BRIEF REPORT

Easy-to-hard effects in perceptual learning depend upon the degree to which initial trials are "easy"

Matthew G. Wisniewski[1] · Barbara A. Church[2] · Eduardo Mercado III[3] · Milen L. Radell[4] · Alexandria C. Zakrzewski[1]

Abstract

Starting perceptual training at easy levels before progressing to difficult levels generally produces better learning outcomes than constantly difficult training does. However, little is known about how "easy" these initial levels should be in order to yield easy-to-hard effects. We compared five levels of initial training block difficulty varying from very easy to hard in two auditory-discrimination learning tasks—a frequency modulation rate discrimination (Experiment 1) and a frequency range discrimination (Experiment 2). The degree of difficulty was based on individualized pretraining ~71% correct discrimination thresholds. Both experiments revealed a sweet spot for easy-to-hard effects. Conditions where initial blocks were either too easy or too difficult produced less benefit than did blocks of intermediate difficulty. Results challenge assumptions that sequencing effects in learning are related to attentional spotlighting of task-relevant dimensions. Rather, they support incremental learning models that account for easy-to-hard effects. Further, the results have implications for how perceptual training regimens should be designed to maximize the benefits of rehabilitative perceptual training.

Keywords Fading · Progressive training · Transfer along a continuum · Adaptive training

Learning can affect the way we experience the world by changing how we process the features of stimuli that reside in it (Goldstone, 1998). Examples of this perceptual learning include how nonnative accents become easier to understand after exposure (Van Engen & Peelle, 2014) and how training to become a wine expert allows one to notice details that a novice misses (James, 1890). Learning is generally more effective if training starts off at an easy level of difficulty before moving to hard perceptual problems. For instance, both rats and humans learn a difficult auditory-identification task better if they start with easy trials than if they complete difficult trials throughout training (Liu, Mercado, Church, & Orduña, 2008). This phenomenon has been referred to as the easy-to-hard

effect. It occurs in several modalities and species (for review, see Wisniewski, Radell, Church, & Mercado, 2017).

One account of easy-to-hard effects is that initial easy trials serve to direct learners' attention to relevant dimensions. Once the "attentional spotlight" is placed on the most relevant dimension, learners perceive that dimension more minutely. Pashler and Mozer (2013) found easy-to-hard effects with participants categorizing face-like "demon" stimuli by horn height. If participants were told that horn height was critical, the easy-to-hard effect disappeared. A single exposure that directs attention to aspects of an image needed to perceive an object (e.g., a face masked by noise) can also yield stable and long-lasting changes in perception (Ahissar & Hochstein, 2004). Such "eureka" effects have been used as support for theories that explain perceptual learning with attentional-spotlighting mechanisms (Ahissar & Hochstein, 2004). Formal models that allow "attentional stretching" of dimensions can account for many instances of easy-to-hard effects as well as other effects of stimulus sequencing (e.g., Carvalho & Goldstone, 2015; Krushcke, 1992; Sutherland & Mackintosh, 1971).

Incremental associative and/or representational modification-based learning processes may additionally drive easy-to-hard effects. The classic gradient interaction theory of Spence (1937)

✉ Matthew G. Wisniewski
   mgwisniewski@ksu.edu

1   Kansas State University, Manhattan, KS, USA

2   Georgia State University, Atlanta, GA, USA

3   University at Buffalo, The State University of New York,
     Buffalo, NY, USA

4   Niagara University, Niagara University, NY, USA

proposes that positive excitatory gradients of generalization develop around reinforced stimuli, while negative inhibitory gradients surround nonreinforced stimuli. An individual's ability to discriminate is governed by the summation of these gradients. If the gradients are too overlapping (as when a reinforced S+ stimulus is difficult to discriminate from a nonreinforced S−), they will mostly cancel each other out, and learning will proceed slowly. If the gradients are more separated, but still overlapping, their summation will produce a stronger difference between S+ and S− that will generalize from the easy to the hard version of the task (also see McLaren & Mackintosh, 2002). Nonassociative representational modification-based models account for easy-to-hard effects by taking into consideration how stimulus sequencing affects the plasticity of stimulus representations. When two stimuli are very similar, they will compete to change the same representational space (e.g., neurons coding for the same frequencies of sound). In this case, learning proceeds slowly. When two stimuli are easier to discriminate, competition is reduced, and thus a more accurate representation of each stimulus can be developed and later refined on the hard version of the task (Saksida, 1999).

