

Perceptual learning without signal

Nicolas Dupuis-Roy *, Frédéric Gosselin

Département de psychologie, Université de Montréal, C.P. 6128, Succ. Centre-Ville, Montréal, Qué., Canada H3C 3J7

Received 8 May 2006; received in revised form 12 October 2006

Abstract

Perceptual learning is characterized by an improvement in a perceptual task following practice. Several studies have demonstrated that top-down processes, such as attention and task-related expectations, can be necessary components of perceptual learning [Ahissar & Hochstein, 1993, 2000, 2002; Fahle & Morgan, 1996; Seitz, Lefebvre, Watanabe, & Jolicoeur, 2005; Seitz, Nanez, Holloway, Koyama, & Watanabe, 2005; Seitz & Watanabe, 2003; Shiu & Pashler, 1992]. Here, we report an experiment that isolated top-down processes in perceptual learning, using a variant of the Gosselin and Schyns (1992) no-signal procedure. Results indicate that top-down processes can be sufficient to produce substantial, possibly long-lasting and rotation-invariant perceptual learning.

© 2006 Elsevier Ltd. All rights reserved.

Keywords: Perceptual learning; Expectations; Memory representation; No-signal procedure; Top-down processing; Classification image analysis

1. Introduction

The ability of the nervous system to improve in a perceptual task after training is known as perceptual learning (PL).¹ The prevailing view among researchers in the field is that top-down processes are necessary to produce PL (Ahissar & Hochstein, 1993, 2000, 2002; Fahle & Morgan, 1996; Jiang & Chun, 2001; Seitz & Watanabe, 2005; but see Fahle, Edelman, & Poggio, 1995; Godde, Stauffenberg, Spengler, & Dinse, 2000; Hodzic, Veit, Karim, Erb, & Godde, 2004; Vaina, Sundaeswaran, & Harris, 1995). Shiu and Pashler (1992) and Ahissar and Hochstein (1993) discovered that practicing one task did not necessarily improve performance in an alternative task, even though both tasks used exactly the same visual stimuli but depended on different stimulus attributes (e.g., either orientation of local elements or global shape). Seitz, Watanabe, and their colleagues (Seitz, Lefebvre, Watanabe, & Jolicoeur, 2005; Seitz, Nanez, Holloway, Koyama, & Watanabe, 2005; Seitz & Watanabe, 2003), who once thought that

bottom-up processes were sufficient to produce PL (Watanabe et al., 2002; Watanabe, Nanez, & Sasaki, 2001), have demonstrated that attention sometimes needs to be temporally paired with – but not focused on – a subliminal and task-irrelevant feature for this feature to be learned.

Here, we show that top-down processes can be sufficient to produce PL in a detection task. To isolate the top-down contribution of PL, we contrasted the learning associated with two previously unknown target textures. Participants were exposed equally often to these two textures but did 24 times more no-signal trials with one than with the other. A no-signal trial consisted in the presentation of a pair of noise fields (Gosselin & Schyns, 2003; see also Gosselin, Bacon, & Mamassian, 2004). Although these fields never contained any consistent shape, we observed a greater improvement following more no-signal practice. We will argue that this can only be attributed to top-down processes.

2. Methods

2.1. Participants

Eleven psychology students aged between 22 and 29 were allotted either to the experimental group ($n=8$) or to the control group ($n=3$). All participants were naïve to the purpose of the experiment but they knew that they were only presented noise fields during the

* Corresponding author.

E-mail addresses: nicolas@dupuis.ca (N. Dupuis-Roy), frederic.gosselin@umontreal.ca (F. Gosselin).

¹ *Abbreviations used:* PL, perceptual learning; CI, classification image; T1, texture one; T2, texture two.

no-signal trials. Only F.L., from the experimental group, had previous experience with the no-signal procedure. All participants had normal or corrected-to-normal vision. Informed consent was obtained from all participants before the beginning of the experiment and a monetary compensation was provided. This study conformed to the tenants of the Declaration of Helsinki.

2.2. Apparatus

The experimental programs were run on a Macintosh G4 computer in the Matlab environment, using the Psychophysics toolbox (Brainard, 1997; Pelli, 1997). All stimuli were presented on a Sony Trinitron monitor (1024 × 768 pixels at 85 Hz), calibrated using a Samsung SyncMaster 753 df photometer to allow linear manipulation of luminance. The resulting corrected table contained 137 luminance levels, ranging from 0.31 to 107 cd/m². The background luminance was equal to 52.85 cd/m². Participants were tested binocularly their heads immobilized by a chin rest at a distance of 57 cm from the computer monitor.

2.3. Stimuli

Eight random textures (32 × 32 pixels, spanning 1 × 1 deg of visual angle) were created by filtering white Gaussian noise fields with an isotropic $1/f^2$ filter. We assigned a pair of such textures to each participant (Fig. 1a). These pairs of textures were (Pearson) uncorrelated and the histogram of their contrasts was Gaussian with a mean equal to zero. Thousands of pairs of white Gaussian noise fields (32 × 32 pixels, spanning 1 × 1 deg of visual angle) were also produced for the no-signal trials. The histogram of the (Pearson) correlations between each of the $1/f^2$ textures and these noise fields had a mean of zero and a standard deviation of 0.03.

