A Computational Model of Auditory Perceptual Learning:

Predicting Learning Interference Across Multiple Tasks

Technical Report: NWU-EECS-11-01

David Little (d-little@u.northwestern.edu)
Bryan Pardo (pardo@northwestern.edu)
Beverly Wright (b-wright@northwestern.edu)

May 5, 2011

Abstract

In this work we build a computational model of several auditory perceptual learning experiments. The modeled experiments show a pattern of learning interference which may help shed light on the structure of both short and long term stores of perceptual memory. It is our hypothesis that the observed interference patterns can be explained by the relationship of stimuli across tasks and how these relationships interact with the limits of human memory. We account for the fact that information is shared across tasks in our model through use of methodology from the machine learning community on transfer learning. When we introduce a set of plausible limits on memory, such a model demonstrates the same pattern of learning interference observed in the human experiments.

Keywords: Perceptual Learning; Perceptual Memory; Consolidation; Acquisition; Learning Interference; Transfer Learning

1 Introduction

Note to reader: This is an extend technical report of a conference paper accepted to 33rd Annual Conference of the Cognitive Science Society in Boston, MA. It includes the full body of the conference paper along with some additional implementation details in the appendices.

With sufficient practice, human beings are able to enhance the acuity of their sensory systems. This is known in the literature as perceptual learning. Recent work in perceptual learning (e.g. Banai et al., 2009; Yotsumoto et al., 2008), has shown that learning on one task (which we call the *target*) may be prevented when a second task (which we call the *distractor*) is practiced either during or shortly after practice of the target: this is called *learning interference*. These results suggest distinct properties of short and long term stores of perceptual memory because what interfered with learning during practice was distinct from what interfered after practice (see the Human Data section for more detail).

Our working hypothesis is that the learning interference observed in these experiments is a consequence of how information is shared across tasks and the limits of human memory. We have built a computational model in an effort towards fully specifying and testing this hypothesis (see the Modeling section for details). An ideal observer would only benefit from sharing information across tasks. However, with the introduction of limited memory, sharing information can also lead to learning interference.

Such sharing of information across tasks is used to accomplish *transfer learning* in the machine learning community. We call a computational technique intended to accomplish transfer learning, *computational* transfer learning. If a system (living or machine) can be seen to have better performance on one task after experience on some prior task, we call this *observable* transfer learning. Prior computational models of perceptual learning, though they have considered observable transfer learning, have ignored matters of computational transfer learning, either by modeling only a single task (e.g. Jacobs, 2009) or by treating learning across several tasks as a single monolithic learning problem (e.g. Petrov et al., 2005). Because of this, none of these models provide an account of how people appropriately segregate and share information across tasks. There are computational models concerned with human memory that can be understood to have some form of computational transfer learning (e.g. McClelland et al., 1995; Anderson, 2002), but these systems do not provide the detail needed to model the current experiments.

In this paper we model one set of learning interference experiments (Wright et al., 2009; Banai et al., 2009) using an ideal observer (Geisler, 2003). We do this by incorporating a method used for computational transfer learning (Roy & Kaelbling, 2007) (see the method section for details). On top of this ideal observer, we introduce a plausible set of memory limits. This approach has the merit of avoiding conflation between task constraints (which both humans and the ideal observer are subject to) and psychological constraints (which only humans are subject to). We hypothesize memory limits that a.) affect the number of distinct stimuli that could be remembered and that b.) introduce a process of consolidation, meaning that over a period of time memories move from a labile, short term form to a stable long term form. We found that when introducing all (and only all) of our limits, our model demonstrated the same pattern of learning interference observed in humans (see the evaluation for details).

2 Human Data

The experiments in Banai et al. (2009) and Wright et al. (2009) suggest two functionally distinct stages of perceptual learning. The first stage occurs *during* practice of a task. We call this stage *acquisition*. The second stage occurs *after* practice is complete and is called *consolidation*. This is supported by the way the *target* task, T1 (see Figure 1(a)) was interfered with. One task (T2) interfered *during* practice of the target but not afterwards, and the task (F1) interfered *after* practice of the target but not during. This dissociation between acquisition and consolidation makes the experiments interesting to model: straightforward interpretations for one half of the data can lead to contradictory predictions of the remaining data.

A point of clarity: throughout our work here we use the word *learning* to mean that some form of memory (either human or computational) is updated to reflect a new experience. When we need to make a distinction we use the term *observable learning* to mean a behaviorally observable improvement in task performance.

The conditions in this experiment involved three tasks: one target (T1) and two distractors (T2 and F). The target is an interval discrimination task. By the term *task* we mean a specific set of stimuli, and the responses expected for these stimuli (in perceptual learning T1 and T2 would often be referred to as the same task). In the target task (T1), the participant had to make a two interval forced choice, indicating which of the two presented stimuli contains a longer temporal interval: the stimuli for the task are shown on the first row of Figure 1(a). Participants heard the two stimuli in a randomized order and received feedback after each trial. The stimuli in T1 each contained two short sinusoidal tones at 1000 Hz, separated by a temporal interval that varied in length. One stimulus (called the *standard*) always contained a 100ms interval. The other stimulus (called the *comparison*) varied in length. The difference between the standard and the comparison is called the *delta*. Over the course of a block (60 trials), the delta was adjusted so that a subject's *threshold* was found. The threshold is the delta at which a person gets 79% of their responses correct (Levitt, 1971).

