

**Low-level sensory plasticity during task-irrelevant perceptual learning:
Evidence from conventional and double training procedures**

Praveen K. Pilly¹, Stephen Grossberg¹, and Aaron R. Seitz²

¹ Department of Cognitive and Neural Systems
Center for Adaptive Systems
Center of Excellence for Learning in Education, Science, and Technology
Boston University
677 Beacon St
Boston, MA 02215
(advait@cns.bu.edu; steve@bu.edu)

² Department of Psychology
University of California Riverside
900 University Ave
Riverside, CA 92521
(aseitz@ucr.edu)

To appear in *Vision Research*, Special Issue on Perceptual Learning

Technical Report BU CAS/CNS-TR-09-007

Abstract

Studies of perceptual learning have focused on aspects of learning that are related to early stages of sensory processing. However, conclusions that perceptual learning results in low-level sensory plasticity are controversial, since such learning may also be attributed to plasticity in later stages of sensory processing or in readout from sensory to decision stages, or to changes in high-level central processing. To address this controversy, we developed a novel random dot motion (RDM) stimulus to target motion cells selective to contrast polarity by ensuring the motion direction information arises only from signal dot onsets and not their offsets, and used these stimuli in the paradigm of task-irrelevant perceptual learning (TIPL). In TIPL, learning is achieved in response to a stimulus by subliminally pairing that stimulus with the targets of an unrelated training task. In this manner, we are able to probe learning for an aspect of motion processing thought to be a function of directional V1 simple cells with a learning procedure that dissociates the learned stimulus from the decision processes relevant to the training task. Our results show direction-selective learning for the designated contrast polarity that does not transfer to the opposite contrast polarity. This polarity specificity was replicated in a double training procedure in which subjects were additionally exposed to the opposite polarity. Taken together, these results suggest that TIPL for motion stimuli may occur at the stage of directional V1 simple cells. Finally, a theoretical explanation is provided to understand the data.

1. Introduction

The stages of brain processing at which perpetual learning takes place is a topic of ongoing controversy. While some studies have argued that perceptual learning can take place at early stages of sensory processing (Furmanski, Schluppeck, & Engel, 2004; Pourtois et al., 2008; Schoups et al., 2001; Watanabe et al., 2002), other studies have proposed that most of the learning takes place in the readout from sensory to decision stages (Doshier & Lu, 1998; Law & Gold, 2008; Smirnakis et al., 2005) or in higher non-retinotopic, central stages (Xiao et al., 2008; Zhang et al., 2009). A likely resolution of this long-standing controversy is that the degree to which learning occurs at various processing stages may depend on details of the procedure, such as training task difficulty (Ahissar & Hochstein, 1997; Liu & Weinshall, 2000), transfer task precision (Jeter et al., 2009), and training procedure (Xiao et al., 2008; Zhang et al., 2009).

A paradigm that has been suggested to lead to a greater degree of low-level sensory plasticity is that of task-irrelevant perceptual learning (TIPL); see Seitz and Watanabe (2009), in this issue, for a detailed review of TIPL. Studies of TIPL have made the remarkable discovery that the brain can refine its sensitivity to parathreshold stimulus features in the environment even when they are presented without awareness or focused attention (Watanabe, Sasaki, & Nanez, 2001; Tsushima, Seitz, & Watanabe, 2008). Moreover, human lateral prefrontal cortex, which is known to subserve high-level functions such as cognitive control and decision-making, has been found to be inactive in conditions that promote TIPL (Tsushima, Sasaki, & Watanabe, 2006). Furthermore, the paradigm of TIPL dissociates the learning aspect of interest (i.e., the task-irrelevant stimulus) from the decision processes that correspond to the explicit training task that the subjects are required to perform with full attention. In fact, TIPL has been found to occur through stimulus-reward conjunctions in the absence of any task for the subjects (Seitz, Kim, & Watanabe, 2009). Given these factors, plasticity in the decision stages seems an unlikely account for TIPL.

While studies of TIPL have generated much attention, they have to date mostly focused on *when* learning occurs and have revealed that for TIPL to occur, the task-irrelevant stimuli should be perceptually neither too weak nor too strong (Tsushima et al., 2008), and either attention needs to be engaged on task-relevant stimuli (Seitz et al., 2005) or external rewards need to be delivered (Seitz et al., 2009) so that reinforcement learning signals can spill over to spatiotemporally nearby task-irrelevant stimuli (Nishina et al., 2007; Seitz & Watanabe, 2003). In this present study, however, we are interested in questions regarding *what* can be learned and *where* in the brain does plasticity occur during TIPL.

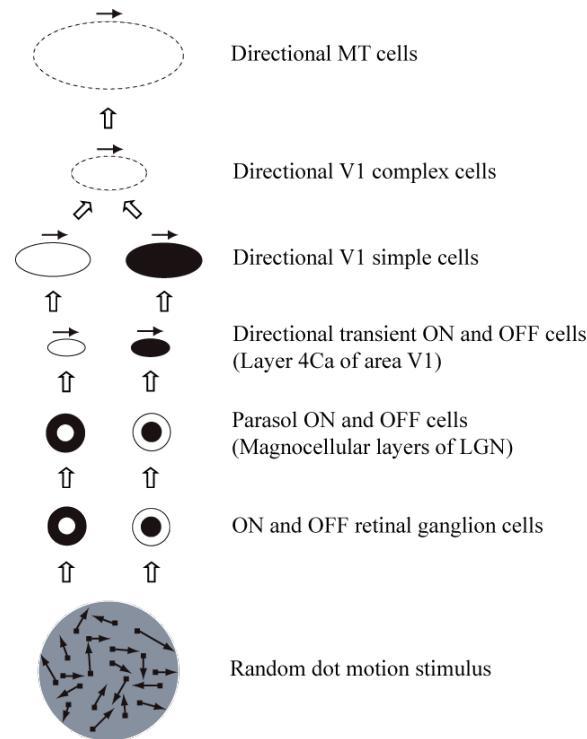


Figure 1. Cartoon of how early motion processing is segregated into ON- and OFF-channels until they summate at the stage of directional V1 complex cells. Several physiological (Alonso, Usrey, & Reid, 2001; Conway & Livingstone, 2003; Goodwin & Henry, 1975; Livingstone, 1998; Movshon & Newsome, 1996; Reid & Alonso, 1995; Schiller, 1982, 1992) and psychophysical (Croner & Albright, 1997; Edwards & Badcock, 1994; van der Smagt & van de Grind, 1999; Wehrhahn & Rapf, 1992) studies support this segregation in both the form and motion pathways.

