Training child learners on non-native vowel contrasts: the role of talker variability

Gwen Brekelmans,^{1,2} Bronwen G. Evans,³ Elizabeth Wonnacott^{2,4}

¹ *Department of Biological and Experimental Psychology, Queen Mary University of London, Fogg Building, Mile End Road, London E1 4NS, United Kingdom*

² *Department of Language and Cognition, University College London, Chandler House, 2 Wakefield Street, London WC1N 1PF, United Kingdom*

³ *Department of Speech, Hearing & Phonetic Sciences, University College London, Chandler House, 2 Wakefield Street, London WC1N 1PF, United Kingdom*

⁴ *Department of Education, University of Oxford, 15 Norham Gardens, Oxford OX2 6PY, United Kingdom*

Running title: Variability in phonetic training for children

Corresponding author: Gwen Brekelmans, g.brekelmans@qmul.ac.uk

[This preprinted manuscript is currently under review].

ABSTRACT

Acquiring non-native speech contrasts can be difficult. In High Variability Phonetic Training, a well-established paradigm used to train learners on non-native phoneme contrasts, learners hear multiple (High Variability – HV) rather than single talkers (Low Variability – LV). While many studies using this paradigm test adults, much less work has looked at children. The current study further investigates the effect of input variability on phonetic training for children, examining whether they show a HV benefit.

In a two-week training study, 58 Dutch 7-year-olds and 51 Dutch 11-year-olds were trained on English vowel contrasts using either HV or LV input. Both groups improved in training, however, only 11-year-olds showed evidence of generalization, with Bayesian evidence for the null for 7-year-olds. Critically, there was no evidence that 11-year-olds showed greater generalization following HV input. Results suggest task demands might interact with input variability, so different age groups might benefit from different L2 learning strategies.

INTRODUCTION

Acquiring non-native speech sounds can be challenging; however, a body of research suggests that learners can benefit from phonetic training i.e., targeted practice in distinguishing minimal pairs, with feedback. A seminal study by Logan et al., (1991) trained Japanese learners of English on /l/-/r/, using natural speech to exemplify the contrast in multiple phonetic contexts with tokens spoken by multiple speakers. Participants were tested on their ability to identify /r/ and /l/ both pre- and post-training: they showed improvement both for trained stimuli and, critically, generalization items with novel voices and items. A follow-up experiment (Lively et al., 1993) demonstrated that generalisation to new talker stimuli was specifically dependent on training including stimuli spoken by multiple talkers.

This benefit of high variability stimuli for phonetic generalisation is theoretically and intuitively appealing: encountering variability helps the listener to recognise which acoustic cues are irrelevant and focus on key cues to identifying the phonemes. Since these seminal papers by Logan, Lively and colleagues, phonetic training with high variability materials (HPVT) has become a standard methodology in the field and the efficacy of such training has been replicated many times for different types of contrast in perception (e.g. *vowels*: Fuhrmeister & Myers, 2020; *consonants*: Sadakata & McQueen, 2013; *lexical tones*: Antoniou & Wong, 2015) and production (e.g. Bradlow et al., 1997), with effects shown to hold over time (e.g. Bradlow et al., 1999).

However, the specific advantage of high (HV) over low variability (LV) input in phonetic training is less well-established: It has only been reported in three additional studies. Wong (2012) trained Cantonese learners on English vowel contrasts for 10 HV or LV sessions, manipulating the amount of variation in both talker and phonetic context. Learners in both conditions improved after training, however those in the HV condition outperformed those in the LV condition. Additionally, they showed greater generalisation to novel talkers and items. Sadakata & McQueen (2013) trained Dutch adults on a Japanese geminate consonant contrast and

manipulated talker and context in 5 HV or LV training sessions. Change from pre- to post-test in an identification task indicated more learning and generalisation to new items and talker after HV training than LV training, although there was no such difference between HV and LV training in a discrimination task. In production, Kartushina & Martin (2019) used articulatory training to investigate whether Spanish learners benefitted from HV when acquiring a French vowel contrast. Although production accuracy improved for both HV and LV, only HV participants showed generalisation when repeating items spoken in a novel voice. Other studies, however, have failed to find an overall high variability advantage: Perrachione et al. (2011) and Sadatoka & McQueen (2014) find HV input detrimental for learners of lower perceptual aptitude, while Zhang et al., (2018), Dong et al., (2019), Wiener et al., (2020); and Deng et al., (2018) find no variability differences. A recent large scale (N=166) replication of the Logan et al. (1991) study showed that participants showed substantial learning and generalization after both HV and LV training, with no evidence for a benefit of multiple talkers on generalisation (Brekelmans et al., 2022).

Beyond second language (L2) speech learning, a benefit of variability for generalisation has been shown in related tasks training voice (Lavan et al., 2019) and dialect identification (Clopper & Pisoni, 2004). There is also substantial experimental evidence that encountering multiple talkers benefits vocabulary learning (see e.g., Sommers & Barcroft, 2011).

Most HVPT studies have employed adult participants, however some recent studies have included children. These have generally predicted that younger learners show improved learning from HVPT due to having greater neural plasticity and/or less well-established L1 categories. However, a complication is that *older* learners generally perform better at pre-test due, at least in part, to longer prior L2 experience [though Wang & Kuhl (2003) also found older learners to be better at differentiating wholly novel lexical tones]. Results are mixed for how much children benefit from training (i.e. improve across training sessions or from pre- to post-test). Some studies find similar amounts of improvement in perceptual performance across adults and

children (e.g. Wang & Kuhl, 2003; Giannakopoulou et al., 2017; Heeren & Schouten, 2010).

Others have found age-based differences: Giannakopoulou et al., (2010; 2013) found greater pre- to post-improvement in Greek-speaking 7/8-year-olds than in adults trained on English /i:/- /ɪ/, which they attributed to greater plasticity. Shinohara and Iverson (2021) found a complex pattern of age differences when comparing Japanese speaking 6/8-year-olds, 8/12-year-olds, 15/18-year-olds, and adults trained on the English /r/-/l/ contrast. Consistent with a plasticity benefit, 8/12 and 15/18-year-olds improved more than adults. However, the youngest group did not improve at all. The authors suggest that 6/8-year-olds may still be developing *phonemic awareness*, meaning they have difficulty hearing the different sounds within the words, and that they may be less familiar with the concept that letters can consistently map to those sounds. This indicates that learning may differ in children and adults for different tasks, and in particular, that children may find it harder to benefit from orthography-based training.

While HVPT has been shown to be generally beneficial, the degree of improvement is, however, not necessarily uniform even for adults. Two studies – Perrachione et al. (2011) and Sadakata & McQueen (2014) – found that for phonetic training of lexical tones, HV input was only beneficial for generalisation to novel talkers for learners who had stronger perceptual abilities. This could be due to a trade-off between the benefits of variability, and the greater initial processing cost adapting to multi-talker input. This burden could be greater in learners with weaker perceptual abilities, offsetting any variability benefit for generalization. This would be in line with literature showing that adapting to multiple talkers in L1 speech processing is initially detrimental to word recall and processing (Martin et al., 1989; Mullennix et al., 1989). Note however, that such effects of aptitude were not replicated in similar studies on adults by Dong et al. (2019) and Wong (2014). Dong et al. (2019) not only found no overall variability benefit but also did not find any evidence of an interaction with individual aptitude. Similarly, Wong (2014) found no modulating effect of learner proficiency in her study training Cantonese speakers on English /e/- /æ/, although in Wong's study *both* groups (rather than *neither*) showed an effect of variability.

Nevertheless, this literature suggests that we should not assume variability benefits will hold across learner groups.

Given the mixed results with regard to the efficacy of HVPT in children, an important question is whether this paradigm is well-suited to younger learners. There are theoretical reasons why variability should boost learning of general phonemic categories, and there is some (mixed) evidence for this in adults, but is this also true for younger learners? So far, only two previous studies have directly compared HV/LV phonetic training, with contradictory results. Evans & Martín-Alvarez (2016) trained 44 Spanish 9/12-year-olds on an English vowel contrast in either an HV (four talkers) or LV (one talker) condition using a picture identification task, testing improvement in discrimination and transfer to production. Only children trained on LV improved in production, while an HV benefit was found in vowel discrimination, although only in test items involving both a novel talker *and* a novel item. In contrast, Giannakopoulou et al., (2017) trained Greek-speaking 8-year-olds and adults using the same stimuli and similar methods (though not testing production) and did *not* find an HV benefit in the discrimination test for any generalization items for either age-group. Instead, while both groups improved in discrimination, adults improved to the same extent in HV and LV while children showed a reverse LV benefit even for generalisation to novel talkers. For adults, the authors noted, that an HV benefit could be masked by near-ceiling performance. For children, they suggested that the LV benefit could be due to greater difficulty in adapting to the multi-talker HV training material (cf. Perrachione et al., 2011; Sadakata & McQueen, 2014), perhaps exacerbated by lower working memory capacity (e.g. Alloway, Gathercole, & Pickering, 2006).