The two distracting tasks are related to the target in distinct ways. Task F (second row of Fig. 1(a)) is a frequency discrimination task meaning that instead of varying the interval of the comparison its frequency was varied over the course of learning. Task T2 (last row of Fig. 1(a)) was a second temporal interval discrimination task, where the standard was 350ms. All stimuli in task F have the *same* temporal interval as the standard of T1 (shown on the second row of Fig. 1(a)), and all stimuli in task T2 contain *distinct* temporal intervals from those present in T1. This is shown in Figure 1(a): a box is drawn around all stimuli that contain a 100ms temporal interval. All other stimuli contain a distinct interval.

There were four conditions in which one of the two distracting tasks was introduced either during acquisition of T1—by interleaving practice with T1—or during consolidation of T1—by presenting it in a block after T1 (also called *blocked* presentation). Figure 1(b) shows that task T2 interfered with the observable learning of T1 during *acquisition* and task F interfered with the observable learning during *consolidation*. Observable learning was said to occur if a subject showed a significantly greater improvement in their threshold, when compared to controls. Controls perform only a pre- and post-test. Participants performed a pre-test, at least 6 days of practice, and then a post-test.

Our hypothesis is that F prevents observable learning on T1 during consolidation because stimuli in T1 and F contain the *same* temporal interval. T2 prevents observable learning during acquisition because T1 and T2 have *distinct* temporal intervals. F1 and T2 place distinct strains on human memory which manifest as a different pattern of learning interference.

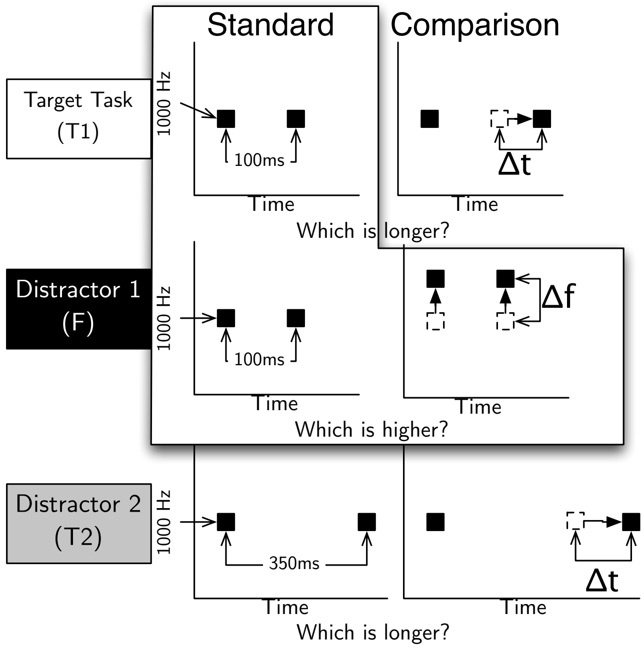

(a) The tasks performed by participants. Interval discrimination at 100 and 350 ms, and frequency discrimination at 1 kHz. The stimuli surrounded by a box all consist of the same interval. The other types of stimuli have distinct intervals.

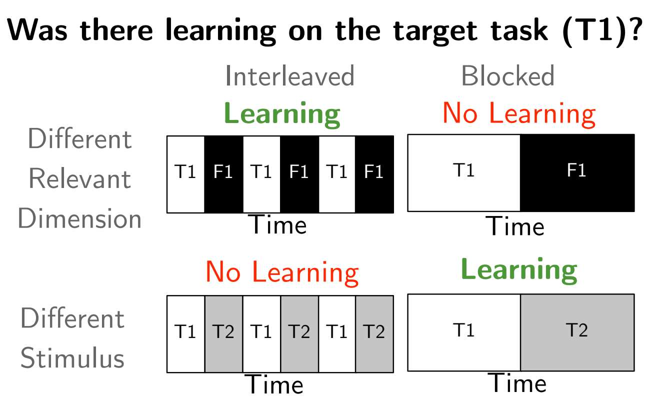

(b) A diagram of the results. Blocks represent number and type of tasks presented on a single day of practice. T2 only interfered with observable learning of T1 during acquisition, and F only interfered with observable learning of T1 during consolidation.

Figure 1: A summary of the results from Wright et al. (2009) and Banai et al. (2009)

3 Modeling

Our model provides one explanation for why *distinct* temporal intervals across tasks would lead to interference during acquisition, and why having the *same* temporal intervals would interfere during consolidation. The idea is that some part of our memory cares solely about intervals, and this is the locus of learning. During acquisition having too many distinct temporal intervals means that there is too much to keep track of; during consolidation having the same temporal interval across tasks prevents consolidation of the first task because the memories are too similar.

It is certainly possible that features other than temporal interval are relevant to the observed interference. However, our model is a demonstration that by using a set of plausible limits on human memory, these features are not *necessary* to explain the human data. To show this we built an ideal observer (Geisler, 2003) of our tasks but made use of only the temporal information in the stimuli. An ideal observer defines what "optimal" behavior is given the same information that humans have to perform a task. Our observer is ideal in the sense that it makes optimal use of the temporal information available in the stimuli. The ideal observer is useful as a baseline to compare to human performance. It is not intended to be psychologically plausible. On top of this ideal observer we introduce a set of memory limits.

