

Simulating the Temporal Dynamics of Learning-Related Shifts in Generalization Gradients with a Single-Layer Perceptron

Matthew G. Wisniewski (mgw@buffalo.edu)

Department of Psychology, University at Buffalo, SUNY
Buffalo, NY 14260 USA

Milen L. Radell (mlradell@buffalo.edu)

Department of Psychology, University at Buffalo, SUNY
Buffalo, NY 14260 USA

Lauren M. Guillette (guillett@ualberta.ca)

Department of Psychology, University of Alberta
Edmonton, AB T6G 2E9 Canada

Christopher B. Sturdy (csturdy@ualberta.ca)

Department of Psychology, University of Alberta
Edmonton, AB T6G 2E9 Canada

Eduardo Mercado III (emiii@buffalo.edu)

Department of Psychology, University at Buffalo, SUNY
Buffalo, NY 14260 USA

Abstract

Neural network models have been used extensively to model perceptual learning and the effects of discrimination training on generalization, as well as to explore natural classification mechanisms. Here we assess the ability of existing models to account for the time course of generalization shifts that occur when individuals learn to distinguish sounds. A set of simulations demonstrates that commonly used single-layer networks do not predict transitory shifts in generalization over the course of training, but that such dynamics can be accounted for when the output functions of these networks are modified to mimic the properties of cortical tuning curves. The simulations further suggest that prudent selection of training criteria can allow for more precise predictions of learning-related shifts in generalization gradients in behavioral experiments.

Keywords: discrimination learning, representation, similarity, perceptual learning, neural network, peak shift

Introduction

When an organism learns that a stimulus results in some consequence, it will often generalize that learning to similar novel stimuli (Shepard, 1987). For instance, Watson and Raynor (1920) famously demonstrated in experiments with Little Albert that fear associated with a white rat can generalize to other stimuli such as a rabbit, a fur coat, or even a piece of cotton. Numerous theoretical efforts have focused on explaining and predicting generalization patterns, with varying degrees of success. Computational models of discrimination learning recently have proven to be adept at simulating many of the empirically observed differences in generalization that relate to differences in stimulus similarity and training variability (e.g., Ghirlanda & Enquist, 1998; Livesey, Pearson, & McLaren, 2005; Saksida, 1999; Suret & McLaren, 2002). Such computational models are becoming increasingly useful tools for generating hypotheses about the mechanisms and cues that participants use during learning and generalization. In the current study, we assessed the ability of existing

perceptron models of discrimination learning and generalization to account for recent observations of learning-related shifts in generalization observed during an auditory learning task.

A phenomenon commonly observed after discrimination training is that the highest levels of responding may occur for stimuli other than those experienced during training. Generally, *peak shift* results when a participant is trained to respond to one stimulus (S+) and not to some other stimulus (S-) that varies along a common dimension. When generalization is measured after training, responding is strongest not to the trained S+, but to a stimulus that is shifted along the dimension even further from S-. In a classic example of this Hanson (1959) trained pigeons to peck a key when presented with a 560 nm light (S+), but not when presented with a 570 nm light (S-). During generalization tests, pigeons responded most strongly to wavelengths other than 560 (such as 540 nm) that were farther along the dimension from the trained S-, 570 nm light.

Most of the experimental studies of peak shift have focused on the presence, size, or generality of the effect. For instance, it is known that the size of peak shift (i.e., the degree of displacement of maximal responding away from the S+) depends upon the similarity of the S+ and S- during training such that the more similar the training stimuli are, the further the learner will shift (Purtle, 1973). Additionally, work shows that peak shift occurs along simple, single dimensions like wavelength (Hanson, 1959), and sine wave frequency (Baron, 1973), as well as along more complex acoustic (Guillette et al., 2010; Verzijden et al., 2007; Wisniewski et al., 2009; 2010), and visual dimensions (Spetch, Cheng, & Clifford, 2004; Livesey et al., 2005).