2.4. Procedure

Fig. 1b outlines the time course of the experiment for participant F.L. on three time dimensions: within-trial, within-session, and between-session time. This figure will serve as the backbone of our general explanation of the experimental procedure because the other participants from the experimental group received a treatment similar to that received by F.L.. Consider the first T1-session. On the within-session time dimension, you have a sample of all the events that occurred during the session. These within-session events were of four types: a 10-s presentation of either the first (T1-pres., in Fig. 1b) or the second target texture (T2-pres.); and a no-signal practice trial associated with either target texture T1 (T1-trial) or target texture T2 (T2-trial). In a T1-session, there were 240 T1-trials and 10 T2-trials whereas, in a T2-session, there were 240 T2-trials and 10 T1-trials. All sessions included two T1-presentations and two T2-presentations. Within-session events were randomly ordered as the crossing arrows indicate.

The within-trial time dimension details what was presented to participants on the computer monitor during the within-session events: during a T1- and T2-presentation, a written cue – either ‘T1’ or ‘T2’ – was displayed at the center of the screen for 1 s; immediately followed by the presentation of the corresponding texture – either T1 or T2 – at the center of the screen for 10 s. During a T1- and T2-trial, a written cue – either ‘T1’ or ‘T2’ – appeared at the center of the screen for 1 s; immediately followed by a white Gaussian noise field presented at the center of the screen for 500 ms; immediately followed by a uniform mid-luminance screen presented 100 ms; immediately followed by another white Gaussian noise field presented at the center of the screen for 500 ms; and immediately followed by a uniform mid-luminance screen presented 100 ms. Participants had to decide which one of the two white Gaussian noise fields was more similar to the cued texture. They knew they were presented noise fields with no signal. No feedback was provided to them.

Participants were paired on the targets to detect and randomization (boxes in Fig. 1a). The randomization determined the order of the texture presentation, the sequence of T1- and T2-trials, and the noise fields. In fact, the only physical difference between paired participants was the

written cues – ‘T1’ and ‘T2’ – attached to the texture presentations. In other words, paired participants differed essentially in the number of times they attempted to detect the two textures in noise fields.

On the between-session time dimension, you can see that participant did no more than one experimental session per day. Two of the eight observers – F.L. being one of them – also did a T2-session and a rotated-T1-session and, after a one-year interruption, ten more T1-sessions. The rationale for these additional tests will be given in Section 3. The participants from the control group were submitted, at a maximum rate of one session per day, over a period of about two months, to two T1-sessions, one before and one after 34 blank sessions; a blank session was identical to a T1-session (or T2-session) except that mid-luminance screens presented for 2.2 s replaced all T1- and T2-trials. Control participants were instructed to carefully attend to the presented textures and to remember them as well as possible in preparation for a second and final T1-session. The rationale for this control condition will also be provided in Section 3.

3. Results and discussion

3.1. Experimental condition: Average results

Gosselin and Schyns (2003; see also Gosselin et al., 2004) have induced the perception of simple objects in noise fields containing no signal. Their no-signal procedure did not allow for the measurement of performance. In our version of the procedure, the cued target texture is always more (Pearson) correlated with one white Gaussian noise field than with the other.² Hence, we calculated performance as the percentage of no-signal trials for which the noise field more correlated with the cued target was chosen. Fig. 2 shows the performance on the first 25 T1- and T2-sessions,³ averaged across the eight experimental participants. We computed least-square linear regressions on a log–log transformation of the T1- and T2-performance data. The best fits are represented as solid black curves and their equations are displayed at the top-left corner of the panels. The dotted lines show the statistical threshold calculated from a Bonferroni-corrected, one-tail binomial test ($p < .05$). The average T1-performance increased, starting at approximately 53.7% (ns) and finishing at approximately 59.4% ($p < .05$). There is no difference between the average of T1-performance on the first 10 T1-trials (.513) and the average T2-performance on the first T1-session (.613, $N = 80$, $Z = .45$, ns). Fine and Jacobs (2002) defined the L index as $d'_{\text{post-treatment}}/d'_{\text{pre-treatment}}$ to quantify learning. Assuming no response bias, $d' = 2Z(P(\text{correct}))$. The L index for the average T1-performance was thus equal to 2.56 ($d'_{\text{post-treatment}} = 0.48/d'_{\text{pre-treatment}} = 0.19$). In contrast, no improvement was observed in T2-performance.

Any actual difference between the T1- and T2-performance can be ascribed to the difference in the number of no-signal trials because everything else was the same. However, our failure to find an improvement in T2-performance might reflect a lack of statistical sensitivity rather than no

² No special precaution was taken to avoid ties but they never actually happened during the experiment.

³ Twenty-five sessions is the least completed by an experimental participant.