Key to the observed learning inteference is our model's *item limit* and *recall limit*. During acquisition the number of distinct stimuli that can be represented in memory is limited (the item limit), this limitation leads to inteference during acquisition when there are many distinct temporal intervals in the stimuli. During consolidation stimuli that resemble each other can cause a memory previously marked for long term storage to be returned to short term memory (the recall limit) leading to interference during consolidation when stimuli have similar temporal intervals across tasks. More details of our memory limits, and their justification are discussed in the subsection Hypothesized Psychological Limits.

We center our discussion of the model around the concept of a *stimulus model*. We start by describing the input provided to our model. We define the meaning of a stimulus model and how it relates to the input during decision making. Then we discuss how the model input is used to learn a better stimulus model, and how the psychological limits affect the results of learning. For full implementation details of our model we refer the reader to the appendices.

3.1 Model Input

As shown in Figure 2, each auditory stimulus, s, presented to the model is transformed to an internal representation $\mathbf{x}$ by the function $R(s)$. The ideal observer is meant to find the best possible decision, given the same information people have. Thus $R(s)$ should be consistent with our understanding of the pertinent information people have to make a decision. We assume for modeling purposes that the data can be explained solely in terms of the intervals present in a stimulus, so this is the only information present in $R(s)$.

The input to $R(s)$ is an audio file and the output is a 32 term vector describing the

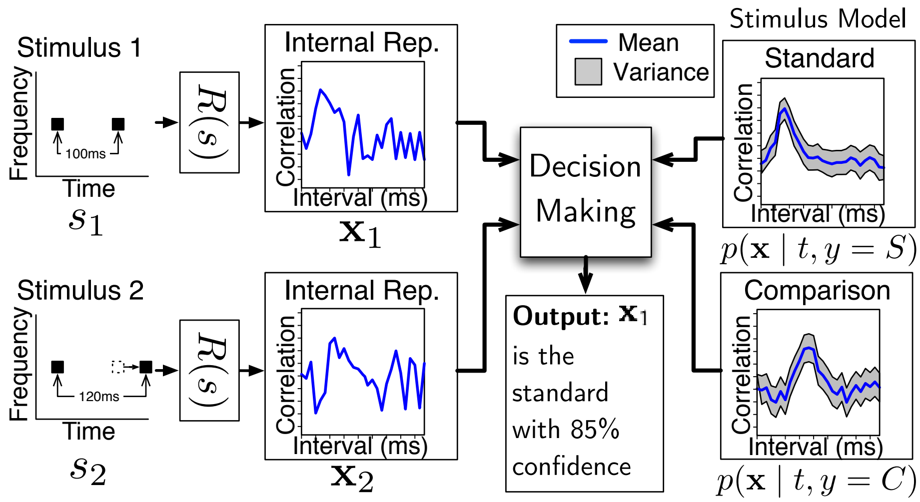

Figure 2: The input to and the output from the ideal decision maker. Arrows towards a box represent input, arrows away from a box represent output. The decision maker is presented two stimuli (s_1 and s_2), transformed according to $R(s)$. Along the x-axis of each $\mathbf{x}$ are the intervals from 10 to 1000ms on a log scale. Along the y-axis is the correlation of onsets in the stimulus to a particular interval. The decision maker is given the distribution of the standard ($p(\mathbf{x} \mid t, y = S)$) and the comparison ($p(\mathbf{x} \mid t, y = C)$) as determined by the ideal learner. These distribuions are called stimulus models and are depicted along the same axes as the input.

temporal intervals present in the stimulus[1]. $R(s)$ applies a windowed auto-correlation function over the onsets in the audio file s, where the window is always proportional to the length of the interval in question. The result of this process is consistent with the model of human interval perception presented in (Buonomano, 2000), in the case where the input contains a single interval. The use of this representation is also supported by the fact that learning on temporal intervals does not generalize to untrained intervals, or to other tasks using the same standard (Wright et al., 1997).

A Gaussian random value is added to each term of the representation $\mathbf{x}$, with an experimentally determined standard deviation σ. This reflects the noise present in sensory systems.

In Figure 2 we represent an observed stimulus with a graph showing all 32 terms of $\mathbf{x}$. Each term corresponds to a time interval from 10 to 1000ms along a log scale (shown along the x-axis of $\mathbf{x}_1$ in Fig. 2). The value at each term of the vector (shown along the y-axis of $\mathbf{x}_1$) corresponds to the correlation in the stimulus to that particular interval: the highest peak in $\mathbf{x}_1$ is near 100ms, because the original stimulus s_1 has a 100ms interval in it.

[1]Note that for reasons of speed, the number of terms in $R(s)$ (32) was chosen to be the smallest number that clearly prevented quantization error from being a limiting factor of model performance.

3.2 Stimulus Models and Decision Making

Intuitively, a stimulus model can be understood as a perceptual template. During decision making, each stimulus ($\mathbf{x}$) is compared to these templates to determine which observed stimulus is most like the standard (e.g. the shorter interval in task T1) and which is most like the comparison (e.g. the longer interval in task T1). Formally, a *stimulus model* is a probability distribution over an internal representation of a stimulus ($\mathbf{x}$) conditioned on a particular task (t = T1,T2 or F) and stimulus type (y = Standard(S) or Comparison(C)). Bayes rule can be applied to these distributions to find the probability that the first observation ($\mathbf{x}_1$) is the standard. A decision is then made by choosing the most probable answer (i.e. "The standard was first." or "The comparison was first.").

