Contrast and stimulus information effects in rapid learning of a visual task

Craig K. Abbey

Department of Psychology, University of California,
Santa Barbara, CA, USA, &
Department of Biomedical Engineering,
University of California, Davis, CA, USA

Binh T. Pham

Department of Psychology, University of California,
Santa Barbara, CA, USA

Steven S. Shimozaki

School of Psychology, University of Leicester, UK

Miguel P. Eckstein

Department of Psychology, University of California,
Santa Barbara, CA, USA

We have previously described a psychophysical paradigm for investigating rapid learning of relevant visual information in detection tasks (M. P. Eckstein, C. K. Abbey, B. T. Pham, & S. S. Shimozaki, 2004). This paradigm uses blocked trials with a set of possible target profiles, and it has demonstrated learning effects after a single trial. When targets are masked by Gaussian luminance noise, there exists a Bayesian ideal observer that also exhibits learning effects over the trials within a block. In this work, we investigate the effect of target contrast and the effect of the information to be learned in the target profile set. Absolute efficiency tracks target contrast closely and ranges from approximately 10% to 25% in these experiments. To disambiguate learning from other effects contributing to absolute efficiency, we define a measure of learning efficiency that measures the observed improvement over a block of trials against the total improvement expected in the ideal observer. We find significant positive trends in learning efficiency both over contrast and the within-block trial number. We find that a two-feature profile set containing orientation and polarity differences leads to a greater within-block gain in performance than a one-feature profile set that contains only orientation differences. However, this apparent difference disappears when efficiency is compared. Lastly, we show that the disparity between task performance and accumulated knowledge of the target profile can be largely explained by a model that only allows learning to occur in trials the observer performs correctly.

Introduction

Many real-world visual tasks are facilitated by familiarity arising through prior experience, and this process is widely recognized as a form of learning (Fahle, 2005; Fine & Jacobs, 2002; Gibson, 1969, 2000; Goldstone, 1998; Seitz & Watanabe, 2005). For example, knowing that someone is wearing a bright blue jacket can make finding that person in a crowd much easier. It is not hard to imagine that knowledge of the jacket could have come from having seen that person at an earlier point in time. In this case, the previous experience allows one to selectively attend to a more relevant feature of the scene—people in bright blue jackets—during the search. Knowledge of the relevant feature may be acquired from very limited experience: A single previous sighting may be sufficient to realize that someone is wearing a blue jacket. Underlying this example is the idea that even relatively simple tasks can have a substantial degree of uncertainty associated with them, and considerable gains in task performance can be derived from learning relevant features from recent experience.

We have recently described an experimental paradigm for studying this sort of rapid learning process in visual tasks (Abbey, Eckstein, & Shimozaki, 2001; Abbey, Pham, Shimozaki, & Eckstein, 2005; Eckstein, Abbey, & Shimozaki, 2002; Eckstein, Abbey, Pham, & Shimozaki, 2004; Rasche, Pham, & Eckstein, 2003). A series of blocked psychophysical trials (explained in more detail below) are used to estimate ensemble performance within each block. Subjects can use their experience in earlier trials of each block to resolve some of the intrinsic uncertainty in the task and thereby improve performance in later trials of the same block. In this way, the learning observed within a block has similarities to fast perceptual

learning described by Fahle (2004) and Poggio, Fahle, and Edelman (1992). However, since this paradigm uses many repeated blocks, subjects will likely be less naive at the onset of a block than they would in one of these more traditional visual learning paradigms. Nonetheless, performance improvements can be assessed with as little as a single previous exposure to a relevant stimulus (Eckstein et al., 2004).

In this paradigm, the intrinsic uncertainty in the task is characterized by a set of possible target profiles, one of which is used throughout the trials constituting a block. When the task includes masking from luminance noise in the stimulus, we have derived the optimal Bayesian ideal observer (Green & Swets, 1966; Peterson, Birdsall, & Fox, 1954)—referred to simply as the ideal observer—for this task (Eckstein et al., 2004). Having an ideal observer allows us to measure the absolute statistical efficiency of human–observer performance (Tanner & Birdsall, 1958). A defining characteristic of the paradigm used in this paper is that the ideal observer exhibits learning within a block. A consequence of an ideal observer that exhibits learning is potential ambiguity between learning effects and absolute performance. As we shall see below, this dichotomy motivates us to formulate a distinct definition of learning efficiency in contrast to absolute efficiency in order to distinguish how effectively subjects learn from the amount of information available to be learned in the task. This definition also provides an attractive alternative to measures based on accuracy, which tend to achieve the largest effects near 50% accuracy because of confounds with floor or ceiling effects. The ideal observer specifies a model of learning through the information gained from prior experience. The learning efficiency measure we propose quantifies (with some caveats) the proportion of this information incorporated by the human observer.

The basic mechanism of learning in this paradigm—which is explicit in the ideal observer and presumed in humans—is adaptive weighting of visual information. In the context of learning, several investigators have used feature-integration models to describe mechanisms of learning (see for example, Beard & Ahumada, 1999; Dosher & Lu, 1998; Gold, 2003; Gold, Bennett, & Sekuler, 1999; Hurlbert, 2000). Li, Levi, and Klein (2004) show a striking example of this retuning in a Vernier acuity task where learning changes both weighting and sampling of the stimulus. In a recent review of perceptual learning experiments, Fine and Jacobs (2002) propose "…that learning might be a consequence of selective reweighting of the neurons that contribute to the psychophysical response…" They use this approach to learning to explain a broad and a diverse set of findings in the perceptual learning literature. We note that all of the perceptual learning experiments reviewed in this work measure perceptual learning in psychophysical procedures carried out over many sessions. The blocked paradigm used in this work functions somewhat differently, but is nonetheless consistent with the idea that learning modifies the visual information used to perform a task.

Previous results using the blocked-target paradigm show that a learning effect can be measured after a single trial within a block (Eckstein et al., 2004). This learning is subject to experimental control through manipulations such as changing the form of feedback (Abbey et al., 2001; Eckstein et al., 2002). In this work, we investigate the role of two important parameters related to learning: the information to be learned in a task and the level of task difficulty. In the context of our experiments, the information to be learned resolves which of a set of possible target profiles is present in a block of trials, and the amount of this information is related to the magnitude of differences between the various profiles. Information places the relevant features within a larger context. Going back to the blue-jacket example, knowledge about the jacket is not particularly useful if everyone in the crowd is wearing the same jacket. We compare learning effects observed on one profile set that contains changes in orientation to another that changes both in orientation and polarity. Stimulus information effects have been shown previously to influence visual processing (Liu, Kersten, & Knill, 1995; Shimozaki, Eckstein, & Abbey, 2002; Tjan & Legge, 1998). Learning in a visual task presumably involves similar processes (feature combination, etc.), motivating the investigation here.

Similarly, task difficulty can mediate our ability to monitor relevant features. In a darkened environment, we may be unable to ascertain the color of jackets being worn, and we may not have realized that the relevant color was blue in the first place. Task difficulty has been recognized as an important component of learning in visual tasks as shown by a number of previous studies (see for example Ahissar & Hochstein, 1997; Liu & Weinshall, 2000; Thompson & Liu, 2006), with larger effects being shown for easier tasks. Having an ideal observer for this experimental paradigm allows us to dissociate differences in performance associated with learning from properties inherent to the task and stimuli.