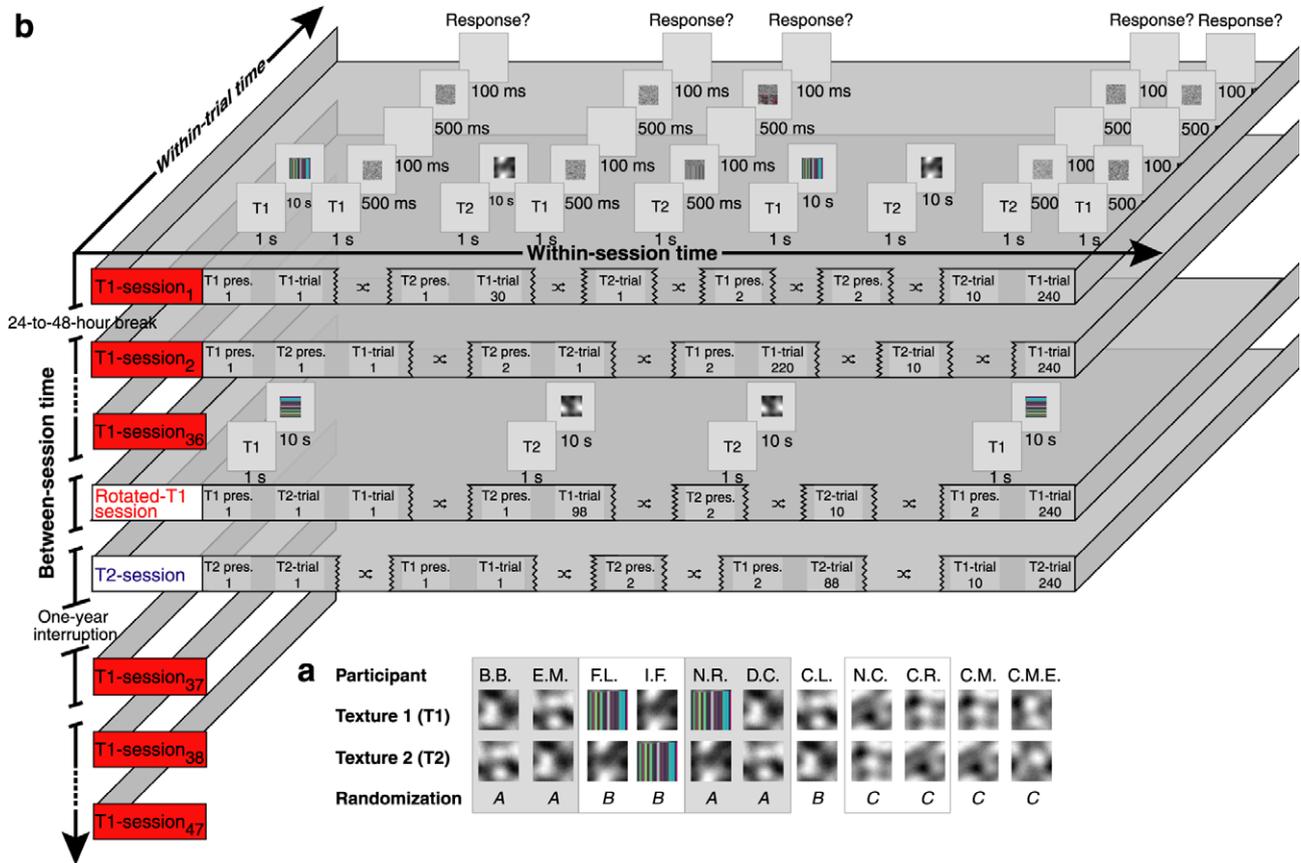


Fig. 1. (a) All $1/f^2$ textures and their assignment to the eight experimental participants (left) and the three control participants (right). Paired participants (boxes) were presented with the same sequence of noise fields; they saw the same two target textures the same number of times; they only differed in the number of times they attempted to detect each of these textures. (b) Outlines the time course of the experiment for participant F.L. on three-dimensions: within-trial, within-session, and between-session time. Exposure to the T1- and T2-textures were identical. In a T1-session, there were 240 T1-trials and 10 T2-trials. T1- and T2-trials contained no signal whatsoever. Therefore, the T1- and T2-textures differed only in no-signal practice.

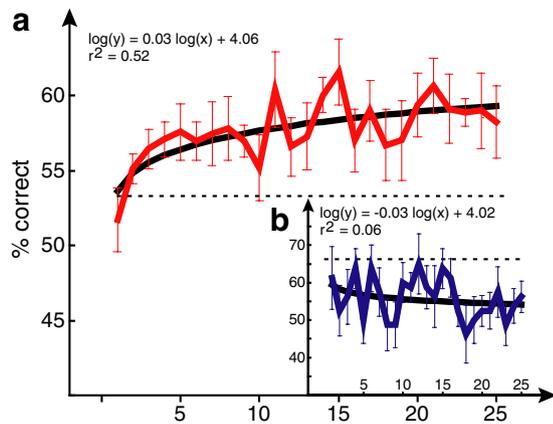


Fig. 2. Percentage of correct responses averaged across the eight participants for the first 25 T1-sessions. (a) T1-performance is in red and (b) T2-performance is in blue. The best fits are in black. The dotted lines indicate the statistical threshold. The length of the error bars is equal to two standard errors.

actual learning. Indeed T1- and T2-performance are not directly comparable due to difference in sample sizes.

To resolve this problem, we estimated the distribution of T1-slopes (the slopes of the least-square linear log-log fits

to the T1-performance) calculated with a reduced sample size via Bootstrap (e.g., Efron & Tibshirani, 1986).⁴ Here, T1-slopes are estimates of the magnitude of PL. The resulting distribution had a mean of 0.03 and a standard deviation of 0.02. The probability of observing the T2-slope of the experimental group (-0.03) given this distribution is less than 0.56%. Thus part of the PL observed in the average data at least is due solely to differences in the no-signal treatments.

3.2. Experimental condition: Individual results

We performed similar Bootstrap analyses on individual data. Results are shown in Table 1. Only the data of I.F. ($p < .0006$) and of F.L. ($p < .0088$) reached statistical significance.

The T1- and T2-slope differences of F.L. and I.F. cannot be explained by intrinsic features of the learned textures or of the randomization procedure because these were paired

⁴ More explicitly, we randomly drew a sample of 10 T1-trials responses from each of the first 25 T1-sessions of each participant, fitted a log-log line on the means of each sampled sessions to extract a T1-slope and repeated the previous steps 10,000 times.