Theories of discrimination learning and peak shift have focused on explaining basic experimental effects and species differences. For instance, several associative learning theories (Blough, 1975; Spence, 1937; McLaren & Mackintosh, 2002) adequately explain the direction of shift and the changes to the size of the effect that result from

variations in stimulus similarity. Other theories have posited that people learn rules during discrimination and that this is what causes peak shift in humans (Thomas, Mood, Morrison, & Wiertelak, 1991). Computational models have been used to test how well different theories explain peak shift and related generalization phenomena across species and conditions.

Although there has been much work exploring the basic predictions of current theories regarding peak shift, relatively few studies, experimental or theoretical, have looked at the dynamics of generalization over time. A few studies have shown that extinction can occur during testing, such that shifts in generalization dissipate as more non-differentially reinforced testing trials are experienced (Cheng et al., 1997; Purtle, 1973). For instance, pigeons that exhibit a strong peak shift in the first block of testing show a reduction in the strength of that shift in the following test blocks (Cheng et al., 1997). Other studies have shown that in both humans (Wisniewski et al., 2009) and nonhumans (Moye & Thomas, 1982), peak shift is stable over time in that it lasts at least 24 hours post-discrimination training. Very few studies, however, have looked at how the extent of training impacts peak shift.

Recent work suggests that peak shift can be a transitory effect related to differential amounts of discrimination learning. Wisniewski et al. (2010) trained humans for 60, 100, 140, 180, 220, or 260 trials on a task requiring the discrimination of two complex sounds that varied in the rate of periodic frequency modulation. Participants who were trained with the fewest trials (60 or 100) or the most trials (220 or 260) did not show a peak shift effect. However, participants who were trained with an intermediate number of trials (140 or 180) did exhibit peak shift. These results suggest that in at least some training conditions, peak shift may only occur at intermediate levels of learning.

The dynamics of generalization have been analyzed in past computer simulations of generalization. For example the shape of generalization gradients produced by connectionist networks shift from being Gaussian to exponential as more training iterations are experienced (Shepard, 1990; Staddon & Reid, 1990). However, to our knowledge, simulations investigating the dynamics of shifts in generalization gradients have not been closely analyzed.

The current study assessed whether existing connectionist approaches can capture the quadratic trend in peak shift that occurs over the course of auditory discrimination training in humans (Wisniewski et al., 2010). Toward this goal, variants of a previously developed perceptron model (Dawson, 2004; 2005) of auditory perception in chickadees that is known to exhibit peak shift effects (Guillette et al., 2010), were used to model human learning. The hypothesis was that the neural networks would show transitory peak-shifts in generalization, comparable to those seen previously in humans (Wisniewski et al., 2010). Given that this model was not originally developed to account for human discrimination, similarities between the simulations and experimental data would

provide support for general learning mechanisms mediating the peak shift effect in humans (Ghirlanda & Enquist, 2006; Mercado, 2008; Spetch et al., 2004; Wisniewski et al., 2009), as opposed to rule-based mechanisms not used by non-humans (Thomas et al., 1991). Also, because past studies of generalization shifts have used different amounts of training (Bizo & McMahon, 2007; Baron, 1973; Derenne, 2010; Galizio, 1980; Lewis & Johnston, 1999; Newlin et al., 1979; Thomas, Mood, Morrison, & Wiertelak, 1991; Wisniewski et al., 2009; 2010), and because there can be large individual differences in improvements during training and generalization (Nicholson & Grey, 1972; Withagen & van Wermeskerken, 2009), we assessed how different criteria for ending training affected the variability of generalization gradients.