Theory

The experimental paradigm we use for studying rapid learning (Eckstein et al., 2004) is shown schematically in Figure 1. The paradigm requires specifying a set of possible target profiles, which we refer to as the target profile set. An experiment is broken into blocks, in which a subject gives a series of localization responses followed by an identification response. For each block, one profile is chosen from the set at random, and used as the target profile for all localization trials in the block. In the identification task, the subject indicates which of the
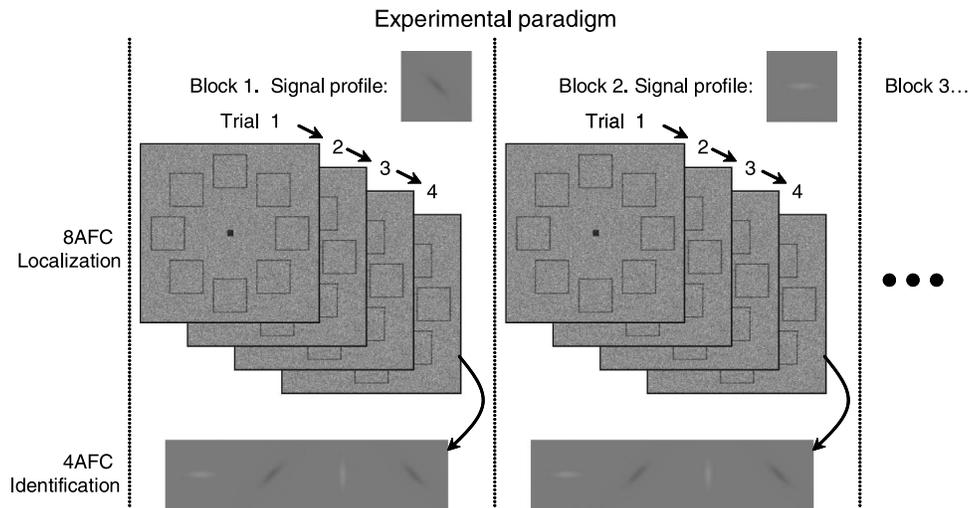
Figure 1. Learning experiments. This diagram shows the sequence of decisions in a learning experiment. The experimental data consists of blocks of four localization trials and a subsequent identification trial. All localization trials in a block have the same signal profile with position randomized between eight possible locations from trial to trial. The signal profile is then randomized between blocks. The identification trial asks the subject to identify which of the possible signal profiles is present in the four localization trials. Each experiment consists of 800 such blocks per condition per subject.

possible target profiles was present in the localization trials for that block.

Because the target profile does not change across the block of 4 localization trials, the ideal observer and potentially the human observer may use that information to improve performance in the later trials within the block. In the first trial of a block, the observer has total uncertainty about which of the target profiles is actually present. This target uncertainty is known to reduce localization performance (Pelli, 1985) since the observer presumably monitors sensors tuned to visual features of all possible targets. Because some of these sensors will be tuned to features of targets other than the one actually present, they will convey little or no relevant information. They may nonetheless respond to noise in the system, thereby reducing task performance. However, in the second localization trial, the subject can use his/her experience in the first trial to reduce uncertainty about which profile is the target. In essence, the subject can learn about the target profile from the localization tasks, and this learning can improve the probability of a correct localization response (PC_{Loc}) by reducing the influence of irrelevant sensory input. Therefore, improved performance in later trials is a sign of learning.

In addition to improved performance in later trials, evidence learning about the target profile is also found in the identification response at the end of the block. At the outset of a block, the subject has no knowledge of which target profile is present, and therefore the probability of a correct identification response (PC_{Id}) before any localization trials have occurred must be at chance levels. If PC_{Id} is significantly above chance performance at some later trial, then the subject must have acquired some knowledge of the target profile during the localization process.

It is worth mentioning an important caveat in associating learning effects to improvement in PC_{Loc}. While learning the identity of the target is sufficient to cause improved performance in later localization trials, it is not necessary. It is possible to imagine an observer that shows improved performance without actually learning about the target profile. For example, imagine the following extreme case in which an observer simply ignores the stimulus in the first trial of a block—responding by guessing—and then does use the stimuli in the second trial. Performance in the first trial will be at chance levels, and the observer will have no rational basis for updating the priors. As a result, the observer will be just as uncertain in the second trial as the first. However, in the second trial the observer uses the stimulus and hence will show improved performance. This performance improvement occurs without any updating of priors between the two trials.

By contrast, it is difficult to imagine a mechanism for performance above chance in the identification task that does not meet a reasonable definition of learning. We therefore contend that learning is necessary for greater-than-chance performance in the identification task, and observed values of PC_{Id} that are significantly greater than chance demonstrate learning that occurs over the localization trials of a block. We use identification performance after the 4th learning trial to establish the presence of a learning effect and then attribute the changes in localization performance over the 4 trials in a block to this effect.

Ideal observer responses

One of the strengths of the learning paradigm we use us that there is a well-defined ideal observer that performs

the task optimally provided that the image is masked by noise with a defined probability distribution. We use white Gaussian luminance (pixel) noise in this work. A notable feature of the ideal observer in this paradigm is that it incorporates experience with previous stimuli in the block to refine a prior distribution on target profiles. The localization and identification responses of the ideal observer have been developed for this paradigm previously (Eckstein et al., 2004). Here, we briefly review them for clarity in this work.

For a given trial t within a block (here $t = 1,...,4$), let the image vector $\mathbf{g}_t$ (256×256 pixels) contain the intensities of all pixels in a localization stimulus with the mean background intensity subtracted, and let $\mathbf{g}_{t,m}$ be independent sub-regions (50×50 pixels) centered on each possible location (here, $m = 1,..., 8$). The ideal observer selects the location that maximizes the posterior probability of being the target location. For forced-choice tasks with equi-probable randomized locations, an equivalent strategy is to choose the target location, $\hat{m}_t$, that maximizes the likelihood of the stimulus, $p(\mathbf{g}_t | m)$. This in turn is equivalent to maximizing the likelihood of the sub-region around the target location

$$\hat{m}_t = \arg \max_m \left(p\left(\mathbf{g}_{t,m} | m\right)\right). \tag{1}$$

Note that the likelihood term in Equation 1 simply specifies the probability of the stimulus vector given a target appears at a given location, it does not yet specify which target profile is present. This likelihood is therefore a sum of the likelihood of each possible target profile, weighted by the prior probability of that target profile. Let j index the possible target profiles ($j = 1,..., N$, with $N = 4$ in this work), then $\pi_{t,j}$ is the prior probability of target j in trial t. In the first trial, all target profiles are equally probable, and hence $\pi_{1,j} = 1/N$. Incorporating the sum over target profiles gives a likelihood defined by

$$p\left(\mathbf{g}_{t,m} | m\right) = \sum_{j=1}^{N} \pi_{t,j} p\left(\mathbf{g}_{t,m} | m, j\right), \tag{2}$$

where $p(\mathbf{g}_{t,m} | m, j)$ is the likelihood of $\mathbf{g}_{t,m}$ given target profile j at location m.

After a given trial, the prior is updated using the stimulus just presented. With location feedback, the ideal observer can update its prior probabilities based on target locations. Let $m_{t,\text{true}}$ be the index denoting the actual location of the target in the tth trial. For $t > 1$, the ideal observer weights the various target profiles according to

$$\pi_{t,j} = Q \prod_{t'=1}^{t-1} p\left(\mathbf{g}_{t,m_{t,\text{true}}} | m_{t,\text{true}}, j\right), \tag{3}$$

where Q is a normalization constant necessary for the priors to sum to unity across j. With this rule, the prior

probability of trial $t + 1$ is the posterior probability from trial t. The ideal observer performs the identification task after t trials by choosing the target index with the largest posterior probability,

$$\hat{j}_t = \arg \max_j \left(\pi_{t+1,j}\right). \tag{4}$$

We note that the development here has been in the context of a single block of trials, but the actual experiments involve many blocks. This can be incorporated into the development here by the addition of another subscript $b = 1,.., B$, where B is the number of blocks. We have omitted this subscript for clarity.

The likelihood term, $p(\mathbf{g}_{t,m} | m, j)$, in Equations 2 and 3 has so far been left general. In this work, we use Gaussian white noise with variance σ^2 to mask the appearance of the target in the stimulus. Let us denote by the vector $\mathbf{s}_j$ the profile of the jth target, which serves as the conditional mean of the observed sub-region. This gives a multivariate Gaussian function for the likelihood term,

$$p\left(\mathbf{g}_{t,m} | m, j\right) = \frac{1}{(2\pi\sigma^2)^{P/2}} e^{-\frac{1}{2\sigma^2} \|\mathbf{g}_{t,m} - \mathbf{s}_j\|^2}, \tag{5}$$

where P is the number of pixels in each sub-region and the $\| \; \|$ indicates the vector magnitude of its argument.

