Decision-making processes in perceptual learning depend on effectors

Vladyslav Ivanov[1*^], Giorgio Manenti[1,2*], Sandrin S. Sudmann[1,2,3], Igor Kagan[2,4], Caspar M. Schwiedrzik[1,2,3†]

[1]Neural Circuits and Cognition Lab, European Neuroscience Institute Göttingen – A Joint Initiative of the University Medical Center Göttingen and the Max Planck Society, Grisebachstraße 5, 37077 Göttingen, Germany; [2]Leibniz ScienceCampus Primate Cognition, Göttingen, Germany; [3]Perception and Plasticity Group, German Primate Center – Leibniz Institute for Primate Research, Kellnerweg 4, 37077 Göttingen, Germany; [4]Decision and Awareness Group, German Primate Center – Leibniz Institute for Primate Research, Kellnerweg 4, 37077 Göttingen, Germany

*These authors contributed equally to this work

^Current address: Sensorimotor Group, German Primate Center – Leibniz Institute for Primate Research, Kellnerweg 4, 37077 Göttingen, Germany

[†]To whom correspondence should be addressed: c.schwiedrzik@eni-g.de

Abstract

Visual perceptual learning is traditionally thought to arise in visual cortex. However, typical perceptual learning tasks also involve systematic mapping of visual information onto motor actions. Because the motor system contains both effector-specific and effector-unspecific representations, the question arises whether visual perceptual learning is effector-specific itself, or not. Here, we study this question in an orientation discrimination task. Subjects learn to indicate their choices either with joystick movements or with manual reaches. After training, we challenge them to perform the same task with eye movements. We dissect the decision-making process using the drift diffusion model. We find that learning effects on the rate of evidence accumulation and the decision criteria depend on effectors, albeit not fully. This suggests that during perceptual learning, visual information is mapped onto effector-specific integrators. Overlap of the populations of neurons encoding motor plans for these effectors may explain partial generalization. Taken together, visual perceptual learning is not limited to visual cortex, but also affects sensorimotor mapping at the interface of visual processing and decision making.

Keywords: visual perceptual learning, effector specificity, orientation, sensorimotor mapping

Introduction

A core characteristic of the brain is its ability to learn. Learning is not limited to acquiring new motor skills or a foreign language – visual perception is also highly plastic. Training enables us to improve our visual abilities well into the hyperacuity range (Poggio, Fahle, & Edelman, 1992) and to learn to see what is initially invisible (Schwiedrzik, Singer, & Melloni, 2011). Such training effects are known as visual perceptual learning (VPL). VPL occurs not only during critical periods of development, but also in adult brains. This has profound implications for how we think about brain organization and function, yet little is known about the principles and neural mechanisms underlying adult VPL.

Classical theories suggest that VPL occurs in visual cortex (Ahissar & Hochstein, 2004; Fahle, 2004; Sagi, 2011). This idea rests upon the striking correspondence between the known selectivities of neurons in early visual areas and the selectivity of behavioral learning effects that are often found in VPL studies. For example, VPL has been found to be specific to the orientation, location (Schoups, Vogels, & Orban, 1995), polarity (Grieco, Casco, & Roncato, 2006), and even eye of origin (Karni & Sagi, 1991) of the training material, resembling the tuning properties of V1 neurons. Accordingly, many classical theories argue that VPL increases the selectivity of neurons in early sensory areas.

While VPL-induced increased selectivity of neurons in early visual cortex has indeed been observed through single cell recordings in monkeys (reviewed in Vogels, 2010), results are somewhat inconsistent across studies (Ghose, Yang, & Maunsell, 2002), the observed changes in selectivity are too small to explain the behavioral benefits (Petrov, Dosher, & Lu, 2005), and modelling points to limitations on how much information can be gained by sharpening tuning curves (Series, Latham, & Pouget, 2004). These results suggest that sharpening of neural tuning in visual cortex is only one of the possible mechanisms of VPL.

More recent studies have started to investigate a previously neglected aspect of VPL, sensorimotor mapping. In a typical VPL task, two stimulus alternatives, e.g., clockwise or counterclockwise orientation of lines, have to be mapped onto two specific motor actions, e.g., a leftward or rightward saccade. This implies that visual information needs to be brought together with motor plans that correspond to the trainee's choice. This suggests selective strengthening of connections between visual and motor neurons as another candidate mechanism of VPL. Indeed, a groundbreaking study found that VPL changes the activity of neurons in the lateral intraparietal area (LIP), an area in posterior parietal cortex that is involved in oculomotor decision-making and where sensorimotor mappings can take place (Law & Gold, 2008). This resonates with theories that frame VPL as a change in the readout of visual information by decision-making areas (Dosher & Lu, 2017). A role of motor areas in VPL is further suggested by several behavioral findings, e.g., that motor tasks improve VPL consolidation

(Klorfeld-Auslender & Censor, 2019), and, reversely, that VPL can improve oculomotor performance (Szpiro, Spering, & Carrasco, 2014).

If VPL affects sensory-motor mapping between visual cortex and the motor system, this would suggest another form of specificity in VPL, namely effector specificity. This is because the motor system contains neurons that code motor plans for specific effectors. These are highly prevalent in the primary motor cortex, but also in parietal cortex: many neurons in the so-called parietal reach region (PRR) code for arm movements, while many neurons in LIP are saccade specific (Christopoulos, Bonaiuto, Kagan, & Andersen, 2015; Christopoulos, Kagan, & Andersen, 2018; de Lafuente, Jazayeri, & Shadlen, 2015; Dickinson, Calton, & Snyder, 2003). During VPL, connections between stimulus-specific visual neurons and effector-specific motor neurons could thus be strengthened, resulting in stimulus- and effector-specific behavioral learning effects. However, areas like LIP also encode components that are effector-independent (Chang, Calton, Lawrence, Dickinson, & Snyder, 2016; Christopoulos et al., 2018; Colby & Goldberg, 1999; Gottlieb & Goldberg, 1999; Snyder, Batista, & Andersen, 1997). This raises the possibility that VPL is effector-independent, too, even if sensorimotor mapping is involved.

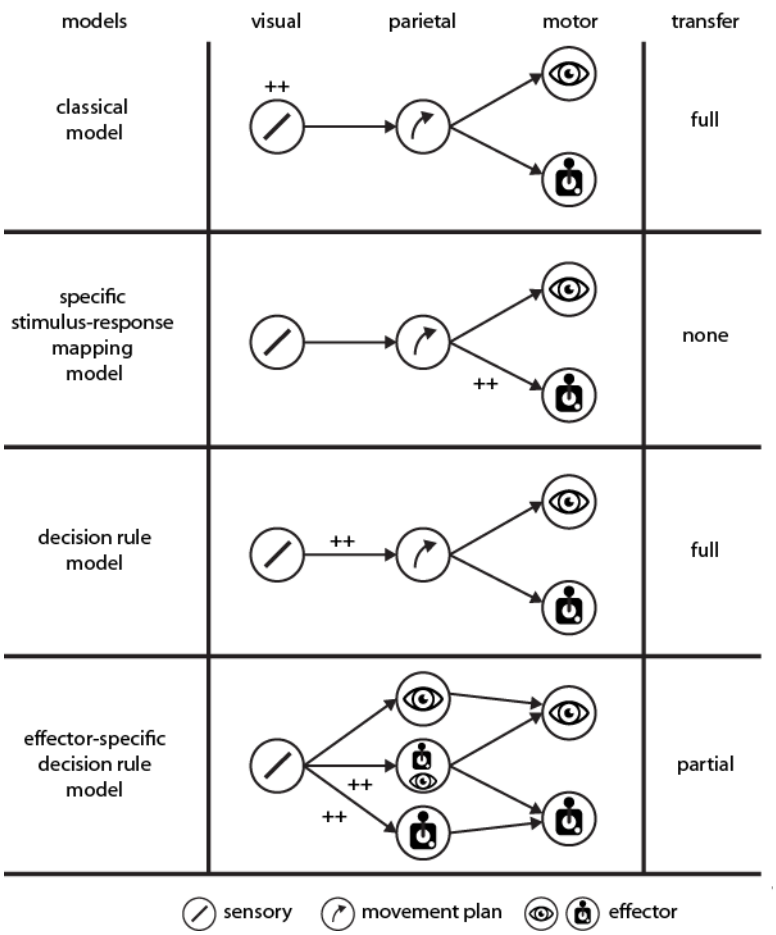

Figure 1: Possible models for effector specificity in visual perceptual learning. *Classical model: learning occurs in visual cortex and results in stimulus- but not effector-specific learning effects. Changing the effector after training results in full transfer of learning effects to the new effector. Specific stimulus-response mapping model: visual information is mapped onto effector-specific motor neurons. This would result in stimulus- and effector-specific learning effects. Decision rule model: stimulus-response mapping takes place between visual neurons and effector-independent neurons involved in decision making and movement planning. This results in stimulus-specific learning effects that generalize across effectors. Effector-specific decision rule model: Sensorimotor mapping occurs between visual neurons and effector-specific integrators. Overlap between effectors at the integration stage results in partial transfer between the trained and untrained effectors.*

A number of studies have started to investigate effector-specificity of VPL, but results are mixed, leaving open several possible models on how sensory and/or motor circuits are affected by training (Fig. 1): in the above-mentioned *classical model* of VPL, learning occurs in visual cortex and results in stimulus- but not effector-specific learning effects. Hence, changing the effector after training would result in full transfer of learning effects to the new effector. Studies that report effector-independent VPL effects can be interpreted this way (Awada, Bakhtiari, & Pack, 2021). In the *specific stimulus-response mapping model*, visual information is mapped onto effector-specific motor neurons in motor cortex (Kazennikov et al., 1999), parietal cortex (Christopoulos et al., 2018; Dickinson et al., 2003), or subcortical structures (B. H. Liu, Huberman, & Scanziani, 2016). This would result in stimulus- and effector-specific learning effects. Evidence for high effector-specificity VPL has been reported in several studies (C. S. Green, Kattner, Siegel, Kersten, & Schrater, 2015; Grzeczkowski, Cretenoud, Herzog, & Mast, 2017). However, these experiments often also involved a change in the task (e.g.,

binary categorization versus continuous adjustment), which possibly requires different readout weights. Hence, these results could also be interpreted to reflect the well-known task-specificity of VPL (Fahle & Morgan, 1996) instead of effector-specificity. Most recently, a *decision rule model* has been put forward (Awada et al., 2021). Here, stimulus-response mapping takes place between visual neurons and effector-independent neurons involved in abstract movement planning and decision making in parietal cortex or subcortical structures like the pulvinar (Wilke, Turchi, Smith, Mishkin, & Leopold, 2010). This would result in stimulus-specific learning effects that generalize well across effectors. This has indeed been reported (Awada et al., 2021).