Notably Watanabe et al. (2002) suggested that in regard to motion stimuli, TIPL occurs at a low-level of motion processing because learning was found only for the local motion directions but not the global direction of a task-irrelevant dynamic dot display in which no dot moved in the perceived global motion direction. However, several stages of cortical motion processing can, in principle, include learning at local directional cells, including directional V1 simple and complex cells, area MT neurons (Majaj, Carandini, & Movshon, 2007), and possibly neurons in other, higher, stages of motion processing. Thus, without integrative theoretical constraints, it is unclear whether the local directional learning observed by Watanabe et al. (2002) resulted from a later stage of motion processing, such as in area MT, changes in the readout from sensory area MT to

decision area LIP (Law & Gold, 2008), or whether it occurred at an earlier stage of processing, such as in directional V1 cells.

In the present study, we attempt to disambiguate the involvement of potential sites in TIPL by testing whether TIPL is specific to the contrast polarity of the task-irrelevant motion stimulus. Coding of local motion specific to a contrast polarity is thought to occur at the short-range motion filtering stage (Anstis & Mather, 1985) of directional V1 simple cells, before pooling of signals between contrast polarity occurs as part of long-range motion filtering process, which feeds into either area MT or beyond, as predicted by the 3D FORMOTION model (Berzhanskaya, Grossberg, & Mingolla, 2007; Chey, Grossberg, & Mingolla, 1997; Grossberg, Mingolla, & Viswanathan, 2001; Grossberg & Pilly, 2008). Thus, if learning is found to not transfer to the opposite contrast polarity, then it is suggestive that TIPL for motion may take place at an early stage of motion processing in area V1.

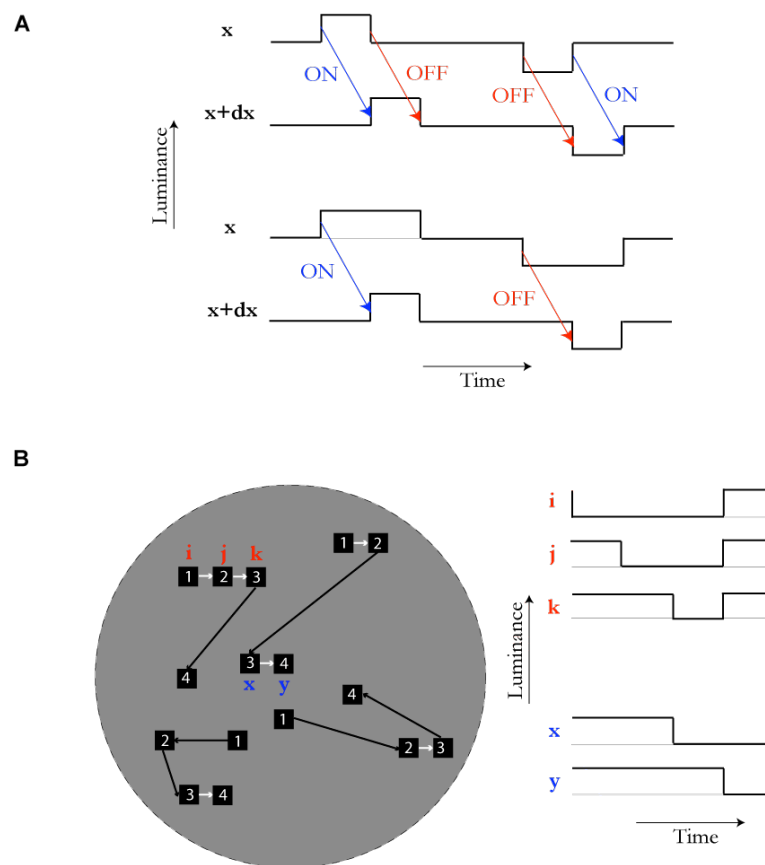


Figure 2. Contrast polarity-specific motion stimuli. (A) A typical apparent motion stimulus, which comprises two brief flashes that occur in different spatial locations at different times, activates both ON- and OFF-cells in early motion processing. However, if the two spatially separate flashes are designed to have the same offset times, then motion cells specific to a single contrast polarity are selectively activated (Wehrhahn & Rapf, 1992). (B) An illustration of a multi-frame variable coherence RDM stimulus that targets OFF motion cells. Here four dots, whose luminance is lower than that of the background, are displayed in each frame with the signal direction being rightward (indicated by white arrow) at 50% coherence, and sample trajectories of the dots are shown for four frames. The number on the dot signifies the frame in which it appears. Temporal luminance profiles are shown at few signal dot locations to illustrate

how our modification for RDM stimuli works. Signal dot lifetime is limited to two frames whenever possible to prevent confounding motion tails in response to high coherences.

2. Experiment 1

In the primate motion processing pathway, neuronal selectivity to motion direction first appears in area V1. Physiological and psychophysical studies have together provided strong evidence for the segregation of early motion processing cells into the ON and OFF channels (see Figure 1). Directional V1 simple cells are sensitive to the contrast polarity of the motion stimuli; however, by the stage of stage of directional V1 complex cells, the ON and OFF channels are combined, and these cells and those at later stages of motion processing have responses that are invariant to the contrast polarity of motion stimuli (Goodwin & Henry, 1975; Movshon & Newsome, 1996; Pack et al., 2006). Thus evidence that TIPL is specific to the contrast polarity of motion stimuli would be consistent with learning that may occur at the stage of directional V1 simple cells, and evidence that TIPL transfers between contrast polarities would be consistent with learning occurring at a later stage of motion processing.

To test this we designed a polarity-specific RDM stimulus based upon a technique by Wehrhahn and Rapf (1992) that was devised to selectively activate ON or OFF cells by ensuring that the two spatially offset flashes that constitute an apparent motion stimulus have different onset but simultaneous offset times (see Figure 2a). Using this approach, we created polarity-specific multi-frame RDM stimuli of variable coherence. Depending on whether the stimulus is of the ON or OFF type, the dots are shown at a luminance level of either 108 cd/m^2 or $\sim 0 \text{ cd/m}^2$ on a gray background (54 cd/m^2). In our RDM algorithm, whenever a dot from a given frame is chosen to move in the signal direction, it persists at its current location in the next frame too. This persistence lasts until the dot is chosen to become noise; that is, relocated to a random position (see Figure 2b). Given that high coherences promote longer signal dot lifetimes, the above modification may create strong motion tails, which can encourage signal dot tracking and confound bottom-up processes involved in direction discrimination. To reduce this effect, the lifetime of signal dots is limited to two frames whenever possible. For coherences used in our study ($\leq 50\%$), this can always be ensured.




Figure 3. Schematic of a testing session trial. In this example, a positive polarity-specific RDM stimulus at some coherence in some direction is shown for 400 ms. Following a delay period of 500 ms, the perceived direction can be chosen within 4000 ms from four alternatives that are displayed on the screen. Note that the aperture boundary and arrows in the figure are only for illustrative purposes and are not actually shown.