Equation 2 gives the localization response of the ideal observer for a given stimulus $\mathbf{g}_t$, and Equation 3 gives the identification response. These responses can be used in Monte Carlo simulations to evaluate ideal observer performance and determine target energy thresholds needed to equate the ideal observer to human–observer performance. This process is used frequently in this work to compute the efficiency of human observer performance (Barlow, 1980; Tanner & Birdsall, 1958). If we let E_{exp} be target energy in a blocked learning experiment and let $E_{\text{IO}}(\text{PC}, i)$ be the target energy required for the ideal observer to achieve a performance of PC in learning trial i, then the observer efficiency is computed as

$$\eta_i = 100\% \times \frac{E_{\text{IO}}(\text{PC}_i, i)}{E_{\text{exp}}}, \tag{6}$$

where PC_i is the observed PC in learning trial i. Implementation of the ideal observer through lookup tables derived from Monte Carlo simulations is described in the Methods section.

Learning efficiency

While the efficiency given in Equation 6 is useful for placing the performance of human observers in the context of optimal performance, it has limited use for isolating the effects of learning or comparisons across

different tasks. For example, if we know that an observer achieves 30% efficiency in learning Trial 3, it is still not immediately clear whether a learning effect has been observed or not. As we shall see below, human–observer efficiency generally goes down as the learning trials increase, even though accuracy improves. In order to quantify a learning effect, we need an appropriate measure of learning relative to the ideal observer. This motivates us to propose a definition of learning efficiency that is linked to the experimental observations we can make.

Let us begin by considering one possible way to define the magnitude of a learning effect, by computing the increase in average localization PC (i.e., accuracy) over learning trials (Eckstein et al., 2004). Since target contrast is constant over all trials, increasing PC is evidence of an improved visual strategy for performing the localization task and hence implies some learning on the part of the observer (with the caveat above).

However, the PC difference is problematic because of ceiling and floor effects. For example, an observer that starts with localization performance on the first trial (PC_1) of 0.98 can only achieve a PC difference of 0.02 or less regardless of how effective learning is for the task. Similarly, imagine target contrast is manipulated so that average localization PC is 0.145 when there is no uncertainty about signal type. In this case, again, the maximum PC difference achievable is 0.02 (recall that chance performance for 8AFC is PC = 0.125). Mid-range values of PC_1 are less constrained by these effects and can be considerably higher.

A reduced measure of learning due to ceiling effect seems particularly counterintuitive for easy tasks. For high values of PC_1, the observer will often go from a state of total uncertainty about the identity of the target before the first trial to almost total certainty after it. Thus, the observer has effectively resolved nearly all uncertainty about the target profile in one trial. The PC difference assigns a modest score to what is effectively near-total learning of target identity.

To generate a more meaningful measure, we propose a definition of "learning efficiency" that utilizes the ideal observer and is independent of the traditional "absolute" efficiency defined in Equation 6. The foundation of this definition is use of energy thresholds from the ideal observer to represent the relevant stimulus information—in units of signal energy—needed to perform the task at a given level of PC. In the first learning trial, when there are no previous images to use to update prior weights, the ideal-observer threshold represents the information accessed from the localization stimuli. For an observer that does not learn (i.e., maintains constant PC over learning trials), the threshold of the ideal observer drops over subsequent trials. This reflects the fact that the total information available to the ideal observer now includes both the current image and previous images in the block. Matching the non-learning observer's performance now requires less target energy because the ideal

observer can use the prior images to more effectively perform the task. The difference in threshold energy thus represents the contribution of prior images to the ideal observer. We define this quantity as

$$\Delta E_{IO}(i) = E_{IO}(PC_1, 1) - E_{IO}(PC_1, i). \tag{7}$$

For a human observer who may have improved performance after the first trial, the ideal observer energy threshold will be elevated somewhat from the non-learning observer. This elevation of the threshold represents the contribution of learning to that observer's performance. We define this quantity as

$$\Delta E_{Obs}(i) = E_{IO}(PC_i, i) - E_{IO}(PC_1, i). \tag{8}$$

The ratio of these two threshold elevations constitutes the definition of learning efficiency used here,

$$LE_i = 100\% \times \frac{\Delta E_{Obs}(i)}{\Delta E_{IO}(i)}. \tag{9}$$

This metric uses the threshold elevation of the ideal observer as a yardstick to measure the threshold elevation of the observer of interest. Note that this definition of learning efficiency is indeterminate in Trial 1, which reflects the fact that the learning process cannot be assessed by this approach from responses to a single trial.

A graphic interpretation of our definition is given in Figure 2. In Figure 2A, performance of a hypothetical human observer (with error bars left off for simplicity here) shows a gradual rise in PC over the four trials of an experiment. Ideal and non-learning observers for whom target contrast has been reduced to match human performance in the first trial are also shown. Ideal observer energy thresholds in Figure 2B match the corresponding performance values given in Figure 2A. The ideal observer, by definition, has flat threshold energy across learning trials. Figure 2C shows the learning efficiency computed according to Equation 9. As expected, the non-learning observer has a learning efficiency of 0%, and the ideal observer achieves 100% learning efficiency. The hypothetical human data achieves a mid range learning efficiency of 45–50%.

As further examples of how learning efficiency functions, Figures 3A and 3B show threshold energies in two hypothetical tasks. Thresholds for the human–observer data are identical. Assuming that target energy is held constant in the blocked trials, absolute efficiency for the two tasks would be identical as well. However, in the first task the energy threshold represents a greater increase from the non-learner, and hence it achieves considerably higher learning efficiency (Figure 3C).
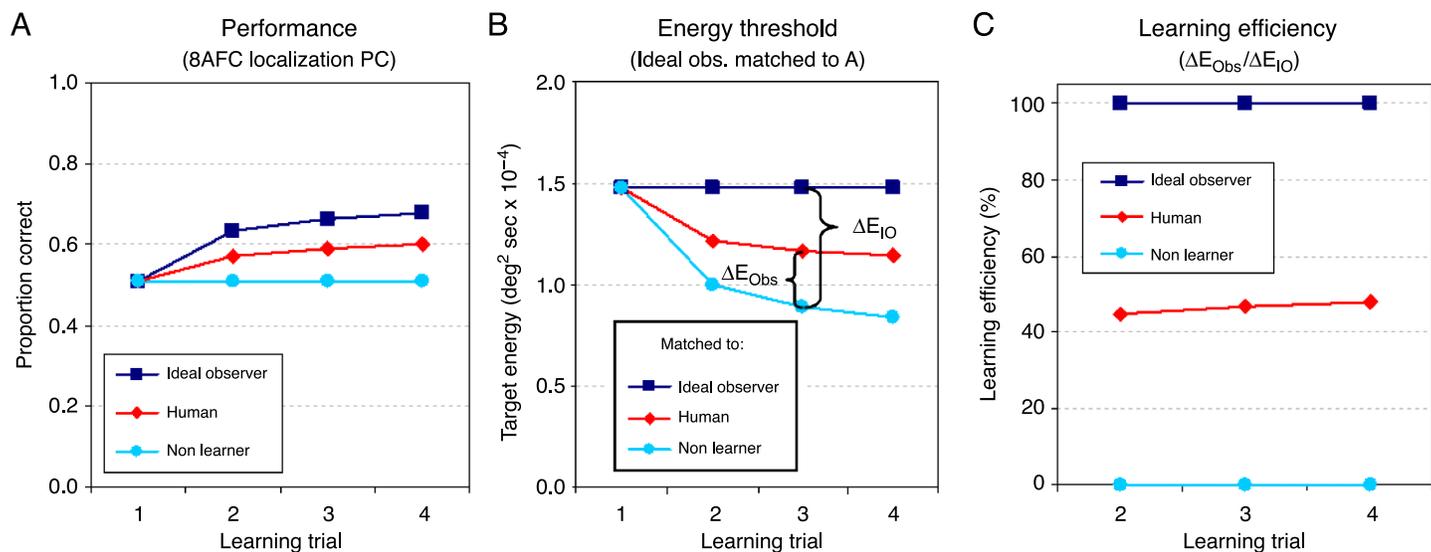
Figure 2. Definition of learning efficiency. Hypothetical observer performance is plotted in panel A with ideal and non-learning observers contrast-matched to human performance in the first trial. Threshold energy for the three observers is plotted in panel B. At learning Trial 3, components defining the learning efficiency are shown. Learning efficiency computed from Equation 9 is shown in panel C. Note that only learning Trials 2–4 are considered since the definition is indeterminate at Trial 1.