3.3 Learning Stimulus Models

Each stimulus model is learned by a processes that can be understood as an averaging over many observations of $\mathbf{x}$. As more stimuli are observed, this average becomes more accurate, leading to more accurate decisions on the part of the model. If there have been no observations presented to the model, then the response given is random. In Figure 2 a standard and comparison stimulus model are shown. The graphs of these stimulus models are along the same axes as the input, and show the mean and variance of the distribution of $\mathbf{x}$ for the given stimulus type.

For the ideal observer, learning occurs after each trial. The input consists of a series of observations. Each trial of a task is a two interval forced choice, meaning that there are two stimuli. Each observation corresponds to one of the two stimuli in a trial, and includes the stimulus ($\mathbf{x}$), the correct label for the stimulus (y)—either standard or comparison—and the task (t) that the stimulus was presented during—T1, T2 or F. The correct label is determined by the feedback provided at the end of a trial. The output of learning is a stimulus model for each task and stimulus type. The goal of learning is to determine how to update stimulus models such that they accurately reflect future observations, leading to better decisions.

To accomplish computational transfer learning for the experiments in question, our ideal observer *learns* which stimuli are drawn from the same distribution. This appears to be the only way in which tasks are relevant to each other in this set of experiments. As noted in the Human Data section, stimuli in task F have the same 100ms interval that the standard has in T1. Because only interval information is represented, our ideal observer represents these three sets of stimuli with the same distribution. The remaining types of stimuli (the standard and comparison for T2 and the comparison for T1) follow their own distinct distribution. Note that although the length of the comparison within each task varies over the course of learning, all comparisons within a task can be represented with the same stimulus model. In the ideal observer the sharing of distributions across tasks only improves learning (at least for the modeled tasks).

Due to considerations of space we do not provide a detailed explanation of how our model learns which stimuli across tasks share a distribution (see the appendices for more details). In short we make use of a Dirichlet processes prior to cluster similar stimuli. This basic approach to transfer learning has been considered elsewhere (e.g.

Roy & Kaelbling, 2007). We assume that observations follow a Dirichlet process prior with a base distribution where each input $\mathbf{x}$ is distributed according to a multivariate Normal distribution, t a Bernoulli distribution and y a Bernoulli conditioned on t.

Prior to observing any trials of a given task humans are capable of above chance performance on the tasks. To represent this prior knowledge we initialize our model by presenting it an experimentally-determined number of trials. Psychological limits are not introduced until after these initial trials, meaning that learning during initialization is optimal. The trials presented during initialization had a comparison whose delta varied around a mean m and standard deviation s, both of which were determined empirically.

3.4 Hypothesized Psychological Limits

We hypothesize four limits on top of our ideal observer for the modeled tasks. To distinguish it from the ideal observer, we refer to the full model as the *Limited Memory Model*. All versions of our model made an ideal decision given the stimulus models they were provided, but the way these stimulus models were learned was not always optimal.

Our choice to express limits in terms of stimulus models means that we are assuming people have something like a stimulus model in their brain: this is a reasonable assumption because to learn anything about a task, the stimuli from the task must be remembered, and a stimulus model is simply a compact representation of previously observed stimuli.

3.4.1 Single Task Limits

There are two limits that apply during the learning of a single task.

The first limit we call the *volatility limit*. It states that during acquisition, trials are represented in a short term store. Stimulus models in this store are said to be *volatile*. Volatile stimulus models decay according to a loss parameter L. Thus, instead of being an average, a volatile stimulus models is more like a moving average. Because of this decay the *effective* number of trials that a voltile stimulus model represents will depend on the rate at which stimuli are presented. The more time that passes without observing more trials, the fewer effective number of stimuli a volatile stimulus model represents. There is evidence suggesting a distinction between short and long term stores of memory and that this short term store is transient (e.g. Izquierdo, 1999; Cowan, 2008). Recent work has shown that when trials are separated this appears to affect the effective number of trials a subject has observed (Zhang & Wright, 2010).

The second limit we call the *consolidation limit*. It states that after a short period of time (15 simulated minutes) during which trials for a task have not been observed, all volatile stimuli with a sufficient *effective* number of trials T are copied to a long term store. (In the full model this not an instantaneous processes, see the *recall limit*). Stimulus models in the long term store are said to be *consolidated*, and do not decay anymore. There is evidence both for a period of memory consolidation (McGaugh, 2000) and that this consolidation does not occur unless enough trials within each day are observed (Wright & Sabin, 2007). During decision making the stimulus models present in the long term store (not the short term store) are used. The model works this way because there is no *observable* learning within a day of practice for the modeled tasks (Wright & Sabin, 2007).