The TIPL procedure, based on Seitz and Watanabe (2003), was as follows. We first assigned each of our human subjects ($n = 7$) to a designated contrast polarity and motion direction (from

one of four non-cardinal directions; 22.5° , 112.5° , 202.5° , 292.5°). The experiment comprised five RSVP sessions and two testing sessions, one before and the other after the exposure stage. The testing sessions measured the subjects' ability to discriminate the direction of random dot motion stimuli shown for 400 ms in various coherence (3%, 5%, 10%, 15%, 25%), signal direction, and contrast polarity conditions, allowing the assessment of TIPL (see Figure 3). In each RSVP session, subjects were repetitively exposed to their designated contrast polarity-specific RDM stimuli at 10% coherence in the periphery while they were required to perform an attentionally demanding RSVP task in the fovea. The targets in each RSVP task trial were temporally paired with dots moving in the designated direction; the distracters co-occurred with those in other directions (see Figure 4). The detailed description of the experimental methods is provided in Section 2.1 below.

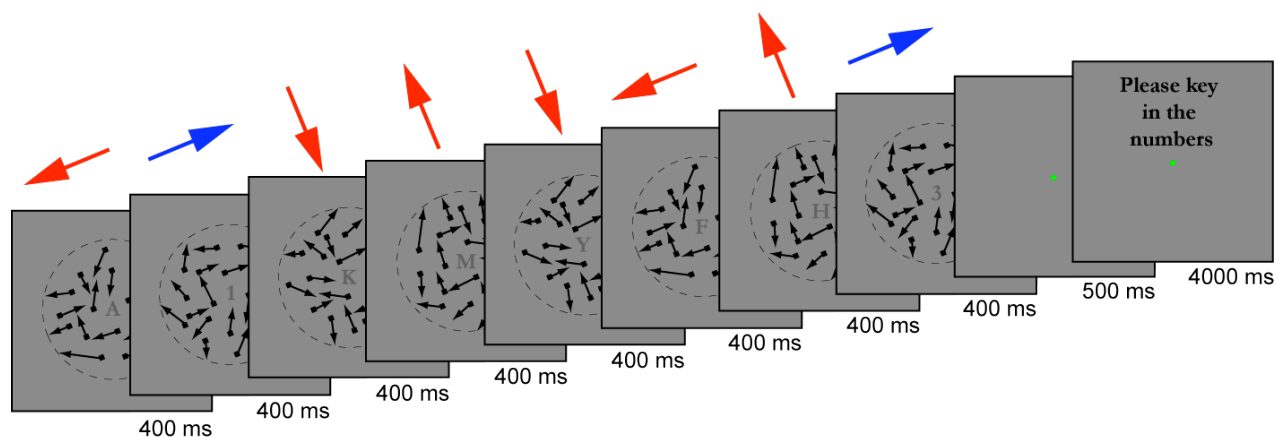


Figure 4. Schematic of an RSVP session trial in Experiment 1. Subjects are instructed to focus on a serially presented sequence of characters that consists of six alphabets and two numbers in random order. Their task is to identify the two numbers, while ignoring the task-irrelevant polarity-specific (OFF in this case) RDM stimuli in the periphery. For each subject, the two number presentations coincide with dots that have net motion in a given direction (22.5° in this case); the six alphabets were randomly paired with motion in other directions equally. RDM stimuli were shown for 400 ms each without an interstimulus interval, and RSVP characters were shown for 275 ms each with an intercharacter interval of 125 ms. A delay period of 500 ms succeeds the sequential presentation of these stimuli, after which the two numbers can be entered within 4000 ms in the order in which they were observed during the trial. Note that the various aperture boundaries and arrows in the figure are only for illustrative purposes and are not actually shown.

2.1. Methods

2.1.1. Subjects

Seven human subjects (18-30 years; four female, three male) were recruited from amongst students of Boston University. They were naïve to the purpose of the experiment and gave written informed consent to their participation for which they received compensation. The methods of the study were approved by the Institutional Review Board of Boston University Charles River Campus.

2.1.2. Apparatus

All stimuli were viewed binocularly at a distance of 60 cm on a 36 cm horizontally wide Dell M992 monitor that was set to a resolution of 1024 x 768 and a refresh rate of 85 Hz. A chin-rest was used to stabilize the head, and the monitor center was ensured to be approximately in the same horizontal plane as that of the subjects' eyes. The experiment was controlled using Psychtoolbox Version 2 (Brainard, 1997; Pelli, 1997) from MATLAB 5.2.1 (The MathWorks, Inc.) on a Macintosh G4 machine running OS 9.

2.1.3. Rapid serial visual presentation (RSVP) task

The RSVP task required subjects to focus on a temporally changing sequence of characters consisting of two numbers and six alphabets in random order and to report the two numbers (or targets) at the end of each trial. The two numbers/targets were randomly chosen without replacement from the set: ('1', '2', '3', '4') and the six alphabets/distracters from the set: ('A', 'E', 'F', 'H', 'K', 'M', 'N', 'P', 'V', 'W', 'X', 'Y'). The possible alphabets had been screened for confusability with any digit (0-9). Each character at the screen center subtended an angle of $\sim 0.8^\circ$.

In a typical trial, the eight characters were shown successively in the center of the screen with each character displayed for 275 ms and followed by a blank period of 125 ms. Then following a delay period of 500 ms during which a green 0.2° fixation point appeared in the center, the subject had 4000 ms within which to press the two numbers, observed during the sequence presentation, on a keyboard. The trial was considered to be correct only if the two numbers were entered in the order in which they appeared. The next trial began after an interval of 400 ms.

2.1.4. Direction discrimination task

In the direction discrimination task, subjects on each trial were asked to choose which of four alternative arrows matched the coherent direction of an RDM stimulus. In each trial, random dot motion (RDM) in some signal direction at some coherence level and specific to a contrast polarity (either ON or OFF) was shown at $12^\circ/\text{s}$ signal dot speed in an invisible 12° diameter aperture centered on the screen. Each dot at the screen center subtended an angle of $\sim 0.1^\circ$. Dot density was fixed at $16.7 \text{ dots deg}^{-2} \text{ s}^{-1}$. ON (OFF) dots were shown at a luminance level of 108 cd/m^2 ($\sim 0 \text{ cd/m}^2$) on a gray background (54 cd/m^2). Given the dot speed and monitor refresh rate values, the spatial and temporal displacements between consecutive signal dot flashes in the stimuli were 0.14° and 11.76 ms, which are both below the classical D_{max} (0.25°) and T_{max} (100 ms) values from Braddick's initial studies (Braddick, 1974) for the directional short-range process. Note, however, that several later studies have shown how these spatial and temporal limits are fluid given different stimulus parameters and task conditions; see Pilly and Seitz (2009).

In a typical trial, the dots were shown around a 0.2° green fixation point for 400 ms followed by a delay period of 500 ms. Then the directional response was recorded by clicking within 4000 ms near the appropriate one of four directional bars that appeared on the screen projecting from the fixation point (see Figure 3). The next trial began after an interval of 400 ms. In order to minimize top-down contributions to performance, subjects were specifically instructed to not engage in individual dot tracking, and instead make the directional decision while fixating the point in the screen center.