Methods

Networks were single layer perceptrons with an input layer consisting of 54 units. Each unit in the input layer was connected to a single output unit. Either sigmoid or value output units were used (see Dawson, 2004; 2005). Sigmoid units are often used in connectionist models and have been used extensively to model generalization and peak shift (Dawson, 2004; 2005; Ghirlanda & Enquist, 1998; 2006; Guillette et al., 2010; Livesey et al., 2005, Suret & McLaren, 2002; Tanaka & Simon, 1996). Previous results, however, showed that using such units in discrimination training can yield gradients that are biased to one side of the generalization distribution, especially when networks are trained extensively (Tanaka & Simon, 1996). We also tested networks using value units. Value units use a Gaussian rather than a sigmoid activation function to convert the sum of the weighted values from each input unit into an output value that ranges between 0 and 1. Like a dose response curve, both very low and very high sums produce smaller outputs than intermediate sums. This allows units to become selective to a range of input values, as is seen in the receptive fields of many cortical neurons (e.g., Elhilali, Fritz, Chi, Shamma, 2007). In contrast, the sigmoid activation function is monotonic. Testing networks using units with different activation functions can thus yield insights into how stimuli are represented and/or what types of receptive fields are important in learning discriminations of complex sounds (Enquist & Ghirlanda, 2005). A detailed description of both the sigmoid and value unit types can be found in Dawson (2004). A 2 (output unit type) x 6 (training criteria) design was used to test how networks with different output unit types and levels of training performance would generalize. Five networks were trained per group.

Stimulus Representations

Wisniewski et al. (2010) tested participants with a set of 8 repetition rates of frequency-modulated sweeps, rank ordered 1-8 (from slow to fast), with rank 5 used as the S+ and rank 4 used as the S-. Two additional sounds, that were not part of the generalization distribution, were used during pre-training. Here, we use overlapping patterns of Gaussian

shaped inputs to represent these stimuli. The inputs had a variance of 5 and a maximum value of 1. The inputs are distributed representations of time-varying sounds in which each input value can be viewed as a population of neurons with selectivity to a particular modulation rate (Gluck, 1992). Similar representations have been used previously in models of peak shift for stimuli that are complex in nature (Livesey et al., 2005; Suret & McLaren, 2002). The input stimulus sets used are shown in Figure 1.




Figure 1. Representations of stimuli used in a) pre-training; b) S+/S- discrimination and c) generalization testing. X-axes depict the unit in the input layer. Y-axes depict the activation in each unit for each pattern. The thicker solid line corresponds to the target stimulus, dotted lines are non-targets, thinner solid lines are novel stimuli.

Training and Testing

All networks were trained using the Rosenblatt learning algorithm (Dawson, 2004; 2005). Initial network weights were set at random between -0.1 and 0.1. Networks were pre-trained with the desired response to S+ as 1 and the response to stimuli displaced on either side of the generalization distribution as 0. Network pre-training continued until the sum of the squared error (SSE) for the single output unit was less than 0.05. SSE provided a measure of the difference between the model's output and the desired output. The pre-training procedure was analogous to pre-training used previously in experimental studies (Spetch et al., 2004; Wisniewski et al., 2009; 2010) and served mainly to initialize the perceptron (Fernández-Redondo & Hernández-Espinosa, 2001; Li, Alnuweiri, & Wu, 1993). This initialization enabled the networks to perform the S+/S- discrimination at levels above chance early in training, as is seen experimentally, rather than starting from a completely naïve state.

After pre-training, all networks were given S+/S- discrimination training. The desired output for networks

was set at 1 for the S+ and 0 for the S-. In order to compare each model's generalization after different levels of training experience, we trained groups of networks to 6 different criteria that were defined by (SSE) in the output unit. The 6 criteria were SSEs of 0.4, 0.3, 0.2, 0.1, 0.05, and 0.02. Network training was stopped after the respective SSE level was reached. SSE was used instead of the number of training trials because networks with different unit types learn at different rates (Dawson, 2004; 2005), and we wanted to make sure that networks in different unit conditions reached similar levels of performance on the S+/S- discrimination.