The studies investigating effector-specificity in VPL have so far exclusively focused on accuracy as a measure of learning and transfer. However, this only considers the end result of a decision-making process and disregards subcomponents involved in the full sensorimotor processing arc. Here, we use a model-based approach, the well-established drift diffusion model (DDM) of decision making (Ratcliff, 1978), to dissociate different aspects of this arc and learning effects upon them. The DDM relates accuracy to reaction times and allows investigation of the rate of evidence accumulation (the drift rate parameter), the decision bound (i.e., when a decision is reached), non-decision time (related to encoding of sensory information and execution of motor actions), and decision bias for one or the other stimulus alternative. This approach has previously been shown to reveal subcomponents of VPL (Dutilh, Krypotos, & Wagenmakers, 2011; C. C. Liu & Watanabe, 2012; Petrov, Van Horn, & Ratcliff, 2011; Zhang & Rowe, 2014), but has so far not been applied to the question of effector-specificity in VPL.

We train subjects in an orientation discrimination task. We focus on orientation VPL since stimulus-specific VPL effects are well documented for this feature (Schoups et al., 1995). Subjects practice for several days to discriminate orientations and to indicate their choice with a joystick movement (Experiment 1). After the last training session, we require them to carry out the same task but to indicate their choice with an eye movement. Our transfer task thus differs from training only in the effectors, allowing us to dissociate the above-mentioned models. In Experiment 2, we train subjects to indicate their choice with a pointing movement instead that closely matches the spatial transformations required for directing saccades during the transfer task. This allows us to further differentiate whether any possible specificities are due to differences in spatial transformations or the effectors themselves.

To preview, across the two experiments, we find that VPL depends on the effector. The DDM taking reaction times into account shows that the specificity arises in the rate of sensory evidence accumulation (i.e., the drift rate) and the decision bound, which reflects the stopping criterion for

choice, two integral components of decision-making. This is most compatible with a new model, where visual information is mapped onto effector-specific integrators.

Methods

Participants

A total of 44 healthy human volunteers (28 female, 1 left-handed, mean age 25.61 yrs, *SD* 6.94 yrs) participated in this study. All subjects had normal or corrected-to-normal vision, reported no history of neurological or psychiatric disease, and gave written informed consent before participation. No sample size estimate was performed, but sample size was selected based on previous studies. Because most comparisons were within subjects, we used convenience sampling. 22 subjects (12 female, 1 left-handed, mean age 23.5 yrs, *SD* 2.32 yrs) participated in Experiment 1 involving joystick and eye movements. 3 subjects had to be excluded from data analysis because lack of significant learning (final *n*=19, 12 female, 1 left-handed, mean age 23.68 yrs, *SD* 2.4 yrs). 12 subjects participated in Experiment 2 involving reaches on a touchscreen and eye movements (*n*=12, 10 female, mean age 26.8 yrs, *SD* 11.78 yrs). 2 subjects were excluded from data analysis because lack of significant learning (final *n*=10, 9 female, mean age 25.8 yrs, *SD* 1.86 yrs). 10 subjects participated in a control experiment for baseline differences between joystick and eye movements (*n*=10, 6 female, mean age 27.3 yrs, *SD* 6.21 yrs). All procedures were in accordance with the Declaration of Helsinki and approved by the Ethics Committee of the University Medical Center Göttingen (protocol number 29/8/17).

Stimuli and procedure

In all experiments, subjects were trained on an orientation discrimination task. Training took place over four to six consecutive days with one training session per day. On the day after the last training session, subjects were instructed to perform the same task, but with a different effector (see below). Stimuli were presented on an LCD monitor (ViewPixx EEG, refresh rate 120 Hz, resolution 1920 × 1080 pixel, viewing distance 65 cm) in a darkened, sound-attenuating booth (Desone Modular Acoustics). Stimulus delivery and response collection were controlled using Psychtoolbox (Brainard, 1997) running in Matlab (The Mathworks, Inc.) on Windows 10. During all experiments, we continuously acquired pupil and gaze measurements using a high-speed, video-based eye tracker (SR Research Eyelink 1000+). Data were sampled at 1000 Hz from both eyes. Subjects in Experiments 1 and 2 were paid €8 per hour. To assure constant motivation over the training sessions, subjects received a bonus of €2 if

they improved by 10% from the previous training session. Subjects in the control experiment were paid 12€ per hour to assure a similarly high level of motivation as in Experiments 1 and 2.

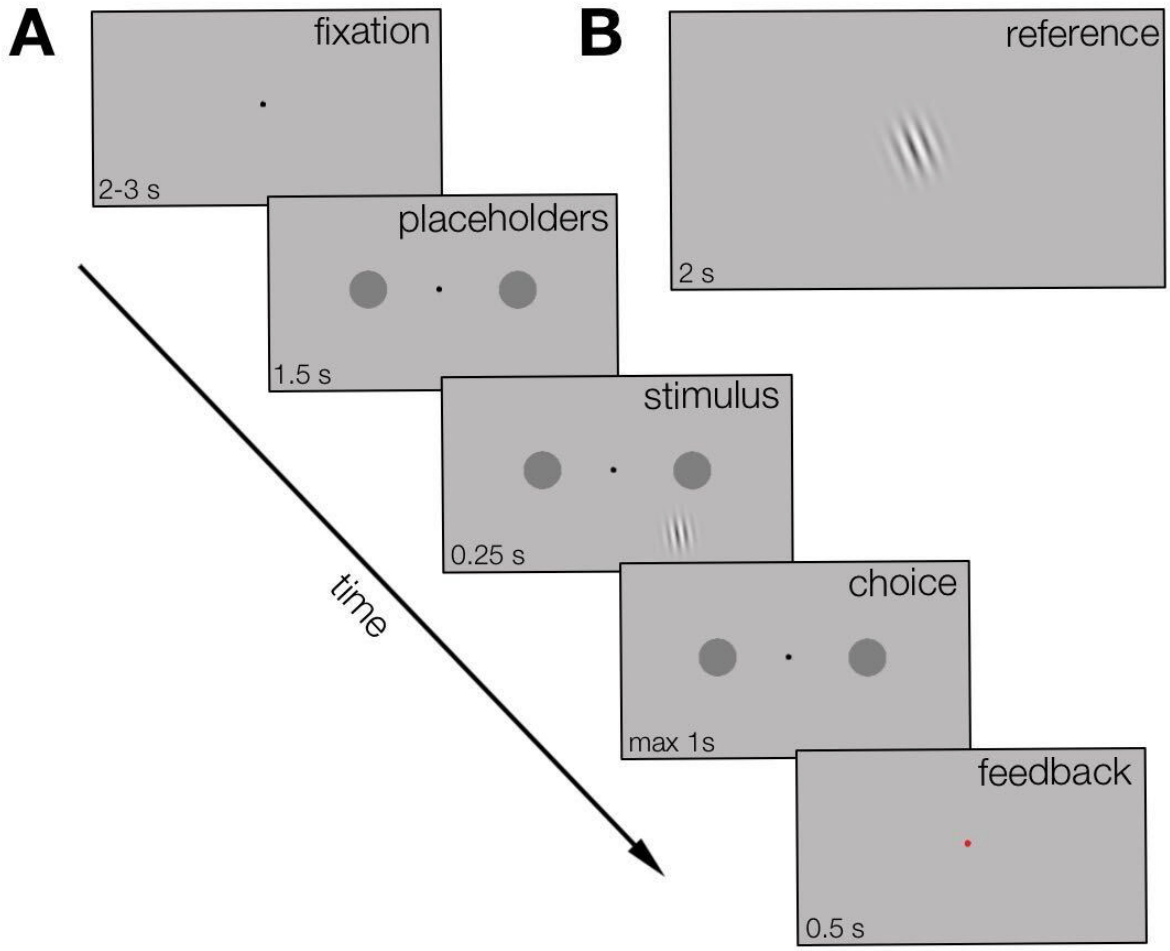

Figure 2: Design of Experiment 1. A. *Each trial started with a fixation phase (2-3 s). Then, placeholders for the saccade version of the task appeared (1.5 s), followed by the presentation of the Gabor grating (0.25 s). Once the Gabor grating disappeared, subjects had 1 s to indicate their choice (with a joystick movement during the training phase or with a saccade during the transfer session). Feedback on accuracy was presented by changing the color of the fixation dot to red or green (0.5 s).* ***B.*** *The reference grating was shown for 2 s in the center of the screen before and in the middle of each block (in total 8 times per session).*

Experiment 1

In Experiment 1, subjects were trained in orientation discrimination and responded with a joystick movement during training, while they had to respond with a saccade in the transfer session. Subjects had to decide whether a Gabor grating (size 2.4 dva, spatial frequency 3.3 cpd, luminance 43.4 cd/m^2) was tilted clockwise or counterclockwise with respect to the reference stimulus. The reference stimulus was a Gabor grating with identical parameters tilted 20° counterclockwise from the vertical meridian presented at the center of gaze. The task stimuli were presented in six logarithmically spaced difficulty levels (±Δ = [0°, 0.31°, 0.56°, 1°, 1.77°, 3.16°]) with 20 trials for each of the orientations,

respectively. The stimuli were presented in pseudo-random order at 8.65° eccentricity in the lower right quadrant against a grey background (39.5 cd/m^2).

Each trial started with a fixation phase of 2-3 s duration. Then, placeholders for the saccade version of the task appeared for 1.5 s at 8° eccentricity (distance from task stimulus 6.7 and 14.5 dva, respectively), followed by the presentation of the Gabor grating for 250 ms. Subjects then had 1 s to respond. The start of the response period was indicated by the disappearance of the Gabor grating. Feedback on accuracy was presented by changing the color of the fixation dot to red (incorrect) or green (correct) for 0.5 s. The next trial started 1.5 s later. If subjects did not respond within 1 s, a message "too late" appeared on the screen and the trial was repeated later during the block. Similarly, if subjects broke fixation throughout the trial, the trial was aborted and repeated at the end of the block.

Each training session included 220 trials evenly distributed among four blocks. Subjects were free to take short breaks between blocks. There were additional breaks in the middle of each block to display feedback about performance and reaction times. The reference was shown for 2 s in the center of the screen before and in the middle of each block (in total 8 times per session).