2.1.5. Procedure

The procedure was based on that used in Seitz and Watanabe (2003). The experiment comprised eight one-hour sessions that were conducted each on a different day. The first session was a practice session to familiarize subjects with the tasks and stimuli. The second and eighth sessions used the direction discrimination task to measure performance on the motion stimuli. The third to seventh sessions were training sessions employing the RSVP task. The session days were as consecutive as they could be scheduled. All sessions were conducted in a dark room.

2.1.5.1. Practice session (Day 1)

The first session was primarily for the subjects to acclimate to the experimental conditions. It comprised of two practice tasks: an RSVP task and a direction discrimination task.

In the practice RSVP task, the luminance of the characters was randomly varied from trial to trial. Five luminance levels with respect to the gray background (54 cd/m^2), ranging from 47.4 cd/m^2 to 52.4 cd/m^2 , were tested at 10 trials per level. A second purpose of this task was to find a level of task difficulty in the RSVP task that would ensure that subjects maintain fixation and attend to the character stimuli. Based on the obtained data, 51 cd/m^2 was chosen as the luminance level for the RSVP task characters, at which initial performance is neither too low nor too high ($\sim 75\%$), in the RSVP sessions.

In the practice direction discrimination task, the RDM stimulus direction and coherence were randomly varied from trial to trial within two randomly ordered blocks corresponding to the two polarities. Four non-cardinal directions (22.5° , 112.5° , 202.5° , 292.5°), five coherence levels (5%, 10%, 15%, 25%, 50%) and the two contrast polarities were used at five trials per condition. Only in this practice task, visual and auditory feedback was given for each trial as to whether the response was right or wrong. A second purpose of this task was to roughly determine the coherence level from the resulting overall psychometric functions where the performance was neither too weak nor too strong. Based on the obtained data, 10% coherence level was chosen to be used for the task-irrelevant motion stimuli during the RSVP sessions; see Section 2.1.5.3 below. This choice was based on the finding that TIPL is best elicited from parathreshold task-irrelevant stimuli (Tsushima et al., 2008).

2.1.5.2. Testing session 1 (Day 2)

This session was used to measure the baseline discrimination performances in response to RDM stimuli corresponding to the four directions and the two contrast polarities, so that the effect of TIPL could be assessed later. A direction discrimination task, similar to the one in the practice session, was employed. The four directions (22.5° , 112.5° , 202.5° , 292.5°), five coherence levels (3%, 5%, 10%, 15%, 25%) and the two contrast polarities were tested at 30 trials per condition. No response feedback was given. The 1200 trials were arranged into 20 blocks. Half of the blocks were randomly assigned to one contrast polarity and the other to the opposite polarity. Within each block, conditions corresponding to various directions and coherences were randomly interleaved. The idea behind the piecewise block design was to reduce any non-stationarities in motion processing that may occur due to arbitrary changes in contrast polarity of the motion signals from trial to trial. Subjects had an opportunity to take a short rest after the completion of every 5 blocks.

2.1.5.3. RSVP sessions (Days 3-7)

The purpose of these sessions was to repeatedly expose the subjects to parathreshold (10% coherence) motion stimuli while their attention is focused elsewhere, on the RSVP task. Of 7 subjects, 3 were randomly selected to be exposed to ON random dot motion stimuli and the other 4 to OFF stimuli. Also, each subject was randomly assigned a (designated) direction. Each RSVP session comprised 440 RSVP task trials. Unlike in the practice session, the RSVP characters (51

cd/m²) in each trial co-occurred temporally with designated contrast polarity-specific motion stimuli (either ~0 or 108 cd/m²), each of which was presented for 400 ms around the character, on the gray background (54 cd/m²). The targets (numbers) co-occurred with motion stimuli in the designated direction, and the distracters (alphabets) co-occurred with motion stimuli in other directions on an equal random basis. The temporal midpoints of the durations for which an RSVP character and its corresponding motion stimulus were presented, respectively, coincided. Other parameters of the RDM stimuli, namely the size of the invisible aperture, signal dot speed, dot size, dot density, and monitor refresh rate, and the dot motion generation algorithm were the same as those used in the testing sessions. The two numbers appeared one each in the two halves of the sequence in order to avoid trials in which both numbers are presented early on, which causes subsequent leakage of attention to the motion stimuli. Subjects were specifically instructed to ignore the motion stimuli surrounding the RSVP characters as they were irrelevant to the RSVP task at hand. They had an opportunity to take a short rest after the completion of every 110 trials.

2.1.5.4. Testing session 2 (Day 8)

The second testing session was identical to the first testing session. This session was used to measure learning that may have occurred in task-irrelevant processing of the polarity-specific motion stimuli.

	Right-tailed paired t-test
(d,e) > (d,u)	0.024
(d,e) > (D,e)	0.024
(d,e) > (D,u)	0.000
(d,u) > (D,u)	0.851
(d,u) > (D,e)	0.905
(D,e) > (D,u)	0.491
(d,e) > 0	0.000
(D,e) > 0	0.011
(d,u) > 0	0.196
(D,u) > 0	0.007

Table 1. Detailed statistical test results of Experiment 1. The legend used in the first column is as follows: [*d*: designated direction; *D*: control directions; *e*: designated polarity; *u*: opposite polarity]. Each entry in the second column specifies the p-value of the right-tailed paired t-test performed on the corresponding relation in the first column. Note that the p-values have been rounded off to three decimal places. The highlighted rows identify the relations that show statistical significance ($p < 0.05$).

2.1.6. Data analysis

A 2-way repeated measures ANOVA was conducted with Direction x Polarity as factors to quantify the obtained learning effects. Right-tailed paired t-tests were conducted on various relations between changes in performance found in conditions of interest; see Table 1.

2.2. Results

As one would expect, the RSVP task performance gradually increased through the five RSVP sessions (see Figure 5). These data show that subjects were generally engaged in the RSVP task and that they underwent task-relevant learning.

In regard to learning for the task-irrelevant motion stimuli, we found that TIPL was specific to the designated contrast polarity (see Figure 6). This was reflected in a significant direction-specific learning effect for the designated direction at the designated contrast polarity ($p < 0.05$, right-tailed paired t-test), but there was no difference in learning between the designated and control directions at the opposite polarity ($p = 0.298$, two-tailed paired t-test). A repeated measures ANOVA found a significant effect of polarity ($p < 0.05$) and a trend for an interaction between polarity and direction ($p = 0.122$). Also for the designated direction, the average improvement in performance was significantly greater at the designated polarity in comparison to the opposite polarity ($p < 0.05$, right-tailed paired t-test) and while highly significant learning was observed at the designated polarity ($p < 0.001$, right-tailed paired t-test), there was no significant learning at the opposite polarity ($p = 0.196$, right-tailed paired t-test). For the control directions, there was significant learning at both contrast polarity conditions ($p = 0.011$ and $p = 0.007$, right-tailed paired t-test vs. 0%, for designated and control directions, respectively) and no difference between them ($p = 0.982$, two-tailed paired t-test). This learning for the opposite polarity can possibly be explained by baseline shifts, that is, consolidation of learning after the first testing session and online learning during the second testing session; see General Discussion for a theoretical explanation. Detailed statistical test results are provided in Table 1. In sum, we found that TIPL was specific both to the direction and the contrast polarity of the motion stimuli that were paired with the targets of the RSVP task.