The debate between attentional spotlighting and incremental learning theories continues. In one recent study, people were trained to discriminate auditory frequency modulation (FM) rates in two different frequency ranges (300–600 Hz or 3000–6000 Hz; Wisniewski, et al., 2017). In one frequency range, difficulty faded from easy to hard. In the other, discriminations remained constantly difficult. Even though the same dimension (FM rate) was relevant in both frequency ranges, participants performed better in the range that received easy-to-hard training. However, other work has failed to find easy-to-hard benefits when attentional-spotlighting effects are kept to a minimum (Pashler & Mozer, 2013). Those authors concluded that easy-to-hard effects should only manifest in relatively high-level perceptual category learning tasks where the learner must first figure out which dimensions are relevant. In some cases, simple exposure to stimuli in an easy-to-hard progression can benefit learning, suggesting that an attention independent nonassociative learning process is at play (Church, Mercado, Wisniewski, & Liu, 2013; Sanjuán, Nelson, & Alonso, 2014). In a similar vein, enhancement of auditory-evoked potentials (AEPs) is larger after easy-to-hard compared with constantly difficult training (Orduña, Liu, Church, Eddins, & Mercado, 2012). This is the case even when participants are asked to ignore sounds during AEP collection (Orduña et al., 2012). However, it could be argued that exposure to easily discriminable contrasts serves as a cue for directing attention to specific dimensions even when there is no task involved, and that participants are paying attention to stimuli during passive AEP collection despite being instructed not to do so.

Here, we pit yet-to-be tested predictions of these two classes of theory against each other in a relatively low-level auditory task. Incremental learning processes are constrained by the degree to which initial trials are "easy." If initial trials contain very distinct nonoverlapping representations, generalization from those easy trials to harder trials will be minimal. For instance, gradients of excitation and inhibition will be too disimilar to yield strong gradient interaction at the points of the hard S+/S− discrimination. From the view of nonassociative representational modification accounts, extremely easy trials make it less likely that representations will be refined that can be usefully modified on a harder version of the task. On the contrary, attentional-spotlighting accounts predict that easy trials should facilitate performance as long as the discrimination-relevant dimension is made obvious.

Experiment 1

Five groups of individuals were trained to discriminate rates of frequency modulation (FM). Groups were either trained constantly at their predetermined 71% correct threshold (hard: H) or faded from easier contrasts to their threshold (easy-to-hard: EH). There were four different EH conditions differing in the degree to which initial easy trials were easy: EH1 (most difficult)–EH4 (easiest). If easy-to-hard effects are exclusively related to learning the discrimination-relevant dimension, then there should be a monotonic trend of increasing posttraining performance from H to EH4 or no effect of training due to a "eureka" experience in the thresholding phase. However, incremental theories predict that there should be a sweet spot where fading benefits are observed somewhere between the H and EH4 conditions.

Method

We designated data collection and analysis plans prior to running experiments. The a priori plans, an annotated postexperiment file, and the raw data can be downloaded from www.alclaboratory.com/opendata.

Participants Seventy individuals at Kansas State University participated in exchange for course credit. All signed a consent form approved by the local Institutional Review Board and reported normal hearing. Fourteen participants were randomly assigned to each of the conditions.

Stimuli and apparatus Single upward directed FM sweeps (1.5–3 kHz) served as stimuli (see Fig. 1a for spectrograms of select stimuli). FM sweeps were rendered online in MATLAB 2018b (The MathWorks, Natick, MA) along with experimental procedures. A 5 octaves/s sweep served as a standard "slow" rate. All other rates were faster.