Because learning efficiency is defined by a ratio of ideal observer threshold energy increments and is made relative to the threshold decrease of the non-learning ideal observer, the measure is independent of stimulus parameters such as task difficulty and the effect of learning for a given profile set. A learning efficiency of 50%, for example, gives no indication of whether the task was easy or difficult, or whether the target profile set supported more or less learning. It also gives no indication of an observer's absolute efficiency for the task. However, all of these can influence the precision with which learning efficiency can be measured. This is particularly true of the profile set effects. For profile sets that support only a small amount of learning, such as the one-polarity set described below, learning efficiency is determined by smaller changes in PC. This makes the measurement of learning efficiency much more sensitive to the measurement errors in PC values.

A final note about learning efficiency as we have defined it is worth making at this point. The measure is not necessarily confined to the range 0–100%, as absolute efficiency is. If (for whatever reason) an observer performs worse in the second trial of an experiment than the first, learning efficiency for that trial will be negative.
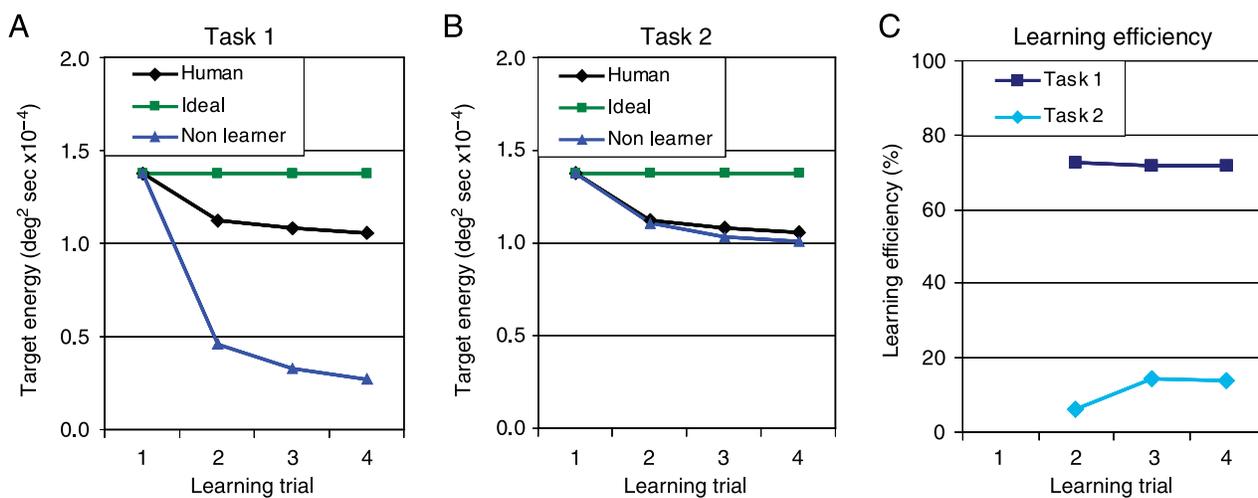

Figure 3. Signatures of efficient and inefficient learning. The plots show hypothetical ideal observer threshold energies matched to achieve equivalent performance in two tasks (A and B). The absolute efficiencies of both tasks are identical at all learning trials since the human–observer thresholds follow the same decreasing pattern in both. However, learning efficiencies (C) of the human observer in the two tasks show a marked difference in magnitude because of the differences with respect to the non-learning observer.

If an observer's PC in the second trial is greater than the ideal observer's (with contrast matched to the first trial), then learning efficiency will be greater than 100%, representing super-efficient learning by this definition. Both cases are likely to be diagnostic of external factors influencing learning effects. For example, a significantly negative learning efficiency may mean that attention or vigilance is somehow being lost in the later learning trials. Super-efficient learning may indicate that the first trial of a learning block is somehow disadvantaged as described for PC_{loc} above.

Methods

In this section, we discuss the design, experimental protocol, and ideal observer analysis of our experiments to investigate the role of the profile set and target contrast in influencing learning.

Design of experiments

The experiments described below investigate learning effects using sets of target profiles shown in Figure 4. All
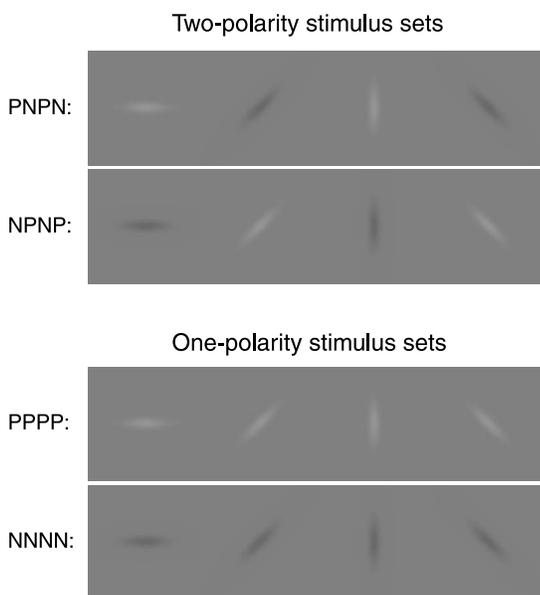

Figure 4. Stimulus sets. The images show the mean signal profiles for the four stimulus sets used in this work. The sets consist of two-polarity stimulus sets in which the profiles have alternating positive (P) and negative (N) contrast and one-polarity stimulus sets where all signal profiles have the same positive or negative contrast. Two contrast reversed stimulus sets are used for each polarity to control for detectability differences between positive and negative contrast targets.

profile sets were comprised of four elongated Gaussian profiles (aspect ratio 4:1) at different orientations (0°, 45°, 90°, 135°). The profile sets differed based on the polarity of the signal profiles. These could be positive (P), indicating a target with increased luminance over the background, or negative (N), indicating targets with decreased luminance. There were two sets of one-polarity target profiles, which consisted of four positive polarity profiles (PPPP) and four negative polarity profiles (NNNN). There were also two sets of two-polarity target profiles, consisting of alternating positive and negative polarity profiles (PNPN) or alternating negative and positive polarity profiles (NPNP). In all experiments, targets were embedded in Gaussian white noise with an RMS contrast of 20.3%.

Contrast experiment

The PNPN profile set was chosen to investigate the effect of contrast on the learning process. Pilot studies were used to find a group of 5 target contrasts that gave a reasonable span of performance in the localization task (7.0%, 9.4%, 11.7%, 14.8%, and 18.0%). Note that for the negative polarity targets, these contrasts should be considered to be negative.

Polarity set experiment

The polarity experiment was designed to test the effect of different signal profiles in the profile set while at the same time controlling for the different stimuli used. This experiment used all four profile sets pictured in Figure 4 at a single target contrast of 14.8% (peak profile luminance/ mean background luminance). Results of the one-polarity profile sets were averaged, as were results of the two-polarity profile sets. Having the average of opposite polarity profile sets meant that the combined results involved exactly the same target profiles, thereby counter-balancing any subject detection asymmetries in positive and negative polarity targets.