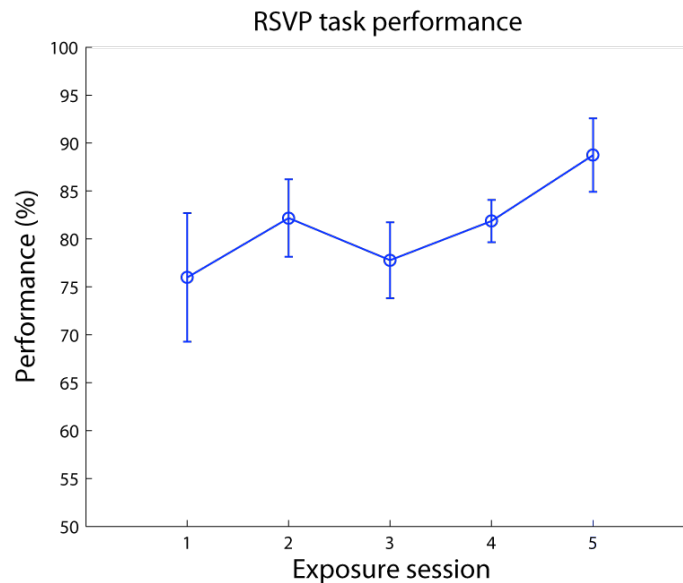


Figure 5. In Experiment 1, RSVP task performance gradually increased through the five RSVP sessions. Error bars represent standard error of mean.

2.3. Discussion

Our results confirm that task-irrelevant perceptual learning is specific to the contrast polarity of the designated parathreshold motion stimuli. These results are difficult to explain by changes either in later stages of sensory processing, because motion-selective neurons higher in the processing hierarchy than directional V1 simple cells are contrast polarity-invariant, or in

decision stages, because the decisions for the training (RSVP) task were dissociated from the decisions required during testing. Moreover, a human imaging study found that lateral prefrontal cortex, a decision-gating area, fails to respond to weak motion stimuli, such as those used during training in the present study (Tsushima et al., 2006). Thus, the results suggest the intriguing conclusion that plasticity for TIPL may occur in directional V1 simple cells.

These data support the prediction of the 3D FORMOTION model that short-range directional filters at directional V1 simple cells are sensitive to motion direction and contrast polarity, before projecting to area MT or beyond via the long-range directional filtering process, which pools between opposite contrast polarities and begins the process of global motion summation and capture (Berzhanskaya et al., 2007; Chey et al., 1997; Grossberg et al., 2001; Grossberg & Pilly, 2008), which was not found to adapt during TIPL (Watanabe et al., 2002).




Figure 6. Contrast polarity-specificity of task-irrelevant learning? Experiment 1 results show that TIPL is specific to the designated contrast polarity of motion stimuli. Significantly more learning was found in response to designated direction stimuli at the designated polarity (blue) when compared to the other polarity (red). Data are averaged across coherences and for the control directions are further averaged across the three unpaired directions. Error bars represent standard error of mean. See Table 1 for detailed statistical test results from Experiment 1.

3. Experiment 2

Some recent studies have brought into question the view that featural specificity is a good indicator of early sensory plasticity (Xiao et al., 2008; Zhang et al., 2009; Doshier & Lu 1998) and explored alternative hypotheses to explain featural specificity. Any perceptual task involves the processing of a perceptual stimulus, which can implicitly encode many individual features such as contrast, retinal location, and orientation. Mollon and Danilova (1996) proposed that specificity to the trained features may simply be due to a central site in the brain learning the “idiosyncrasies” related to the restricted stimulus space used in the training task. Within this framework, one would expect that if the stimulus space is expanded, either using the same task or different tasks, then a broader extent of learning transfer would be found.

Along these lines, an influential study by Xiao et al. (2008) found that “double training” allowed a surprising amount of transfer across feature dimensions. Consistent with previous

studies, they first showed that practicing contrast discrimination of a vertically-oriented Gabor did not transfer to a new retinal location. However, when subjects further practiced at the new location an unrelated task with or without a new stimulus, or the same task with a new stimulus, such as orientation discrimination of a horizontally-oriented Gabor, then there was complete transfer of contrast learning at vertical orientation to the new location. The results of this “double training” are difficult to explain in terms of plasticity in retinotopic early visual areas as the basis for featurally-specific perceptual learning. Further, Zhang et al. (2009) showed that the unrelated task in the second training could be very coarse identification of a stimulus comprising the transfer feature, such as discriminating if the presented stimulus is a Gabor exposed at the transfer orientation or the letter C, and concluded that stimulus exposure was a sufficient form of double training to produce transfer of learning.

While in Experiment 1 we found that task-irrelevant learning of direction discrimination is specific to the direction and contrast polarity of unattended motion stimuli spatiotemporally paired with targets of an RSVP task, in light of the task-relevant double training results (Xiao et al., 2008; Zhang et al., 2009), it is possible that our observed lack of transfer to the opposite polarity could just simply be due to the lack of priming that polarity during the RSVP sessions. To address this possibility, which has not previously been tested in the context of TIPL, we performed Experiment 2.

The procedure for Experiment 2 was the same as that used in Experiment 1 except for the RSVP sessions; see Section 3.1 below. In each RSVP session, subjects were inattentively exposed to moving dots at both polarities in alternating blocks of trials. Similar to Experiment 1, each subject was randomly assigned one of the four directions and also one contrast polarity. The trials in half of the blocks were exactly like those in Experiment 1 RSVP sessions where the RSVP task targets were always paired with RDM in the designated direction and all RDMs were presented at the designated polarity. In the other blocks, the RDMs were shown at the opposite polarity and there was no consistent pairing between the RSVP task targets and motion in any particular direction. The double training hypothesis predicts that this simple exposure to the other contrast polarity should ‘unlock’ transfer of TIPL and thus that the direction-specific learning, observed in Experiment 1, should transfer between contrast polarities in Experiment 2.




Figure 7. Block schedule in each session of the double training procedure employed in Experiment 2. Subjects perform two kinds of trials in alternating blocks in each RSVP session. Trials in the TIPL blocks are similar to those used during the RSVP sessions in Experiment 1. In these trials (shown by blue arrow), the RSVP task targets are temporally paired with random dot motion that is 10% coherent in a designated direction and specific to a designated contrast polarity (negative polarity in this example), whereas the RSVP task distracters temporally co-occur with similar polarity-specific motion in other three directions chosen randomly and equally; see Figure 4 for an example trial of this kind. Trials in the exposure blocks (shown by red arrow) are similar, except that the RSVP task characters are paired with motion specific to

the opposite contrast polarity (positive polarity in this example) and there is no consistent temporal pairing between RSVP task targets and motion in the designated direction.