Table 1
Column 2: slopes of the best-fitted log–log lines on T2-performance

Participants	T2-slope	Parameters of the simulated T1-slopes null distribution		$P_{T2\text{-slope}/T1\text{-trials}}$
		\bar{x}	S	
		B.B.	-0.06	
E.M.	0.08	0.27	0.25	<i>ns</i>
F.L.	-0.46	0.12	0.24	<.0088
I.F.	-0.58	0.21	0.24	<.0006
N.C.	0.38	0.12	0.17	<i>ns</i>
C.R.	-0.23	0.13	0.32	<i>ns</i>
C.M.	-0.05	0.13	0.20	<i>ns</i>
C.M.E.	0.12	-0.08	0.19	<i>ns</i>

Columns 3 and 4: the mean and the standard deviation of the simulated null distribution of T1-slopes ($n = 10,000$) calculated from samples of 10 T1-trials. Column 5: the probability of observing the T2-slope given the simulated population. *ns*, non significant.

subjects; they were presented with exactly the same sequence of events, except that the ‘T1’ and ‘T2’ written cues attached to the texture presentations were swapped (Section 2.4; Fig. 1). However, some intrinsic features of the pair of textures of F.L. and I.F. might have minimized the T1–T2 learning transfer.

Fig. 3 shows the performance of these two participants on the T1-sessions (red curves). The thick solid black curves are the best-fitted log–log lines to the accuracy of the first 36 T1-sessions, and the thick dotted curves are the forecasted accuracies based on these best fits.

The gray areas delimit the confidence interval ($p < .05$). The thin dotted lines show the statistical threshold calculated from a Bonferroni-corrected, one-tail binomial test ($p < .05$). The T1-performance of I.F. and F.L. increased, starting at approximately 55.5% (*ns*) and 57.1% (*ns*), respectively, and attaining approximately 65.4% ($p < .05$) and 65.4% ($p < .05$) after 36 T1-sessions. This corresponds to L indexes of 2.86 ($d'_{\text{post-treatment}} = 0.79/d'_{\text{pre-treatment}} = 0.28$) and of 2.21 ($d'_{\text{post-treatment}} = 0.79/d'_{\text{pre-treatment}} = 0.36$), respectively, for I.F. and F.L.. This improvement is also reflected in the increasing number of T1-sessions above statistical threshold with practice. In addition, each of these two participants completed one T2-session (240 T2-trials and 10 T1-trials) following the 36 T1-sessions (Section 2.4; Fig. 1b). For F.L., the difference between the performances on this T2-session (57.5%; $N = 240$) and the average T2-performances on the first 36 sessions (56.7%; $N = 360$) was not significant ($Z = 0.20$, *ns*); for

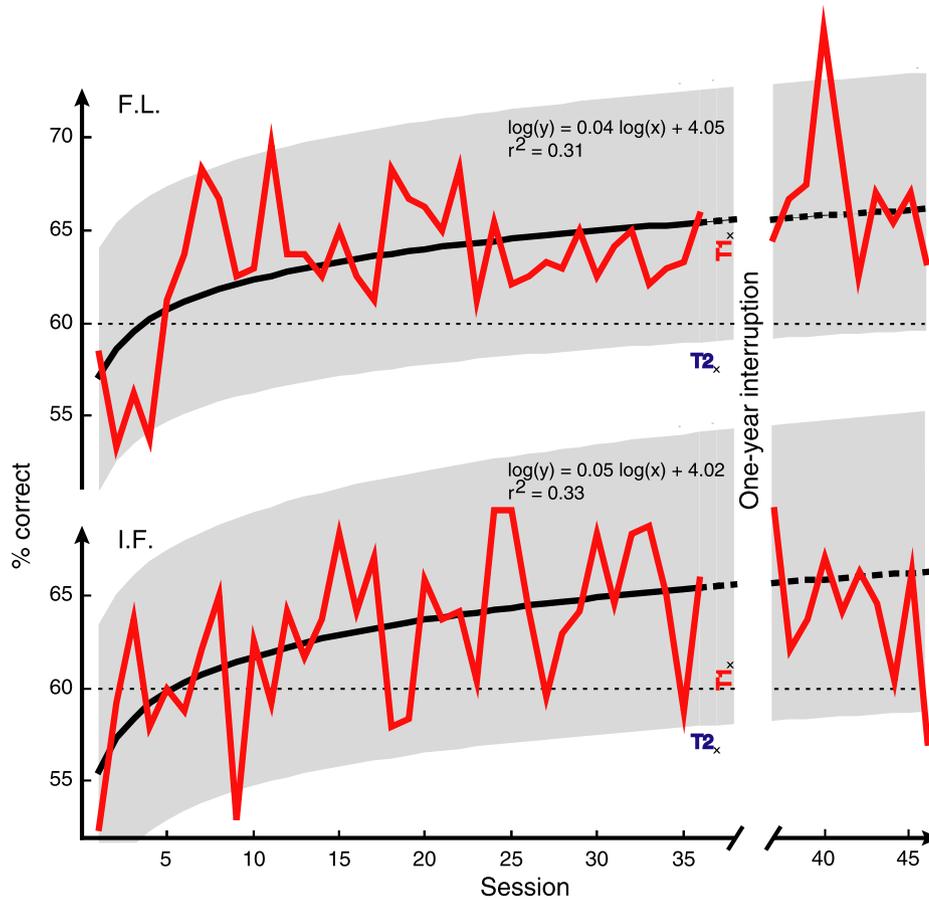


Fig. 3. The performance of I.F. and F.L. in T1-sessions is displayed as red curves and the best-fitted log–log lines as solid black curves. The bold dotted black curves indicate the T1-performance predicted from the first 36 T1-sessions. The performance of I.F. and F.L. in the T2-transfer session is represented by a blue ‘T2’ and that in the orientation transfer test by a rotated red ‘T1’. The right panel depicts T1-performance of I.F. and F.L. after a one-year break. The gray area is the prediction interval. The thin dotted lines are the statistical threshold.