While the above discussion focussed on generalisation after training, many of these studies also investigated the difference between HV and LV input *during* training itself (i.e., tracking trial-by-trial performance). Studies often find an LV advantage in both adults and children (Perrachione et al., 2011; Dong et al., 2019; Giannakopoulou et al., 2017; Evans & Martín-Alvarez, 2016), indicating greater ease in processing single-talker input. In another, single-session, training

study comparing HV and LV materials for L2 vocabulary learning, Sinkevičiūtė et al., (2019) show a difference in adult and child learning. In this study, English speaking adults, 7/8-year-olds and 10/11-year-olds, all without previous knowledge of Lithuanian, completed a single L2 vocabulary training session learning 12 Lithuanian words using a picture identification game, with words produced by either a single-talker (LV), or multiple talkers (HV). Adults showed the predicted HV benefit in production recall (cf. Barcroft & Sommers, 2005; Sommers & Barcroft, 2011), however neither of the child groups showed evidence of a HV benefit in either production or comprehension. On the other hand, 7/8-year-olds – but not adults – *did* show a difference between variability conditions during training itself; specifically they showed a *low* variability advantage. Importantly, this study used Bayes Factor analyses and demonstrated that – at least in some of the tests – there was evidence for the null, i.e. *against* a high variability effect at test for both adults and children.

Given the clear results showing improvements following HVPT in adults and the mixed results with regard to the efficacy of HVPT in children, further work is required to assess whether HVPT is effective in training the perception of L2 contrasts for younger learners, and whether it is specifically advantageous over LV training. The current study aimed to investigate this.

Participants were two groups of native Dutch primary school children: 7/8-year-olds and 11/12-year-olds. This older age-group was chosen in light of the already strong knowledge of English in Dutch adolescents and adults¹, and also allowed us to look at differences in age groups whilst recruiting children from within the same school (11/12-year-olds are the oldest age group in Dutch primary schools). Using a two-week training paradigm, children were taught four English vowel contrasts known to be hard for Dutch learners. Half of the participants were trained on multi-talker (HV) and half on single-talker (LV) input. Data was collected during training itself, and from a battery of pre/post generalization tests which all tested generalisation to novel talkers. In addition, we also aimed to look for age-based differences both in terms of overall improvement resulting from training, and effects of variability. We tested a set of pre-registered

hypotheses based on previous literature (see Table 1). Note that in some cases we test two contrasting hypotheses, based on different literature.

Prediction	Supporting literature
Improvement in both age groups	Giannakopoulou et al., 2017; Evans & Martín-Alvares, 2016; Shinohara, 2014; Wang & Kuhl, 2003; Giannakopoulou, Uther, & Ylinen, 2010, 2013; Heeren & Schouten, 2010
a) older children will improve more than younger children	a) Shinohara, 2014 (comparison 6/7-year-olds versus 8/18-year-olds)
b) younger children will improve more than older children	b) Giannakopoulou, 2013 (comparison 7/8-year-olds versus adults)
There will be strongest performance during training in the LV group.	Giannakopoulou et al., 2017; Evans & Martín-Alvarez, 2016 (training data with adults and 7/8-year-olds)
a) Generalization (i.e., improvement from pre- to post-test) will be greater in HV condition.	Barcroft & Sommers, 2005 (vocabulary); Lively et al., 1993
b) Generalization (i.e., improvement from pre- to post-test) will be greater in LV condition.	Giannakopoulou et al., 2017 (result in discrimination test with 7-year-olds)
Any effect of HV will be stronger in older than younger children	Giannakopoulou et al., 2017

Any effect of HV will be stronger in novel
items

Evans & Martín-Alvarez, 2016

Table 1. Pre-registered hypotheses and supporting literature.

Where appropriate, we also tested whether overall performance was above chance. The study was preregistered² at https://osf.io/bgdxp/?view_only=209600aca2cf46a89f0dfd402ab46769.

Where analyses deviated from the pre-registration, this is explained.

METHODS

Participants

Participants were recruited from two year groups in a school in the south of the Netherlands: *Groep 4* (7/8-years-olds), and *Groep 8* (11/12-year-olds). Opt-out parental consent was used; children could decide not to participate in a session at any time, and parents/carers could withdraw them at any point.

Each child participated in 10 sessions, each on separate days over a maximum of 12 consecutive school days, receiving daily stickers and a certificate on completion. Our preregistered target sample was 48 7/8-year-olds and 48 11/12-year-olds. However, as anticipated, our actual sample was affected by practicalities of testing in school where whole classes wished to participate. In fact, a total of 109 children was tested: 51 11/12-year-olds (29 HV, 29 LV) and 58 7/8-year-olds (25 HV, 26 LV). In line with the preregistration, data from 20 of these were excluded due to: not completing at least 60% of training; having dyslexia/hearing impairments/other language disorders; having English-speaking parents/carers; having lived in an English-speaking country for more than two months. This left 50 7/8-year-olds (mean age = 7;6, SD = 4 months, range 6;10-8;3 years) and 39 11/12-year-olds (mean age = 11;7, SD = 4 months, range 10;9-12;0 years). Additionally, for each task, participants were excluded from

analyses if either their pre-test or post-test data was missing. Full details on participants' language background and participant numbers included in each task can be found on OSF. Both age groups had relatively little exposure to English. The two age groups differed slightly in previous English exposure in school: 7-8 year-olds had not yet started English lessons, while 11/12-year-olds had received one English lesson per week for about a school year (i.e. 40 weeks). Crucially, these English lessons were focussed on communicative content rather than pronunciation, and were taught by native Dutch speakers who had no specific pronunciation training and will likely not have contrasted the phonemes of interest. Additionally, children reported their English exposure outside of school was generally low, and while 11/12-year-olds had a little more exposure to English than 7/8-year-olds, this was matched for the variability conditions within the groups (see Supplementary information A2 on OSF).

Stimuli

Stimuli were recorded by 7 native speakers of Standard Southern British English in a sound-attenuated booth at a sampling rate of 44.1 kHz 16-bit samples per second. They were later downsampled to 22.05kHz and band-pass filtered from 60-20,000 Hz with a smoothing factor of 10. Finally, intensity was scaled to 70 dB. All processing was carried out in Praat (Boersma & Weenink, 2015). Stimuli were presented to speakers using ProRec (Huckvale, 2016) and can be found on OSF: https://osf.io/bgdxp/?view_only=209600aca2cf46a89f0dfd402ab46769.

Trained stimuli

Trained stimuli consisted of 16 monosyllabic minimal pairs. All were English CVC-words with pairs differing in the vowel. Consonants were chosen such that they were present in the consonant inventory of both Dutch and English. There were 3 vowel contrasts which were expected to be difficult (Collins & Mees, 2003): GOOSE-FOOT /u:/-/ʊ/, STRUT-LOT /ʌ/-/ɒ/, DRESS-TRAP /e/-/æ/, and a control contrast FLEECE-THOUGHT /i:/-/ɔ:/ which should be easy for Dutch learners to distinguish without training as it has a close native counterpart. The latter provides a

way to check whether participants understood the task, and to boost confidence/task appeal by having easier items.

Each vowel contrast was exemplified by 4 minimal pair items (8 words), resulting in a total of 32 trained items, or 16 trained pairs (full stimuli list in Supplementary information A3 on OSF).

Words were chosen on the basis of imageability and, as far as possible, transparency and consistency of orthography (since pictures and written representations were used in training). It was not possible to control for participants' initial familiarity with the words due to the limited number of feasible minimal pairs for the contrasts of interest. However, all words were deemed appropriate for primary school children, and all concepts were deemed familiar to both age groups after initial pilot tests.

Novel stimuli

In addition to the 32 trained items, several of the pre/post-test tasks included novel items to test generalisation (see Supplementary information A3 on OSF). The Discrimination task, in addition to the trained items, also used 32 novel items (16 pairs). All items used were real English minimal pair CVC words differing only in the vowel. The Orthography Identification task used 32 novel items (16 pairs), all of which were minimal pairs of English pseudowords differing in the vowel only. Since English orthography is known to be opaque, the novel pseudoword items were created to be as transparent as possible and kept vowel orthography identical to that used in the trained items.