Subjects were instructed to respond as accurately as possible, however, they also had to consider the time limitation. To respond, subjects had to tilt a joystick (Logitech Extreme 3D Pro, either in standard configuration or in a customized configuration for left-handed subjects) with their dominant hand to the left if a stimulus was counterclockwise rotated and to the right if clockwise (while maintaining fixation). Hand dominance was assessed at the beginning of the experiment with the Edinburgh Handedness Inventory (Oldfield, 1971). In the transfer session, subjects were instructed to instead direct their gaze to the left or right placeholder, respectively. Subjects thus had to learn an arbitrary mapping of a visual stimulus to a response alternative, and potentially remap this association to a new effector in the transfer session. Previous research has shown that effectors do not affect accuracy pre training (Jaskowski & Sobieralska, 2004; Szul, Bompas, Sumner, & Zhang, 2020; Szumska, van der Lubbe, Grzeczkowski, & Herzog, 2016). We focus on effector-specificity since stimulus-specificity does not differentiate the different models for sensorimotor mapping in VPL.

Participant had to reach at least 75% correct responses before we transitioned to the transfer session. All but two subjects reached this performance level within four days; two subjects required five and six days of training, respectively.

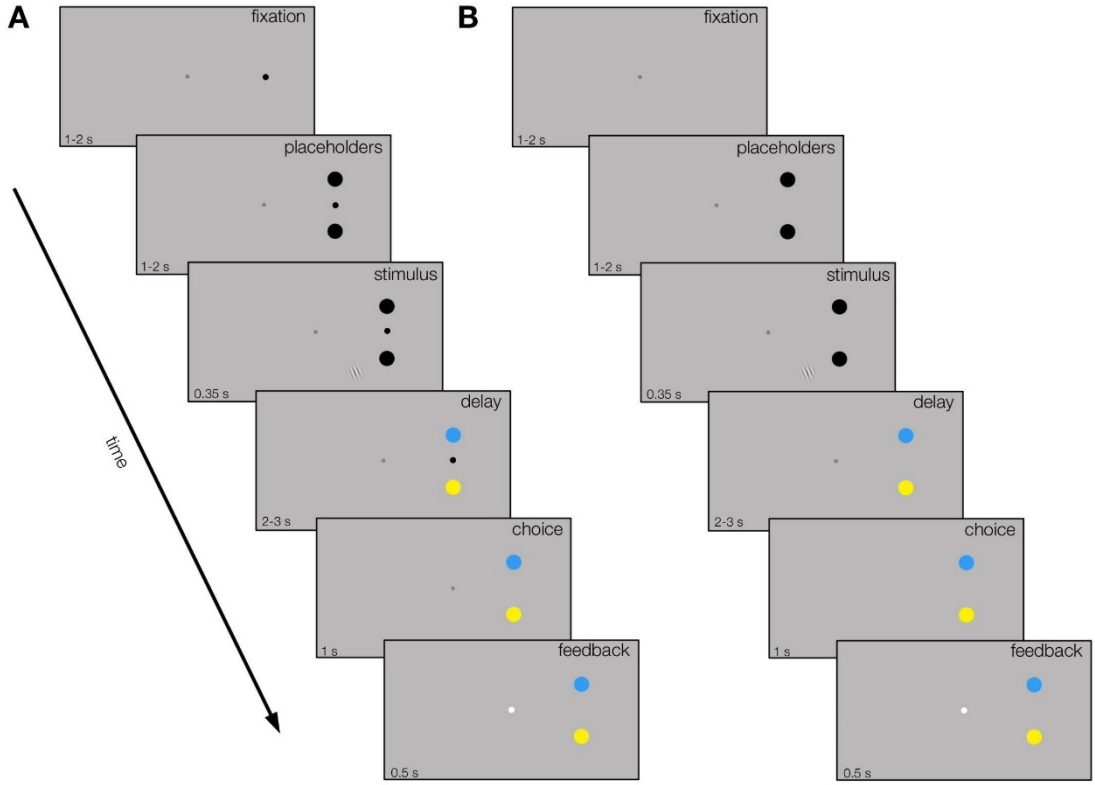

Figure 3: Design of Experiment 2. A. *In the touchscreen version of the task, each trial started with a finger and eye fixation phase of (1-2 s). Subjects had to maintain central fixation with their gaze and simultaneously touched a finger fixation dot between the two target locations on the screen. Then, black placeholders appeared (1-2 s), followed by the presentation of the Gabor grating (0.35 s). A delay phase (2-3 s) started with a switch of the placeholders' color to blue and yellow. Subjects had to respond within 1 s once the finger fixation dot disappeared. Feedback on accuracy was presented by changing the color of the fixation dot to blue or white (0.5 s).* ***B.*** *The trial structure of the saccade version of the task was identical to the touchscreen version, with the exception that the finger fixation dot was not shown. Instead, the disappearance of the central fixation dot indicated the start of the response period.*

Experiment 2

In Experiment 2, subjects were trained on orientation discrimination and responded with a finger movement during training, while they had to respond with a saccade in the transfer session. We recorded finger movements by means of a transparent touch screen (Magic Touch, Tyco Touch Inc.), modelled after Pilacinsky and Lindner (2019), which was mounted in front of the screen. Subjects were provided with an arm support (ErgoRest) and a finger stylus (TrueTip) to ensure ease of use.

As in Experiment 1, subjects had to decide whether a Gabor grating (size 2.4 dva, spatial frequency 3.3 cpd, luminance 10.9 cd/m^2) was tilted clockwise or counterclockwise with respect to the reference stimulus. The reference stimulus was a Gabor grating with identical parameters tilted 20° counterclockwise from the vertical meridian. The task stimuli were presented in six logarithmically spaced difficulty levels (±Δ = [0°, 0.31°, 0.56°, 1°, 1.77°, 3.16°]) with 20 trials for each of the orientations, respectively. The stimuli were presented in a pseudo-random order at 8.65° eccentricity in the lower right quadrant against a dark grey background (6.4 cd/m^2).

Each trial started with a finger and eye fixation phase of 1 to 2 s duration. Subjects had to maintain central fixation with their gaze and simultaneously touched a finger fixation dot between the two target locations on the screen. Then, black placeholders appeared for 1 to 2 s at 11° eccentricity (distance from task stimulus 5 and 11.4 dva, respectively), followed by the presentation of the Gabor grating for 350 ms, a delay of 2 to 3 s during which the placeholders' color switched from to blue and yellow, respectively. Subjects had 1 s to respond. The start of the response period was indicated by change of fixation dot. The next trial started 0.5 to 2 s later. If subjects did not respond within 1 s, the trial was repeated later during the block. Similarly, if subjects broke fixation throughout the trial, the trial was aborted and repeated at the end of the block. Feedback about correct and incorrect responses was provided with a blue or white dot, respectively, at fixation.

Each training session included 220 trials evenly distributed among four blocks. Subjects were free to take short breaks between blocks. The reference was shown for 2 s in the center of the screen before and in the middle of each block (in total 8 times per session).

Subjects were instructed to respond as accurately as possible. To respond, subjects had to move their finger from the finger fixation spot between the two targets to the blue target if the stimulus was tilted clockwise or to the yellow target if the stimulus was tilted counterclockwise (while maintaining central fixation). Hand dominance was assessed at the beginning of the experiment with the Edinburgh Handedness Inventory (Oldfield, 1971). In the transfer session, subjects were instructed to instead direct their gaze to the blue or yellow placeholder, respectively. For all subjects, training lasted four days.

<u>Control experiment for differences between effectors</u>

To determine whether any differences between effectors were the result of training or already existed before training began, we conducted a control experiment comparing orientation discrimination performance between effectors. This was done in a separate group of subjects (n=10, 6 female, mean age 27.3 yrs, *SD* 6.21 yrs) to avoid possible effects of pretesting on generalization performance. The task was identical to Experiment 1, and subjects were asked to complete 220 trials with one effector and, after a break, 220 trials with the other effector. The order of effectors was counterbalanced across subjects.

Analyses

The Signal Detection Theory parameter *d'* (D. M. Green & Swets, 1966) was computed per subject and session. *d'* were entered into repeated measures analyses of variance (rmANOVA) with factors session

(4 levels for training session effects, or 2 levels for last training session to transfer effects), difficulty (5 levels), and their interaction. Reaction times were entered into a rmANOVA with factors session and correctness (levels correct, incorrect). Two subjects required more than four days of training (5 and 6) in Experiment 1; therefore, we restrict our analyses to the last four training sessions. Due to technical issues, we lost the data from the third training session for one subject in Experiment 1. Thus, to be able to compute rmANOVA across training sessions (1-4), we excluded the subject (total *n*=18) from these analyses, but we included the subject in the transfer test (last training versus transfer, total *n*=19). For *t*-tests and rmANOVAs, we compute Hedges' *g* and partial η^2, respectively, as effect sizes, using the Measures of Effect Size Toolbox (Hentschke & Stüttgen, 2011). In addition, to quantify transfer performance, a Specificity Index (Petrov et al., 2011) was calculated, as follows:

$$SI = \frac{d'_{last\ training} - d'_{transfer}}{d'_{last\ training} - d'_{first\ training}}$$

The index reflects the training session used to compute the median *d'*. All behavioral data are publicly available on Figshare (link will be share upon journal acceptance).

Drift Diffusion Model

Sequential sampling models like the DDM assume that decision making is a stochastic process of evidence accumulation about a stimulus in time (Ratcliff, 1978). Decisions are made when the accumulation process reaches an evidence threshold (decision criterion) for one of the alternatives. The rate of accumulation and the decision criteria determine the reaction time (RT), and fluctuations in the process are the reason for errors and the distribution of RTs. The DDM can account for effects of different experimental conditions from accuracy, mean RT for correct and error responses, and the shape of RT distributions for correct and error responses. The full DDM includes seven parameters: drift rate (v) and its trial-by-trial variability (η), non-decision time (T_{er}) and its trial by-trial variability (s_t), boundary separation (a), the starting point of the decision process (z) and its trial-by-trial variability (s_z). v is the key component of the model, as it describes the rate of evidence accumulation and positively correlates with the accuracy level of subjects. a is the distance between decision criteria for both alternatives, and the bigger the distance is, the longer it takes to make a decision. The evidence accumulation starts from z, which lies between 0 and a. z can be positioned closer to one of the alternatives, which biases the decision process toward it. T_{er} accounts for the non-decision phases: stimulus encoding and response output. The DDM states that v and T_{er} are independent, and hence, that manipulations of the effector only affect T_{er} (Gomez, Ratcliff, & Childers, 2015). This assumption has been previously empirically validated (Gomez et al., 2015), and it is thus common in the field to

compare DDM parameters directly between effectors (e.g., MacInnes, 2017). Here, we follow the same logic.