Neutral valence silent videos were downloaded from videos.pexels.com under a creative commons zero license. A
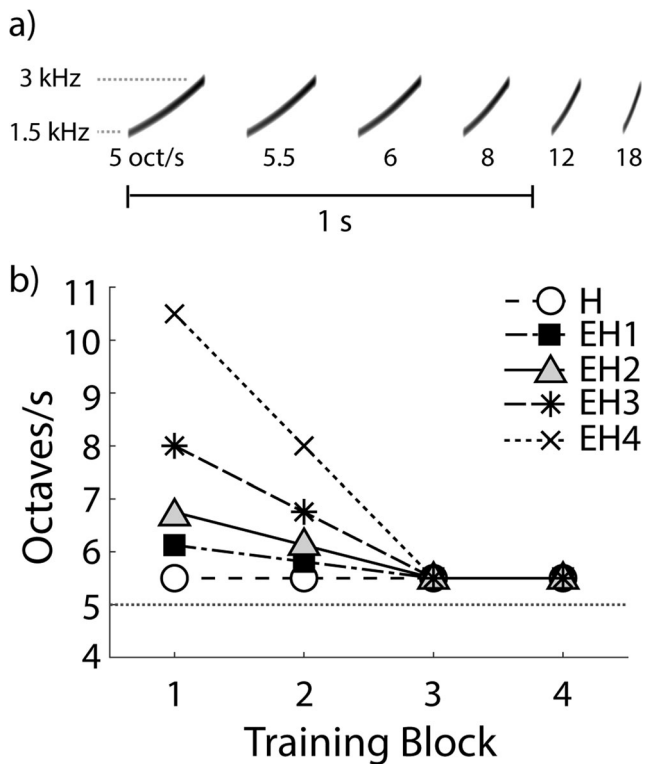
a)

b)

Fig. 1 **a** Spectrograms of select FM sweep stimuli at rates from 5 octaves/s to 18 octaves/s. **b** Depiction of fast FM rates that would be presented under each condition for a participant with a threshold of .5 octaves/s

random video was selected after each block and played during a forced 30-s break. This served to mitigate overwhelming listeners with constant auditory stimulation.

Participants sat in sound-attenuating booths (WhisperRoom, Knoxville, TN) and heard stimuli over Sennheiser HD-280 closed-back headphones (Sennheiser, Germany) connected to Focusrite Scarlett external sound cards (Focusrite, UK). Stimuli were presented at ~81 dB SPL. Responses were made on custom keypads (P.I., Engineering, Williamston, MI).

Procedures A two-interval two-alternative forced-choice (2i–2afc) task was used to identify ~71% correct thresholds in an initial 60 trial block. On each trial, two FM sounds were played, separated by a 500-ms interstimulus interval. One was always the standard. The other was faster (order selected at random on each trial). Participants' task was to indicate which sound was slower. The FM rate difference started at 4 octaves/s on Trial 1. That is, the "slow" rate was 5 octaves/s, and the "fast" was 9 octaves/s. Using a one-up, two-down procedure (Levitt, 1970), the difference was made larger by dividing by 0.9 after every incorrect response and was made smaller by multiplying by 0.9 after every two consecutive correct responses. The mean octaves/s difference on the last 10 trials of the block was taken as a participant's threshold.

Participants then proceeded to 1i–2afc training. Participants were asked to indicate whether a single FM sweep was "fast" or "slow." Participants were informed that 50% of sounds would be "fast" and 50% would be "slow." Responses were not registered until the sound completed playing. Feedback of correctness was presented after responding. If a response was not made within a 5-s window following sound offset, a missing response was recorded. These trials were not included in the analysis. A pseudorandom trial order was used such that the same stimulus could not occur more than 4 times consecutively.

Group H completed four blocks of training, all at threshold. Group EH1 had the first block of training at 2.25 times threshold, and the second training block at 1.625 times threshold. First and second block training levels for Groups EH2, EH3, and EH4 were 3.5 and 2.25, 6 and 3.5, and 11 and 6, respectively. Training Blocks 3 and 4 were at threshold for EH Groups 1–4, just as the H group. Figure 1b shows an example of the progressions of fast FM rates that would be used for Training Blocks 1–4 for a hypothetical individual with a 0.5 octaves/s threshold under each of the conditions. Training blocks were all 32 trials long.