3.4.2 Multiple Task Limits

The third limit we call the *item limit*. It limits the effective number of stimulus models (or items) allowed in the short term store. Specifically it states that decay (L) is proportional to the total number of volatile stimulus models. This item limit is consistent with the notion that short term memory can only effectively store a limited number of items (e.g. Cowan, 2008). This limit explains why learning fails during the interleaved practice of T1 and T2, but not during interleaved practice of T1 and F. There are four distinct stimulus models when practicing T1 and T2 (the standard and the comparison for both tasks), all of which are volatile during interleaved practice: this means stimulus models decay too quickly and so the effective number of trials is never large enough for consolidation to occur. There are only two distinct stimulus models during interleaved practice of T1 and F (since there are two distinct intervals across these tasks), and so much less decay occurs, allowing consolidation.

The fourth limit we call the *recall limit*. It states that there is a period of time before models become fully consolidated when a stimulus model is being moved from the short to long term store. During this period, in which the model is said to be *transferring*, the stimulus models can be *recalled*, meaning they return to a volatile state. At this point they will only be consolidated for the same reasons that any volatile stimulus model is consolidated. This recall occurs when a newly observed stimulus belongs to one of the transferring stimulus models. In our model stimulus models move from a *transferring* to a *consolidated* state at the end of a simulated day. This limit is consistent with the idea that consolidation is not an instantaneous process: more permanent memories are formed over extended periods of time, and before consolidation is complete, it can be interrupted (e.g. McGaugh, 2000).

The recall limit explains why learning is interfered with during blocked practice of T1 and F, but not T1 and T2. When task F begins, T1 begins to be consolidated, and so T1's stimulus models are transferring. However, task F shared a stimulus model with T1 and so all the transferring stimulus models are recalled. During blocked practice of T1 and T2, T2 shares no stimulus models with T1, and so the stimulus models of T1 can safely transfer from the short to the long term store. Note that T1 and F must also be consolidated during interleaved practice, and so the reader might view the recall limit as preventing learning in this case: however, because the consolidation limit states that consolidation begins shortly after a task is complete, and consolidates *all* stimulus models with sufficient trials, both tasks' stimulus models are consolidated as a single unit in this case.

4 Evaluation

The purpose of our evaluation was to demonstrate that our limited memory model qualitatively matched the learning interference patterns observed in Wright et al. (2009) and

Banai et al. (2009) and that this behavior of the model was due to all of our hypothesized limits.

To evaluate the hypothesized psychological limits we compared six different models: the ideal observer, the limited memory model—which included all hypothesized limits—and four more versions, each with one of the limits removed. If all limits are necessary to explain the data then all but the full model should fail to predict when learning interference will occur for humans.

We simulated the experiments from Wright et al. (2009) and Banai et al. (2009) in the following way. For each task there were 60 trials per block and 6 blocks per day of practice. For each condition we ran 11 simulations (to simulate 11 participants). There was noise present in every stimulus, which meant each simulation of the experiment was different. We used 11 simulations for each condition because this is the maximum number of subjects for any condition used in Wright et al. (2009) and Banai et al. (2009). For each simulation we presented the stimuli to the model, following the same adaptive tracking procedure (to find the model's threshold). The model provided the response it predicted to most likely be the correct response. After all 360 trials for each task were presented for a "day" the model was allowed to "sleep". During this period of the simulation the system finished consolidation of any stimulus models still transferring, and all volatile stimulus models were fully forgotten if any decay was present. In this way the simulation of the trained conditions was made as parallel as possible to the human experiment.

	Condition			
	Interleaved		Blocked	
Learner	T1/F	T1/T2	T1/F	T1/T2
Human	X	-	-	X
Ideal	X	X	X	X
LMM	**X**	**-**	**-**	**X**
LMM - volatile	X	X	X	X
LMM - consolidated	-	-	-	-
LMM - item	X	X	-	X
LMM - recall	X	-	X	X

Table 1: Qualitative results across all learners. An X in a column indicates that the given learner showed observable learning on task T1 when interleaved or blocked with the specified task. LMM stands for the limited memory model, and LMM - *L* indicates that limit *L* was removed from the LMM model.

We simulated 11 control subjects by running two blocks (60 trials each), where no learning step was performed. This differed from the procedure used for control subjects in Wright et al. (2009) and Banai et al. (2009) in that some learning may have occurred during the pre- and post-test. This was because our model only simulated behavior during days of learning, not the pre- and post-test behavior. Model parameters (e.g. input noise σ) were held constant across all computational models.

If there was a significantly greater difference from pre- to post-test of a model, compared to control subjects, for a given condition, the model was said to have learned on this condition. This was determined by performing a two time (day 1 to day 6) by two group (trained vs. control) ANOVA, with time as a repeated measure. Table 1 summarizes the results for all models. For all simulations marked with an X, p values were below 0.013, and all dashes where above 0.18. The results for the human data are taken from the prior analysis in Wright et al. (2009) and Banai et al. (2009). This table shows that, among the models we tested, only the limited memory model shows the same pattern of learning interference that humans showed.

5 Conclusions

In closing, we have presented a framework from which a variety of learning interference experiments might be modeled and studied, and have shown that this framework is capable of predicting the qualitative results of one challenging set of human data. Our work was grounded in the hypothesis that learning interference was an effect of how information is shared across tasks and the limits of human memory.