Experimental protocol

Four naive observers participated in all experiments reported here. All had normal or corrected-to-normal vision. Experiments were performed in a darkened room on a monochrome monitor (Monitor M17LMAX, Image Systems, Minnetonka MN) with a video board (Planar systems, INC. Beaverton OR) that allows for control of the monitor lookup table to match predetermined specifications. Pixel size was 0.3 mm. The luminance of the monitor was calibrated through photometer measurements to a linear scale with a mean monitor luminance of 20 cd/ m². Viewing distance was not constrained, but observers

maintained a comfortable viewing distance of approximately 40 cm.

Human observer data were collected using software developed in-house with the IDL programming environment (Research Inc. Boulder, CO). All stimuli were generated and stored before the experiments began. The software would read the appropriate data, display the stimuli, and record observer responses. The experiments followed closely the scheme laid out in Figure 1. An observer would proceed through four localization trials followed by an identification trial in a single block of the experiment. In a localization trial, the observer would initiate the trial with a mouse click. After a delay of 300 ms, an 8AFC stimulus would appear for a duration of 200 ms, followed by a blanking screen which would query the observer for the location containing the target. The eight possible target locations were located on a circle of radius 2.25° from fixation at the center of the screen. These locations were cued with a dark box (2 deg width) appearing at each possible location (see localization stimuli in Figure 1). After presentation of the stimulus and blanking screen, the subject responded with a mouse click in one of 8 boxes appearing at each of the eight locations. Feedback was given by changing the polarity of the box cuing the correct location. At the end of four learning trials, a screen containing the four possible target signals appeared and queried the subject for a mouse click on the profile representing the signal in the four previous localization trials. Feedback was given for the identification trial as well, in the form of a box appearing around the correct signal profile. This had no direct influence on learning since the block finished with the identification trial. It was used because we have found that in some cases, feedback can help with subject vigilance.

Before beginning any of the profile sets, each subject performed 50 trials of task-specific training in each experimental condition. All subjects had prior experience with the experimental protocol, and they completed experiments in sessions of 50 blocks. Each condition of an experiment consisted of 800 blocks broken into 16 sessions, and subjects cycled through one session in each condition (order of the sessions was randomized) before progressing.

Ideal observer analysis

We evaluate the ideal observer using pre-computed lookup tables derived from Monte Carlo simulations to capture the performance of the ideal observer in the one- and two-polarity profile sets. Note that since the ideal observer is equally sensitive to positive or negative contrasts, the same lookup table can be used for both the PNPN and NPNP profile sets. This is also true for the PPPP and NNNN sets. Ideal observer performance is evaluated for each localization and identification trial at 26 signal energies corresponding to contrasts of 0.0–11.7% in steps of 0.78% contrast. This gives a reasonable sampling of PC from chance performance

(0.125 for localization and 0.25 for identification) to over 0.999. Each point in the lookup table is determined from 100,000 Monte Carlo samples, which allows for effectively negligible estimation error in the table ($SE < 0.0016$).

The contrast needed to achieve a targeted proportion correct for a given trial is computed from the lookup table using linear interpolation. Let PC_{targ} be the targeted value of proportion correct, and let PC_{lo} and PC_{hi} be values from the lookup table that bracket it, with corresponding contrasts C_{lo} and C_{hi}. Linear interpolation yields a threshold contrast of

$$C_{targ} = \frac{\left(PC_{hi} - PC_{targ}\right)}{\left(PC_{hi} - PC_{lo}\right)} C_{lo} + \frac{\left(PC_{targ} - PC_{lo}\right)}{\left(PC_{hi} - PC_{lo}\right)} C_{hi}. \quad (10)$$

Target contrast can be converted to contrast energy by knowing the area of a displayed pixel, A_{pix}, in units of degrees^2; the stimulus duration, T, in seconds; and the sum of squared pixel intensities averaged across all target profiles at a nominal contrast of 1, SS_{ave}. Conversion to signal energy is given by

$$E_{targ} = SS_{ave} A_{pix} T C_{targ}^2, \quad (11)$$

which specifies target contrast energy in units of deg.^{2}sec. For the experiments reported here, the average target sum of squares is $SS_{ave} = 82.84$, the area of a pixel in degrees is $A_{pix} = 0.001847$ deg^2 (for pixel area of 0.09 mm^2 and 40 cm viewing distance), and the stimulus duration is $T = 0.2$ s. Thus, the conversion to threshold energy is $E = 0.0306 C^2$. Efficiency with respect to the ideal observer is then computed according to Equation 6, and learning efficiency is computed using this approach according to Equation 9.

Results and discussion

Contrast experiment

Figure 5 plots average human observer performance as proportion correct (Figure 5A) and efficiency with respect to the ideal observer (Figure 5B). The five plots in each graph show performance for the five contrast values used in the experiment.

As expected, increasing target contrast increases observer PC in Figure 5A. Of greater interest here is the increase in performance for a given contrast as the number of learning trials increases. For the lower contrast values, there appears to be substantially less improvement in PC going from Trials 1 to 4 than in the higher contrast values.
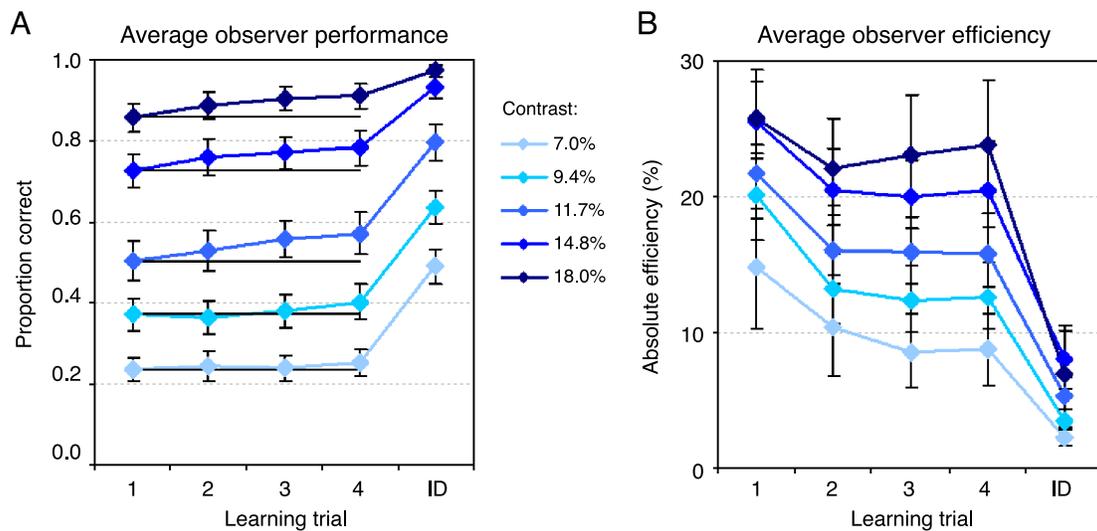
Figure 5. Performance in stimulus contrast experiments. The plots show the average of observer proportion correct (A) and absolute efficiency (B) across the five contrast settings from 7.0% to 18.0%. One observer had perfect identification performance in the highest contrast leading to an undefined efficiency. That datum was excluded from the efficiency average. Error bars represent a 95% confidence interval averaged across observers, and horizontal lines at the level of performance in the first trial are plotted for reference in panel A.

The average efficiency plotted in Figure 5B shows that efficiency is generally increasing as target contrast increases, but decreasing over the trials within a block. At the lowest level of target contrast, efficiency decreases from 14.8% efficiency in Trial 1 to 8.7% in Trial 4. At the highest contrast, efficiency is a factor of 2 higher and ranges from 25.9% to 23.8%. Efficiencies of the identification task are considerably lower and range from a minimum efficiency of 2.2% at low contrast to a maximum efficiency of 8.4% at high contrast. The highest experimental contrast used (18.0%) results in slightly lower measured identification efficiency (7.0%). However, this may be due to excluding the pest performing subject at this data point because of perfect identification performance ($PC_{ID} = 1$) and is also consistent with the larger effects of measurement error at very high levels of PC (0.969 ± 0.011).