We use the Hierarchical Drift Diffusion Model (HDDM 0.8.0) (Wiecki, Sofer, & Frank, 2013), available as a toolbox in Python, to fit the model parameters to our data of Experiment 1. The HDDM employs Bayesian inference to recover the model parameters. First, the HDDM estimates the parameter values for a population (p_{group}), and then it uses the population fit to constrain the parameter values for individual subjects (p_{subj}). The hierarchical Bayesian approach performs significantly better than using χ^2 as the optimization criterion when data per subject is limited (Ratcliff & Childers, 2015).

To identify the combination of model parameters which best explains our data, we fit several models built on theoretically meaningful parameter combinations (see Results) and compare them on the basis of the Deviance Information Criterion (DIC, Spiegelhalter, Best, Carlin, & van der Linde, 2002).

In the best performing model, we estimate parameters for all sessions jointly while allowing *v* and *a* to vary across the sessions and the difficulty levels (see Discussion). The stimuli are encoded by their orientation (clockwise/counterclockwise), and *v i*s estimated for both orientations together. We allow T_{er} and *z* to vary only between sessions because a priori, encoding, muscle response time, and the orientation bias should be the same across difficulty levels. We exclude the variability parameters from the analysis; for proper estimation of these parameters, the HDDM requires noticeably more samples per condition than we have available. We run the fitting algorithm four times with 30000 Markow-Chain-Monte-Carlo samples, 10000 samples burn-in, and discarding every 5th sample.

We utilize two formal diagnostics to assess the convergence of the model: the Gelman-Rubin and Geweke method, respectively. The Gelman-Rubin method compares between-chain and within-chain variances for each model parameter (Brooks & Gelman, 1998). Significant differences (diagnostics value >1.2 or 1.1 under strict conditions) between these variances indicate non-convergence (Gelman & Rubin, 1992). We compute the diagnostic with the help of the kabuki (0.6.3) Python library. The Geweke method evaluates only within-chain fluctuation (Geweke, 1991). It calculates *z*-scores of the difference between the initial and subsequent segments of the chain. If the chain has converged, most points should be within two standard deviations of zero (*z*-score $\in$ [-2, 2]). We chose the first 10% of the chain as the initial segment, and the last 50% divided into 20 segments. To compute the Geweke diagnostic, we use the PyMC (2.3.8) Python library.

To compare model parameters between conditions within subjects, we assess the overlap of the respective posterior distribution. To assess differences between groups, we compare the distribution of differences between conditions between groups and report the probability of one group having a higher estimated difference in parameter value than the other group (Pedersen et al., 2021).

To investigate the relationship between drift rate and boundary separation, we use a 2D minimum energy test (Aslan & Zech, 2005; Székely & Rizzo, 2013), as implemented by B. Lau (https://github.com/brian-lau/multdist). This test assesses via permutations whether two multidimensional samples are drawn from the same distribution.

Results

Learning and transfer effects on accuracy and reaction time

We find that across all five difficulty levels, subjects significantly improve their performance (*d'*) as function of training, with increases ranging between 0.60 and 1.93 (Fig. 4A, mean increase between first and last training session 1.25, $t(18)=-11.78$, $p<0.001$, $g=3.04$; rmANOVA main effect of session, $F(3, 51)=138.94$, $p<0.001$, $\eta_p^2=0.519$). Subjects learn more on easy than on difficult orientation differences (main effect of difficulty, $F(4, 68)=96.18$, $p<0.001$, $\eta_p^2=0.647$, interaction session × difficulty $F(12, 204)=13.18$, $p<0.001$, $\eta_p^2=0.145$). When the effector is changed from a manual joystick response to a saccadic response in the transfer session, we find a significant decrease in performance relative to the last session across all difficulty levels (mean decrease 0.32, $t(18)=6.41$, $p<0.001$, $g=0.98$; rmANOVA main effect of session, $F(1,18)= 21.48$, $p<0.001$, $\eta_p^2=0.54$ main effect of difficulty, $F(4,72)= 290.68$, $p<0.001$, $\eta_p^2=0.94$, interaction session × difficulty $F(4,72)=0.487$, $p=0.745$, $\eta_p^2=0.02$). This is

also evident when we calculate the SI, which has a median of 0.30 (Fig. 4C, p<0.001, sign test, one-sided). In contrast, there are no differences in *d'* when the task is carried out with joystick or saccades in the control group in the absence of training (mean difference in *d'* 0.14, $t(9)=1.10$, $p=0.653$, $g=0.21$). Accordingly, the difference in *d'* between effectors in the trained group is significantly larger than the difference in the control group (mixed ANOVA, main effect of experimental group, $F(1,27)=8.593$, $p=0.007$, $\eta_p^2=0.241$, no main effect of difficulty, $F(4,108)=1.126$, $p=0.348$, $\eta_p^2=0.04$, no interaction experiment × difficulty $F(4,108)=0.96$, $p=0.432$, $\eta_p^2=0.03$). This suggests that the drop in d' we find in in Experiment 1 when effectors change is learning-induced and cannot be explained by preexisting differences.

Mean reaction times for correct and incorrect responses do not differ from each other throughout the sessions and do not change significantly as a function of learning (all p>0.29, all $\eta_p^2 < 0.069$). When the effector is changed, reaction times become uniformly shorter (Fig. 4B, mean difference 0.143 s, main effect of session, $F(1,18)=91.56$, p<0.001, $\eta_p^2 = 0.83$). As there was no significant main effect of correctness or any significant interaction (all p>0.09, all $\eta_p^2 < 0.14$) in reaction times, this is unlikely to constitute a learning effect and can be attributed to generally shorter latencies of saccades as

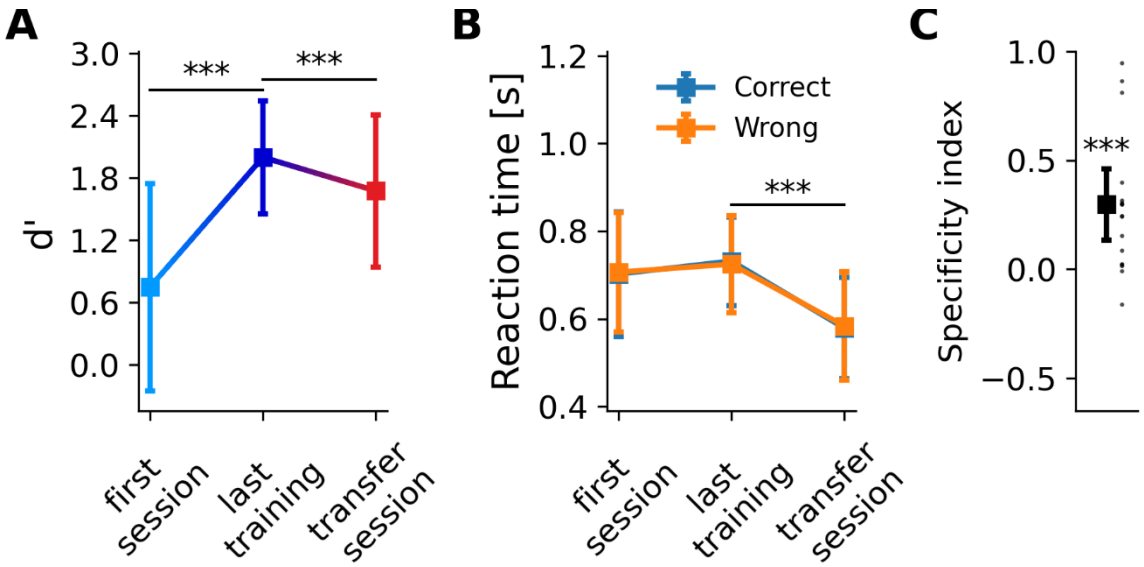

Figure 4: Accuracy and reaction times for Experiment 1. A. *Mean accuracy (in units of d') increases with learning from the first to the last training session (mean difference 1.25, t(18)=-11.78, p<0.001, g=3.04, two-sided) and drops from the level reached in the last training session when the effector is changed from a manual joystick response to a saccadic response (mean difference 0.32, t(18)=6.41, p<0.001, g=0.98, two-sided).* ***B.*** *Mean reaction times for correct and incorrect responses do not differ from each other throughout the sessions and do not change significantly as a function of learning (all p>0.29, all $\eta_p^2 < 0.069$). When the effector is changed, reaction times become uniformly shorter (main effect of session, F(1,18)=91.56, p<0.001, $\eta_p^2 = 0.83$, no main effect of correctness or interaction, all p>0.09, all $\eta_p^2 < 0.14$).* ***C.*** *Learning depends on the effector, as evidenced by a median Specificity Index of 0.30, which is significantly bigger than 0 (p<0.001, sign test, one-sided). Dots reflect individual subjects' Specificity Indices. 17 out of 19 subjects have Specificity Indices >0. In panels A and B, error bars reflect ±2 SDs. In panel C, error bars reflect the standard error of the median (McKean & Schrader, 1984). *** stand for p < 0.001, ** for p < 0.01, and * for p<0.05.*

compared to hand movements (Jaskowski & Sobieralska, 2004) that we also find in our control group (mean difference 0.12 s, rmANOVA main effect of effector, $F(1,9)=6.731$, $p=0.029$, $\eta_p^2 = 0.428$, all other effects $p>0.84$)).

We find very similar results in Experiment 2, where subjects perform the same task with a precise reach on a touch screen. In this smaller sample ($n=10$), *d'* also improves with training over sessions (Fig. 5A, mean increase between first and last training session 0.82, $t(9)=7.76$, $p<0.001$, $g=1.65$; rmANOVA main effect of session, $F(3,27)=24.46$, $p<0.001$, $\eta_p^2 = 0.20$), with larger improvements for easier orientation differences than for difficult orientation differences (main effect of difficulty, $F(4,36)=106.65$, $p<0.001$, $\eta_p^2 =0.71$, interaction session × difficulty $F(12,108)=4.88$, $p<0.001$, $\eta_p^2 =0.12$). When we change the motor component of the task from reaches to saccades in the transfer session, we again find a drop in discrimination across difficulty levels, albeit not statistically significant (mean decrease 0.12, $t(9)=0.99$ $p=0.347$, $g=0.2$; rmANOVA main effect of session, $F(1,9)=0.98$, $p=0.347$, $\eta_p^2 =0.01$, main effect of difficulty, $F(4,36)=66.24$, $p<0.001$, $\eta_p^2 =0.69$, interaction session × difficulty $F(4,36)=1.97$, $p=0.119$, $\eta_p^2 =0.03$). Specificity Indices nevertheless show that training is specific to the effector in 7 out of 10 subjects (Specificity Index >0); expressed in median SI, this is 0.49 (Fig. 5C, $p=0.035$, sign test, one-sided).