The model provides concrete predictions concerning future experiments. It predicts that if two tasks are interleaved they will interfere if there are many distinct stimuli across tasks. It predicts that during blocked presentation interference can occur when there are identical or very similar stimuli used across tasks. This is a consequence of the item limit, which limits how many distinct stimuli can be remembered at one time and the recall limit, which prevents consolidation of one task when a new task contains similar stimuli. These limits in turn have implications for the form and function of short and long term stores of perceptual memory.

Acknowledgments

We would like to thank the anonymous reviewers of our draft manuscript, Mark Cartwright, Zhiyao Duan, Jinyu Han, Eric Hoover, Andrew Lovett, Alex Madjar, Nicole Marrone, Zafar Rafii, Andy Sabin and Matthew Waggenspack for their helpful feedback. This research was supported, in part, by Northwestern University's Cognitive Science program and by US National Science Foundation grant 0643752.

A Internal Representation Implementation

This appendix describes the representation used by our model. This is found using the function $R(s)$. $R(s)$ takes in an audio file and outputs a 32 term vector $\mathbf{x}$ where each term corresponds to an interval from 10ms to 1000ms along a log scale. This is produced from the initial audio file according the following steps.

Initial Representation Each audio file is given to our front end as a PCM audio file at a sampling rate of 8000Hz, which we denote as a.

Auditory Model The audio file is transformed into a spectrogram-like representation according to the model of the peripheral auditory system described in Wang &

Shamma (1994). The model uses a set of logarithmically spaced asymmetric band-pass frequency filters resembling the observed shape of human auditory filters in the cochlea. The resulting output is a matrix $F(a)$, where each row $F(a)_{i,\bullet}$ is the response of a single frequency filter across all points in time, and each column $F(a)_{\bullet,j}$ is a single point in time across all filters. There are several parameters for this model discussed in Wang & Shamma (1994): we use a frame length of 2ms, an integration constant of 4ms, and no compression.

"Onset" Detection To identify intervals onset detection is used. The output is a vector whose ith term is denoted $o[i]$. Given the very simple nature of our audio input we use the following formula:

$$o[i] = \max\left\{0, \sum_j F(a)_{j,i+1} - \sum_j F(a)_{j,i}\right\} \tag{1}$$

Temporal Analysis A temporal analysis over the onsets was performed as an intermediate step to encode the intervals in a stimulus. The analysis resembled an autocorrelation function. For lag y and time t a single real number is found according to the following formula.

$$T[y,t] = \frac{1}{\|W_y\|}(o * W_y)[t] \times (o * W_y)[t-y] \tag{2}$$

In Equation 2 the '$*$' denotes convolution, and $(o * W_x)$ is the resulting function of convolvling the onsets with W_x, which is a Gaussian window. The window was determined using the MATLAB function `gausswin`, with a length equal to $\lfloor y \cdot 0.9 \rfloor$. $W_y[n]$ is defined as follows.

$$W_y[n] = C \cdot \exp\left[-\frac{1}{2}\left(2.5\frac{n - \lfloor y \cdot 0.45 \rfloor}{\lfloor y \cdot 0.45 \rfloor}\right)^2\right] \tag{3}$$

The value n ranges from 0 to $\lfloor y \cdot 0.9 \rfloor$. The constant C is set such that the euclidean norm of W_y is equal to one. This ensured responses were on the same scale across all values of y. We evaluated $T[y,t]$ for 32 values of t evenly spaced along a logarithmic scale from 10ms to 1000ms. Research suggests that intervals from about 10ms to 1000ms are represented qualitatively differently than lengths above and below this range (Buonomano & Karmarkar, 2002). The number 32 was chosen to be the smallest number that clearly prevented quantization error from being a limiting factor of model performance.

Interval Representation The final interval representation was a 32 term vector $\mathbf{x}$, where term $x_y = \max_t \{T[y,t]\}$. A normally distributed zero mean random value was added to each term in this vector, whose standard deviation, σ, was experimentally determined.

B Ideal Observer Learning

We explain learning in the ideal observer in three steps. We begin by describing how learning would work without transfer learning. We then describe a simplified version of transfer learning, and then describe the actual method used for learning. This serves as a gentle introduction to how transfer learning can be accomplished using the Dirichlet process prior: very similar approaches have been described in Roy & Kaelbling (2007) and Xue et al. (2007), for example. Throughout the following description we use a lower case letter to indicate a random variable and a subscripted lower case letter to indicate a specific observation from this random variable.

Recall that each observation during learning consists of a stimulus $\mathbf{x}$, a type y and a task t. We denote the set of all previous observations $\{x_i, y_i, t_i\}$, as D. Without considering transfer between tasks, learning a stimulus model is a matter of finding the following probability distribution over some new observation with index n.

$$p(\mathbf{x}_n \mid y_n, t_n, D) = p(\mathbf{x}_n \mid D_{y_n, t_n}) \tag{4}$$

In Equation 4 $D_{y,t}$ denotes the set of all observations with index i such that $y_i = y$ and $t_i = t$. Assuming that $\mathbf{x}$ is drawn from a Normal distribution, and that our prior over this Normal distribution is conjugate, Equation 4 can be calculated for a given data set using standard results for the exponential family of distributions (Gelman, 2004). The assumption of conjugacy is made merely as a matter of convenience. Whether this assumption is actually sensible will depend on the particular application. In our case we chose to use conjugacy lacking any strong evidence that anything more complicated was necessary.