After training, all participants completed an 80-trial test at threshold. The task was identical to training, except that no feedback was given after responding. Test performance served as the critical variable for our hypotheses.

Results and discussion

The mean difference threshold across all individuals was .56 octaves/s (SD = .27). There were no significant differences in thresholds across conditions, $F < 2$.

We used A' as a measure of accuracy in training and testing, computed using Equation 1 when hit rate (H) was greater than the false-alarm rate (F), and Equation 2 otherwise.

$$A' = .5 + \frac{(H-F)(1+H-F)}{4H(1-F)} \qquad (1)$$

$$A' = .5 - \frac{(F-H)(1+F-H)}{4F(1-H)} \qquad (2)$$

Figure 2a shows A' for the training blocks. Unsurprisingly, the initial two training blocks showed different accuracies across the conditions. The H condition showed the lowest performance as the first two blocks were at ~71% thresholds. For the EH groups, there was higher accuracy for EH4 compared to EH1, with the others were in between. Accuracy in the third and fourth blocks of training was comparable across conditions.

The lowest level of test performance was observed for the H condition, while the highest performance was observed for
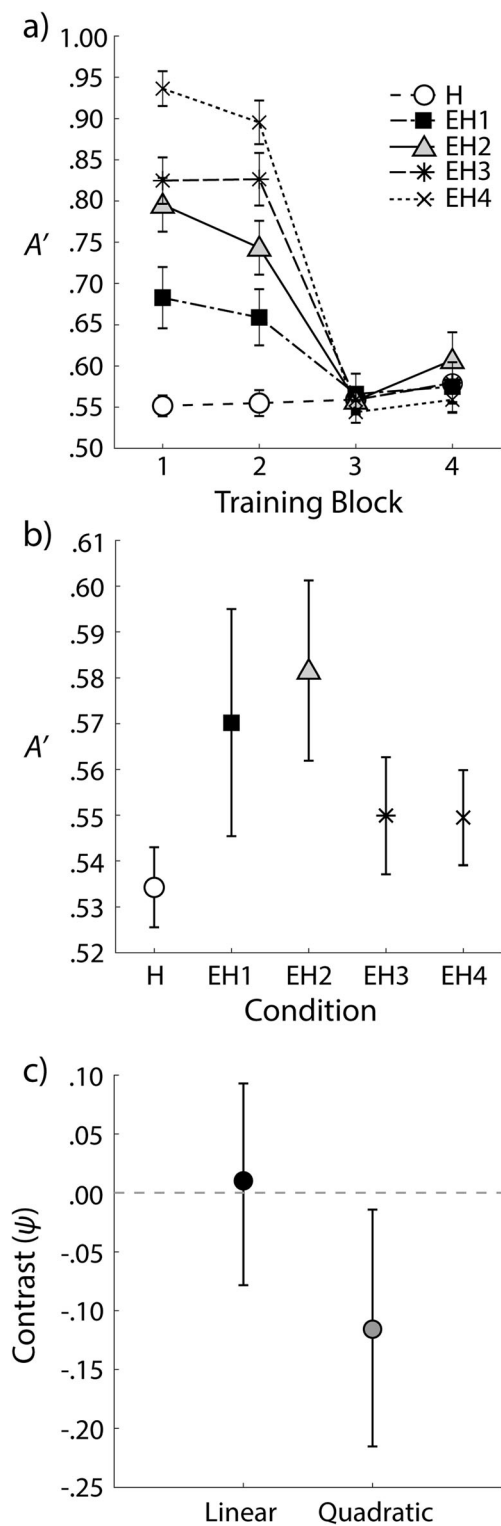
a) Accuracy (A') across training blocks for all groups in Experiment 1.

b) Accuracy (A') in the test for all groups.

c) Linear and quadratic contrasts.

the EH2 condition (see Fig. 2b). Planned orthogonal linear (attentional spotlighting prediction; coefficients = −2, −1, 0, 1, 2) and quadratic contrasts (incremental prediction; coefficients = 2, −1, - 2, −1, 2) were computed for the trends in means across conditions H to EH4. These contrasts were assessed for significance using a nonparametric permutation-based procedure. For 1,000 iterations, condition labels were randomly shuffled and a linear (ψ_{linear}) and quadratic ($\psi_{quadratic}$) contrast statistic was computed. These permuted ψ values created a null hypothesis distribution. A p value was considered to be the proportion of ψ values in the distribution that exceeded that of the actual data (α = .05).