Average reaction times decrease over sessions (Fig. 5B, mean decrease 0.09 s, $t(9)=3.83$, $p=0.004$, $g=-0.97$; rmANOVA main effect of session, $F(3,27)=10.94$, $p<0.001$, $\eta_p^2=0.170$, main effect of correctness, $F(1,9)=1.20$, $p=0.301$, $\eta_p^2=<0.001$, interaction session × correctness $F(3,27)=0.89$, $p=0.45$, $\eta_p^2<0.001$) and further drop significantly when the effector is changed (mean decrease 0.19 s, $t(9)=-6.14$, $p<0.001$; $g=-2.63$; rmANOVA main effect of session, $F(1,9)=37.74$, $p<0.001$, $\eta_p^2 = 0.673$, main effect of correctness, $F(1,9)=0.315$, $p=0.588$, $\eta_p^2<0.001$, interaction session × correctness $F(1,9)=1.80$, $p=0.213$, $\eta_p^2=0.001$). The latter can again be attributed to generally shorter reaction times for saccades than for reaches. Together, the SIs from Experiments 1 and 2 show that learning effects on accuracy in orientation discrimination depend on the effector without any change in the stimulus material or the task.

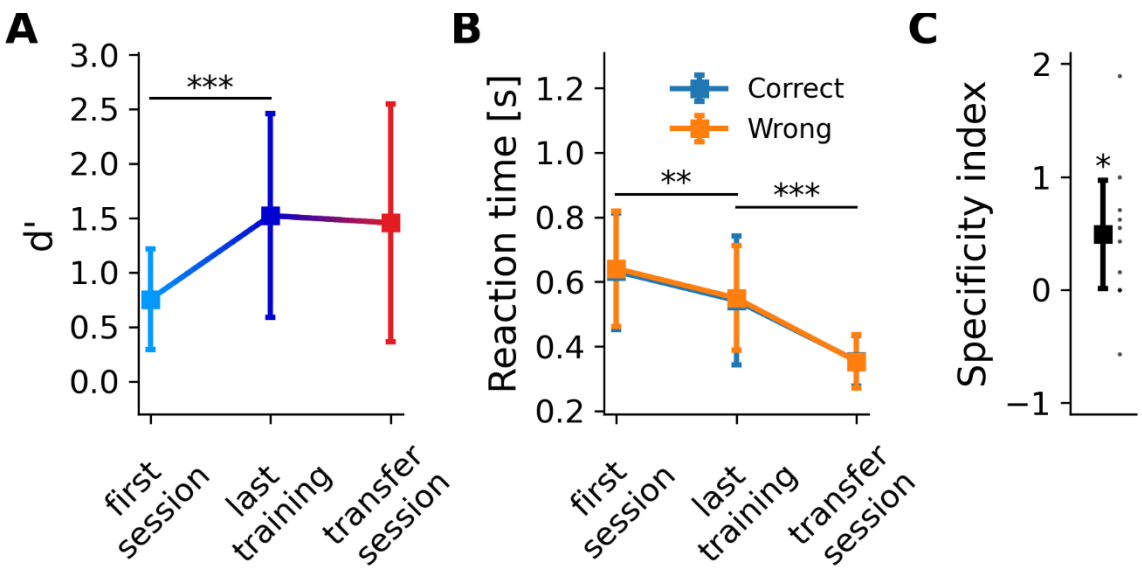

Figure 5: Accuracy and reaction times for Experiment 2. A. *Mean accuracy (in units of d') increases with learning from the first to the last training session (mean difference 0.82, t(9)=7.76, p<0.001, g=1.65, two-sided). There is a numerical, but no statistically significant drop in accuracy when the effector is changed from a manual joystick response to a saccadic response (mean difference 0.12, t(9)=0.99, p=0.345, g=0.2, two-sided).* ***B.*** *Mean reaction times decrease over sessions (mean decrease 0.09 s, t(9)=3.83, p=0.004, g=-0.97) and further drop significantly when the effector is changed (mean decrease 0.19 s, t(9)=-64.1497, p<0.001; g=-2.6325).* ***C.*** *Training is specific to the effector in 7 out of 10 subjects (Specificity Index >0), with a median Specificity Index of 0.49, which is statistically significantly bigger than 0 (p=0.035, sign test, one-sided). Dots reflect individual subjects' Specificity Indices. In panels A and B, error bars reflect ±2 SDs. In panel C, error bars reflect the standard error of the median (McKean & Schrader, 1984). *** stand for p < 0.001, ** for p < 0.01, and * for p<0.05.*

Learning and transfer effects on components of the decision-making process

To dissociate different aspects of the decision-making process and learning effects upon them, we use Hierarchical Drift Diffusion Models (HDDM) to relate accuracy to reaction times. To this end, we consider the data from Experiment 1 (joystick – saccade) in which the task was performed under time

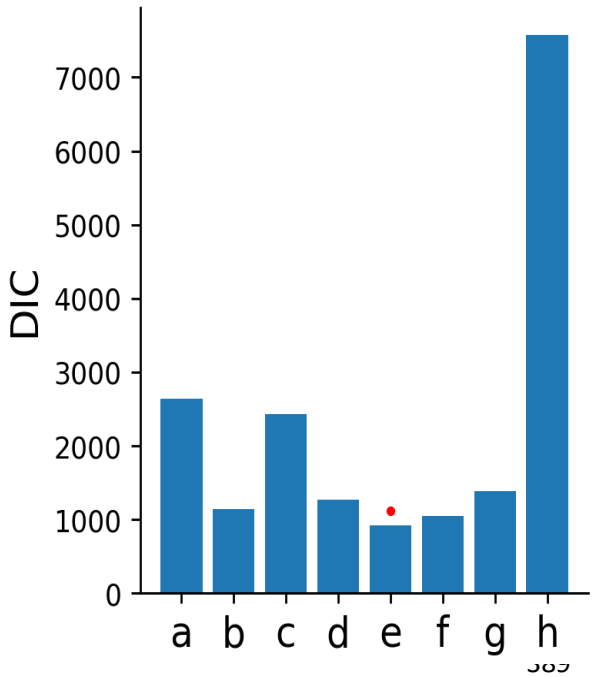

Figure 6: Model comparison. *Description of the models (v = drift, a = decision boundary, t = non-decision time, z = bias; conditions: ses = session, diff = difficulty, const = constant across conditions):*
***a**: v(ses)_a(ses*diff)_t(ses)_z(ses),*
b: *v(ses*diff)_a(ses)_t(ses)_z(ses),*
c: *v(ses)_a(ses)_t(ses)_z(ses),*
d: *v(ses*diff)_a(ses)_t(ses),*
e: *v(ses*diff)_a(ses*diff)_t(ses)_z(ses),*
f: *v(ses*diff)_a(ses*diff)_t(ses),*
g: *v(ses*diff)_a(const)_t(ses)_z(ses),*
***h**: v(ses*diff)_a(ses*diff)_t(const)_z(ses).*
*The best performing model is marked with a red dot. Model **e** has the lowest DIC value and outperforms the closest model by ΔDIC > 100.*

pressure (by means of a fixed response time window of 1 s) and without a delay phase; the DDM is most appropriate for such speeded reaction times (Ratcliff & Rouder, 1998). We consider the parameters drift rate, decision boundary, bias, and non-decision time. To find the combination of model parameters which best explains our data, we fit several models and compare their DIC values (Fig. 6). The model which allows drift and decision boundary to vary across sessions and difficulty levels and non-decision time and bias to vary only across sessions provides the best fit to our data (model e). This model outperforms the closest model (f), which does not account for bias, by a difference in DIC >100. According to Gelman-Rubin method, all parameters of model (e) converged with a value < 1.01 and there is no difference in convergence quality between sessions (Fig. S1). The Geweke method indicates that all group parameters (n=64) have an average z-score ∈ [-1.5, 1], and at worst, 70% of the values are within two standard deviations of zero (Fig. S2). All subject parameters (n=1138) have an average z-score = ±2, which is within acceptable bounds. Only few parameters have 50-70% of the values within two standard deviations of zero. Also, the fit quality is not different across sessions in each of the four independent runs (Fig. S3). Together, this shows that our HDDMs converged well. Additionally, we verify the model through a posterior predictive comparison (Table S1), which demonstrates that the predicted accuracy and reaction time distributions match the empirical data.

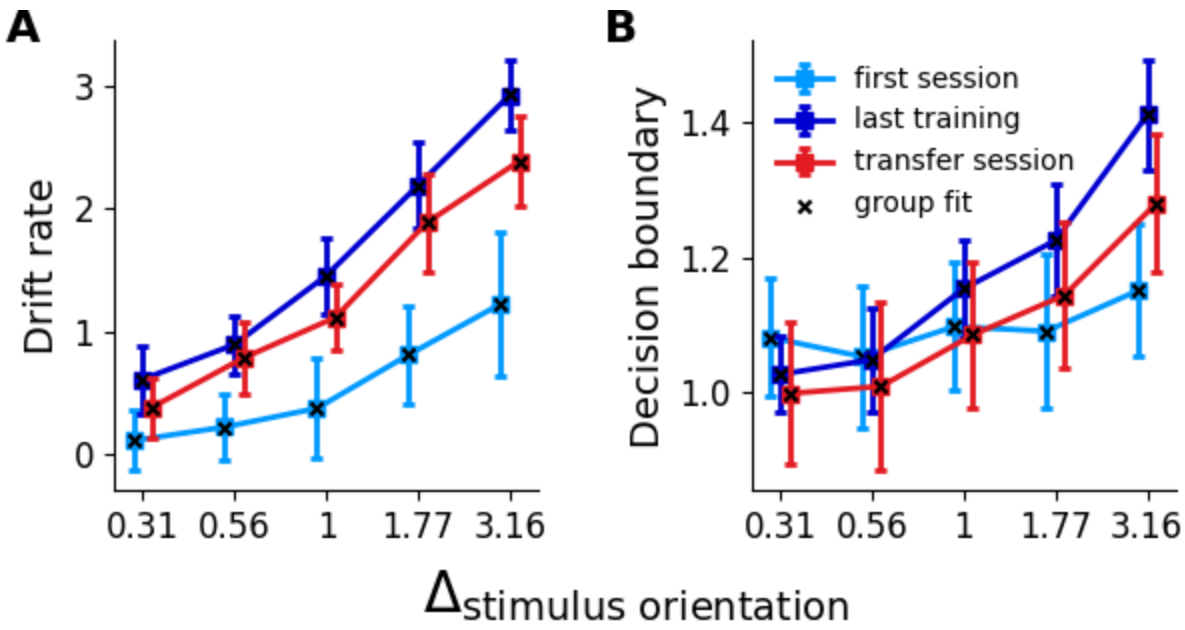

Figure 7: Hierarchical Drift Diffusion Model fit of drift rates and decision boundaries in Experiment 1. A. *Drift rate (v) increases between the first and the last training session across orientations. When the effector changes in the transfer session, the learning-induced gains in drift rate drop.* ***B.*** *The separation between decision boundaries (a) increases between the first and last training session for the easier orientations and decreases or stays the same for the most challenging conditions. Learning effects drop when the effector changes in the transfer session. Squares signify the mean individual fits, errors bars ±2 standard deviations; crosses indicate the group fits. The discrepancy along the abscissa between sessions is intentional to increase visibility of data points.*