Learning efficiency and task difficulty

Average learning efficiency for the contrast study is shown in Figure 6A. Each contrast groups the learning efficiency results for Trials 2 (LT2), 3 (LT3), and 4 (LT4) within the learning block. Learning efficiency values that are significantly different from zero are indicated with an asterisk. These are found starting in the middle contrast (11.7%), corresponding to an 8AFC PC of about 0.5, and going to the highest contrast measured (18.0%). For reference and comparison, absolute efficiency is re-plotted from Figure 5 at the scale of contrast for the learning trial.

Learning efficiency increases with contrast and with learning trial. Linear trends across contrasts are significant for all three learning trials (one-tailed t-test, $df = 3$, $p < 0.022$ (LT2), 0.006 (LT3), and 0.013 (LT4)), and a

linear trend in learning trials averaged across contrasts was also significant (one-tailed t-test, $df = 3$, $p < 0.014$). Significance of the trend in learning trials was greater when averaging was restricted to the three highest contrasts ($p < 0.001$). In contrast, absolute efficiency is significant at all contrasts with no significant learning trial effect.

These results show two important facets of human–observer learning in this paradigm. Significance of the linear trend in learning efficiency across contrast shows that observers become more effective at learning as the
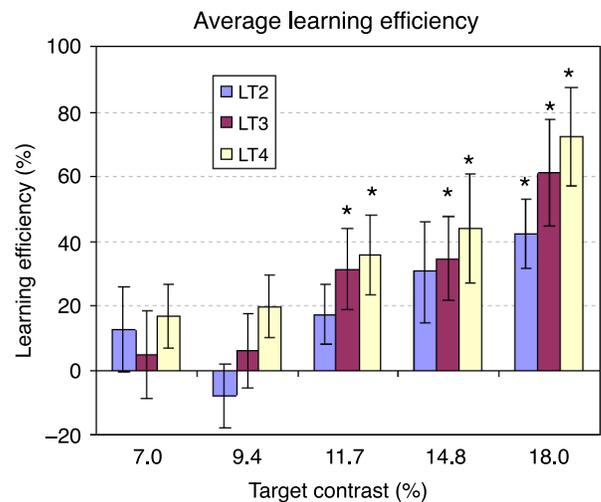

Figure 6. Learning efficiency. Average learning efficiencies in learning Trials 2–4 (LT2–LT4) are plotted at each of the five experimental contrasts. Error bars represent ±1 *SE*. Asterisks indicate measured learning efficiencies significantly different from zero (t-test; $p < 0.05$, $df = 3$). Significant effects for linear trends in contrast and learning trial are also found (see text).

task becomes easier. This finding is somewhat different than would be determined from the increase in accuracy (PC_{Loc}) over learning trials seen in Figure 5A, where the largest effects are found at mid-range target contrast. Observers are most effectively incorporating prior experience at high contrast when the information gleaned from that experience is the highest quality.

The highest average learning efficiency observed (at 18% contrast in Trial 4) is 72%, which indicates that humans are generally quite effective learners, even though the absolute efficiency in this case was less than 25%. The increase in learning efficiency with learning trials implies that human observers are somewhat slower than the ideal observer at incorporating previous experience in a block to improve performance.

Localization and identification efficiency

A finding seen in Figure 5B that was consistent across all measurements of absolute efficiency was the substantial drop in efficiency going from localization to identification. While PC was typically higher in the identification task, we would generally expect a higher value since identification is 4AFC and localization is 8AFC. Following previous work (Eckstein et al., 2004), we will examine a suboptimal model for updating priors to explain this difference in efficiency. The central idea behind this model is that observers only update the prior on trials in which they localize the target correctly. We refer to this as the update on correct (UPCOR) model.

Recall that stimulus duration in the localization experiments was approximately 200 ms, and feedback came after a response was obtained in the form of a highlighted box around the correct location. Thus, incorporating feedback required memory of the stimulus present at a given location. This could be difficult to do for locations that were not deemed to contain the target. As a result, when an incorrect localization was made, we hypothesize that the feedback was effectively unusable and the observer would not be able to update any sort of prior on the target profile. Evidence for this conjecture can be found in Eckstein et al. (2004), who considered performance in the second trial of a learning block after correct or incorrect first trials. Unlike the ideal observer, human observer performance in the second trial was very sensitive to performance in the first trial. Improved localization performance in the second trial was only found when the observer made a correct localization in the first trial. Localization performance in the second trial after an incorrect localization in the first was no different than localization performance in the first trial. These effects are reproduced in the current data as seen in Figure 7. After an incorrect localization, PC in the second trial is significantly less than PC with a correct localization in the first trial (one-tailed *t*-test, $df = 3$, $p < 0.001$). Significant effects from an incorrect localization persist

throughout the learning trials and the identification task ($p < 0.001$ (LT3), 0.002 (LT4), and 0.003 (ID)).

Could an inability to incorporate information from incorrect trials explain the relatively poor identification efficiency we observe in our human–observer data? To investigate this question, we have evaluated the efficiency of the UPCOR model. The prior update term for the UPCOR model is similar to Equation 3, except that the product only includes only the previous trials that were correctly localized instead of all previous trials. Other aspects of the UPCOR model are identical to the ideal observer. UPCOR performance was encapsulated into a lookup table using the same procedure as was used for the ideal observer described above.

Figure 8 plots identification efficiency against localization efficiency in the first learning trial. Under the hypothesis of equal efficiency, we would expect points to lie on the diagonal. Average subject efficiency rests well below the diagonal line, demonstrating reduced efficiency in the identification task with the largest deviations coming at low contrasts. The ideal observer—disadvantaged by reduced contrast to match human performance in the first localization trial—by definition maintains constant efficiency. Therefore, the ideal observer is constrained to lie on the diagonal and cannot exhibit reduced
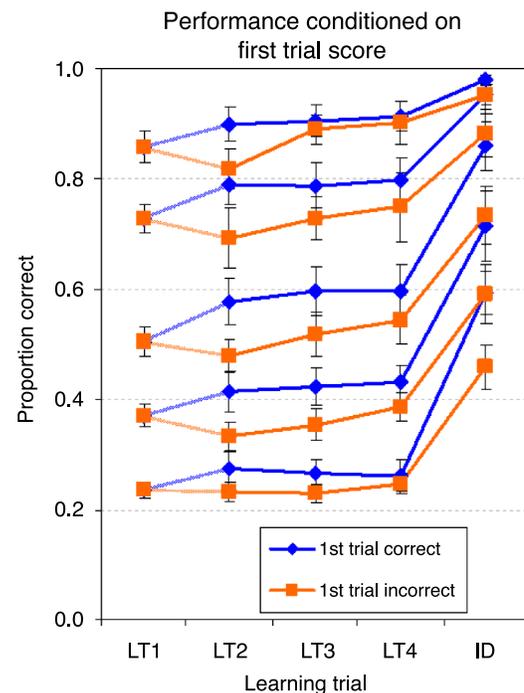
Figure 7. Effect of first trial outcome. Proportion correct is plotted as a function of learning trial for the five contrast experiments. In each experiment, performance is decomposed into blocks for which the first trial is correct and incorrect to show persistence effects from the first trial. For reference, both plots use the average PC in Trial 1, and a horizontal line at this value is drawn across learning trials.
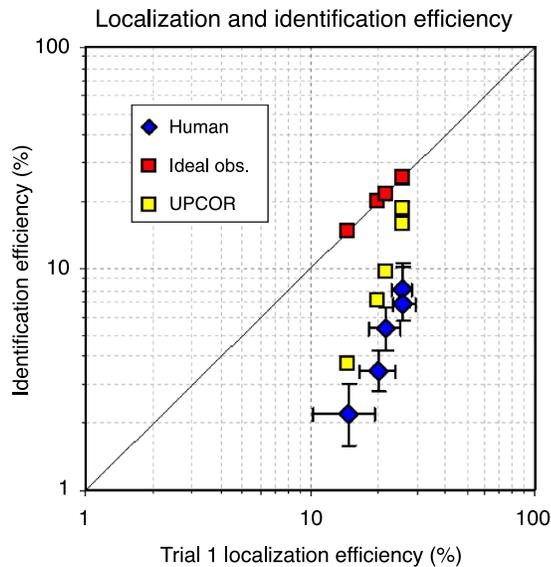
Localization and identification efficiency

Figure 8. Efficiency scatter plots. Average subject efficiency for learning Trial 1 localization is plotted against identification efficiency after learning Trial 4. Deviation from the diagonal line indicates the magnitude of disparity between localization and identification efficiency. The ideal observer and UPCOR model efficiency values (contrast matched to human performance in learning Trial 1) are also plotted.

efficiency in the identification task. The UPCOR model covers a substantial portion of the difference in localization and identification efficiency, particularly at low contrast where the effect is strongest.