Figure 2c depicts the observed linear and quadratic contrast statistics along with 90% confidence intervals for one-sided hypotheses (derived from obtained null hypothesis distribution). The linear contrast was not significant, p = .448. The quadratic contrast was, p = .029, demonstrating that intermediate levels of difficulty for initial easy trials led to better test accuracy than levels that were too difficult (H condition) or too easy (EH4 condition).

Experiment 2

Experiment 2 was similar to Experiment 1, except that stimuli varied on the dimension of frequency range rather than FM rate. The purpose was to test whether or not Experiment 1 effects could be replicated for a different acoustic dimension.

Method

With one exception, the data collection and analysis plans were identical to Experiment 1. Instead of using FM rate difference threshold as a drop criterion, a threshold of two semitones frequency range difference was used.

Participants Seventy-three individuals enrolled in courses at Kansas State University participated in exchange for credit. All signed a consent form approved by Kansas State University's Institutional Review Board. All participants self-reported normal hearing. Fourteen participants were initially randomly assigned to each of the conditions. Three participants were dropped due to high pitch discrimination thresholds (>2 semitones) and were replaced.

Stimuli and apparatus A standard "low" pitched stimulus was identical to the "slow" stimulus in Experiment 1 (5 octaves/s; 1.5–3 kHz frequency range). All other stimuli spanned a higher frequency range.

Procedures The thresholding task was similar to Experiment 1, except that the participant's task was to indicate which sound was lower in pitch. Frequency range difference started

Fig. 2 **a** Accuracy (A') across training blocks for all groups in Experiment 1. Error bars show standard error of the mean (*SEM*). **b** Accuracy (A') in the test for all groups. Error bars show *SEM*. **c** Linear and quadratic contrasts. Error bars show 90% confidence intervals according to a null hypothesis distribution generated using a permutation procedure (see text)

at 2.16 semitones. The standard "low" sound swept from 1.5 to 3 kHz, while the "high" sound swept from 1.7 kHz to 3.4 kHz. The difference was made larger by dividing the difference in Hz by 0.9 after every incorrect response and made smaller by multiplying by 0.9 after every two consecutive correct responses. The mean difference on the last 10 trials of the block was taken as a participant's ~71% correct threshold.

Training and testing procedures were the same as Experiment 1, except that the stimulus differences were based off thresholds for the frequency range of FM sweeps rather than rate. Also, the task was to indicate whether a sound was "high" or "low" rather than "fast" or "slow."

Results and discussion

The mean threshold across participants was .49 semitones (*SD* = .38). There were no significant differences across conditions, $F < 2$. As in Experiment 1, the EH4 group performed best and the H group performed worst for the first two blocks of training (see Fig. 3a). The other EH groups were in between. There also appeared to be a trend in Blocks 3 and 4 for higher accuracy in the EH1–3 groups compared with the others. This is consistent with a nonmonotonic trend in accuracy across degree of initial block easiness. We turn next to the test data, as this data were critical to evaluating the competing hypotheses.

Figure 3b shows A' in the test. Much like Experiment 1, the data show a quadratic trend from H to EH4 such that conditions where initial trials were too easy or too difficult showed lower accuracy than conditions with intermediate levels of difficulty for the first two blocks. The same statistics used in Experiment 1 were employed in the analysis of Experiment 2 (see Fig. 3c). The linear trend was not significant, $p = .311$. The quadratic contrast was significant, $p = .008$, as predicted by incremental theories of learning.