I.F. this difference was significant (56.7%; $N = 240$; 64.7%; $N = 360$; $Z = -1.99$, $p < .025$) albeit in a direction, indicating forgetting rather than learning. The converging evidence presented in this section demonstrates that at least part of the observed PL in at least some participants can be attributed to differences in the number of completed no-signal trials.

Several factors could explain why we did not observe a significant effect in more individual participants. As previously mentioned, we might not have enough statistical sensitivity to detect actual differences. The T2-slope of six out of eight observers is less than the mean of the estimated distribution of T1-slopes. This non-significant trend and the average results presented in the previous section give some weight to this speculation. Another possible factor is individual differences. Participants were not as talented at perceiving the target textures in pure noise. One indication of this is found in Table 2. This table shows individual T1-slopes and their probability given the null hypothesis. The slopes of two participants (C.R. and C.M.E.) did not reach significance. Moreover, participants might have had different levels of motivation. A final possibility is intrinsic properties of the pairs of textures. Some pairs of textures might have allowed more T1–T2 learning transfer. This could be the case for the textures used by paired participants B.B. and E.M.

3.3. Control condition

The experimental participants had to learn previously unknown textures via signal-dependent PL. This signal-dependent learning was sufficient after one T1-session, or equivalently two exposures to the T1-texture, for the performance of subjects to reach statistical significance (see Fig. 2). The textures were presented to participants in all T1-sessions. Thus this signal-dependent PL might have been abrupt – only contaminating learning in the first few T1-sessions – or it might have been gradual – possibly contaminating all T1-sessions. The failure to find any improvement in the T2-performance suggests indirectly that the signal-dependent PL reached an asymptote abruptly or was small in comparison to no-signal PL. In the control condition, we measured the effect of these repeated

texture presentations directly using two T1-sessions. None of the performances at the post-test (55%, 44.17%, and 49.58%) were greater than the performances at the pre-test (55.42%, 55.42%, and 51.67%). This suggests gradual no-signal PL and abrupt or comparatively small, signal-dependent PL. We will thus assume in the remainder of this article that examining T1-learning informs us mostly about the characteristics of no-signal PL.

3.4. Additional tests

We submitted I.F. and F.L. – the two participants who showed reliable differences between their T1- and T2-slopes – to two additional tests. This first one was designed to assess the orientation specificity of the no-signal PL. Participants completed a rotated-T1-session that is, a T1-session in which the two target textures were rotated 90° counterclockwise. Subjects were informed of this alteration to the target textures. Their performance (64.58% and 61.25%; rotated ‘T1’ in Fig. 3) fell within the prediction intervals ($p < .05$). Thus no-signal PL can be invariant to a 90° counterclockwise rotation.

Finally, we assessed the retention of no-signal PL with 10 supplementary T1-sessions completed after a one-year interruption. Only two of the twenty new data points fell outside the prediction area ($p < .05$), one per participant. One of these outliers was above the prediction interval. Thus no-signal PL can be remarkably resistant to the passage of time.

3.5. Classification image analyses

No-signal PL was further examined via classification image analyses (Gosselin & Schyns, 2003). Only the six participants having T1-slopes significantly greater than zero were studied (*family-wise* $p < 0.05$; Table 2).

We performed least-square multiple linear regressions on the pairs of random white Gaussian noise fields (explanatory variables) and accuracy (predictive variable). This amounts to subtracting the sum of all rejected noise fields from the sum of all selected noise fields. The resulting planes of regression coefficients are called Classification Images (CI) (Eckstein & Ahumada, 2002). We computed CIs at three critical stages of learning: the first three sessions (0% of learning), the three sessions at half the performance range⁵ (50% of learning), and the last three sessions (100% of learning). We also performed least-square multiple linear regressions on the pairs of white Gaussian noise fields and the correct responses. These ideal CIs can be understood as benchmarks to which human CIs can be compared.

The colored pixels in Fig. 4a reached statistical significance in a series of Bonferroni-corrected (*family-wise*

Table 2
Column 2: slopes of the best-fitted log–log line on T1-performance

Participants	T1-slope	p
B.B.	0.04	<.0037
E.M.	0.07	<.0000
F.L.	0.04	<.0001
I.F.	0.05	<.0000
N.C.	0.03	<.0017
C.R.	0.02	<i>ns</i>
C.M.	0.03	<.0010
C.M.E.	–0.01	<i>ns</i>

Columns 3: the probability of observing the T1-slope given the null hypothesis. *ns*, non significant.

⁵ We fitted the following bilinear curve to the T1-performance data points: $y = b_1x + a_1$ when $x < x_0$, and $y = a_2$ when $x \geq x_0$, where a_2 is the learning asymptote.