Design

The study comprised three stages: pre-test, training, and post-test. The pre-test and post-test were identical and consisted of a Discrimination, Vocabulary, Picture Identification, and Orthography Identification task. Both vowel discrimination and identification were tested, since differences between these are sometimes reported (e.g. Heeren & Schouten, 2010). In addition, we had separate identification tests using orthography and pictures, bearing in mind the potential challenges of orthography for younger learners. We also included a vocabulary test in light of

variability benefits reported in the vocabulary learning literature (e.g. Sinkevičiūtė et al., 2019). Talker variability was manipulated between-participants with half assigned to HV and half to LV training. Talkers were drawn from a total of 7 native speakers of Standard Southern British English and assignment of talkers was counterbalanced.

In HV training, participants heard four different speakers (two male, two female) while in LV training stimuli were spoken by a single talker (one of the four used in HV, with voices rotated across versions to ensure all talkers were equally represented). In the HV condition, stimuli from the four speakers were presented in separate blocks, as trial-by-trial speaker alteration has been shown to be detrimental to training performance in phonetic training paradigms (e.g. Perrachione et al., 2011). Speaker order was counterbalanced across blocks. Different talkers were used for pre- and post-tests to measure generalisation to novel voices. Pre- and post-tests were identical across conditions, and all involved novel speakers (three females, rotated between tasks) who had not been encountered during Training. Full counterbalancing details are in Supplementary information A4 on OSF.

Procedure

The study was run in Gorilla (www.gorilla.sc, Anwyl-Irvine et al., 2019) on 30 Dell Vostro 15 3000 series laptops, and each child had their own in-ear headphones. Testing took place in the children's classroom, administered in full class groups by the first author. Children first completed the pre-test the day before their first training session. After the pre-test, children undertook 8 training sessions on separate days. Multiple vowel contrasts were used since training on a more extended vowel set has been shown to be beneficial (Nishi & Kewley-Port, 2007; 2008). Training involved a minimal pair identification task using both pictures and orthography. The post-test was completed the day after the last training session, ensuring that any temporary phoneme categories built up during training were no longer maintained (Heeren & Schouten, 2010). Figure 1 gives an overview of the procedure. See Supplementary information A1 on OSF for screenshots of all tasks.

Figure 1. Overview of the Training and pre/post-tests procedure, including the number and type of experimental trials per task. Task numbering shows the order in which tasks were conducted. Pre- and post-test were administered in a separate session.

Training

The training task was a two-alternative forced-choice (2AFC) identification task. Training consisted of 8 sessions, 4 blocks each, with each block comprising 64 trials – twice with each of the 32 training words as target – spoken by a single talker. In total, 256 trials were presented per session. Trial order was randomised within blocks on a participant-by-participant basis. In each trial, participants heard a word and viewed two pictures, with accompanying written words, representing the target and foil (i.e., the minimal pair of the target). The task was to select the correct word. They then received feedback and a reward (an animated coin), and were shown the correct response.

Pre/post-tests

Discrimination task

Discrimination was measured using a three-alternative forced-choice (3AFC) oddity task. In each trial, three animated figures appeared on the screen; a speech bubble appeared when the corresponding word was played. Participants heard three words (ISI 200ms), two of which were different tokens of the same word (e.g., *bed*, *bed*), and one of which was a minimal pair foil containing the other vowel from the trained contrast (e.g., *bad*). After all words had played, participants indicated which was the odd one out by clicking on the corresponding speech bubble. Each token was spoken by a different talker.

There were 16 trained and 16 untrained minimal pairs. Each pair occurred once with each member of the pair as target (64 trials in total). Trial order was pseudo-randomised. No feedback was provided.

2AFC Orthography identification task

This was like the training task except that (1) only written words but no pictures were used (2) there was no feedback. The task comprised 128 trials: 32 minimal pairs from training, each occurring once as the target; and 32 novel minimal pairs each occurring once as the target. Presentation of trained and novel items was blocked, with a randomised order within blocks. All items were spoken by a novel speaker. No feedback was provided.

2AFC Picture identification task

This task was like the training task except that (1) only pictures but no written words were used (2) there was no feedback. There were 128 trials: 32 minimal pairs used in training each occurring once as the target; 32 non-minimal pair combinations that differed in more than just the vowel, each occurring once as the target. Presentation of minimal pairs and non-minimal pairs was blocked, with the order randomised within blocks. Note that (1) it is not possible to include novel *items* in this test, since it depends on knowledge of picture/item correspondences, and (2) including the non-minimal pair items provided an additional test of *vocabulary knowledge* without requiring an ability to discriminate the vowels. No feedback was provided.

Analyses and statistical approach

Data were analysed using both Frequentist and Bayesian statistics.

Model building

Data were analysed using generalised logistic mixed effects models (glmer) (Baayen et al., 2008) using the lme4 package (Bates et al., 2015) for the R environment (R Core Team, 2018). The plyr (Wickham, 2011), knitr (Xie, 2018) and kableExtra packages (Zhu, 2018) were used to present the results. Data was visualised using ggplot2 (Wickham, 2016) and ggpirate (Braginsky, 2018).

The approach was to include all experimentally manipulated variables and their interactions for each task as fixed factors in the models, regardless of whether they contributed to the model. Key fixed effects were *variability* (HV/LV), *session* (1-8 for training, pre/post for all other tasks), and *item novelty* (in tasks where this was manipulated, i.e., trained/novel for Discrimination and Orthography Identification, minimal pair/non-pair for Picture Identification) and all the interactions between them. As preregistered, the two age groups were first analysed separately, followed by targeted comparisons including *age group* as a variable.

The approach to random effects was to automatically include *participant* as a random effect with a full random slope structure, or the most complex structure to converge, as recommended by Barr et al., (2013). For *item*, random intercepts were automatically included, however by-item slopes were only included if they contributed to the model following the approach recommended in Matuschek et al., (2017). The final model structure detailing the random effects structure and included control variables is reported for each model.

A centred coding was used for predicting variables to reduce collinearity effects between main effects and interactions, and so that the intercept corresponded to the grand-mean (in log-odds space) and main effects were evaluated as average effects over all levels of other factors. All reported models converged using the Bound Optimization by Quadratic Approximation (BOBYQA optimization: Powell, 2009).

In addition to the experimental factors, factors to control for *vowel contrast* (DRESS-TRAP, GOOSE-FOOT, STRUT-LOT, and FLEECE-THOUGHT) and *talker* (4 talkers in training; 3 in the tests) were included in the analyses as it was hypothesized that these factors would likely contribute variance to the data (different contrasts are different perceptually, and talker idiosyncrasies contribute to difficulty). Both are included as fixed in the model and are coded as a set of centred contrasts (3 for vowel contrast³; for talker – 4 for training and 3 for the tests) so that the intercept continues to represent the grand mean. Neither were included in any interactions. Conceptually, *talker* is a random effect, but this cannot have fewer than 6 levels (Bolker, 2019) so is entered as a fixed effect. *Talker* was only included when it contributed to the model fit (with $p > .2$ similar to what is suggested for random effects by Matuschek et al., 2017). When included, it was not interpreted and not included in interactions. Full model structures for each analysis are reported on OSF.

Inference criteria - Frequentist and Bayesian analyses

Although models included all experimentally manipulated variables and their interactions, we only interpret and report statistics for fixed effects relevant to our hypotheses. The key inferential statistic used to evaluate these hypotheses is the Bayes Factor (BF), which computes the strength of evidence for the hypothesis (H1) over the null hypothesis (H0), or vice versa. This provides a measure of the strength of evidence for a hypothesis, compared with the null. These differ from p-values, which do not provide the option to evaluate the strength for the H0 ($p > .05$ does not tell us that we should increase our confidence in the null hypothesis, despite this common misinterpretation). Since Bayes Factors might be less familiar, p-values [automatically computed from z-values by the *lme4* package (Bates et al., 2015)] are additionally reported. Bayes Factors were computed using the method advocated by Dienes (2008, 2015), modelling H1 as a half-normal⁴, testing a one-sided prediction in each case since we are testing directional hypotheses. This method of computing Bayes Factors requires three numbers: (i) an estimate of mean difference, (ii) the standard error – we take these from the beta and standard errors of the

relevant fixed effects in the models described above (allowing for working in log odds space, meeting assumptions of normality), and (iii) an estimate of x , the predicted mean difference under H_1 which is then used as the standard deviation of the half-normal modelling H_1 . The key issue is how to determine appropriate values of x . In some cases, this was done by using estimates taken from equivalent analyses in Giannakopoulou et al. (2017), which used a similar task design. Where this was not possible or appropriate, values from within the current data set were used⁵ wherever possible using values from independent participants, i.e., using the beta value for the equivalent fixed effect for 7/8-year-old children to inform that for 11/12-year-old children and vice versa. Our choice of values of H_1 for each test is laid out in Supplementary information A5 on OSF. Justifications are based on the logic laid out in Silvey et al., (2021). This was all as pre-registered except where noted in endnote 6). BF results were reported using the notation $BF_{H(0,x)}$ to denote a Bayes Factor where x is the standard deviation of the half normal used to model H_1 (following Dienes, see <https://osf.io/hzcv6/>).