3.1. Methods

3.1.1. Subjects

Six new human subjects (18-24 years; four female, two male) were recruited from amongst students of Boston University. They were naïve to the purpose of the experiment and gave written informed consent to their participation for which they received compensation. The methods of the study were approved by the Institutional Review Board of Boston University Charles River Campus.




Figure 8. In Experiment 2, RSVP task performance gradually increased through the first five RSVP sessions and then again from sixth to tenth sessions. The contrast between RSVP task characters and the gray background was reduced for the second half of sessions to ensure performance does not reach ceiling, which is important to ensure significant attentional resources are not diverted to the task-irrelevant motion stimuli. Error bars represent standard error of mean.

3.1.2. Procedure

The apparatus, RSVP task, direction discrimination task, practice session, and testing sessions were exactly the same as those in Experiment 1. The only procedural differences were in the RSVP sessions. The 440 RSVP task trials in each RSVP session were divided into eight blocks. The subjects were allowed to take a short rest after the completion of every two blocks. Of six subjects, ON polarity was randomly designated to one half and OFF polarity to the other half. The eight blocks alternatively exposed the subjects to motion selective of either polarity. The order of their presentation was randomly determined on a per-session basis. The blocks of trials corresponding to the designated polarity (TIPL blocks) were the same as the trials during the RSVP sessions in Experiment 1. However in the other blocks of trials (exposure blocks), the possible direction set was randomized on each trial for pairing with the RSVP task characters (see Figure 7), thereby eliminating the advantage of consistent pairing with task targets for irrelevant motion in the designated direction. In order to replicate the total number of trials in Experiment 1 for the designated polarity, the training stage in Experiment 2 consisted of 10 (third to twelfth) sessions. Also, the luminance of the RSVP task characters was increased to

51.6 cd/m² starting in the sixth RSVP session as a precaution against subjects reaching a performance ceiling, which may facilitate some attentional leak to the task-irrelevant stimuli. In any case, it should be noted that the subjects were clearly instructed to completely ignore the RDM stimuli.

3.1.3. Data analysis

A 2-way repeated measures ANOVA was conducted with Direction x Polarity as factors to quantify the obtained learning effects. Right-tailed paired t-tests were conducted on various relations between changes in performance found in conditions of interest; see Table 2.

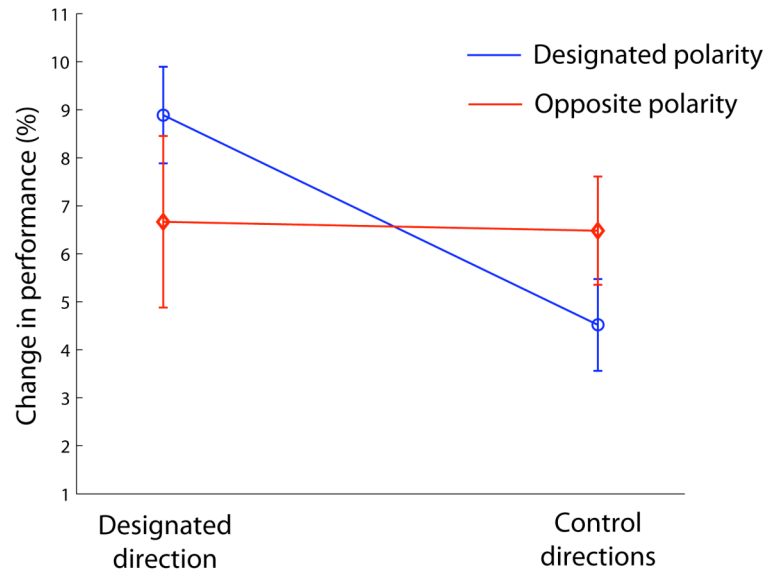


Figure 9. Results from the double training version of the task-irrelevant paradigm for motion stimuli. Data indicates that TIPL is indeed specific to the contrast polarity of motion stimuli. Directionally selective learning occurred only for the designated polarity and not for the opposite polarity, denoting that the polarity-specificity result of Experiment 1 passed the test of double training. Data are averaged across coherences and for the control directions are further averaged across the three control directions. Error bars represent standard error of mean. See Table 2 for detailed statistical test results from Experiment 3.

3.2. Results

As in Experiment 1, the subjects improved their RSVP task performance from the first to fifth RSVP session and then from the sixth to tenth RSVP session, when the contrast between the RSVP task characters and the grey background was slightly reduced (see Figure 8).

In regard to TIPL, for the designated polarity, significantly more learning occurred for the designated direction than for the control directions ($p < 0.05$, right-tailed paired t-test), but no direction-specific learning was found for the opposite polarity ($p = 0.951$, two-tailed t-test). This result replicates the polarity-specific directional learning observed in Experiment 1 and shows that double training failed to elicit complete transfer to the opposite contrast polarity. However, in Experiment 2, significant learning was found in all polarity and direction conditions ($p < 0.05$, right-tailed paired t-tests); see Figure 9, as well as Table 2 for detailed statistical results from Experiment 2. Moreover, the overall learning for the designated (6.704 +/- 0.733%) and opposite (6.574 +/- 0.733%) polarities were about the same, which reflected in the lack of effect of

polarity in the repeated measures ANOVA ($p = 0.939$), although a trend was still found for an interaction between polarity and direction ($p = 0.102$). In sum, we found more overall learning in Experiment 2, than in Experiment 1, but the component of learning from TIPL that was direction-specific did not transfer to the opposite contrast polarity.

	Right-tailed paired t-test
$(d,e) > (d,u)$	0.205
$(d,e) > (D,e)$	0.024
$(d,e) > (D,u)$	0.151
$(d,u) > (D,u)$	0.475
$(d,u) > (D,e)$	0.250
$(D,e) > (D,u)$	0.930
$(d,e) > 0$	0.017
$(D,e) > 0$	0.046
$(d,u) > 0$	0.041
$(D,u) > 0$	0.001

Table 2. Detailed statistical results of Experiment 2. Refer to caption of Table 1 for legend information. As in Table 1, the second column shows the p-value of the corresponding t-test.

3.3. Discussion

The results from Experiment 2 do not support the hypothesis that TIPL is mediated by a high-level central brain site. If it were true, then the *direction learning* from TIPL blocks for the designated polarity and the *opposite polarity learning* from the exposure blocks should have enabled complete transfer of directional learning to the opposite polarity. The lack of this transfer strengthens the conclusion that TIPL for motion stimuli could be based on changes in the brain as early in the hierarchy as directional V1 simple cells.