General discussion

In two experiments, we pitted predictions of attentional spotlighting and incremental learning theories against each other in the context of easy-to-hard effects. We found a sweet spot for easy-to-hard effects in auditory learning such that training protocols where initial blocks are too easy or too difficult produce less benefit than blocks of intermediate difficulty. This result was observed for two different acoustic dimensions, was predicted by incremental accounts of learning, and runs counter to predictions of attentional spotlighting. It is also worth noting that discrimination thresholds on the relevant dimension for every listener were collected before training. Listeners were able to reach reasonable thresholds, suggesting they had knowledge of the relevant dimension
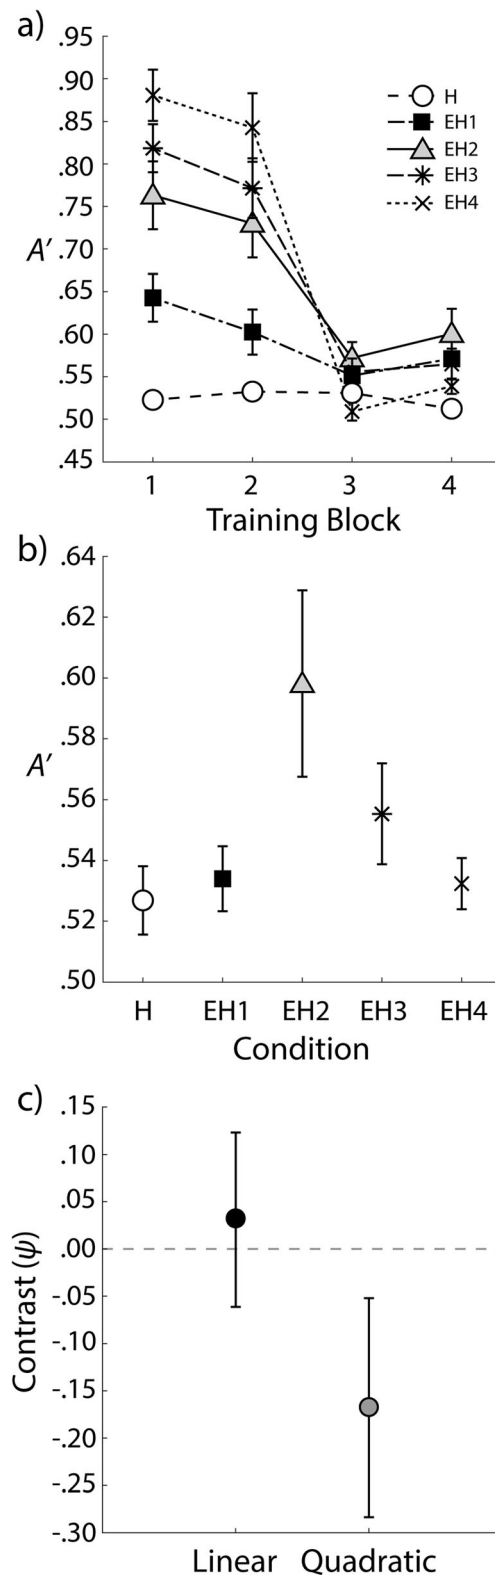

Fig. 3 **a** Accuracy (A') across training blocks for all groups in Experiment 2. Error bars show standard error of the mean (*SEM*). **b** Accuracy (A') in the test for all groups. Error bars show *SEM*. **c** Linear and quadratic contrasts. Error bars show 90% confidence intervals according to a null hypothesis distribution generated using a permutation procedure (see text)

before ever starting training. Easy-to-hard effects observed as a result of training are thus unlikely to come from a "eureka" experience (Ahissar & Hochstein, 2004) or dimensional discovery (Pashler & Mozer, 2013).