Together, Figures 7 and 8 show that correct and incorrect localization judgments have important effects in later trials, and a substantial portion—though not all—of the disparity in observer efficiency going from localization to identification can be explained by an inability to update prior information (i.e., learn) in trials which the observer was unable to correctly localize the target. Despite being given feedback after every localization trial, subjects only appear able to reinforce the learning process when they correctly localized the target.

Learning across sessions

We note that traditional learning effects across multiple sessions were present in this data (not shown). Recall that each observer's performance was determined from 16 sessions of 50 learning trial blocks. In most cases, positive trends in performance were observed over the 16 sessions. Since sessions of different target contrasts were intermixed, it is not clear to what extent this improvement comes from prior experience with the same contrast as opposed to generalization from different contrasts. Restricting analysis to the last 8 sessions improved the

PC and efficiency values in Figures 5 and 6 (largest improvement was 0.067 in PC and 6.3% in efficiency) but did not change their qualitative character or invalidate any significant differences.

Polarity experiment

The previous experiment shows how target contrast influences learning. However, other properties, such as polarity of the profiles, could also play an important role. Conventional interpretations of learning suggest that learning is improved by having more stimulus dimensions to learn (for a review, see Fine & Jacobs, 2002). For example, the two-polarity profile sets can be thought of as varying across orientation and polarity dimensions while the one-polarity sets only vary across orientation. Learning could be assumed to occur in representations of each dimension resulting in a facilitated performance with summation of better representations in later trials.

Average human performance for the polarity experiment is shown in Figure 9A. There is a significant difference in proportion correct going from Trials 1 to 4 for both one- and two-polarity target profiles ($p < 0.042$ for the one-polarity stimuli, $p < 0.001$ for the two-polarity stimuli: paired t-test two-tailed, $df = 3$). However, the two-polarity stimuli show a substantially larger effect. The average difference in proportion correct between Trial 4 and Trial 1 is 0.091 (±0.004) for the two-polarity stimuli as opposed to 0.033 (±0.010) for the one-polarity stimuli. This difference is primarily due to the first trial, where PC for the two-polarity stimuli is significantly lower ($p < 0.001$) than the one-polarity stimuli. There is also a significant difference ($p < 0.027$) in identification performance after Trial 4. In this case, two-polarity stimuli result in higher average performance by 0.079 (±0.019) units of PC.

Some aspects of these data would appear to agree with the conventional interpretation that learning is improved by having two stimulus dimensions to learn (orientation and polarity) as opposed to just one (orientation). The change in PC from learning Trial 1 to learning Trial 4 is larger for the two-polarity stimulus sets. However, this interpretation is problematic here. Summation of independent representations of orientation and polarity predict that localization performance should be better for the two-polarity stimulus sets since they have more stimulus dimensions upon which to represent the target. However, the PC data in Figure 9A show that localization performance is generally higher for the one-polarity stimulus sets. And yet when identification performance is considered, a significant improvement is observed for the two-polarity stimulus sets indicating that the subjects did learn more effectively.

We contend that the greater learning effects can be attributed to a more general property of the stimuli captured by efficiency with respect to the ideal
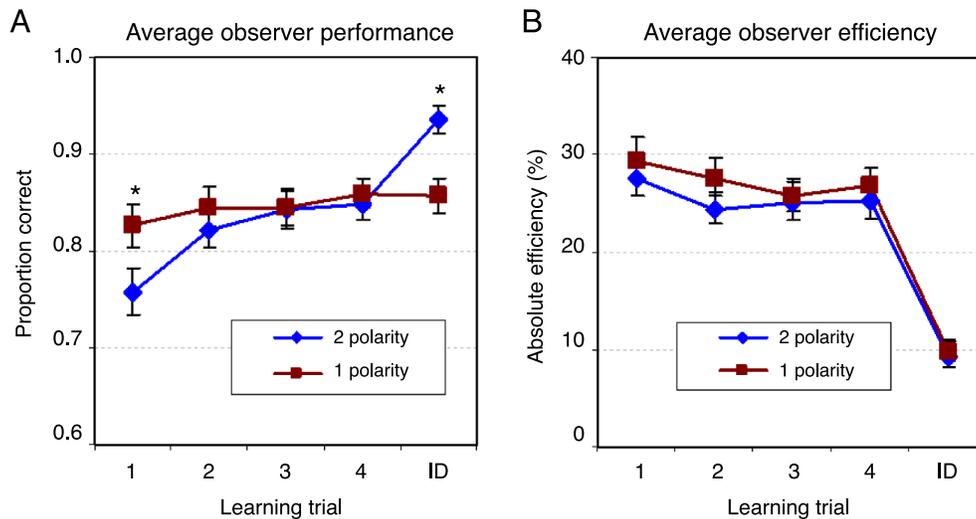
Figure 9. Performance in stimulus polarity experiments. The plots show the average of observer proportion correct (A) and absolute efficiency (B). Error bars represent the 95% confidence interval averaged across observers. Performance for learning trials marked with an asterisk (*) shows significant differences between the stimulus sets.

observer. Previous investigators such as Shimozaki et al. (2002) and Tjan and Legge (1998) have shown how the ideal observer can be used as a more general alternative to representational summation models. Intuitively, the ideal observer approach capitalizes on the notion that the reason the two-polarity targets are more difficult in the first trial is there are more ways for noise to look like a target for these stimuli. If all patterns look approximately the same, then there is only limited opportunity for noise mimic the target. Conversely, when the target profiles are very different, noise in the stimulus has a greater probability of mimicking one of the profiles to mislead the observer thereby reducing performance. As the actual profile is learned through the trials in a block, noise masks that would have suggested a different target are no longer confused with the target and performance improves.

Figure 9B plots the average efficiency of the human observers for the one- and two-polarity target profiles. In contrast to the proportion correct results, average efficiency values are very similar for both one- and two-polarity profile sets, although there appears to be some reduced localization efficiency for the two-polarity profiles (1.8 percentage points on average). The efficiency also very nearly equates performance in the identification task. Learning efficiency as defined in Equation 9 is not shown for this experiment since the measure could not be computed with any meaningful precision for the one-polarity profile sets. However, the absolute efficiency data show that for both profile sets, the observers are equally efficient at extracting task relevant information from the stimuli. In this context, learning appears to be stimulus driven.

Summary and conclusions

We have examined aspects of rapid learning that take place over a few psychophysical trials using experiments with blocked localization and identification responses from stimuli masked by noise. Previous efforts have derived the Bayesian ideal observer for this experimental paradigm. An important component of the ideal observer is that previous trials in a block are used to update prior probabilities to improve performance. Thus, the ideal observer itself is subject to a form of learning. Our purpose in this work has been to investigate the influence of two key controllable experimental parameters, target contrast and information, in the profile set that forms the basis for learning.