The main difference between the results of Experiment 1 and those of Experiment 2 is the large baseline learning effect found in Experiment 2. A possible explanation for this is that the mere exposure of stimuli of the opposite polarity during training allowed for a partial transfer of TIPL, consistent with the hypothesis of double training. However, an alternative explanation for this partial transfer is that learning occurred due to accidental pairings with RSVP task targets that took place for opposite polarity-specific motion in each of the four directions. This alternative explanation would explain why equal learning was found for the designated and control directions at the opposite contrast polarity and is consistent with previous findings of TIPL (cf., Nishina et al., 2007).

4. General Discussion

Our results show for the first time that TIPL is specific to the contrast polarity of motion stimuli. Furthermore, the specificity of TIPL to the designated direction of motion did not transfer between contrast polarities when a double training paradigm was employed in which the other contrast polarity was exposed an equal number of times as the designated polarity. Thus these results show for the first time an example of specificity of perceptual learning that holds up to double training.

This finding of specificity of TIPL to contrast polarity of motion stimuli is consistent with previous TIPL studies that showed specific learning to other low-level visual features such as retinal location and local motion direction (Watanabe et al., 2002), orientation and retinal location (Nishina et al., 2007), and eye of exposure and orientation (Seitz et al., 2009). Furthermore, a recent electrophysiological study of TIPL in awake, behaving macaque monkeys showed that TIPL results in plasticity at or before the level of visual area V4 (Franko, Seitz, & Vogels, in press).




Figure 10. Illustration of the adaptive motion processing hierarchy in the 3D FORMOTION model in response to our experiments. The red asterisk indicates the location of primary plasticity in the model at the short-range motion filter stage, which corresponds to directional V1 simple cells in Figure 1. Directional cells above the level of short-range motion filter stage are depicted by ‘dashed’ directional ellipses to signify their lack of contrast polarity-sensitivity. This long-range motion filter process, which pools local motion signals over opposite polarities, both eyes, multiple orientations, and a larger spatial region, may occur in a sequence of substages from directional V1 complex cells to the deep laminar layers of area MST. The interconnections between any two consecutive stages are adaptive. Short-range cells coding the two polarities are shown opposing each other. The rounded rectangle encompassing the long-range and global motion cells is used to indicate the potential modulatory influence of top-down spatial attention.

We suggest that our results arise at least partially from plasticity in directional V1 simple cells, which correspond to short-range motion filter cells in 3D FORMOTION model. Notably, the 3D FORMOTION model provides a parsimonious explanation of secondary aspects of the data; see Figure 10. In Experiment 1, the designated direction at the opposite polarity showed the lowest performance change and was the only condition that showed insignificant learning; see Figure 6 and Table 1. We posit this occurs due to opponent interactions between the polarity-specific short-range cells (Baloch et al., 1999). As short-range cells tuned to the designated direction

become active, they hyperpolarize the corresponding opposite polarity cells. As a result, anti-learning occurs in the connections between the opposite polarity cells and the long-range filter cells that pool activities from these two kinds of short-range cells, leading to a relatively lower learning for the designated direction at the opposite polarity. In Experiment 2, we observed the lowest degree of learning for the control directions at the designated polarity; see Figure 9. This is also well explained by the 3D FORMOTION model. As discussed in Section 3.3, learning for the opposite polarity may have occurred due to accidental pairings of each of the directions with the targets of the RSVP task. In the same manner that pairings in Experiment 1 resulted in negative learning for the designated direction at the opposite polarity, the accidental pairings for the control directions at the opposite polarity may have caused anti-learning for the control directions at the designated polarity, explaining why the performance change was lowest for this condition.

Our findings illustrate an important difference between task-relevant and task-irrelevant perceptual learning paradigms with respect to the effect of the double training procedure on the transfer of learning to untrained variants of features. Comparison of the current results with those from task-relevant paradigms, namely Xiao et al. (2008) and Zhang et al. (2009), points out that the presence of attention may be an important factor in the recruitment of higher-level central brain sites for perceptual learning. The fact that task-irrelevant learning bypasses some effects of the attentional system may explain why double training has different impact on task-irrelevant than for task-relevant perceptual learning.

Some aspects of how our findings generalize to task-relevant learning paradigms may also be clarified by properties of the 3D FORMOTION model. Notably, previous research using different methods found that task-relevant learning can transfer between contrast polarities (Wehrhahn & Rapf, 2001). Why do not these data conflict with our present explanation of non-transfer between contrast polarities? When subjects attend to the motion stimulus, this may be driving activity at higher stages of processing with polarity-invariant receptive fields (see Figure 10) such that primary plasticity occurs there, accounting for the transfer to the unexposed polarity. A pertinent question for future research is why an additional component of low-level learning was not found by Wehrhahn & Rapf (2001).

It should be noted that, whereas Grieco, Casco, and Roncato (2006) showed texture segregation learning specific to contrast polarity, their learning procedure does not preclude the attribution of underlying neural plasticity to the readout from V1 simple cells to decision stages and their results have not yet been shown to pass the specificity-test of double training. In all, our results suggest that TIPL may differ from task-relevant learning in its degree of featural specificity and that TIPL can lead to a greater proportion of plasticity at early stages of sensory processing.

5. Conclusion

The main contribution of the article is that it makes a case for plasticity in early visual cortex as the basis for task-irrelevant perceptual learning (TIPL) for motion stimuli. While individual behavioral studies of perceptual learning are at best ambiguous regarding the neural locus of the learning effect, we argue that the high-level of specificity to low-level visual features found in this and other studies of TIPL, the finding in this study that specificity to the contrast polarity of motion stimuli occurs in both conventional and double training procedures, and the recent electrophysiological finding that TIPL changes responses of visual cortical local field potentials (Franko et al., in press) are all consistent with the involvement of early visual cortex in TIPL.

Thus as a whole, it seems that TIPL is a procedure that is capable of inducing learning at the earliest levels of visual cortical processing.

Acknowledgements

P.K.P. and S.G. were supported in part by CELEST, an NSF Science of Learning Center (SBE-0354378) and by the DARPA SyNAPSE program (HR0011-09-3-0001, HR001-09-C-0011). A.R.S. was supported in part by NSF (BCS-0549036) and NIH (R21 EY017737).