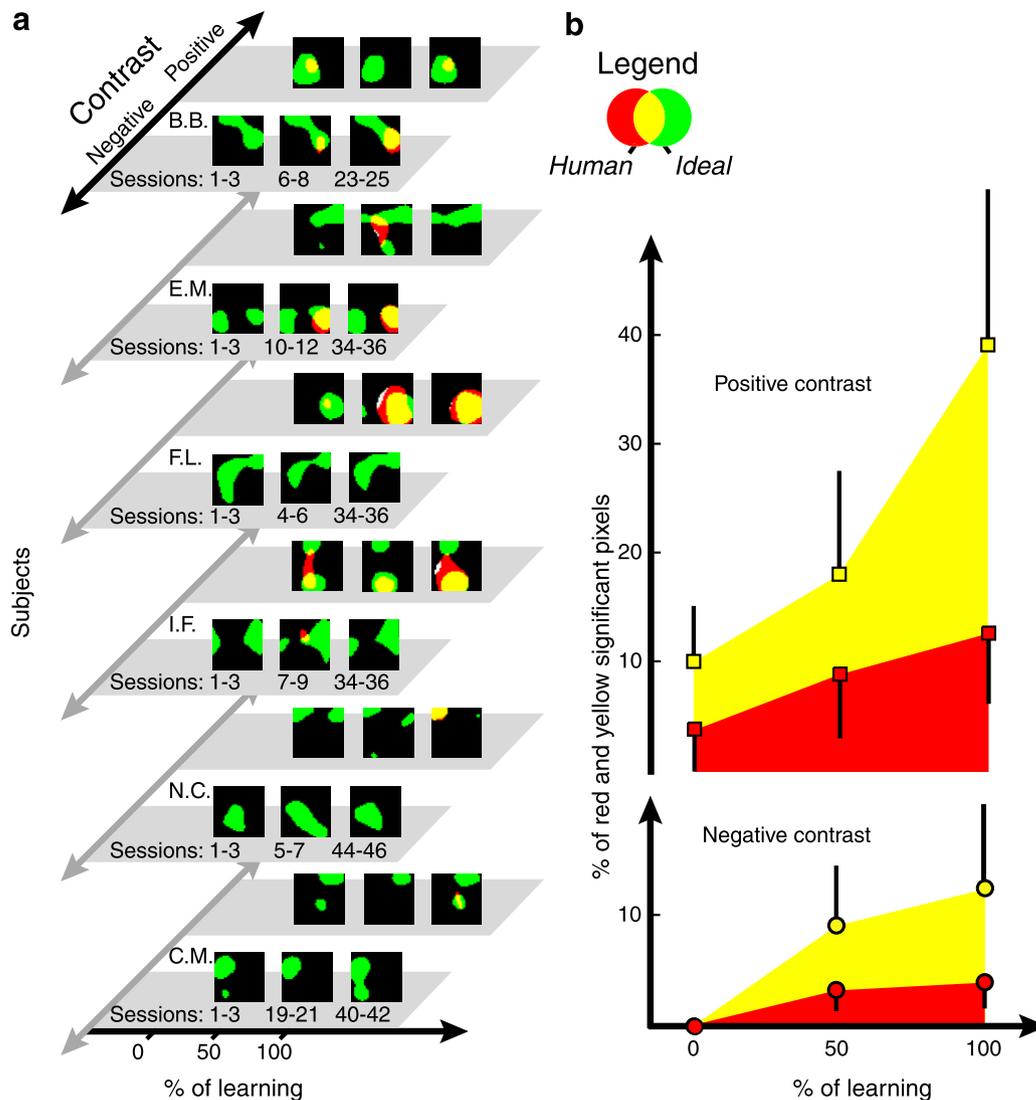


Fig. 4. (a) Positively and negatively contrasted pixels reaching statistical significance in the classification images (CIs) of human observers (red), ideal observers (green), and both classes of observers (yellow) at three critical stages of learning: the first three T1-sessions (or 0% learning); the three T1-sessions around the middle performance range (or 50% learning); and the last three T1-sessions (or 100% learning). (b) The average proportion of significant pixels that are red and yellow at the three critical stages of learning and for positively and negatively contrasted pixels.

$p < 0.05$), two-tail Pixel tests⁶ (Chauvin, Worsley, Schyns, Arguin, & Gosselin, 2005). CI pixels attaining significance at the positive and the negative tail are presented in two images. Green pixels reached significance only for the ideal observers, red ones only for the human observers, and yellow ones for both classes of observers. Fig. 4b illustrates the average proportion of significant pixels that are red and yellow at the three critical stages of learning. The information used by humans only (red pixels) and the information used by both humans and ideals (yellow pixels) increase monotonically. This suggests that no-signal PL is at least partly due to an increase in absolute efficiency (Murray, Bennett, & Sekuler, 2005). On average, human observers showed a bias for positively contrasted regions

of the T1 textures (Fig. 4b, upper half vs. lower half) even though the ideal observer did not. That paired participants B.B. and E.M. did not exhibit this bias could mean that intrinsic properties of paired textures might have been at least partly responsible for the bias.

4. General discussion

We believe that two kinds of PL intervened in our experimental condition. Firstly, high-level representations of the T1- and T2-textures were learned during the texture presentations. This first kind of learning necessarily had to contain a signal-dependent component because the textures were previously unknown to the participants and all participants recognized them above chance level. It might also have had a top-down component; we will briefly come back to this point later. In any case, the results obtained in the

⁶ The CIs were smoothed using a Gaussian kernel with a standard deviation of 3.96 pixels and Z-scored prior to the Pixel tests.

control condition indicate that this first kind of learning was not sufficient to produce an increase in performance measurable by our no-signal instrument. This suggests that signal-dependent PL was abrupt, possibly reaching an asymptote after as few as two exposures to the textures, or of small magnitude.