Drift rate

Both when assessed on the group and on the subject level, we find that drift rate increases between the first and the last training session across difficulty levels (Fig. 7A, $p_{subj}(v_{last_training} > v_{first_training})$ = [0.91, 0.97, 0.99, 0.99, 1.0], $p_{group}(v_{last_training} > v_{first_transfer})$ = [0.99, 1.0, 1.0, 1.0, 1.0] for $\pm\Delta_{orientation}$ = [0.31°, 0.56°, 1°, 1.77°, 3.16°], respectively). This is in keeping with previous reports that show that VPL increases the drift rate in various tasks (e.g., C. C. Liu & Watanabe, 2012; Petrov et al., 2011). When we change the effector in the transfer session, drift rates drop, which is evident in the single subject and, especially, group fit of the HDDM ($p_{subj}(v_{last_training} > v_{transfer})$ = [0.73, 0.61, 0.84, 0.78, 0.91], $p_{group}(v_{last_training} > v_{transfer})$ = [0.95, 0.78, 0.99, 0.97, 0.99], see Table S2 for the comparison between all sessions). We do not find this difference in drift rate between effectors in the untrained control group ($p_{group}(v_{hand} > v_{eyes})$ = [0.45, 0.2 , 0.59, 0.23, 0.36], $p_{group}(\Delta v_{Exp1} > \Delta v_{Control})$ = [0.78, 0.85, 0.79, 0.92, 0.95]). We note that statistical power of the between-groups comparison may be affected by higher variance in the drift estimation (see Fig. S4). In the control group, drift was conditioned on effectors and not sessions, which may lead to an increase in the variance of the MCMC samples. Nevertheless, the results overall suggest that learning effects on the rate of evidence accumulation depend on the effector.

Decision boundaries

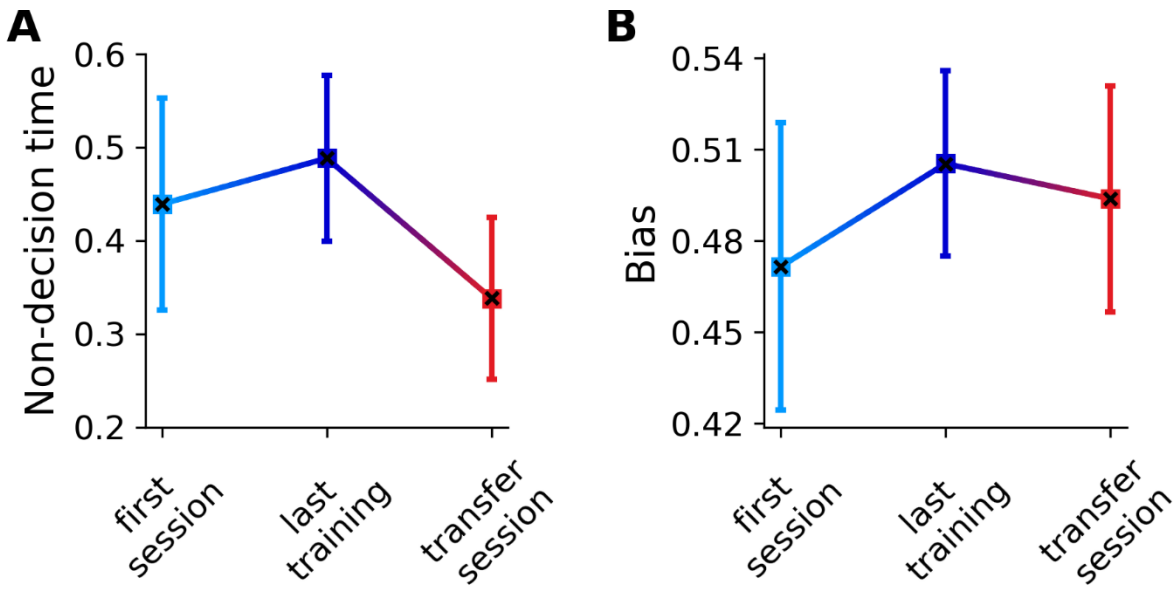

Figure 8: Hierarchical Drift Diffusion Model fit of non-decision times and biases in Experiment 1. A. *Non-decision time (t) increases between the first and the last training session, and then decreases in the transfer session* ***B.*** *Subjects had a slight bias (z) in the clockwise direction during the first training session, which disappeared with the last training and remained low for the transfer session. Squares signify the mean individual fits, errors bars ±2 standard deviations; crosses indicate the group fits.*

We find that the separation between decision boundaries corresponding to clockwise and counterclockwise choices increases between the first and the last training session, most clearly and consistently in the easier conditions, but decreases or stays the same for the two most challenging conditions (Fig. 7B and Table S2, $p_{subj}(a_{last_training} > a_{first_training})$ = [0.25, 0.46, 0.72, 0.89, 0.99], $p_{group}(a_{last_training} > a_{first_training})$ = [0.05, 0.44, 0.93, 0.99, 1.0]] for $\pm\Delta_{orientation}$ = [0.31°, 0.56°, 1°, 1.77°, 3.16°], respectively). The separation is reduced when we change the effector in the transfer session ($p_{subj}(a_{last_training} > a_{transfer})$ = [0.60, 0.63, 0.74, 0.76, 0.83], $p_{group}(a_{last_training} > a_{transfer})$ = [0.8, 0.87, 0.96, 0.97, 0.99]). The difference in decision boundary between effectors in the untrained control group and the trained group is weak ($p_{group}(\Delta a_{Exp1} > \Delta a_{Control})$ = [0.41, 0.55, 0.65, 0.76, 0.79], see Fig. S5) for the same reasons as for drift rates. However, within the control group, the difference between effectors is insignificant ($p_{group}(a_{hand} > a_{eyes})$ = [0.74, 0.64, 0.65, 0.56, 0.67]) and the difference within effectors between groups indicates the presence of learning for the easier conditions ($p(a_{eyes_exp1} - a_{eyes_control})$ = [0.44, 0.41, 0.85, 0.94, 0.94]). Hence, like the drift rate, learning effects on the decision boundary depend on the effector.

<u>Non-decision time</u>

Non-decision time increases between the first training session and the last training session, as is clearly evident on the group level (Fig. 8A, $p_{group}(t_{last_training} > t_{first_training})$ = 0.995), but less consistently on the subject level ($p_{subj}(t_{last_training} > t_{first_training})$ = 0.769). It then decreases massively in the transfer session ($p_{subj}(t_{last_training} > t_{transfer})$ = 1.0, $p_{group}(t_{last_training} > t_{transfer})$ = 1.0). The latter effect can likely be attributed

to generally shorter latencies of saccades as compared to hand movements (Gomez et al., 2015; Jaskowski & Sobieralska, 2004).

Bias

We find a slight bias in clockwise orientation (Fig. 8B, $z_{subj} = 0.471$, $z_{group} = 0.471$) during the first training session, which disappears with the last training ($z_{subj} = 0.505$, $z_{group} = 0.505$; $p_{subj}(z_{last_training} > z_{first_training}) = 0.796$, $p_{group}(z_{last_training} > z_{first_training}) = 0.998$). Hence, subjects seem to optimally adjust their response bias in accordance with the task demands. During the transfer session, the bias is almost eliminated ($z_{subj} = 0.493$, $z_{group} = 0.493$; $p_{subj}(z_{last_training} > z_{transfer}) = 0.633$, $p_{group}(z_{last_training} > z_{transfer}) = 0.851$).

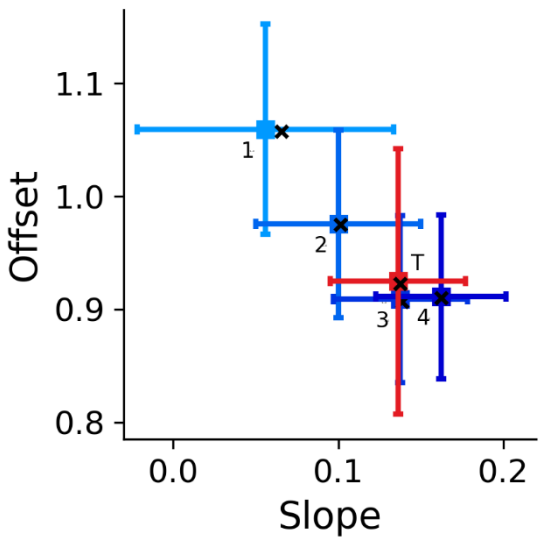

Figure 9: Relationship between drift rate and decision boundary. *Slope and offset of a linear fit of drift rates against decision boundaries across orientations. Throughout the training sessions, the slope and the offset stabilize at higher and lower values, respectively. This is evidenced by a significant difference in the a-v relationship between the first and last training session (2D minimum energy test, e=2.391, p=0.001). During the transfer session, only the slope changes (t(18)=4.28, p=0.0004), while the offset remains the same as in the last training session (t(18)=-0.77, p=0.44). Colored error bars display the estimates based on the per-subject parameters (mean ±2 SDs), the black crosses show the estimates based on group parameters. Numbers indicate training sessions (1-4), transfer session is red and marked with T.*

Strategic adjustment of decision boundary to drift

We observe the emergence of a correlation between drift rate (v) and decision boundary (a) throughout training (Spearman correlation $\bar{\rho}_{first_training}=0.511$, $p=0.380$; $\bar{\rho}_{last_training}=0.947$, $p=0.021$; $t(18)=5.88$, $p<0.001$, $g=1.70$, see Table S3). This correlation remains significant during the transfer session (Spearman correlation $\bar{\rho}_{transfer}=0.947$, $p=0.021$; $t(18)<0.001$, $p>0.999$, $g\approx0$). We can relate this finding to the well-known speed-accuracy trade-off. The subjects, apparently, preferred speed over accuracy for more challenging conditions, and accuracy over speed for easier conditions. They lowered or stabilized a even after the improvement of v for the two most difficult orientations (see Fig. 7A-B). Given the low drift rates, the emphasis on speed reduced the probability of exceeding the response time limit (1.25s after the stimulus onset) and, hence, the overall number of trial attempts. For the three least demanding conditions, subjects strategically used the additional time afforded by an increase in the drift rate to sample additional information towards higher accuracy instead of saving time on task by decreasing their reaction times through tighter decision bounds. To quantify the interaction between v and a, we fit the first-degree

polynomial to the data per subject and session, where *a* depends on *v* over difficulty (orientation) levels. Thus, the slope of the linear fit reflects how *a* changes with growing *v*, and the offset is *a* when *v* equals zero. Throughout the training, the slope and offset converge on higher and lower values, respectively (Fig. 9). This is evident from a significant difference in the *v*-*a* relationship between the first and last training session (2D minimum energy test, $e = 2.391$, $p = 0.001$). During the transfer session, only the slope changes ($t(189) = 4.28$, $p = 0.0004$), while the offset is not different from the last training session ($t(18) = -0.77$, $p = 0.44$). The slight change in the slope can be explained by a non-linear interaction between *v* and *a*. Accordingly, the difference between the final training session and the transfer session is statistically significant in 2D ($e = 0.165$, $p = 0.04$); however, not in comparison with the penultimate training session ($e = 0.040$, $p = 0.65$). Considering the overall faster reaction times (by ca. 150 ms) during the transfer session, subjects could apparently compensate for the reduced *v* by raising *a*, which would be reflected in higher offset values.