We now introduce a simplified form of transfer learning. Recall that for our tasks the key objective of transfer learning is to identify which stimuli follow the same distribution (e.g.. the standard of T1, and the standard and comparison of F all follow the same distribution). We first explain transfer learning as if the distribution that each stimulus comes from is known. In this case each observation includes the value z_i, which is a natural number indicating which distribution the observation x_i belongs to. The the value of z_i can range form 0 to N where N is the number of observations. Learning is now a matter of calculating the following distribution.

$$p(\mathbf{x}_n \mid z_n, y_n, t_n, D) = p(\mathbf{x}_n \mid D_{z_n}) \tag{5}$$

In Equation 5, D_{z_n} denotes all observations with index i such that $z_i = z_n$, and this is calculated, as above, assuming a conjugate prior for the distribution of $\mathbf{x}$.

In the above description z is an observed variable. In reality this is information the ideal observer does not have during learning, and so it must infer the value of each z_i given each observation (x_i, y_i, t_i). In a Bayesian context this means assuming some distribution over the set $Z = \{z_1, \ldots, z_i, z_n\}$, using Bayes rule to find the posterior distribution of Z, and then marginalizing over Z, as follows.

$$p(\mathbf{x}_n \mid y_n, t_n, D) = \sum_Z p(\mathbf{x}_n \mid D_{z_n}) p(z_n \mid y_n, t_n, D, Z \setminus z_n) p(Z \mid D) \tag{6}$$

$$p(z_n \mid y_n, t_n, D, Z \setminus z_n) \propto p(y_n, t_n \mid D_{z_n}) p(z_n \mid Z \setminus z_n) \tag{7}$$

The sum in Equation 6 is a sum over all possible assignments to the variables in Z. To calculate Equation 6 a distribution over Z, y and t, must be defined. These will be given below. The sum in Equation 6 is computationally intractable, so in practice we use a Sequential Monte Carlo technique introduced in Fearnhead (2004). This means that we essentially perform a randomized greedy search over some limited number of clusterings (M): as each new observation is introduced, the possible ways this observation could be added to each clustering becomes a new possible clustering to be considered. This increases the number of possible clusterings considered. To avoid an exponential explosion of possible clusterings, we sample back down to M clusterings at each step, favoring the more probable clusterings.

To calculate Equation 6 we must assume a distribution over Z, y and t, and their priors (even in a non-Bayesian context such an assumption would be made during clustering, it just wouldn't be explicit). Below we include the distribution over $\mathbf{x}$ for completeness.

$$\mathbf{x} \mid z \sim \mathcal{N}(\mu_z, \Sigma_z) \qquad t \mid z \sim \text{Bernoulli}(T_z) \tag{8}$$

$$y \mid z, t \sim \text{Bernoulli}(Y_{z,t}) \qquad z \sim GEM(\alpha) \tag{9}$$

The distribution $GEM(\alpha)$ is a specific formulation of the Dirichlet process. We refer the reader to Neal (2000) for an introduction to Dirichlet processes. The distribution over t is assumed to take only two values (and hence a Bernoulli) because only two tasks are seen in any experimental condition. We chose to condition y on t because the label of a stimulus in one task may change when it is used in another task. For the above parameters to be learned, we must assume some distribution over T, Y, μ and Σ.

We assume the following priors over T, Y, μ and Σ.

$$\mu, \Sigma \sim \mathcal{NIW}(\mathbf{m}, r, S, d) \qquad T \sim W(a) \tag{10}$$

$$Y \sim W(b) \tag{11}$$

$$W(x;a) = (0.5 - a)\delta(x - \varepsilon) + a\delta(x - 0.5) + (0.5 - a)\delta(x - (1 - \varepsilon)) \tag{12}$$

In the above $\mathcal{NIW}$ denotes the Normal inverse Wishart distribution, which is the conjugate prior to the Normal distribution (Gelman, 2004). The distribution $W(x;a)$ is defined using the dirac delta distribution δ. We define ε to be $\exp(10^{-20})$. This prior encodes the assumption that a distribution is either shared across tasks (in which case the distribution of t is evenly divided between tasks) or it is not shared (in which case it is very nearly always), like-wise, distributions either make a distinction between a label (e.g. the data is relevant to the task), or they don't (e.g. the data is not relevant to the task). We found that these hard and fast distinctions lead to improved performance of our ideal observer for the modeled tasks.

The final model includes the following free parameters: a, b, $\mathbf{m}$, r, S, d, and α. In the current experiments these parameters were selected by hand to qualitatively match the human data when all computational limits were present.

C Memory Decay

Here we describe how memory decay is implemented in our limited memory model. This memory decay is used to implement the effects of the *volatility* and *item* limits. In the above model, a stimulus model given z is defined by the distribution of all observations $(\mathbf{x}_i, y_i, t_i, z_i = z)$. Because these distributions come from the exponential family (Gelman, 2004), the posterior over the distributions parameters ϕ can be expressed in the following form.

$$p(\phi \mid D_z) = h(\phi)\exp\{\eta(\phi) \cdot \sum_i T(x_i, y_i, t_i) - nA(\phi)\} \tag{13}$$

In the above equation $T(x_i, y_i, t_i)$ is called the sufficient statistic of the ith observation, and n is the number of observations seen so far. Thus, to find the posterior over ϕ we use the number of observations n and a sum of the sufficient statistics of all observed stimuli from z. A decay is applied the sufficient statistics, to reduce the amount that is remembered.