Secondly, and most importantly, PL occurred during the no-signal trials. We contrasted two learning conditions (T1- and T2-learning), which only differed in the number of no-signal trials they comprised. On average, experimental participants showed a T1-over-T2 learning advantage. Two individual participants – F.L. and I.F. – also showed a reliable T1-over-T2 advantage. This advantage can be ascribed only to top-down processes because no signal was presented to the participants during the no-signal trials.

Each no-signal trial contained two noise fields. At the limit, these noise fields would sum into a homogenous field. This is unlike the stimuli used by Shiu and Pashler (1992), by Ahissar and Hochstein (1993), and by Seitz, Watanabe and colleagues that is, the stimuli used in the main experiments that have established that top-down processing is sometimes necessary for PL. In Watanabe et al. (2001), for example, the subliminal and task-irrelevant coherent motion summed to the learnt direction. The task of our participants was to determine which noise field from each pair of noise fields was most (Pearson) correlated with the cued texture. In 99% of the no-signal trials, the correlation between the textures and the noise fields was less than .07. The noise fields never resembled much the cued textures. This is unlike the Shiu and Pashler (1992) and the Ahissar and Hochstein (1993) stimuli, but somewhat like the Seitz, Watanabe and colleagues "subliminal" stimuli. However, the goal of the later researchers was to minimize top-down processing whereas our goal was to minimize bottom-up, signal-dependent processing. A clear indication that we have succeeded is that participants I.F. and F.L. were submitted to the same unstructured, visual stimulation and nonetheless exhibited orthogonal PL. For all these reasons, the no-signal trials are best understood as probes that minimally bias behavior.

Overall we found no evidence of T1-to-T2 transfer although, as discussed previously, the pair of textures presented to participants B.B. and E.M. might have allowed a minimum amount of T1-to-T2 transfer. This suggests that no-signal PL is not caused by changes resulting from characteristics shared by paired textures. Indeed, it suggests that it cannot be the consequence of global spatial frequency tuning because the two textures had $1/f^2$ energy spectrums. Furthermore, it suggests that it cannot be the consequence of a task-specific improvement such as the capacity to integrate an increasing amount of pixels because the task was the same for T1- and T2-trials. In fact, only the shapes of the textures differed in T1- and T2-trials. Classification image analyses allowed us to visualize the transformation occurring in the shape of the memory representations of these textures. Comparison between human

and ideal classification images showed an enlarging overlap. We believe that during the no-signal trials the high-level representations of the T1-textures (and, to a lesser extent, of T2-textures) activated mid-level cells involved in the coding of visual features in a top-down fashion; that each activation of the high-level representations resulted in a greater tuning of these mid-level cells to the features of these representations possibly via Hebbian learning.

Sigala and Logothetis (2002) have discovered that feature selectivity in the macaque inferior temporal cortex is shaped by categorization of objects. Inferotemporal cells are thus promising candidates for the role of the mid-level cells in the above story. We have gathered some evidence that no-signal PL could be rotation invariant; a hallmark of inferotemporal cells is that they often show robustness to stimulus transformations such as scale and position and possibly to other affine transformations, including rotation (Riesenhuber & Poggio, 2004). Also the preference of the experimental participants for positively contrasted blobs suggests that the no-signal PL occurred after figure-ground segregation; inferotemporal cells have been shown to discriminate between figure and ground (Baylis & Driver, 2001; Kourtzi & Kanwisher, 2001). Finally, we have collected evidence, indicating that no-signal could be long-lasting; Kobatake, Wang, and Tanaka (1998) have found long-lasting effects of PL in inferotemporal neurons.

Such top-down tuning of inferotemporal neurons could accompany bottom-up tuning of low- to mid-level cells in PL with signal. Consistently, the magnitude of no-signal PL (top-down) appears to be somewhat less than that found in comparable signal-dependent PL (bottom-up and possibly top-down). Fine and Jacobs (2002) examined the magnitude of PL in 16 psychophysical studies with signal. The L index observed in the study most similar to ours – a band-passed (2–4 cycles per image) random texture discrimination task – was of about 2.25 after four sessions (Gold, Bennett, & Sekuler, 1999); we observed an average L index of about 1.32 after four sessions (participant I.F. and F.L. exhibited L indexes of 1.71 and 1.45 after four sessions, respectively). The largest L index reported in the Fine and Jacob review is of about 5; we observed an average L index of about 2.56 (participant I.F. and F.L. exhibited total L indexes of 2.86 and 2.21, respectively).

In sum, we have shown that top-down processes can be sufficient to produce substantial, possibly rotation-invariant and long-lasting PL. More work is needed to examine precisely the nature of the top-down component of PL. For example, is it really rotation invariant? Does it show other affine invariance (Riesenhuber & Poggio, 2004)? Is it polarity invariant (Baylis & Driver, 2001)? How long does it last? Is it partly due to internal noise reduction unlike PL with signal (Gold et al., 1999)? Does it show interference effects such as PL with signal (Seitz, Yamagishi, Werner, Kawato, & Watanabe, 2005)? We believe that the no-signal procedure employed in this article will be an important tool for this endeavor.

Acknowledgments

We thank Catherine Mello, France Landry, Cynthia Roy, and Karine Tadros for their helpful corrections. This study was supported by FCAR scholarship awarded to Nicolas Dupuis-Roy and by NSERC and NATEQ grants awarded to Frédéric Gosselin.