Bayes factors were interpreted using the convention in Jeffreys (1961) and expanded by Dienes (2014), where $BF > 3$ indicates substantial evidence for H_1 , a BF between $1/3$ and 3 indicates ambiguous evidence (the data is insensitive to test the hypothesis), and $BF < 1/3$ indicates substantial evidence for H_0 . Since there is subjectivity in the choice of value to inform H_1 , robustness regions are included, showing the range of estimates of H_1 (i.e., the value used as the standard deviation of the half normal) for which our data would support the same conclusion (e.g., if the reported effect is substantial, the range of estimates would also yield a $BF > 3$). The larger the robustness region, the wider the range of values of H_1 estimates that would support the same conclusion. This shows that our conclusions are robust even if there are small but reasonable differences in the potential assumptions we used to inform the H_1 estimates. This robustness region is noted as $[x_1, x_2]$ with x_1 being the smallest standard deviation and x_2 the largest standard deviation (as recommended by Dienes, 2017). To compute the ranges, values were tested in increments of 0.01 from a difference of 0 from chance to the log-odds

score corresponding to the difference between chance and ceiling performance (deemed to be all but one trial correct, i.e., 99.6% (log-odds 5.5156) in Training; 98.4% (log-odds 4.8299) in Discrimination; 99.2% (log-odds 4.84) in Orthography Identification; 98.4% (log-odds 4.1425) in Picture Identification;). Note that there will always be some value which provides evidence for H0. Where this was not found within the range tested, the end point of the tested range is noted as e.g., >4.8299 (Discrimination). For ranges of values already giving evidence for H0, the maximum is infinity. Robustness Regions should be interpreted bearing in mind that larger values of H1 bias evidence for the *null*, whereas smaller values bias in favour of H1.

RESULTS

Results below are presented per task and separately for each age group, starting with initial performance throughout training, and followed by the pre/post-test generalisation tasks⁶

Training

Figure 2. Response accuracy during Training. Error bars indicate 95% CI; the dashed line indicates chance. Note that the dip in improvement visible at Day 6 occurred after a weekend.

7/8-year-olds

As can be seen in Figure 2, there was substantial evidence that 7/8-year-olds performed above chance in the Training task ($\beta = 0.806$, $SE = 0.089$, $z = 9.076$, $p < .001$; $BF_{H(0,2.362)} = 5.45 * 10^{16}$ [RR 0, >5.516]) and that they improved across sessions ($\beta = 0.033$, $SE = 0.010$, $z = 3.370$, $p = .001$; $BF_{H(0,0.166)} = 33.719$ [RR 0, >5.516]). Note that the dip in improvement at Day 6 occurred after a weekend. However, the evidence for LV outperforming HV overall ($\beta = 0.139$, $SE = 0.145$, $z = 0.957$, $p = .339$; $BF_{H(0,0.648)} = 0.557$ [RR 0, 1.116]) and for there being more

improvement in LV than in HV ($\beta = 0.032$, $SE = 0.020$, $z = 1.637$, $p = .102$; $BF_{H(0,0.103)} = 1.3$ [RR 0, 0.316]) was ambiguous.

11/12-year-olds

As seen in Figure 2, there was substantial evidence that 11/12-year-olds performed above chance in the Training task ($\beta = 2.362$, $SE = 0.183$, $z = 12.907$, $p < .001$; $BF_{H(0,0.806)} = 1.11 \times 10^{34}$ [RR 0, >5.516]), that they improved across sessions ($\beta = 0.166$, $SE = 0.038$, $z = 4.378$, $p < .001$; $BF_{H(0,0.033)} = 92.972$ [RR 0, >5.516]), and that LV outperformed HV overall ($\beta = 0.648$, $SE = 0.215$, $z = 3.012$, $p = .003$; $BF_{H(0,1.626)} = 22.596$ [RR 0.1156, >5.516]). The evidence for there being more improvement in LV than in HV was ambiguous ($\beta = 0.039$, $SE = 0.076$, $z = 0.513$, $p = .608$; $BF_{H(0,0.103)} = 0.854$ [RR 0, 0.316]).

Age-group comparisons

There was substantial evidence that 11/12-year-olds outperformed 7/8-year-olds overall across Training ($\beta = 1.370$, $SE = 0.160$, $z = 8.537$, $p < .001$; $BF_{H(0,1.998)} = 8.50 \times 10^{14}$ [RR 0.1156, >5.516]) and that they showed greater improvement across training sessions ($\beta = 0.103$, $SE = 0.023$, $z = 4.494$, $p < .001$; $BF_{H(0,0.083)} = 6328.34$ [RR 0, >5.516])⁷

Discrimination

Figure 3. Response accuracy in the Discrimination task. Violins show kernel density estimation; points are individual participants. Mean performance is shown by the horizontal line with the band around it showing 95% CI; dashed line indicates chance. ‘Pre’ indicates scores from pre-test, while ‘post’ indicates scores from post-test. Data marked ‘Train’ used trained items, while data marked ‘new’ used novel items.

7/8-year-olds

As seen in Figure 3, there was substantial evidence that 7/8-year-olds performed above chance in the Discrimination task ($\beta = 0.833$, $SE = 0.096$, $z = 8.686$, $p < .001$; $BF_{H(0,1.689)} = 2.43 \times 10^{15}$ [RR 0, >4.830]), but substantial evidence that they did not improve overall from pre- to post-test ($\beta = -0.154$, $SE = 0.055$, $z = -2.785$, $p = .005$; $BF_{H(0,0.219)} = 0.063$ [RR 0.130, >4.8299]). Since there was no effect of session, no further interactions with session are reported.

11/12-year-olds

As seen in Figure 3, there was substantial evidence that 11/12-year-olds both performed above chance in the Discrimination task ($\beta = 1.689$, $SE = 0.133$, $z = 12.680$, $p < .001$; $BF_{H(0,0.833)} = 3.49 \times 10^{33}$ [RR 0, >4.830]) and improved from pre- to post-test ($\beta = 0.219$, $SE = 0.070$, $z = 3.143$, $p = .002$; $BF_{H(0,1.689)} = 11.441$ [RR 0, >4.830]). The evidence for there being more improvement in HV than in LV ($\beta = 0.043$, $SE = 0.141$, $z = 0.306$, $p = .760$; $BF_{H(0,0.219)} = 0.672$ [RR 0.030, 0.430]), more improvement in LV than in HV ($\beta = -0.043$, $SE = 0.141$, $z = -0.306$, $p = .760$; $BF_{H(0,0.219)} = 0.445$ [RR 0.030, 0.230]), and greater improvement for HV than LV for novel items only ($\beta = 0.020$, $SE = 0.303$, $z = 0.067$, $p = .946$; $BF_{H(0,0.219)} = 0.835$ [RR 0.030, 0.830]) was all ambiguous.

Orthography Identification

Figure 4. Response accuracy in the Orthography Identification task. Violins show kernel density estimation; points are individual participants. Mean performance is shown by the horizontal line with the band around it showing 95% CI; the dashed line indicates chance. 'Pre' indicates scores

from pre-test, while 'post' indicates scores from post-test. Data marked 'Train' used trained items, while data marked 'new' used novel items.

7/8-year-olds

As seen in Figure 4, there was substantial evidence that 7/8-year-olds performed above chance in the Orthography Identification task ($\beta = 0.228$, $SE = 0.036$, $z = 6.387$, $p < .001$; $BF_{H(0,0.978)} = 5.12 \times 10^7$ [RR 0, >4.840]), but substantial evidence that they did *not* improve from pre- to post-test ($\beta = -0.048$, $SE = 0.038$, $z = -1.273$, $p = .203$; $BF_{H(0,0.228)} = 0.073$ [RR 0.05, >4.840]). Since there was no effect of session, no further interactions with session are reported.