References

- Alonso, J.M., Usrey, W.M., & Reid, R.C. (2001). Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. *Journal of Neuroscience*, *21*, 4002-4015.
- Ahissar, M., & Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature*, *387*, 401-406.
- Anstis, S.M., & Mather, G. (1985). Effects of luminance and contrast on direction of ambiguous apparent motion. *Perception*, *14*, 167-179.
- Baloch, A.A., Grossberg, S., Mingolla, E., & Nogueira, C.A.M. (1999). A neural model of first-order and second-order motion perception and magnocellular dynamics. *Journal of the Optical Society of America A*, *16*, 953-978.
- Berzhanskaya, J., Grossberg, S., & Mingolla, E. (2007). Laminar cortical dynamics of visual form and motion interactions during coherent object motion perception. *Spatial Vision*, *20*, 337-395.
- Braddick, O. (1974). A short-range process in apparent motion. *Vision Research*, *14*, 519-527.
- Brainard, D.H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433-436.
- Chey, J., Grossberg, S., & Mingolla, M. (1997). Neural dynamics of motion grouping: From aperture ambiguity to object speed and direction. *Journal of the Optical Society of America A*, *14*, 2570-2594.
- Conway, B.R., & Livingstone, M.S. (2003). Space-time maps and two-bar interactions of different classes of direction-selective cells in macaque V1. *Journal of Neurophysiology*, *89*, 2726-2742.
- Croner, L.J., & Albright, T.D. (1997). Image segmentation enhances discrimination of motion in visual noise. *Vision Research*, *37*, 1415-1427.
- 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 USA*, *95*, 13988-13993.
- Edwards, M., & Badcock, D.R. (1994). Global motion perception: interaction of the ON and OFF pathways. *Vision Research*, *34*, 2849-2858.
- Franko, Seitz, & Vogels. (2009). Dissociable neural effects of long term stimulus-reward pairing in macaque visual cortex. *Journal of Cognitive Neuroscience*, in press.
- Furmanski, C.S., Schluppeck, D., & Engel, S.A. (2004). Learning strengthens the response of primary visual cortex to simple patterns. *Current Biology*, *14*, 573-578.
- Grieco, A., Casco, C., & Roncato, S. (2006). Texture segregation on the basis of contrast polarity of odd-symmetric filters. *Vision Research*, *46*, 3526-3536.
- Grossberg, S., Mingolla, E., & Viswanathan, L. (2001). Neural dynamics of motion integration and segmentation within and across apertures. *Vision Research*, *41*, 2521-2553.
- Grossberg, S., & Pilly, P.K. (2008). Temporal dynamics of decision-making during motion perception in the visual cortex. *Vision Research*, *48*, 1345-1373.

- Goodwin, A.W., & Henry, G.H. (1975). Direction selectivity of complex cells in a comparison with simple cells. *Journal of Neurophysiology*, 38, 1524-1540.
- Jeter, P.E., Doshier, B.A., Petrov, A., & Lu, Z.-L. (2009). Task precision at transfer determines specificity of perceptual learning. *Journal of Vision*, 9(3), 1-13.
- Law, C.T., & Gold, J.I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience*, 11, 505-513.
- Liu, Z.L., & Weinshall, D. (2000). Mechanisms of generalization in perceptual learning. *Vision Research*, 40, 97-109.
- Livingstone, M.S. (1998). Mechanisms of direction selectivity in macaque V1. *Neuron*, 20, 509-526.
- Majaj, N.J., Carandini, M., & Movshon, J.A. (2007). Motion integration by neurons in macaque MT is local, not global. *Journal of Neuroscience*, 27, 366-370.
- Mollon, J.D., & Danilova, M.V. (1996). Three remarks on perceptual learning. *Spatial Vision*, 10, 51-58.
- Movshon, J.A., & Newsome, W.T. (1996). Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. *Journal of Neuroscience*, 16, 7733-7741.
- Nishina, S., Seitz, A.R., Kawato, M., & Watanabe, T. (2007). Effect of spatial distance to the task stimulus on task-irrelevant perceptual learning of static gabors. *Journal of Vision*, 7, 1-10.
- Pack, C.C., Conway, B.R., Born, R.T., & Livingstone, M.S. (2006). Spatiotemporal structure of nonlinear subunits in macaque visual cortex. *Journal of Neuroscience*, 26, 893-907.
- Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10, 437-442.
- Pilly, P.K., & Seitz, A.R. (2009). What a difference a parameter makes: A psychophysical comparison of random dot motion algorithms. *Vision Research*, 49, 1599-1612.
- Pourtois, G., Rauss, K.S., Vuilleumier, P., & Schwartz, S. (2008). Effects of perceptual learning on primary visual cortex activity in humans. *Vision Research*, 48, 55-62.
- Reid, R.C., & Alonso, J.M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378, 281-284.
- Schiller, P.H. (1982). Central connections of the retinal ON and OFF pathways. *Nature*, 297, 580-583.
- Schiller, P.H. (1992). The ON and OFF channels of the visual system. *Trends in Neurosciences*, 15, 86-92.
- Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, 412, 549-553.
- Seitz, A.R., Kim, D., & Watanabe, T. (2009). Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron*, 61, 700-707.
- Seitz, A.R., Lefebvre, C., Watanabe, T., & Jolicoeur, P. (2005). Requirement for high-level processing in subliminal learning. *Current Biology*, 15, R753-R755.
- Seitz, A.R., & Watanabe, T. (2003). Is subliminal learning really passive? *Nature*, 422, 36.
- Seitz, A.R., & Watanabe, T. (2009). The phenomenon of task-irrelevant perceptual learning. *Vision Research*, in press.
- Smirnakis, S.M., Brewer, A.A., Schmid, M.C., Tolias, A.S., Schuz, A., Augath, M., Inhoffen, W., Wandell, B.A., & Logothetis, N.K. (2005). Lack of long-term cortical reorganization after macaque retinal lesions. *Nature*, 435, 300-307.

- Tsushima, Y., Sasaki, Y., & Watanabe, T. Greater disruption due to failure of inhibitory control on an ambiguous distractor. *Science*, *314*, 1786-1788.
- Tsushima, Y., Seitz, A.R., & Watanabe, T. (2008). Task-irrelevant learning occurs only when the irrelevant feature is weak. *Current Biology*, *18*, R516-R517.
- van der Smagt, M.J., & van de Grind, W.A. (1990). Integration and segregation of local motion signals: the role of contrast polarity. *Vision Research*, *39*, 811-822.
- Watanabe, T., Nanez Sr., J.E., 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., Sasaki, Y., & Nanez, J. (2001). Perceptual learning without perception. *Nature*, *413*, 844-848.
- Wehrhahn, C., & Rapf, D. (1992). ON- and OFF-pathways form separate neural substrates for motion perception: psychophysical evidence. *Journal of Neuroscience*, *12*, 2247-2250.
- Wehrhahn, C., & Rapf, D. (2001). Perceptual learning of apparent motion mediated through ON- and OFF-pathways in human vision. *Vision Research*, *41*, 353-358.
- Xiao, L.Q., Zhang, J.Y., Wang, R., Klein, S.A., Levi, D.M., & Yu, C. (2008). Complete transfer of perceptual learning across retinal locations enabled by double training. *Current Biology*, *18*, 1922-1926.
- Zhang, J.Y., Xiao, L.Q., Klein, S., Levi, D., & Yu, C. (2009). Enabling complete transfer of perceptual learning across orientations in foveal vision through double training. *Vision Sciences Society Abstracts*, 872.