83:707-36.

Johnson, K. 1997. Speech perception without speaker normalization: An exemplar model. In *Talker variability in speech processing*, ed. by K. Johnson and J. Mullennix, 145-165. San Diego: Academic Press.

Kohn, S., J. Melvold and K. Smith. 1995. Consonant harmony as a compensatory mechanism in fluent aphasic speech. *Cortex* 31:747-56.

Legendre, G., Y. Miyata and P. Smolensky. 1990. Can connectionism contribute to syntax? Harmonic Grammar, with an application. In *CLS 25,* ed. by M. Ziolkowski, M. Noske and K. Deaton, 237-252. Chicago: Chicago Linguistic Society.

Li, F., et al. 2011. Language specificity in the perception of voiceless sibilant fricatives in Japanese and English: implications for cross-language differences in speech-sound development. *Journal of the Acoustical Society of America* 129:999-1011.

Macnamara, J. 1982. *Names for things: a study of child language*. Cambridge: Bradford Books / MIT Press.

McAllister, T. 2009. Positional asymmetries in phonological acquisition. Doctoral dissertation, MIT.

McAllister Byun, T., S. Inkelas, and Y. Rose. Transient patterns in phonological development: the A-map model. Ms. New York University, UC Berkeley, and Memorial University.

Munson, B., B. Kurtz and J. Windsor. 2005. The influence of vocabulary size, phonotactic probability, and wordlikeness on nonword repetitions of children with and without Specific Language Impairment. *Journal of Speech, Language, and Hearing Research* 48:1033–1047.

Ota, M. and S. Green. 2013. Input frequency and lexical variability in phonological development: a survival analysis of word-initial cluster production. *Journal of Child Language* 40:1-28.

Pater, J. 1997. Minimal violation and phonological development. *Language Acquisition* 6:201–253.

Pater, J. 2002. Form and substance in phonological development. In *WCCFL 21*, ed. by L. Mikkelsen and C. Potts, 348-372. Somerville: Cascadilla Press.

Pierrehumbert, J. 2001. Exemplar dynamics: Word frequency, lenition and contrast. In *Frequency and the emergence of linguistic structure*, ed. by J. Bybee and P. Hopper, 137-157. Amsterdam: John Benjamins.

Pinker, S. 1984. *Language learnability and language development*. Cambridge: Harvard University Press.

Pouplier, M. 2008. The role of a coda consonant as error trigger in repetition tasks. *Journal of Phonetics* 36:114–140.

Pouplier, M. and L. Goldstein. 2005. Asymmetries in the perception of speech production errors. *Journal of Phonetics* 33:47–75.

Schwartz, M., et al. 1994. Disordered speech production in aphasic and normal speakers. *Brain and Language* 47:52-88.

Shattuck-Hufnagel, S. and D. Klatt. 1979. The limited use of distinctive features and markedness in speech production: Evidence from speech errors. *Journal of Verbal Learning and Verbal Behaviour* 18:41–55.

Steriade, D. 2001. Directional asymmetries in place assimilation: A perceptual account. In *The role of speech perception in phonology*, ed. by E. Hume and K. Johnson, 219-50. New York: Academic Press.

Tessier, A.-M. 2012. USELISTEDERROR: a grammatical account of lexical exceptions in phonological acquisition. In *NELS 39*, ed. by S. Lima, K. Mullins, and B. Smith, 813-827. Amherst: GLSA.

Vihman, M. and W. Croft. 2007. Phonological development: toward a "radical" templatic phonology. *Linguistics* 45:683–725.

Werker, J. and S. Curtin. 2005. PRIMIR: A developmental framework of infant speech processing. *Language Learning and Development* 1:197–234.

Tara McAllister Byun
Deparment of Communicative Sciences and Disorders
New York University — Steinhardt
New York, NY 10012

tara.byun@nyu.edu

Sharon Inkelas
Department of Linguistics
UC Berkeley
Berkeley, CA 94720

inkelas@berkeley.edu