First, let us re-express the sum of sufficient statistics recursively:

$$S_0 = 0 \tag{14}$$

$$S_n = T(x_i, y_i, t_i) + S_{n-1} \tag{15}$$

In the volatile memory store we add a decay term L, which is a function of the number of distinct values in Z, as follows.

$$S_n = T(x_i, y_i, t_i) + S_{n-1} \cdot L(Z) \tag{16}$$

$L(Z)$ can be appplied to n in a similar way. The function $L(Z)$ determines the amount of decay and depends on the number of clusters present in the short term memory store.

$$L(Z) = [1 + \exp(-L_s \cdot (\text{count}(Z) - L_t))]^{-1} \tag{17}$$

In Eqution 17 count(Z) is the number of distinct values in Z. Z comes from the volatile memory store. The parameters L_s and L_t were chosen by hand to qualitatively match the human data. A sigmoid function was chosen because this constrained the decay to be within 0 to 1. A decay of 0 means the model learns optimally and a decay of 1 means the model learns nothing.

References

Anderson, J. (2002). ACT: A Simple Theory of Complex Cognition John R. Anderson. *Cognitive modeling*, 49. Available from `http://act-r.psy.cmu.edu/publications/pubinfo.php?id=97`

Banai, K., Ortiz, J. A., Oppenheimer, J. D., & Wright, B. A. (2009). Learning two things at once: differential constraints on the acquisition and consolidation of perceptual learning. *Neuroscience*.

Buonomano, D. V. (2000). Decoding temporal information: a model based on short-term synaptic plasticity. *Journal of Neuroscience*, *20*(3), 1129.

Buonomano, D. V., & Karmarkar, U. R. (2002). How do we tell time? *The Neuroscientist*, *8*(1), 42.

Cowan, N. (2008). What are the differences between long-term, short-term, and working memory? *Progress in brain research*, *169*, 323.

Fearnhead, P. (2004). Particle filters for mixture models with an unknown number of components. *Statistics and Computing*, *14*(1), 11–21.

Geisler, W. S. (2003). Ideal observer analysis. *The visual neurosciences*, 825–837.

Gelman, A. (2004). *Bayesian data analysis*. CRC press.

Izquierdo, I. (1999, August). Separate mechanisms for short- and long-term memory. *Behavioural Brain Research*, *103*(1), 1–11. Available from `http://dx.doi.org/10.1016/S0166-4328(99)00036-4`

Jacobs, R. A. (2009). Adaptive precision pooling of model neuron activities predicts the efficiency of human visual learning. *Journal of Vision*, *9*(4), 22.

Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *Journal of the Acoustical Society of America*, *49*(2), 467–477.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995, July). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, *102*(3), 419–57. Available from `http://www.ncbi.nlm.nih.gov/pubmed/7624455`

McGaugh, J. L. (2000). Memory–a Century of Consolidation. *Science*, *287*(5451), 248–251.

Neal, R. M. (2000, June). Markov Chain Sampling Methods for Dirichlet Process Mixture Models. *Journal of Computational and Graphical Statistics*, *9*(2), 249–265.

Petrov, A. A., Dosher, B. A., & Lu, Z. (2005). The Dynamics of Perceptual Learning: An Incremental Reweighting Model. *Psychological Review*, *112*(4), 715.

Roy, D. M., & Kaelbling, L. P. (2007). Efficient Bayesian task-level transfer learning. In *Proceedings of the twentieth international joint conference on artificial intelligence, hyderabad, india.*

Wang, K., & Shamma, S. (1994). Self-Normalization and Noise-Robustness in Early Auditory Reprersentations. *IEEE Trans. On Speech and Audio Processing*, *2*(3), 421–435.

Wright, B. A., Banai, K., Sabin, A. T., & Zhang, Y. (2009). *Distinct phases of auditory learning identified by differences in vulnerability to intervening events.* Poster Presentation at ARO 32nd MidWinter Meeting.

Wright, B. A., Buonomano, D. V., Mahncke, H. W., & Merzenich, M. M. (1997). Learning and generalization of auditory temporal-interval discrimination in humans. *Journal of Neuroscience*, *17*(10), 3956–3963.

Wright, B. A., & Sabin, A. (2007). Perceptual learning: how much daily training is enough? *Experimental Brain Research*, *180*(4), 727–736.

Xue, Y., Liao, X., Carin, L., & Krishnapuram, B. (2007). Multi-Task Learning for Classification with Dirichlet Process Priors. *The Journal of Machine Learning Research*, 35. Available from `http://portal.acm.org/citation.cfm?id=1248661`

Yotsumoto, Y., Watanabe, T., & Sasaki, Y. (2008, March). Different dynamics of performance and brain activation in the time course of perceptual learning. *Neuron*, *57*(6), 827–33. Available from `http://www.ncbi.nlm.nih.gov/pubmed/18367084`

Zhang, Y., & Wright, B. (2010). Disruption of frequency-discrimination learning by a 30 minute break. In *Association for research in otolaryngology abstracts* (pp. 340–341). Baltimore, Maryland. Available from `http://www.aro.org/archives/2010/2010\_992\_1287522000.html`