It appears that attentional spotlighting is not a sufficient explanation of easy-to-hard effects on its own. Others have come to similar conclusions based on recent work in category learning. Lee and Livesey (2018) trained participants to categorize circles that varied in color and size such that paying attention to a single dimension was encouraged. Afterwards, participants were tested either in a condition where they could consistently apply a rule based on the relevant dimension (e.g., small circles are Category A, but large circles are Category B), or in a condition where such a rule led to inconsistent performance. Participants tested in the rule-consistent condition showed generalization patterns that were indicative of rule use on the relevant dimension (e.g., size). However, the inconsistent group showed generalization patterns more consistent with incremental learning processes (also, see Wisniewski, Church, & Mercado, 2014a; Wisniewski, Radell, Guillette, Sturdy, & Mercado, 2012).

The recent work of Lee and Livesey (2018) combined with the results reported here suggest an interesting possibility for why attentional spotlighting has remained such a favored account of easy-to-hard effects: Strong attentional spotlighting effects have the potential to mask effects of other learning processes. Easy-to-hard effects unlikely to be related to attentional spotlighting are generally smaller (e.g., Cohen's d = 0.90 for Experiment 1; Liu et al., 2008) compared with the effect of knowing the relevant dimension(s) (e.g., Cohen's d = 2.58 for Experiment 5; Pashler & Mozer, 2013). At the same time, selection of suboptimal "easy" levels may lead to weak easy-to-hard effects generated by incremental learning processes, making this discrepancy even larger. Many models account for sequencing effects through adjustments to a parameter that creates "attentional stretching" of a dimension by reducing stimulus similarity (sequential attention theory: Carvalho & Goldstone, 2015; ALCOVE: Kruschke, 1992; analyzer theory: Sutherland & Mackintosh, 1971). Often, other models (e.g., McLaren & Mackintosh, 2002; Saksida, 1999) or model components that do not allude to attention are ignored or quickly abandoned as explanations of sequencing. We do not protest the idea that attentional spotlighting contributes to sequencing effects in learning, only to the notion that attentional spotlighting is a sufficient explanation on its own. Models that take into consideration attentional states, along with incremental learning processes, will do a better job of simulating easy-to-hard effects and perceptual learning in general.

Largely based on the assumption that initial easy trials help learners discover relevant dimensions, adaptive training procedures that start at very easy levels are used extensively in perceptual learning studies. Some have suggested that the most effective way to train individuals should be to use initial easy trials until an individual discovers the relevant dimension, then switch to constantly difficult trials that allow for most accurate tweaking of category boundaries (e.g., Pashler & Mozer, 2013). This is somewhat in line with how adaptive procedures work, where the majority of trials are at an individual's asymptotic performance. The current data highlight the need for empirically based selection of perceptual training regimens. Though comparisons of different adaptive procedures sometimes show little to no effects of varying conditions (Amitay, Irwin, Hawkey, Cowan, & Moore, 2006), future work should compare the effectiveness of adaptive training procedures to those of fixed levels of progression selected to optimize learning (e.g., by designing initial trials to be in between too easy and too difficult). Since incremental learning processes are also affected by sequencing, it is further important to examine stimulus sequencing effects under conditions of exposure (cf. Church et al., 2013; Sanjuán et al., 2014; Wright, Sabin, Zhang, Marrone, & Fitzgerald, 2010), latent learning between blocks or sessions (Molloy, Moore, Sohoglu, & Amitay, 2012), and when performance depends on generalizing across training-irrelevant dimensions and/or stimulus characteristics (Wisniewski, Liu, Church, & Mercado, 2014b; Wisniewski, Mantell, & Pfordresher, 2013). It will also be necessary to examine whether initial trial ease affects learning in other modalities the same way.

Consistent with incremental based learning theory explanations of easy-to-hard effects, we found a sweet spot for initial trial ease in two different auditory tasks. Considering incremental learning as well as attentional-spotlighting explanations of sequencing effects in training will help advance theory, as well as lead to more effective training procedures.

Compliance with ethical standards

Open practices statement A preexperiment plan, annotated postexperiment file, and the raw data are available at www.alclaboratory.com/opendata.

References

Ahissar, M., & Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. *Trends in Cognitive Sciences*, *8*(10), 457–464.

Amitay, S., Irwin, A., Hawkey, D. J., Cowan, J. A., & Moore, D. R. (2006). A comparison of adaptive procedures for rapid and reliable threshold assessment and training in naïve listeners. *The Journal of the Acoustical Society of America*, *119*(3), 1616–1625.