References

- Ahissar, M., & Hochstein, S. (1993). Attentional control of early perceptual learning. *Proceedings of the National Academy of Science USA*, *90*, 5718–5722.
- Ahissar, M., & Hochstein, S. (2000). The spread of attention and learning in feature search: effects of target distribution and task difficulty. *Vision Research*, *40*, 1349–1364.
- Ahissar, M., & Hochstein, S. (2002). The role of attention in learning simple visual task. In M. Fahle & T. Poggio (Eds.), *Perceptual learning* (pp. 367–380). Cambridge: MIT Press.
- Baylis, G. C., & Driver, J. (2001). Shape-coding in IT cells generalizes over contrast and mirror reversal, but not figure-ground reversal. *Nature Neuroscience*, *4*, 937–942.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.
- Chauvin, A., Worsley, K. J., Schyns, P. G., Arguin, M., & Gosselin, F. (2005). Accurate statistical tests for smooth classification images. *Journal of Vision*, *5*(9), 659–667.
- Eckstein, M.P. & Ahumada, A.J., Jr. (2002). Classification images: a tool to analyze visual strategies. *Journal of Vision*, *2*(1), i–i. <<http://journalofvision.org/2/1/i/>>. doi:10.1167/2.1.i.
- Efron, B., & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, *1*, 54–75.
- Fahle, M., Edelman, S., & Poggio, T. (1995). Fast perceptual learning in hyperacuity. *Vision Research*, *35*, 3003–3013.
- Fahle, M., & Morgan, M. (1996). No transfer of perceptual learning between similar stimuli in the same retinal position. *Current Biology*, *6*, 292–297.
- Fine, I. & Jacobs, R. A. (2002). Comparing perceptual learning tasks: a review. *Journal of Vision*, *2*(2), 190–203. <<http://journalofvision.org/2/2/5/>>. doi:10.1167/2.2.5.
- Godde, B., Stauffenberg, B., Spengler, F., & Dinse, H. R. (2000). Tactile co-activation-induced changes in spatial discrimination performance. *Journal of Neuroscience*, *20*, 1597–1604.
- Gold, J., Bennett, P., & Sekuler, A. (1999). Signal but not noise changes with perceptual learning. *Nature*, *402*, 176–178.
- Gosselin, F., Bacon, B. A., & Mamassian, P. (2004). Internal surface representations approximated by reverse correlation. *Vision Research*, *44*, 2515–2520.
- Gosselin, F., & Schyns, P. G. (2003). Superstitious perceptions reveal properties of internal representations. *Psychological Science*, *14*, 505–509.
- Hodjic, A., Veit, R., Karim, A. A., Erb, M., & Godde, B. (2004). Improvement and decline in tactile discrimination behavior after cortical plasticity induced by passive tactile co-activation. *Journal of Neuroscience*, *24*, 442–446.
- Jiang, Y., & Chun, M. M. (2001). Selective attention modulates implicit learning. *Quarterly Journal of Experimental Psychology*, *54*(A), 1105–1124.
- Kobatake, E., Wang, G., & Tanaka, K. (1998). Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. *The Journal of Neurophysiology*, *80*, 324–330.
- Kourtzi, Z., & Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science*, *293*, 1506–1509.
- Murray, R. F., Bennett, P. J., & Sekuler, A. B. (2005). Classification images predict absolute efficiency. *Journal of Vision*, *5*(2), 139–149. <<http://journalofvision.org/5/2/5/>>. doi:10.1167/5.2.5.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.
- Riesenhuber, M., & Poggio, T. (2004). In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences* (pp. 1640–1653). Cambridge: MIT Press.
- Seitz, A., Lefebvre, C., Watanabe, T., & Jolicoeur, P. (2005). Requirement for high-level processing in subliminal learning. *Current Biology*, *15*, R753–R755.
- Seitz, A., Nanez, J. E., Holloway, S. R., Koyama, S., & Watanabe, T. (2005). Seeing what is not there shows the costs of perceptual learning. *Proceedings of the National Academy of Science USA*, *102*, 9080–9085.
- Seitz, A. R., & Watanabe, T. (2003). Is subliminal learning really passive. *Nature*, *422*, 36.
- Seitz, A., & Watanabe, T. (2005). A unified model for perceptual learning. *Trends in Cognitive Science*, *9*, 329–334.
- Seitz, A., Yamagishi, N., Werner, B., Kawato, M., & Watanabe, T. (2005). Task-specific disruption of perceptual learning. *Proceedings of the National Academy of Sciences*, *102*, 14895–14900.
- Shiu, L.-P., & Pashler, H. (1992). Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Perception and Psychophysics*, *52*, 582–588.
- Sigala, N., & Logothetis, N. K. (2002). Visual categorization shapes feature selectivity in primate temporal cortex. *Nature*, *415*, 318–320.
- Vaina, L. M., Sundaeswaran, V., & Harris, J. G. (1995). Learning to ignore: psychophysics and computational modeling of fast learning of direction in noisy motion stimuli. *Cognitive Brain Research*, *2*, 155–163.
- Watanabe, T., Nanez, J. E., Sr., Koyama, S., Mukai, I., Liederman, J., & Sasaki, Y. (2002). Greater plasticity in lower-level than higher-level visual motion processing in a passive perceptual learning task. *Nature Neuroscience*, *5*, 1003–1009.
- Watanabe, T., Nanez, J. E., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature*, *413*, 844–848.