After training, generalization was assessed by presenting networks with the S+, the S-, and 6 novel stimuli with no feedback. The output unit's activations to presentations of test stimuli in each group of networks are reported. Since the quadratic trend in shift over the course of learning has not been modeled before, our main focus was on examining how the performance of different models qualitatively fit the data. The behavioral results from Wisniewski et al. (2010) are shown in Figure 2.




Figure 2. The generalization gradients reported by Wisniewski et al. (2010) for groups of participants trained for different amounts of trials. The y-axis depicts the proportion of times participants indicated a stimulus was the S+. The x-axis shows each stimulus in the generalization distribution. Solid vertical line is the S+; dashed line is the trained S-.

Results

The mean generalization gradients for networks in each group are shown in Figure 3.

Sigmoid Output Units: Perceptrons with sigmoid units showed an increase in shift as SSE decreased on the S+/S- discrimination. The peak activation of the output unit shifted away from the S+, even in the most extensively trained group of networks.

Value Output Units: Conversely, the single layer perceptrons made up of value units that were trained to an SSE of 0.05 and 0.02 were most strongly activated by the S+. Value unit networks that were trained to the criterions

of 0.2 and 0.10 SSE showed shifts in peak activation to stimulus 6.




Figure 3. Shown are (top) the generalization gradients for sigmoid unit networks, and (bottom) gradients for networks trained with value units. Solid vertical line is the S+; dashed line is the trained S-.

Training to criteria based on number of trials vs. performance levels

Many studies of generalization and peak shift involve training participants for a specific number of trials (Bizo & McMahon, 2007; Baron, 1973; Galizio, 1980; Lewis & Johnston, 1999; Newlin et al., 1979; Thomas et al., 1991), rather than to specific discrimination performance criteria. Recent experimental data (Wisniewski et al., 2010) and the simulations presented here, however, showed that generalization differs strongly and nonlinearly with changes in discrimination performance. Therefore, training to a specific trial criterion may lead to more variability between participants in generalization than training to a performance criterion. To examine this possibility, single layer networks of value units were trained to the 6 criteria levels, rank-ordered from lowest to highest, and defined by the previously used SSEs or for a pre-specified number of training trials (10, 160, 320, 640, 1280, or 2560 trials). Learning rates of 0.01, 0.02, 0.04, and 0.08 were used to simulate individual differences in learning capacity. Three networks were trained per learning rate and per criterion level defined by SSE or number of training trials. The standard deviation of generalization gradient means was higher for networks that were trained for a specified number of trials than for networks that were trained to criterion

levels 1-4. Criterion levels 5 and 6 showed similar standard deviations for networks trained to a criterion defined by SSE and number of trials.

Discussion

The current simulations demonstrate that perceptron models of discrimination learning can replicate the quadratic trend for shifts in generalization gradients reported by Wisniewski et al. (2010), in which a peak shift effect emerges and then dissipates during the course of learning. However, networks constructed with a sigmoid output unit were not successful in capturing this trend. Sigmoid units have been popular in previous perceptron models of peak shift, which is understandable given that they predict peak shift and are consistent with theories of how stimulus characteristics are associatively reweighted after learning (more informative characteristics gain more weight) (Ghirlanda & Enquist, 1998; 2006; Livesey et al., 2005; Suret & McLaren, 2002; Tanaka & Simon, 1996). Here, however, they appear to result in stronger shifts with greater learning on the S+/S-discrimination. In contrast, networks trained with value units qualitatively replicated the quadratic trend in gradient shifts. Shifts in gradients only occurred after intermediate levels of performance were reached on the S+/S-discrimination. Discrimination performance on the S+/S-discrimination that was too poor, or too good, led to little or no shift in the generalization gradients of networks. Because the value, but not sigmoid units were successful in capturing the quadratic trend in generalization with perceptrons, it could be the case that, during learning, elements of a stimulus are not always reweighted in the manner proposed by previous theory. A connectionist architecture more complicated than a single-layer perceptron may be capable of simulating the behavioral data using only sigmoid units, but previously proposed perceptron models are not.