Carvalho, P. F., & Goldstone, R. L. (2015). The benefits of interleaved and blocked study: Different tasks benefit from different schedules of study. *Psychonomic Bulletin & Review*, *22*(1), 281–288.

Church, B. A., Mercado, E., III, Wisniewski, M. G., & Liu, E. H. (2013). Temporal dynamics in auditory perceptual learning: Impact of sequencing and incidental learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(1), 270–276.

Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology, 49*, 585-612

James, W. J. (1890). The principles of psychology. New York: Dover.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99(1), 22–44.

Lee, J. C., & Livesey, E. J. (2018). Rule-based generalization and peak shift in the presence of simple relational rules. *PLOS ONE*, 13(9), e0203805.

Levitt, H. (1970). Transformed up-down methods in psychoacoustics. *Journal of the Acoustical Society of America, 49*, 467–477.

Liu, E. H., Mercado, E., III, Church, B. A., & Orduña, I. (2008). The easy-to-hard effect in human (Homo sapiens) and rat (Rattus norvegicus) auditory identification. *Journal of Comparative Psychology, 122*, 132–145.

McLaren, I. P. L., & Mackintosh, N. J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. *Animal Learning & Behavior*, 30(3), 177–200.

Molloy, K., Moore, D. R., Sohoglu, E., & Amitay, S. (2012). Less is more: Latent learning is maximized by shorter training sessions in auditory perceptual learning. *PLOS ONE*, 7(5), e36929.

Orduña, I., Liu, E. H., Church, B. A., Eddins, A. C., & Mercado, E., III. (2012). Evoked-potential changes following discrimination learning involving complex sounds. *Clinical Neurophysiology*, 123(4), 711–719.

Pashler, H., & Mozer, M. C. (2013). When does fading enhance perceptual category learning?. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(4), 1162–1173.

Saksida, L. M. (1999). Effects of similarity and experience on discrimination learning: A nonassociative connectionist model of perceptual learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 25(3), 308–323.

Sanjuán, M. C., Nelson, J. B., & Alonso, G. (2014). An easy-to-hard effect after nonreinforced preexposure in a sweetness discrimination. *Learning & Behavior*, 42(3), 209–214.

Spence, K. W. (1937). The differential response in animals to stimuli varying within a single dimension. *Psychological Review, 44*, 430–444.

Sutherland, N. S., & Mackintosh, N. J. (1971). Mechanisms of animal discrimination learning. New York: Academic Press.

Van Engen, K. J., & Peelle, J. E. (2014). Listening effort and accented speech. *Frontiers in Human Neuroscience, 8*, 577.

Wisniewski, M. G., Church, B. A., & Mercado III, E. (2014a). Individual differences during acquisition predict shifts in generalization. *Behavioural Processes*, 104, 26–34.

Wisniewski, M. G., Liu, E. H., Church, B. A., & Mercado, E., III. (2014b). Learning to discriminate frequency modulation rate can benefit and worsen pitch acuity. *The Journal of the Acoustical Society of America*, 135(2), EL55–EL60.

Wisniewski, M. G., Mantell, J. T., & Pfordresher, P. Q. (2013). Transfer effects in the vocal imitation of speech and song. *Psychomusicology: Music, Mind, and Brain, 23*(2), 82–99.

Wisniewski, M. G., Radell, M. L., Church, B. A., & Mercado, E., III. (2017). Benefits of fading in perceptual learning are driven by more than dimensional attention. *PLOS ONE*, 12(7), e0180959.

Wisniewski, M. G., Radell, M. L., Guillette, L. M., Sturdy, C. B., & Mercado, E. (2012). Predicting shifts in generalization gradients with perceptrons. *Learning & Behavior*, 40(2), 128–144.

Wright, B. A., Sabin, A. T., Zhang, Y., Marrone, N., & Fitzgerald, M. B. (2010). Enhancing perceptual learning by combining practice with periods of additional sensory stimulation. *Journal of Neuroscience, 30*, 12868–12877.