11/12-year olds

As seen in Figure 4, there was substantial evidence that 11/12-year-olds performed above chance in Orthography Identification ($\beta = 0.978$, $SE = 0.109$, $z = 9.012$, $p < .001$; $BF_{H(0,0.228)} = 2.04 \times 10^{14}$ [RR 0, >4.840]), but substantial evidence that they did *not* improve from pre- to post-test ($\beta = 0.035$, $SE = 0.048$, $z = 0.726$, $p = .468$; $BF_{H(0,0.978)} = 0.099$ [RR 0.28, >4.840]). Since there was no effect of session, no further interactions with session are reported.

Picture Identification: Minimal Pair Items⁸

Figure 5. Response accuracy in the Picture Identification task. Violins show kernel density estimation; points are individual participants. Mean performance is shown by the horizontal line with the band around it showing 95% CI; the dashed line indicates chance. ‘Pre’ indicates scores from pre-test, while ‘post’ indicates scores from post-test. ‘MP’ are the minimal pair items, while ‘NP’ are the non-pair items.

7/8-year-olds

As seen in Figure 5, there was substantial evidence that 7/8-year-olds performed above chance on the minimal pair items ($\beta = 0.224$, $SE = 0.050$, $z = 4.507$, $p < .001$; $BF_{H(0,1.373)} = 1836.957$ [RR 0, >4.143]) but ambiguous evidence for improvement from pre- to post-test ($\beta = 0.140$, $SE = 0.075$, $z = 1.877$, $p = .061$; $BF_{H(0,0.767)} = 1.075$ [RR 0.223, 2.463]). Since there was not substantial evidence for an effect of session, no further interactions with session are reported.

11/12-year-olds

As seen in Figure 5, there was substantial evidence that 11/12-year-olds performed above chance on the non-pair items ($\beta = 1.373$, $SE = 0.137$, $z = 10.003$, $p < .001$; $BF_{H(0,0.224)} = 5.46 \times 10^{15}$ [RR 0, >4.143]), and improved from pre- to post-test ($\beta = 0.767$, $SE = 0.115$, $z = 76.663$, $p < .001$; $BF_{H(0,0.14)} = 7.05 \times 10^5$ [RR 0, >4.143]). There was substantial evidence against there being more improvement after HV than LV training ($\beta = -0.081$, $SE = 0.219$, $z = -0.370$, $p = .712$; $BF_{H(0,0.767)} = 0.214$ [RR 0.463, >4.143]), and ambiguous evidence as to there being more improvement in LV than HV training ($\beta = 0.081$, $SE = 0.219$, $z = 0.370$, $p = .712$; $BF_{H(0,0.767)} = 0.376$ [RR 0, 0.873]).

Picture Identification: Non-Pair Items (testing vocabulary knowledge)

7/8-year-olds

As seen in Figure 5 there was substantial evidence that 7/8-year-olds performed above chance on the non-pair items ($\beta = 0.572$, $SE = 0.079$, $z = 7.253$, $p < .001$; $BF_{H(0,3.277)} = 1.25 \times 10^{10}$ [RR 0, >4.143]) but substantial evidence against pre- to post-test improvement ($\beta = -0.299$, $SE = 0.110$, $z = -2.720$, $p = .007$; $BF_{H(0,0.572)} = 0.053$ [RR 0.083, >4.143]). Since there was no effect of session, no further interactions with session are reported.

11/12-year-olds

As seen in Figure 5, there was substantial evidence that 11/12-year-olds performed above chance on the non-pair items ($\beta = 3.277$, $SE = 0.164$, $z = 19.950$, $p < .001$; $BF_{H(0,0.224)} = 2.14 \times 10^{42}$ [RR 0, >4.143]). There was substantial evidence against pre- to post-test improvement ($\beta = -0.067$, $SE = 0.284$, $z = -0.235$, $p = .815$; $BF_{H(0,3.277)} = 0.07$ [RR 0.663, >4.143]). Since there was no effect of session, no further interactions with session are reported.

DISCUSSION

The present study investigated the role of input variability in children acquiring L2 vowels via phonetic training. Our hypotheses concerned the role of input variability and age in the ability to

improve phonetic category identification both in training and on a range of identification and discrimination tasks conducted at pre- and post-test. Overall, this study found evidence that both age groups improved during the training task itself, although improvement was greater in 11/12-year-olds. However, although both age groups performed above chance in all of the generalisation test tasks, surprisingly 7/8-year-olds did *not* show substantial evidence of pre- to post- improvement in *any* test, with evidence for the null in every case but one. This does not mean that the 7/8-year-olds did not learn from training, as performance during training clearly increased. Instead, it means that the youngest children were unable to apply what they learnt to novel talkers in the pre/post-tests. In contrast, 11/12-year-olds improved in three out of five tasks. Since all tests involve novel voices, this improvement indicates generalisation. Importantly, there was no evidence that the amount of generalisation improvement for 11/12-year-olds was greater after HV training than LV training, with three tests providing evidence for the null for this hypothesis. In fact, the only task that showed any evidence in terms of a variability effect was Training itself, where 11/12-year-olds did better with LV than HV input, suggesting adaptation to the single talker in the input. In short, no HV benefit effects were found in any task. The following discussion will first focus on why no pre- to post improvement was found in any task with 7/8-year-olds and was only found in some tasks with 11/12-year-olds, and will consider the implications for our hypotheses about age-based differences in learning. We then turn to consider the findings with regards to the role of talker variability.

Learning in 7/8-year-olds

The most surprising finding in this data is that although the 7/8-year-olds improve during training, indicating learning, they do *not* show improvement - and thus generalisation - in any pre/post-test. Note that they are above chance *overall* (i.e., including at pre-test) on the Discrimination, Picture Identification and Orthography Identification tasks, indicating that, at least to some extent, they are able to distinguish the non-native vowels, and have some knowledge of the associations between English words and their meanings (the pictures), as well as some

knowledge of the mappings between the vowels and their orthography. However, there is *no* evidence to suggest that they improve in these when tested with generalization items, with evidence for the null in all cases except the minimal pair Picture Identification task - the task most like Training - where the evidence is ambiguous. This contrasts with previous studies where 7/8-year-olds did show pre- to post- improvement and generalisation after similar amounts of training (Giannakopoulou et al., 2013, 2017; Shinohara & Iverson, 2013).

One thing to consider in interpreting this lack of generalisation improvement is whether the task demands at test were too high for 7/8-year-olds. This could be the case for Orthography Identification as 7/8-year-olds are still developing readers and writers. However, task difficulty cannot explain why we see no pre- to post-test improvement in the Picture Identification with minimal pairs and Discrimination tasks, since these have been successfully used in other research with this age group (e.g., Giannakopoulou et al. 2017). An alternative is that the training task itself made it harder to transfer acquired knowledge to generalisation beyond the trained materials. The current task trained learners on four different vowel contrasts at the same time, whereas a majority of studies using the HVPT paradigm, and all those conducted with children, trained learners on a single phoneme contrast. Although there is evidence that training multiple contrasts is helpful for adult learners, (Nishi & Kewley-Port, 2007), it is possible that this may increase complexity in a way that is detrimental for child learners. In this particular case, it seems likely that the 7/8-year-olds were unable to generalise what they had learnt during training to novel talkers and items as the knowledge that was being acquired in training itself was not yet sufficiently developed. Further consolidation through extended training sessions or training on fewer contrasts so that the overall amount of training per contrast increased might have resulted in increased generalisation for our 7/8-year-olds.

Learning in 11/12-year-olds

This age group showed pre- to post-test improvement in most tasks with the exception of the non-minimal pair Picture Identification and Orthography Identification tasks where there was

evidence for the null. For the non-minimal pair Picture Identification task, this lack of improvement looks likely to be a result of ceiling effects (especially since this is primarily a measure of vocabulary learning, and improvement is found for this age group in the Vocabulary test). However, the lack of improvement for the Orthography Identification task suggests that 11/12-year-olds may not actually use the orthography presented during training. If so, this contrasts with existing findings showing that for Dutch adult learners of English, spelling is especially helpful initially when learning words relying on non-native phoneme contrasts (Escudero et al., 2008). Theoretically, it seems sensible that mapping to orthographic representations could help with developing phoneme categories, since they provide an invariant representation of those categories across instances. However, it is an open question whether this is actually helpful for children with less reading experience. Indeed, Shinohara & Iverson (2021) also found a lack of improvement in their youngest age group (6/8-year-olds) and attributed this to difficulty in using orthography in training. That said, it is possible that orthography was helpful in our study, even though children did not show learning when tested on orthography in isolation: it might have an additive benefit when presented in combination with the pictures. To shed light on this, future research could include a pre/post-test which *combines* orthography and pictures to check for better performance in comparison to picture/orthography-only tasks. On the other hand, the inclusion of orthography could actually *harm* learning, for example if orthography distracts from more useful cues. This could be tested by comparing the learning in the current study with a condition with pictures-only training, to check whether adding orthography actively *harms* learning.