Value units may have been better for simulating the empirical results of Wisniewski et al. (2010) for a couple of reasons. First, the receptive fields of many neurons in cortex that code for the features of sound selectively respond to specific features of an input with response decreasing to properties that are dissimilar to those features (e.g., Elhilali et al., 2007; Linden et al., 2003). It could be that Gaussian activation functions simulate these types of receptive fields and that these receptive fields are important for discrimination training. Second, Wisniewski et al. (2010), and many others studying generalization (Baron, 1973; Bizo & McMahon, 2007; Lewis & Johnston, 1999; Spetch et al. 2004; Thomas et al., 1991; Wisniewski et al., 2009), used a task for which participants were told to make a response only to the trained stimulus and nothing else. In single-layer perceptron models Gaussian activation functions may be best for modeling types of tasks for which responses should be withheld in the presence of stimuli that are different from the trained stimulus in both directions on the dimension. In contrast, sigmoid functions may be best at modeling tasks for which participants are instructed to

generalize responses to novel stimuli as done in studies of the caricature effect (Tanaka & Simon, 1996).

We also found that generalization was more variable when networks were trained to a criterion based on the number of training trials than to a criterion based on discrimination performance. This finding suggests that when investigating trajectories of changes in generalization, the method of training for a certain number of trials (Bizo & McMahon, 2007; Baron, 1973; Derenne, 2010; Galizio, 1980; Lewis & Johnston, 1999; Newlin et al., 1979; Thomas et al., 1991; Wisniewski et al., 2009; 2010) may be less effective than training to a performance criterion (Spetch et al., 2004; Thomas, Svinicki, & Vogt, 1973; Wills & Mackintosh, 1998). In addition, there are large individual differences in how participants generalize (Landau, 1968; Guttman & Kalish, 1956; Nicholson & Gray, 1972). Some of the previously reported differences in generalization may stem from participants not reaching similar levels of performance.

Finally, the fact that we were able to simulate the temporal dynamics of human generalization with a model that can model songbird perceptual discriminations of naturally occurring stimuli strongly suggests that there are similar mechanisms for human and nonhuman generalization and learning (Ghirlanda & Enquist, 2006; Mercado, 2008; Spetch et al., 2004; Wisniewski et al., 2009). Some have argued that different mechanisms account for the peak shift effect in humans versus nonhumans (Bizo & McMahon, 2007; Thomas et al., 1991), whereas others have challenged this idea (Ghirlanda & Enquist, 2006; Livesey et al., 2005; Spetch et al., 2004; Suret & McLaren, 2002; Wisniewski et al., 2009). The fact that simple perceptron models can adequately model learning and generalization across species is consistent with the involvement of common underlying mechanisms.

In conclusion, the simulations presented here suggest that: 1) it is important to consider how stimuli are neurally coded as well as task design when simulating behavioral data from discrimination learning experiments; 2) equalizing subjects/participants on performance leads to less variable generalization than equalizing trial numbers; and 3) common learning mechanisms, shared between human and nonhuman species, likely mediate the peak shift effect.

Acknowledgments

This work was supported in part by NSF grant #SBE 0542013 to the Temporal Dynamics of Learning Center, an NSF Science of Learning Center. LMG is supported by an Izaak Walton Killam Memorial Scholarship at the University of Alberta. CBS was supported by a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, an Alberta Ingenuity Fund (AIF) New Faculty Grant, a Canada Foundation for Innovation (CFI) New Opportunities Grant along with start-up funding and CFI partner funding from the University of Alberta, Edmonton, Alberta, Canada.