Discussion

We find that VPL of orientation discrimination depends on the effector with which the task is trained. Analyses with the HDDM reveal that this dependency arises from effector specificity of two subcomponents of the decision-making process, namely the rate of evidence accumulation and the decision bound. Our results thus confirm that sensorimotor mapping is affected by VPL but suggests a new model on how and where this mapping may take place.

Previous models of VPL suggest either full or no transfer between effectors. Full transfer would arise if learning took place entirely in visual cortex (Ahissar & Hochstein, 2004; Fahle, 2004; Sagi, 2011), as originally conceived, or if sensorimotor mapping occurs between visual and decision making neurons, e.g., in parietal cortex, in the form of learning an abstract “rule” (Awada et al., 2021). No transfer between effectors would be predicted if sensorimotor mapping would take place between visual and effector-specific neurons in motor or parietal cortex (Grzeczkowski et al., 2017). In theory, learning abstract rules leads to excellent generalization capacity, which has for example been suggested to be the basis of visual training from action video games (C. S. Green, Pouget, & Bavelier, 2010). On the other hand, decomposing tasks into effector-specific components greatly simplifies reinforcement learning and effector-specific reward prediction errors supporting such learning have been found in humans using fMRI (Gershman, Pesaran, & Daw, 2009). Effector specificity has been found in several other forms of learning involving sensorimotor mapping (e.g., Feulner, Postin, Schwiedrzik, & Pooresmaeili, 2021), but generalization across effectors is also not uncommon (reviewed in Halsband & Lange, 2006). Our results can be interpreted as an *effector-specific decision rule*. Here, sensorimotor

mapping would occur between visual neurons and effector-specific integrators, e.g., in parietal cortex, similar to the *specific stimulus-response mapping model*.

However, the degree of specificity we find falls between the two previously reported extremes, as transfer is neither full nor none. Given the co-localization of saccade-specific, reach-specific, and effector-invariant motor neurons in parietal cortex (Calton, Dickinson, & Snyder, 2002; Diomedi, Vaccari, Filippini, Fattori, & Galletti, 2020), it is possible that there is overlap of the populations of neurons encoding motor plans for these two effectors. The degree of overlap would then determine the degree of transfer between the trained and untrained effectors in the transfer session. A possible reason for the colocalization of saccade- and reach-selective neurons in parietal cortex is the need for eye-hand coordination in many real-world tasks. Saccades may always be planned along limb movements for visual guidance, but not executed if fixation is required by the task (Levy, Schluppeck, Heeger, & Glimcher, 2007). Such consistent coactivation could also account for partial transfer results, as saccades would be trained along hand movements in an implicit fashion, perhaps through Hebbian plasticity.

In Experiment 2, we used a spatially-localized visually-guided pointing movement, as compared to the simple left/right joystick tilt response in Experiment 1, to match as well as possible the spatial contingencies of the saccade task. This was done to ensure that the finding of partial transfer in the first experiment was not due to very different spatial transformation between joystick movements and saccades but indeed dependent on the effector: there are considerable differences in parietal and occipital cortex activity associated with simple non-visual actions such as joystick tilt or a button press, compared to visually-guided actions where the effector (the hand, or a joystick cursor) is visible and needs to be directed to a spatial target (e.g., Bernier, Whittingstall, & Grafton, 2017; Breveglieri, Galletti, Dal Bo, Hadjidimitrakis, & Fattori, 2014). In line with this, lesions of parietal cortex influence spatially-contingent reach-and-point responses but do not affect arbitrary stimulus to simple non-spatial motor response transformations (Rushworth, Nixon, & Passingham, 1997). At the same, index finger pointing and saccades activate largely overlapping areas in parietal and prefrontal cortex (Hagler, Riecke, & Sereno, 2007), making the learning transfer question very relevant. Since we find significant SIs in both experiments, this indicates that VPL is indeed effector-specific; however, given that only Experiment 2 included a delay phase and that there was only a numerical but not a statistically significant drop in accuracy when the effector was changed suggests that further research is required to illuminate the role of spatial mapping in this form of specificity.

Classical VPL models in which learning takes place only in visual cortex predict full transfer and cannot easily explain partial specificity of learning effects to effectors. However, visual cortex activity has been shown to be differentially modulated by saccades versus reaches in humans using fMRI (Gallivan,

542 Chapman, Gale, Flanagan, & Culham, 2019; Levy et al., 2007). In rodents, there is evidence that motor
543 effects in visual cortex increase information about the orientation of visual stimuli (reviewed in Pakan,
544 Francioni, & Rochefort, 2018). This suggests that interactions between visual and motor components
545 of VPL could hypothetically already arise in early visual areas but still points to sensorimotor
546 interactions as a relevant contributor to VPL.

547 Because neurons in parietal cortex have been shown to integrate visual information towards a decision
548 bound, the parameters of the DDM are often interpreted to arise in parietal cortex as well (Smith &
549 Ratcliff, 2004). This would fit with the known existence of effector-specific integrators and their
550 colocalization in areas like LIP that from the basis of the newly suggested effector-specific but
551 overlapping integrators model. An increased drift rate is usually interpreted as better sensory input to
552 decision-making process (Petrov et al., 2011; Ratcliff, 2014). However, this interpretation cannot fully
553 account for the effector specificity of the drift rate parameter that we found, because the sensory
554 input does not change when the effector changes. Rather, our results suggest that the drift rate (also)
555 reflects the integration process itself, i.e., the action of the integrator. For example, it is possible that
556 integrators in parietal cortex increase their sensitivity to sensory information, which could further
557 contribute to a fast integration of information towards the decision bound.

558 We find that VPL also has an effect on the decision bound, and that this effect is also (partially) effector-
559 specific. Trainees seem to adjust their decision bound such that faster integration is in a way
560 counteracted by a more conservative decision bound. This speed-accuracy policy results in higher
561 accuracy because more information can be sampled but forgoes the opportunity to save time on task
562 by reaching a decision earlier on the basis of a higher drift rate alone. Previous studies have reported
563 a decreased decision bound (C. C. Liu & Watanabe, 2012) and less variability in the decision bound
564 parameter (Petrov et al., 2011) with VPL. However, differences in task design and instructions, e.g., an
565 emphasis of speed over accuracy, may explain this discrepancy.

566 Our best performing model included decision bounds that varied with the difficulty level of the task.
567 Often, DDMs are instead fitted with a single decision bound that spans all difficulty levels and is
568 thought to reflect a strategic criterion to optimize speed versus accuracy. Balci et al. (2011) and Evans
569 & Brown (2017) suggest that with sufficient practice, subjects can learn to adjust their decision bound
570 to difficulty levels. This is indeed in line with our results, as decision bounds are initially, i.e., in the first
571 session, almost identical across difficulty levels, but differentiate only in later sessions (Fig. 9), speaking
572 to a learning effect. Given this and that the model with decision bounds per difficulty level substantially
573 outperformed the model with a single, fixed decision bound per session (Fig. 6), subject may have
574 developed a sense of time to exploit the available response time window at maximum. Interestingly,
575 the adjustment of *a* to *v* could be transferred to the new effector, suggesting that task strategies are

not effector specific. However, we also note that the effector specificity of the drift rate parameter fully holds in the model with a single decision bound across difficulty levels that only varies between training sessions (see Table S4).

Neurally, narrower decision bounds correlate with an increase in activity at the outset of the integration process in area LIP (Hanks, Kiani, & Shadlen, 2014). Since many neurons in LIP are effector specific, it is possible that the same neurons that may account for the learning effects and the specificity of the drift rate also account for the effects on the decision bound. Our model would thus predict a reduction in spontaneous choice-predictive activity in areas like LIP (Shadlen & Newsome, 2001) with VPL. Alternatively, a central decision bound circuit, e.g., in the basal ganglia (Zhang, 2012), where correlates of VPL have also been found (Guggenmos, Wilbertz, Hebart, & Sterzer, 2016), could account for our results.

Our results are in partial conflict with previous reports on full transfer (Awada et al., 2021) and no transfer between effectors (C. S. Green et al., 2015; Grzeczkowski et al., 2017). As mentioned in the introduction, the latter results can conceivably be explained by the task-specificity of VPL. Several possibilities could account for the discrepancy with the study by Awada and colleagues: the study by Awada and colleagues tested a much smaller sample of subjects (6 and 4 subjects in their Experiment 1 and 2, respectively), which may have been insufficient to reveal partial transfer effects. In fact, the results from Awada and colleagues' Experiment 1 show a small trend towards partial transfer that did not reach statistical significance. However, it is also possible that differences in the stimuli and task account for the divergent results. In contrast to our study on orientation discrimination, Awada and colleagues investigated motion direction discrimination VPL. Motion direction is processed in dorsal stream areas, which are anatomically closer to parietal decision-making areas than neurons encoding orientation, which primarily reside in the ventral stream. Furthermore, parietal areas such as LIP contain intrinsically motion direction sensitive (i.e., sensory) neurons (Fanini & Assad, 2009). Both these factors could ease sensorimotor mapping to multiple colocalized effectors. Further studies with high statistical power explicitly comparing motion direction and orientation discrimination tasks will be able to resolve these alternative interpretations.