Age-based differences in learning

It is clear that there was no “plasticity” effect: younger learners did not benefit more from training and did not show improvement during training. In contrast, we find clear evidence that not only do older children perform better at the start of training (likely as a result of their greater experience with English), but that they also learn more during training. This contrasts with

previous findings that have shown younger is better (e.g., Giannakopoulou et al., 2013).

Shinohara (2014) also found a plasticity benefit for 8/12-year-olds compared to adults, however 6/8-year-olds did not show any pre- to post-test improvement, a pattern similar to the current study. Older learners outperforming younger learners is not unprecedented in language learning research: Huyck & Wright (2010) found that adolescents showed much less, or even no, improvement in auditory perceptual learning tasks compared to adults who did improve, despite all adolescents starting out within the adult range of scores at the start of training. Moreover, Snow & Hoefnagel-Höhle (1978) found that younger learners consistently showed less improvement in various areas of natural L2 learning compared to older children and adults. This suggests that there may be more to the plasticity benefit, and that it may depend on more than just age factors.

It is possible that any inherent plasticity benefit may have been offset by task difficulties. We also cannot rule out that our 7/8-year-olds might catch up and eventually overtake 11/12-year-olds given more training (note that even after 8 training sessions, their performance did not reach the starting point of the 11/12-year-olds in this study). On the other hand, it is notable that although starting younger appears to be beneficial for long-term phonetic learning in naturalistic environments (Flege, 1995), there is actually relatively little direct experimental evidence in the literature that younger is better for phonetic training. One possibility is that older children benefit more from training because they have a better starting point: if phonetic training leads to more accurate mappings from an existing L1 to an L2 category (cf. Iverson & Evans, 2009), then young children's L1 categories might not be well-developed enough to support learning in this way. As children grow older and their L1 categories become more stable, their learning strategies might change to more closely resemble those of adults such that they show a benefit from phonetic training.

Another possibility relates to differences in the amount of exposure to English between the two age groups. If the younger children start out having less exposure than the older children, then

this may influence how well they do in certain tasks. While older children might be finetuning their knowledge of the words, the younger children with no prior knowledge to access, would be unable to do so. Note that prior knowledge of the words in and of itself would not necessarily be helpful. Given that the children are in a Dutch context, part of their exposure to English will have been in a Dutch-accented context where the phoneme contrasts of interest might not have been correctly or consistently distinguished. This could potentially be detrimental if the mental representation of the words were inaccurate for the vowel contrasts in question. Theoretically, this could put older learners at a disadvantage as they would need to adjust fixed representations. In contrast, younger learners, without pre-existing lexical representations, could start “from scratch”. One possibility in the current study is that the relatively small amount of learning and familiarity in the 11/12-year-olds is enough to make the learning tasks easier but has not led to inaccurate representations which are too entrenched.

On the other hand, we also point out that there is a relatively small difference in the amount of English learning by our participants - although 7/8-year-olds had no formal training, 11/12-year-olds had had only one session a week in basic English for a year. In many published studies the gap is greater, however our reading of the literature is that there is no clear pattern as to whether learners with more L2 experience show greater or lesser improvement. For example, Giannakopoulou et al. (2017) found that adults with on average over 10 years of exposure showed less improvement than children with 1-2 years English education, but Shinohara & Iverson (2021) found that adults with many years of exposure through schooling improved more than 8/10-year-olds with no formal English training, but less than 16/18-year-olds with four to six years. It is thus difficult to draw strong conclusions about the role of previous L2 exposure or to pull this apart from age at the time of testing.

When considering the impact of possible age-based differences in the phonetic training paradigm, it is worth considering that others might have found no younger-is-better effect, but that due to the file drawer problem these null findings have not been published (Franco et al.,

2014). In addition, traditional statistical methods make it difficult to interpret null effects where these are reported [e.g., Heeren & Schouten (2010) who found no effect of age on learning]. The current work makes a contribution in using Bayes factors to clearly show evidence *against* a “younger is better” advantage for learning in training. While the study aimed to investigate similar effects in the pre/post-test tasks, the lack of improvement across these tasks for the younger age group means this comparison would not have been statistically informative beyond confirming that one group did indeed not learn at all.. Future studies should compare different age-groups using a variety of training and test tasks, to explore if and when a plasticity effect is apparent, and whether or not this is affected by the types of task and materials.

Variability

There are theoretical reasons to expect a variability benefit in the current paradigm, since encountering variability should help listeners to learn to ignore speaker-specific cues and tune in to the relevant cues for the phonetic contrasts. This should help participants discriminate and identify the vowels in the novel-talker stimuli used in all of our tests. However, a high variability benefit was *not* found for 11/12-year-olds (the only age-group who showed any evidence of generalisation and where we can thus draw conclusions about variability) in the current study, with evidence for the null in several tests (Discrimination, Orthography Identification, Non-Pair Picture Identification). Note that due to a lack of pre-to-post improvement we did not test for a benefit specifically for test stimuli involving novel words (cf. Evans & Martín-Alvarez, 2016). One possible explanation for the lack of variability benefit is that the task demands were too high even for our 11/12-year-olds. Note that we saw strong evidence for a low variability benefit during training itself. One possibility is that the added complexity of dealing with multiple talkers in training, in addition to the complexity of training multiple vowels simultaneously as discussed above, combined to make the HV training task too cognitively demanding for the children and that this offset any potential benefit of input variability. This explanation is similar to that given in Sinkeviciute et al. (2019), where younger children also showed a LV benefit in (vocabulary)

training, and no HV advantage for generalization, in contrast to adults who showed an HV advantage for generalization. On the other hand, we must acknowledge that in the phonetic training paradigm HV benefits in generalization are not always found, even in adults (Brekelmans et al., 2022). We also note that although the current study found no evidence for an HV benefit, we also found no LV benefit for generalization (with evidence for the null in almost all tasks), failing to replicate Giannakopoulou et al. (2017), who found this in 7-year-olds in both training and generalization. Future research should further probe the trade-off between benefits of simplicity versus variability helping with generalisation in learners of different ages.

CONCLUSION

The current study trained native Dutch 7/8-year-olds and 11/12-year-olds on four English vowel contrasts using a two-week phonetic training paradigm that used both meaningful pictures and familiar orthography in training. While both age groups showed improvement during training, the 7/8-year-olds were not able to generalise this acquired knowledge to any of the post-test tasks using novel speakers. In contrast, 11/12-year-olds were able to generalise beyond the trained materials in most pre/post-tasks. Crucially, this study did not find that talker variability in the input stimuli had any benefit for generalization. This suggests that different ages groups might be using different strategies for L2 learning, and might benefit most from different types of input and approaches as a result.

NOTES

¹ A pilot study with these materials with 27 16-year-olds found that they were at ceiling even at pre-test.

² Note that the pre-registered plan also includes a production task, a phonological working memory test, and a category boundary task. These tasks are not reported as part of the current study due to data loss. A Vocabulary task was also pre-registered; for space reasons full methods and results for this task are reported on OSF.

³As preregistered, for each set of analyses, before any models were run, performance was plotted by vowel contrast to ensure performance on the control contrast FLEECE-THOUGHT was not near ceiling (in which case the data would have been removed) but this was never the case. No further analyses were done for this factor though data are available on OSF. Full analyses were also run without the FLEECE-THOUGHT control contrast in the data and can be found on OSF- the overall pattern of results remained the same.

⁴ Note that using a half-normal rather than a uniform distribution has the advantage of favouring smaller values (Dienes, 2014).

⁵ In some cases standard deviations were computed using knowledge of constraints on the likely maximum value, such as a maximum of the scale (see also Dienes & McLatchie, 2018). Some of these calculations require an estimation of minimal performance. For most tests this was chance.