References

- Baron, A. (1973). Postdiscrimination gradients of human subjects on a tone continuum. *Journal of Experimental Psychology*, 101(2), 337-342.
- Bizo, L. A., & McMahon, C. V. (2007). Temporal generalization and peak shift in humans. *Learning & Behavior*, 35(2), 123-130.
- Blough, D. S. (1975). Steady state data and a quantitative model of operant generalization and discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*, 1(1), 3-21.
- Cheng, K., Spetch, M. L., & Johnston, M. (1997). Spatial peak shift and generalization in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 23(4), 469-481.
- Dawson, M. R. W. (2004). *Minds and Machines: Connectionism and Psychological Modeling*. Malden, MA: Blackwell Publishing.
- Dawson, M. R. W. (2005). *Connectionism: A Hands-On Approach*. Malden, MA: Blackwell Publishing.
- Derenne, A. (2010). Shifts in postdiscrimination gradients within a stimulus dimension based on bilateral facial symmetry. *Journal of the Experimental Analysis of Behavior*, 93, 485-494.
- Elhilali, M., Fritz, J. B., Chi, T., Shamma, S. A. (2007). Auditory cortical receptive fields: Stable entities with plastic abilities. *Journal of Neuroscience*, 27(39), 10372-10382.
- Enquist, M., & Ghirlanda, S. (2005). *Neural Networks and Animal Behavior*. Princeton, NJ: Princeton University Press.
- Fernández-Redondo, M., & Hernández-Espinosa, C. (2001). Weight initialization methods for multilayer feedforward. *ESANN'2001 proceedings - European Symposium on Artificial Neural Networks* (pp. 119-124) Bruges, Belgium.
- Galizio, M. (1980). Conditional gradient displacements: The effects of conditional discrimination training on human auditory frequency generalization. *Journal of Experimental Psychology: Human Learning and Memory*, 6(6), 766-777.
- Ghirlanda, S., & Enquist, M. (1998). Artificial neural networks as models of stimulus control. *Animal Behaviour*, 56, 1383-1389.
- Ghirlanda, S., & Enquist, M. (2006). How training and testing histories affect generalization: a test of simple neural networks. *Philosophical Transactions of the Royal Society B*, B362, 449-454.
- Gluck, M. A., (1992). Stimulus sampling and distributed representations in adaptive network theories of learning. In A. Healy, S. Kosslyn, & R. Shiffrin (Editors), *From Learning Theory to Connectionist Theory: Essays in Honor of William K. Estes* (pp. 169-199). New Jersey: Lawrence Erlbaum Associates.
- Guillette, L. M., Farrell, T. M., Hoeschele, M., Nickerson, C. M., Dawson, M. R. W., & Sturdy, C. B. (2010). Mechanisms of call note-type perception in black-capped