For target contrast, efficiency relative to the ideal observer generally increases with target contrast across learning trials and in the identification task. Similar increases in efficiency have been found for simple detection tasks in the past (for example, Burgess, Wagner, Jennings, & Barlow, 1981). However, for a fixed contrast, absolute efficiency relative to the ideal observer decreases with learning trials. In the context studied here, it is of interest to know whether these efficiency effects simply reflect general changes in ability to perform the task or if increased contrast in some way facilitates learning specifically. To disambiguate a learning effect from more general changes in performance, we have defined a quantity we term the learning efficiency. This ratio of threshold energy differences measures changes in threshold relative to the ideal observer and a non-learning observer that is otherwise ideal. Learning efficiency is independent of absolute efficiency, although absolute

efficiency and the profile set may strongly influence how accurately learning efficiency can be measured. We do not find evidence for any significant learning efficiency until signal contrast reaches the range where performance in the 8 alternative localization task is approximately 50% correct, with significant learning efficiency effects for higher contrasts. Learning efficiency also improves significantly with more learning trials. Our data shows learning efficiencies as high as 70% in a task where absolute (performance) efficiency was only 25%. This analysis shows that humans are not only more efficient at performing visual detection tasks at high contrast, but that we are more efficient at learning them as well.

We also show that a suboptimal model of learning can explain much of the difference in efficiency going from the localization task to the identification task. Typically efficiency relative to the ideal observer drops by a factor of more than three between these two tasks. Modifying the ideal observer so that prior probabilities are updated only on correct trials explains over much of that difference.

The information to be learned in a given stimulus profile set strongly influences the performance of subjects in both the localization task and in the subsequent identification task. Single polarity profile sets show significantly less of a learning effect than the two-polarity profile sets in the localization tasks, but this is because performance in early learning trials is significantly higher for the one-polarity profile sets. In the identification task after the fourth learning trial, performance is significantly higher in the two-polarity profile set. However, these differences all disappear when human observer performance is made relative to an ideal observer by computing efficiency. In this case, the ideal observer shows that humans are extracting roughly the same amount of task relevant information for both profile sets. The differences in performance are a consequence of how much information there is to learn in each profile set and not related the summation of putative features such as polarity and orientation.

Acknowledgments

Commercial relationships: none.
Corresponding author: Craig K. Abbey.
Email: abbey@psych.ucsb.edu.
Address: Department of Psychology, University of California, Santa Barbara, CA 93106.

References

Abbey, C. K., Eckstein, M. P., & Shimozaki, S. S. (2001). The efficiency of perceptual learning in a visual detection task [Abstract]. *Journal of Vision, 1*(3):28, 28a, http://journalofvision.org/1/3/28/, doi:10.1167/1.3.28.

Abbey, C. K., Pham, B. T., Shimozaki, S. S., & Eckstein, M. P. (2005). Contrast effects in rapid learning of a visual detection task [Abstract]. *Journal of Vision, 5*(8):1051, 1051a, http://journalofvision.org/5/8/1051/, doi:10.1167/5.8.1051.

Ahissar, M., & Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature, 387,* 401–406. [PubMed]

Barlow, H. B. (1980). The absolute efficiency of perceptual decisions. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 290,* 71–82. [PubMed]

Beard, B. L., & Ahumada, A. J., Jr. (1999). Detection in fixed and random noise in foveal and parafoveal vision explained by template learning. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 16,* 755–763. [PubMed]

Burgess, A. E., Wagner, R. F., Jennings, R. J., & Barlow, H. B. (1981). Efficiency of human visual signal discrimination. *Science, 2,* 93–94. [PubMed]

Dosher, B. A., & Lu, Z. L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proceedings of the National Academy of Sciences of the United States of America, 95,* 13988–13993. [PubMed] [Article]

Eckstein, M. P., Abbey, C. K., Pham, B. T., & Shimozaki, S. S. (2004). Perceptual learning through optimization of attentional weighting: Human versus optimal Bayesian learner. *Journal of Vision, 4*(12):3, 1006–1019, http://journalofvision.org/4/12/3/, doi:10.1167/4.12.3. [PubMed] [Article]

Eckstein, M. P., Abbey, C. K., & Shimozaki, S. S. (2002). Short term negative learning produced by monitoring erroneous templates [Abstract]. *Journal of Vision, 2*(7):560, 560a, http://journalofvision.org/2/7/560/, doi:10.1167/2.7.560.

Fahle, M. (2004). Perceptual learning: A case for early selection. *Journal of Vision, 4*(10):4, 879–890, http://journalofvision.org/4/10/4/, doi:10.1167/4.10.4. [PubMed] [Article]

Fahle, M. (2005). Perceptual learning: Specificity versus generalization. *Current Opinion in Neurobiology, 15,* 154–160. [PubMed]

Fine, I., & Jacobs, R. A. (2002). Comparing perceptual learning tasks: A review. *Journal of Vision, 2*(2):5, 190–203, http://journalofvision.org/2/2/5/, doi:10.1167/2.2.5. [PubMed] [Article]

Gibson, E. J. (1969). *Principles of perceptual learning and development*. Englewood Cliffs, NJ: Prentice-Hall.

Gibson, E. J. (2000). *An ecological approach to perceptual learning and development*. Oxford, NY: Oxford University Press.

Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Signal but not noise changes with perceptual learning. *Nature, 402,* 176–178. [PubMed]

Gold, J. M. (2003). Dynamic classification images reveal the effects of perceptual learning in a hyperacuity task [Abstract]. *Journal of Vision, 3*(9):162, 162a, http://journalofvision.org/3/9/162/, doi:10.1167/3.9.162.

Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology, 49,* 585–612. [PubMed]

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and Psychophysics*. New York: Wiley.

Hurlbert, A. (2000). Visual perception: Learning to see through noise. *Current Biology, 10,* R231–R233. [PubMed] [Article]

Li, R. W., Levi, D. M., & Klein, S. A. (2004). Perceptual learning improves efficiency by retuning the decision 'template' for position discrimination. *Nature Neuroscience, 7,* 178–183. [PubMed]

Liu, Z., Kersten, D., & Knill, D. C. (1995). Dissociating stimulus information from internal representation-a case study in object recognition. *Vision Research, 39,* 603–612. [PubMed]

Liu, Z., & Weinshall, D. (2000). Mechanisms of generalization in perceptual learning. *Vision Research, 40,* 97–109. [PubMed]

Pelli, D. G. (1985). Uncertainty explains many aspects of visual contrast detection and discrimination. *Journal of the Optical Society of America A, Optics and Image Science, 2,* 1508–1532. [PubMed]

Peterson, W. W., Birdsall, T. G., & Fox, W. C. (1954). The theory of signal detectability. *Transactions of the IRE Professional Group on Industrial Electronics, 4,* 171–212.

Poggio, T., Fahle, M., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science, 256,* 1018–1021. [PubMed]

Rasche, C., Pham, B. T., & Eckstein, M. P. (2003). The influence of stimulus information on human perceptual learning: An ideal observer analysis [Abstract]. *Journal of Vision, 3*(9):668, 668a, http://journalofvision.org/3/9/668/, doi:10.1167/3.9.668.

Seitz, A., & Watanabe, T. (2005). A unified model for perceptual learning. *Trends in Cognitive Sciences, 9,* 329–334. [PubMed]

Shimozaki, S. S., Eckstein, M. P., & Abbey, C. K. (2002). Stimulus information contaminates summation tests of independent neural representations of features. *Journal of Vision, 2*(5):1, 354–370, http://journalofvision.org/2/5/1/, doi:10.1167/2.5.1. [Pubmed] [Article]

Tanner, W. P., & Birdsall, T. G. (1958). Definitions of d' and η as psychophysical measures. *Journal of the Acoustical Society of America, 30,* 922–928.

Thompson, B., & Liu, Z. (2006). Learning motion discrimination with suppressed and un-suppressed MT. *Vision Research, 46,* 2110–2121. [PubMed]

Tjan, B. S., & Legge, G. E. (1998). The viewpoint complexity of an object-recognition task. *Vision Research, 38,* 2335–2350. [PubMed]