⁶ Note our analyses deviate from the preregistration due to unexpected null effects in our results. While the proposed analyses planned to test for interactions with *session* and *variability* in all pre/post-tests, we did not anticipate null effects of learning. If there is no effect of learning, it makes no sense to see if the effect of learning differs depending on the variability condition, as no evidence for an effect was found in the first place. Since these analyses were pre-registered, we did conduct these and they are reported on OSF

(https://osf.io/bgdxp/?view_only=209600aca2cf46a89f0dfd402ab46769) for completeness. In

almost every case, they yielded ambiguous evidence, and where they did not, they yielded substantial evidence for the null. In the current paper, if no substantial evidence was found for an effect of session (i.e. no improvement from pre- to post-test was seen), no further follow-up interactions were reported. The preregistration further proposed age comparisons between the 7/8-year-olds and 11/12-year-olds on every pre/post-test, to test for differences in the amount of generalisation shown in the age groups. However, given that the 7/8-year-olds consistently show a lack of improvement on pre/post-tests, we do not think it sensible to report age comparisons looking at the amount of improvement in generalisation if one of the two groups showed none to start with. For this reason, age comparisons are only reported for training, where learning across sessions was found in both age groups. Since these analyses were pre-registered, they were again conducted and are reported on OSF for completeness. In almost every case, there was evidence for older learners outperforming younger learners, and often also for older learners improving more than younger learners. Where this was not the case there was often evidence *against* younger learners outperforming older learners but ambiguous evidence in the opposite direction.

⁷ We also tested the reverse hypothesis that younger learners improved more than older learners and there was substantial evidence for the null ($\beta = -0.103$, $SE = 0.023$, $z = -4.494$, $p < .001$; $BF_{H(0,0.083)} = 0.045$ [RR 0, >5.516]).

⁸ The preregistration document proposed analysing the data in the same way as other pre/post-tests, with item novelty as a factor. However, in hindsight, the contrast between the minimal pair trials and the non-minimal pair trials is quite different from the item novelty manipulation, as all words are familiar in the non-pair condition. This difference from the item novelty manipulation is reflected in rather different performances on the minimal pair items compared to the non-pair items. For this reason, we felt it was more appropriate to analyse the minimal pair items separately from the non-pair items for each age group.

REFERENCES

- Alloway, T. P., Gathercole, S. E., & Pickering, S. J. (2006). Verbal and visuospatial short-term and working memory in children: Are they separable?. *Child development*, 77(6), 1698-1716. <https://doi.org/10.1111/j.1467-8624.2006.00968.x>
- Antoniou, M., & Wong, P. C. (2015). Poor phonetic perceivers are affected by cognitive load when resolving talker variability. *The Journal of the Acoustical Society of America*, 138(2), 571-574. <https://doi.org/10.1121/1.4923362>
- Anwyl-Irvine, A. L., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. K. (2019). Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods*. <https://doi.org/https://doi.org/10.1101/438242>
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390–412. <https://doi.org/10.1016/j.jml.2007.12.005>
- Barcroft, J., & Sommers, M. S. (2005). Effects of acoustic variability on second language vocabulary learning. *Studies in Second Language Acquisition*, 27(3), 387-414. <https://doi.org/10.1017/S0272263105050175>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Boersma, P., & Weenink, D. (2015). *Praat: doing phonetics by computer* (Version 5.4.17). <http://www.praat.org/>
- Bolker, B. (2019). *GLMM FAQ*. <http://bbolker.github.io/mixedmodels-misc/glmmFAQ.html#should-i-treat-factor-xxx-as-fixed-or-random>

- Bradlow, A. R., Akahane-Yamada, R., Pisoni, D. B., & Tohkura, Y. I. (1999). Training Japanese listeners to identify English/r/and/l: Long-term retention of learning in perception and production. *Perception & psychophysics*, 61(5), 977-985.
<https://doi.org/10.3758/BF03206911>
- Bradlow, A. R., Pisoni, D. B., Akahane-Yamada, R., & Tohkura, Y. I. (1997). Training Japanese listeners to identify English/r/and/l: IV. Some effects of perceptual learning on speech production. *The Journal of the Acoustical Society of America*, 101(4), 2299-2310.
<https://doi.org/10.1121/1.418276>
- Braginsky, M. (2018). *ggpirate: pirate plots for ggplot2*. <https://github.com/mikabr/ggpirate>
- Brekelmans, G., Lavan, N., Saito, H., Clayards, M. & Wonnacott, E. (2022). Does high variability training improve the learning of non-native phoneme contrasts over low variability training? A replication. *Journal of Memory and Language*, 126. 104352.
<https://doi.org/10.31234/osf.io/gs7y9>
- Collins, B., & Mees, I. M. (2003). *The Phonetics of English and Dutch* (Fifth revised edition). Leiden: Brill.
- Clopper, C. G., & Pisoni, D. B. (2004). Effects of talker variability on perceptual learning of dialects. *Language and Speech*, 47(3), 207-238.
<https://doi.org/10.1177/00238309040470030101>
- Deng, Z., Chandrasekaran, B., Wang, S., Wong, P.C.M. (2019). Training-induced brain activation and functional connectivity differentiate multi-talker and single-talker speech training. *Neurobiology of Learning and Memory*, 151. 1-9.
<https://doi.org/10.1016/j.nlm.2018.03.009>
- Dienes, Z. (2008). *Understanding Psychology as a Science: An Introduction to Scientific and Statistical Inference*. Palgrave Macmillan.
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5(July), 1–17. <https://doi.org/10.3389/fpsyg.2014.00781>

- Dienes, Z. (2015). How Bayesian statistics are needed to determine whether mental states are unconscious. In M. Overgaard (Ed.), *Behavioral Methods in Consciousness Research* (pp. 199–220). Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199688890.003.0012>
- Dienes, Z. (2017). *Bayes Factors: Context, Principles, Criticism*. Retrieved from <https://www.youtube.com/watch?v=g9blfZ4KqCQ>
- Dienes, Z., & McLatchie, N. (2018). Four reasons to prefer Bayesian analyses over significance testing. *Psychonomic Bulletin & Review*, 25(1), 207–218. <https://doi.org/10.3758/s13423-017-1266-z>
- Dong, H., Clayards, M., Brown, H., & Wonnacott, E. (2019). The effects of high versus low talker variability and individual aptitude on phonetic training of Mandarin lexical tones. *PeerJ*, 7, e7191. <https://doi.org/10.7717/peerj.7191>
- Escudero, P., Hayes-Harb, R., & Mitterer, H. (2008). Novel second-language words and asymmetric lexical access. *Journal of Phonetics*, 36(2), 345–360.
<https://doi.org/10.1016/j.wocn.2007.11.002>
- Evans, B. G., & Martín-Alvarez, L. (2016). Age-related differences in second-language learning? A comparison of high and low variability perceptual training for the acquisition of English /i/-/ɪ/ by Spanish adults and children. *Proceedings of New Sounds.: 8th International Conference on Second Language Speech*. Aarhus, Denmark. Retrieved from <https://conferences.au.dk/newsounds2016/>
- Flege, J. E. (1995). Second Language Speech Learning: Theory, Findings, and Problems. In W. Strange (Ed.), *Speech Perception and Linguistic Experience: Issues in Cross-Language Research* (pp. 233–277). Timonium, MD: York Press.
- Franco, A., Malhotra, N., & Simonovits, G. (2014). Publication bias in the social sciences: Unlocking the file drawer. *Science*, 345(6203), 1502-1505.
<https://doi.org/10.1126/science.1255484>

- Fuhrmeister, P., & Myers, E. B. (2020). Desirable and undesirable difficulties: Influences of variability, training schedule, and aptitude on nonnative phonetic learning. *Attention, Perception, & Psychophysics*, 1-17. <https://doi.org/10.3758/s13414-019-01925-y>
- Giannakopoulou, A., Brown, H., Clayards, M., & Wonnacott, E. (2017). High or low? Comparing high and low-variability phonetic training in adult and child second language learners. *PeerJ*, 5, e3209. <https://doi.org/10.7717/peerj.3209>
- Giannakopoulou, A., Uther, M., & Ylinen, S. (2010). Phonetic cue-weighting in the acquisition of a second language (L2): evidence from Greek speakers of English. *Achievements and Perspectives in SLA of Speech: New Sounds 2010*, 1(1), 91–201.
- Giannakopoulou, A., Uther, M., & Ylinen, S. (2013). Enhanced plasticity in spoken language acquisition for child learners: Evidence from phonetic training studies in child and adult learners of English. *Child Language Teaching and Therapy*, 29(2), 201–218. <https://doi.org/10.1177/0265659012467473>
- Heeren, W. F. L., & Schouten, M. E. H. (2010). Perceptual development of the Finnish /t-t:/ distinction in Dutch 12-year-old children: A training study. *Journal of Phonetics*, 38(4), 594–603. <https://doi.org/10.1016/j.wocn.2010.08.005>
- Huckvale, M. (2016). *ProRec: Prompt & Record* (2.3). <http://www.phon.ucl.ac.uk/resource/prorec/>
- Huyck, J. J. & B. A. Wright. (2010). Late maturation of auditory perceptual learning. *Developmental Science*, 14(3), 614-621. <https://doi.org/10.1111/j.1467-7687.2010.01009.x>
- Iverson, P., & Evans, B. G. (2009). Learning English vowels with different first-language vowel systems II: Auditory training for native Spanish and German speakers. *The Journal of the Acoustical Society of America*, 126(2), 866–877. <https://doi.org/10.1121/1.3148196>
- Jeffreys, H. (1961). *Theory of probability*. Oxford: Oxford University Press.