- chickadees (*Poecile atricapillus*): Peak shift in a note-type continuum. *Journal of Comparative Psychology*, 124(1), 109-115.
- Guttman, N., & Kalish, H. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology*, 51, 79-88.
- Hanson, H. M. (1959). Effects of discrimination training on stimulus generalization. *Journal of Experimental Psychology*, 58, 321-334.
- Landau, J. S. (1968). Postdiscrimination generalization in human subjects of two different ages. *Journal of Experimental Psychology*, 76(4), 656-663.
- Lewis, M. B., & Johnston, R. A. (1999). Are caricatures special? Evidence of peak shift in face recognition. *European Journal of Cognitive Psychology*, 11(1), 105-117.
- Li, G., Alnuweiri, H., Wu, Y. (1993). Acceleration of backpropagations through initial weight pre-training with delta rule. *Proceedings of the IEEE International Conference on Neural Networks* (pp. 580-585). IEEE.
- Linden, J. F., Liu, R. C., Sahani, M., Schreiner, C. E. & Merzenich, M. M. (2003). Spectrotemporal structure of receptive fields in areas A1 and AAF of mouse auditory cortex. *Journal of Neurophysiology*, 90, 2660-2675.
- Livesey, E. J., Pearson, L. S., & McLaren, I. P. L. (2005). Spatial variability and peak shift: A challenge for elemental associative learning? *Proceedings of the XXVIIth Annual Convention of the Cognitive Science Society*. (pp. 1302-1307) Hillsdale, NJ: Lawrence Erlbaum Associates.
- McLaren, I. P. L., & Mackintosh, N. J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. *Animal Learning & Behavior*, 30(3), 177-200.
- Mercado, E., III (2008). Neural and cognitive plasticity: From maps to minds. *Psychological Bulletin*, 134(1), 109-137.
- Moye, T. B., & Thomas, D. R. (1982). Effects of memory reactivation treatments on postdiscrimination generalization performance in pigeons. *Animal Learning & Behavior*, 10(2), 159-166.
- Newlin, R. J., Rodgers, J. P., & Thomas, D. R. (1979). Two determinants of the peak shift in human voluntary stimulus generalization. *Perception & Psychophysics*, 25(6), 478-486.
- Nicholson, J. N., & Gray, J. A. (1972). Peak shift, behavioral contrast, and stimulus generalization as related to personality and development in children. *British Journal of Psychology*, 63(1), 47-62.
- Purtle, R. B. (1973). Peak shift: A review. *Psychological Bulletin*, 80(5), 408-421.
- 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.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237, 1317-1323.
- Shepard, R. N. (1990). Neural nets for generalization and classification: Comment on Staddon and Reid (1990). *Psychological Review*, 97(4), 579-580.
- Spence, K. W. (1937). The differential response in animals to stimuli varying within a single dimension. *Psychological Review*, 44, 430-444.
- Spetch, M. L., Cheng, K., & Clifford, C. W. G. (2004). Peak shift but not range effects in recognition of faces. *Learning and Motivation*, 35(3), 221-241.
- Staddon, J. E. R., & Reid, A. K. (1990). On the dynamics of generalization. *Psychological Review*, 97(4), 576-578.
- Suret, M., & McLaren, I. P. L. (2002). An associative model of human learning on an artificial dimension. *Proceedings of the 2002 International Joint Conference on Neural Networks* (pp. 806-811). Piscataway, NJ: IEEE.
- Tanaka, J. W., & Simon, V. B. (1996). Caricature recognition in a neural network. *Visual Cognition*, 3(4), 305-324.
- Thomas, D. R., Mood, K., Morrison, S., & Wiertelak, E. (1991). Peak shift revisited: A test of alternative interpretations. *Journal of Experimental Psychology: Animal Behavior Processes*, 17(2), 130-140.
- Thomas, D. R., Svinicki, M., & Vogt, J. (1973). Adaptation level as a factor in human discrimination learning and stimulus generalization. *Journal of Experimental Psychology*, 97(2), 210-219.
- Verzijden, M. N., Etman, E., van Heijningen, C., van der Linden, M., & ten Cate, C. (2007). Song discrimination learning in zebra finches induces highly divergent responses to novel songs. *Proceedings of the Royal Society B*, 274, 295-301.
- Watson, J. B., & Raynor, R. (1920). Conditioned emotional reactions. *Journal of Experimental Psychology*, 3(1), 1-14.
- Wills, S., & Mackintosh, N. J. (1998). Peak shift on an artificial dimension. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, 51(1), 1-31.
- Wisniewski, M. G., Church, B. A., & Mercado, E., III (2009). Learning-related shifts in generalization gradients for complex sounds. *Learning & Behavior*, 37(4), 325-335.
- Wisniewski, M. G., Church, B. A., & Mercado, E., III (2010). Temporal dynamics of generalization and representational distortion. *Psychonomic Bulletin & Review*, 17(6), 809-814.
- Withagen, R., & van Wermeskerken, M. (2009). Individual differences in learning to perceive length by dynamic touch: Evidence for variation in perceptual learning capacities. *Attention, Perception, & Psychophysics*, 71(1), 64-75.