- Kartushina, N., & Martin, C. D. (2019). Talker and acoustic variability in learning to produce nonnative sounds: evidence from articulatory training. *Language Learning*, 69(1), 71-105. <https://doi.org/10.1111/lang.12315>
- Keuleers, E., Brysbaert, M., & New, B. (2010). SUBTLEX-NL: A new measure for Dutch word frequency based on film subtitles. *Behavior Research Methods*, 42(3), 643–650. <https://doi.org/10.3758/BRM.42.3.643>
- Lavan, N., Knight, S., Hazan, V., & McGettigan, C. (2019). The effects of high variability training on voice identity learning. *Cognition*, 193, 104026. <https://doi.org/10.1016/j.cognition.2019.104026>
- Lively, S. E., Logan, J. S., & Pisoni, D. B. (1993). Training Japanese listeners to identify English/r/and/l/. II: The role of phonetic environment and talker variability in learning new perceptual categories. *The Journal of the acoustical society of America*, 94(3), 1242-1255. <https://doi.org/10.1121/1.408177>
- Logan, J. S., Lively, S. E., & Pisoni, D. B. (1991). Training Japanese listeners to identify English/r/and/l: A first report. *The Journal of the Acoustical Society of America*, 89(2), 874-886. <https://doi.org/10.1121/1.1894649>
- Maguire, M. J., Hirsh-Pasek, K., Michnick Golinkoff, R., Brandone, A.C. (2008). Focusing on the relation: fewer exemplars facilitate children's initial verb learning and extension. *Developmental science*, 11(4), 628-634. <https://doi.org/10.1111/j.1467-7687.2008.00707.x>
- Martin, C. S., Mullennix, J. W., Pisoni, D. B., & Summers, W. V. (1989). Effects of talker variability on recall of spoken word lists. *Journal of experimental psychology. Learning, memory, and cognition*, 15(4), 676–684. <https://doi.org/10.1037//0278-7393.15.4.676>
- Matuschek, H., Kliegl, R., Vasishth, S., Baayen, R. H., & Bates, D. (2017). Balancing Type I error and power in linear mixed models. *Journal of Memory and Language*, 94, 305–315. <https://doi.org/10.1016/j.jml.2017.01.001>

- Mullennix, J. W., Pisoni, D. B., & Martin, C. S. (1989). Some effects of talker variability on spoken word recognition. *The Journal of the Acoustical Society of America*, 85(1), 365–378. <https://doi.org/10.1121/1.397688>
- Nishi, K., & Kewley-Port, D. (2007). Training Japanese Listeners to Perceive American English Vowels: Influence of Training Sets. *Journal of Speech, Language, and Hearing Research : JSLHR*, 50(6), 1496–1509. [https://doi.org/10.1044/1092-4388\(2007/103\)](https://doi.org/10.1044/1092-4388(2007/103))
- Nishi, K., & Kewley-Port, D. (2008). Nonnative Speech Perception Training Using Vowel Subsets: Effects of Vowels in Sets and Order of Training. *Journal of Speech, Language, and Hearing Research*, 51(6), 1480–1493. [https://doi.org/10.1044/1092-4388\(2008/07-0109\)](https://doi.org/10.1044/1092-4388(2008/07-0109))
- Perrachione, T. K., Lee, J., Ha, L. Y. Y., & Wong, P. C. M. (2011). Learning a novel phonological contrast depends on interactions between individual differences and training paradigm design. *The Journal of the Acoustical Society of America*, 130(1), 461–472. <https://doi.org/10.1121/1.3593366>
- Powell, M. J. D. (2009). The BOBYQA algorithm for bound constrained optimization without derivatives. *NA Report NA2009/06*, 39. <https://doi.org/10.1.1.443.7693>
- R Core Team. (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. <https://www.r-project.org>
- Sadakata, M., & McQueen, J. M. (2013). High stimulus variability in nonnative speech learning supports formation of abstract categories: Evidence from Japanese geminates. *The Journal of the Acoustical Society of America*, 134(2), 1324-1335. <https://doi.org/10.1121/1.4812767>
- Sadakata, M., & McQueen, J. M. (2014). Individual aptitude in Mandarin lexical tone perception predicts effectiveness of high-variability training. *Frontiers in psychology*, 5, 1318. <https://doi.org/10.3389/fpsyg.2014.01318>

- Shinohara, Y., & Iverson, P. (2021). The effect of age on English/r/-/l/perceptual training outcomes for Japanese speakers. *Journal of Phonetics*, 89, 101108.
<https://doi.org/10.1016/j.wocn.2021.101108>
- Sinkevičiūtė, R., Brown, H., Brekelmans, G., & Wonnacott, E. (2019). The role of input variability and learner age in second language vocabulary learning. *Studies in Second Language Acquisition*, 1–26.
<https://doi.org/10.1017/S0272263119000263>
- Silvey, C., Dienes, Z., & Wonnacott, E. (2021). Bayes factors for mixed-effects models. *PsyArXiv*. <https://doi.org/10.31234/osf.io/m4hju>
- Snow, C. E. & M. Hoefnagel-Höhle. (1978). The Critical Period for Language Acquisition: Evidence from Second Language Learning. *Child Development*, 49(4), 1114-1128.
<https://doi.org/10.2307/1128751>
- Sommers, M. S., & Barcroft, J. (2011). Indexical information, encoding difficulty, and second language vocabulary learning. *Applied Psycholinguistics*, 32(2), 417. <https://doi.org/10.1017/S0142716410000469>
- Wang, Y., & Kuhl, P. K. (2003). Evaluating the “critical period” hypothesis: perceptual learning of Mandarin tones in American adults and American children at 6, 10 and 14 years of age,”. In Proceedings of the 15th International Congress of Phonetic Sciences (pp. 1537-1540). Barcelona, Spain: Futurgraphic. Retrieved from
https://www.internationalphoneticassociation.org/icphs-proceedings/ICPhS2003/p15_1537.html
- Wickham, H. (2011). The Split-Apply-Combine Strategy for Data Analysis. *Journal of Statistical Software*, 40(1). <https://doi.org/10.18637/jss.v040.i01>
- Wickham, H. (2016). ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag.
- Wiener, S., Chan, M.K.M. and Ito, K. (2020), Do Explicit Instruction and High Variability Phonetic Training Improve Nonnative Speakers’ Mandarin Tone Productions?. *The Modern Language Journal*, 104: 152-168. <https://doi.org/10.1111/modl.12619>

- Wong, J. (2012). Training the Perception and Production of English/e/and/æ/ of Cantonese ESL Learners: a comparison of low vs. high variability phonetic training. In *Proceedings of the 14th Australasian International Conference on Speech Science and Technology. Sydney, Australia* (pp. 3-6). Retrieved from <https://assta.org/proceedings/sst/SST-12/SST2012/PDF/AUTHOR/ST120021.PDF>
- Wong, J. W. S. (2014). The Effects of High and Low Variability Phonetic Training on the Perception and Production of English Vowels / e / - / æ / by Cantonese ESL Learners with High and Low L2 Proficiency Levels. *Proceedings of the 15 Th Annual Conference of the International Speech Communication Association*, 524–528. Retrieved from https://repository.hkbu.edu.hk/hkbu_staff_publication/6234
- Xie, Y. (2018). knitr: A General-Purpose Package for Dynamic Report Generation in R (1.20).
- Zhu, H. (2018). *kableExtra: Construct Complex Table with “kable” and Pipe Syntax* (0.9.0). <https://cran.r-project.org/package=kableExtra>