Taken together, we find that VPL is not as specific, but also not as generalizable as previously thought. Taking into account different components of the sensorimotor arc revealed additional aspects of VPL beyond the traditional focus on (early) visual areas. Our results also suggest an expanded interpretation of the drift rate as reflecting not only the amount of accumulable information, but also the sensitivity of the integrator itself. Finally, in terms of application of VPL (Deveau & Seitz, 2014), where generalization is key, it may be beneficial to counteract overtraining by varying effectors. Given the new effector-specific learning rule model we propose, more effector variability during training may

enable more generalization of training outcomes (also see Grzeczkowski, Cretenoud, Mast, & Herzog, 2019; Manenti, Dizaji, & Schwiedrzik, 2023).

Acknowledgements

This project was supported by a seed fund from the Leibniz ScienceCampus 'Primate Cognition', Göttingen, Germany (to CMS and IK), and has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 802482, to CMS). CMS is supported by the German Research Foundation Emmy Noether Program (SCHW1683/2-1). The funders had no role in study design, data collection and interpretation, decision to publish, or preparation of the manuscript. The authors declare no competing financial interests.

References

Ahissar, M., & Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. *Trends Cogn Sci, 8*(10), 457-464. doi:10.1016/j.tics.2004.08.011

Aslan, B., & Zech, G. (2005). Statistical energy as a tool for binning-free, multivariate goodness-of-fit tests, two-sample comparison and unfolding. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 537*(3), 626-636. doi:10.1016/j.nima.2004.08.071

Awada, A., Bakhtiari, S., & Pack, C. C. (2021). Visual perceptual learning generalizes to untrained effectors. *J Vis, 21*(3), 10. doi:10.1167/jov.21.3.10

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., & Cohen, J. D. (2011). Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Atten Percept Psychophys, 73*(2), 640-657. doi:10.3758/s13414-010-0049-7

Bernier, P. M., Whittingstall, K., & Grafton, S. T. (2017). Differential Recruitment of Parietal Cortex during Spatial and Non-spatial Reach Planning. *Front Hum Neurosci, 11*, 249. doi:10.3389/fnhum.2017.00249

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spat Vis, 10*(4), 433-436.

Breveglieri, R., Galletti, C., Dal Bo, G., Hadjidimitrakis, K., & Fattori, P. (2014). Multiple aspects of neural activity during reaching preparation in the medial posterior parietal area V6A. *J Cogn Neurosci, 26*(4), 878-895. doi:10.1162/jocn_a_00510

Brooks, S. P., & Gelman, A. (1998). General Methods for Monitoring Convergence of Iterative Simulations. *Journal of Computational and Graphical Statistics, 7*(4), 434-455. doi:10.1080/10618600.1998.10474787

Calton, J. L., Dickinson, A. R., & Snyder, L. H. (2002). Non-spatial, motor-specific activation in posterior parietal cortex. *Nat Neurosci, 5*(6), 580-588. doi:10.1038/nn0602-862

Chang, S. W., Calton, J. L., Lawrence, B. M., Dickinson, A. R., & Snyder, L. H. (2016). Region-Specific Summation Patterns Inform the Role of Cortical Areas in Selecting Motor Plans. *Cereb Cortex, 26*(5), 2154-2166. doi:10.1093/cercor/bhv047

Christopoulos, V. N., Bonaiuto, J., Kagan, I., & Andersen, R. A. (2015). Inactivation of Parietal Reach Region Affects Reaching But Not Saccade Choices in Internally Guided Decisions. *J Neurosci, 35*(33), 11719-11728. doi:10.1523/JNEUROSCI.1068-15.2015

Christopoulos, V. N., Kagan, I., & Andersen, R. A. (2018). Lateral intraparietal area (LIP) is largely effector-specific in free-choice decisions. *Sci Rep, 8*(1), 8611. doi:10.1038/s41598-018-26366-9

Colby, C. L., & Goldberg, M. E. (1999). Space and attention in parietal cortex. *Annu Rev Neurosci, 22*, 319-349. doi:10.1146/annurev.neuro.22.1.319

de Lafuente, V., Jazayeri, M., & Shadlen, M. N. (2015). Representation of accumulating evidence for a decision in two parietal areas. *J Neurosci, 35*(10), 4306-4318. doi:10.1523/JNEUROSCI.2451-14.2015

Deveau, J., & Seitz, A. R. (2014). Applying perceptual learning to achieve practical changes in vision. *Front Psychol, 5*, 1166. doi:10.3389/fpsyg.2014.01166

Dickinson, A. R., Calton, J. L., & Snyder, L. H. (2003). Nonspatial saccade-specific activation in area LIP of monkey parietal cortex. *J Neurophysiol, 90*(4), 2460-2464. doi:10.1152/jn.00788.2002

Diomedi, S., Vaccari, F. E., Filippini, M., Fattori, P., & Galletti, C. (2020). Mixed Selectivity in Macaque Medial Parietal Cortex during Eye-Hand Reaching. *iScience, 23*(10), 101616. doi:10.1016/j.isci.2020.101616

Dosher, B., & Lu, Z. L. (2017). Visual Perceptual Learning and Models. *Annu Rev Vis Sci, 3*, 343-363. doi:10.1146/annurev-vision-102016-061249

Dutilh, G., Krypotos, A. M., & Wagenmakers, E. J. (2011). Task-related versus stimulus-specific practice. *Exp Psychol, 58*(6), 434-442. doi:10.1027/1618-3169/a000111

Evans, N. J., & Brown, S. D. (2017). People adopt optimal policies in simple decision-making, after practice and guidance. *Psychon Bull Rev, 24*(2), 597-606. doi:10.3758/s13423-016-1135-1

Fahle, M. (2004). Perceptual learning: a case for early selection. *J Vis, 4*(10), 879-890. doi:10.1167/4.10.4

Fahle, M., & Morgan, M. (1996). No transfer of perceptual learning between similar stimuli in the same retinal position. *Curr Biol, 6*(3), 292-297. doi:10.1016/s0960-9822(02)00479-7

Fanini, A., & Assad, J. A. (2009). Direction selectivity of neurons in the macaque lateral intraparietal area. *J Neurophysiol, 101*(1), 289-305. doi:10.1152/jn.00400.2007

Feulner, B., Postin, D., Schwiedrzik, C. M., & Pooresmaeili, A. (2021). Previous Motor Actions Outweigh Sensory Information in Sensorimotor Statistical Learning. *eNeuro, 8*(5). doi:10.1523/ENEURO.0032-21.2021

Gallivan, J. P., Chapman, C. S., Gale, D. J., Flanagan, J. R., & Culham, J. C. (2019). Selective Modulation of Early Visual Cortical Activity by Movement Intention. *Cereb Cortex, 29*(11), 4662-4678. doi:10.1093/cercor/bhy345

Gelman, A., & Rubin, D. B. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science, 7*(4), 457-472.

Gershman, S. J., Pesaran, B., & Daw, N. D. (2009). Human reinforcement learning subdivides structured action spaces by learning effector-specific values. *J Neurosci, 29*(43), 13524-13531. doi:10.1523/JNEUROSCI.2469-09.2009

Geweke, J. F. (1991). *Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments*. Retrieved from https://ideas.repec.org/p/fip/fedmsr/148.html

Ghose, G. M., Yang, T., & Maunsell, J. H. (2002). Physiological correlates of perceptual learning in monkey V1 and V2. *J Neurophysiol, 87*(4), 1867-1888. doi:10.1152/jn.00690.2001

Gomez, P., Ratcliff, R., & Childers, R. (2015). Pointing, looking at, and pressing keys: A diffusion model account of response modality. *J Exp Psychol Hum Percept Perform, 41*(6), 1515-1523. doi:10.1037/a0039653

Gottlieb, J., & Goldberg, M. E. (1999). Activity of neurons in the lateral intraparietal area of the monkey during an antisaccade task. *Nat Neurosci, 2*(10), 906-912. doi:10.1038/13209

Green, C. S., Kattner, F., Siegel, M. H., Kersten, D., & Schrater, P. R. (2015). Differences in perceptual learning transfer as a function of training task. *J Vis, 15*(10), 5. doi:10.1167/15.10.5

Green, C. S., Pouget, A., & Bavelier, D. (2010). Improved probabilistic inference as a general learning mechanism with action video games. *Curr Biol, 20*(17), 1573-1579. doi:10.1016/j.cub.2010.07.040

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

Grieco, A., Casco, C., & Roncato, S. (2006). Texture segregation on the basis of contrast polarity of odd-symmetric filters. *Vision Res, 46*(20), 3526-3536. doi:10.1016/j.visres.2006.05.002

Grzeczkowski, L., Cretenoud, A., Herzog, M. H., & Mast, F. W. (2017). Perceptual learning is specific beyond vision and decision making. *J Vis, 17*(6), 6. doi:10.1167/17.6.6

Grzeczkowski, L., Cretenoud, A. F., Mast, F. W., & Herzog, M. H. (2019). Motor response specificity in perceptual learning and its release by double training. *J Vis, 19*(6), 4. doi:10.1167/19.6.4

Guggenmos, M., Wilbertz, G., Hebart, M. N., & Sterzer, P. (2016). Mesolimbic confidence signals guide perceptual learning in the absence of external feedback. *Elife, 5*. doi:10.7554/eLife.13388

Hagler, D. J., Jr., Riecke, L., & Sereno, M. I. (2007). Parietal and superior frontal visuospatial maps activated by pointing and saccades. *Neuroimage, 35*(4), 1562-1577. doi:10.1016/j.neuroimage.2007.01.033

Halsband, U., & Lange, R. K. (2006). Motor learning in man: a review of functional and clinical studies. *J Physiol Paris, 99*(4-6), 414-424. doi:10.1016/j.jphysparis.2006.03.007

Hanks, T., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *Elife, 3*. doi:10.7554/eLife.02260

Hentschke, H., & Stüttgen, M. C. (2011). Computation of measures of effect size for neuroscience data sets. *Eur J Neurosci, 34*(12), 1887-1894. doi:10.1111/j.1460-9568.2011.07902.x

Jaskowski, P., & Sobieralska, K. (2004). Effect of stimulus intensity on manual and saccadic reaction time. *Percept Psychophys, 66*(4), 535-544. doi:10.3758/bf03194899

Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proc Natl Acad Sci U S A, 88*(11), 4966-4970. doi:10.1073/pnas.88.11.4966

Kazennikov, O., Hyland, B., Corboz, M., Babalian, A., Rouiller, E. M., & Wiesendanger, M. (1999). Neural activity of supplementary and primary motor areas in monkeys and its relation to bimanual and unimanual movement sequences. *Neuroscience, 89*(3), 661-674. doi:10.1016/s0306-4522(98)00348-0

Klorfeld-Auslender, S., & Censor, N. (2019). Visual-oculomotor interactions facilitate consolidation of perceptual learning. *J Vis, 19*(6), 11. doi:10.1167/19.6.11

Law, C. T., & Gold, J. I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat Neurosci, 11*(4), 505-513. doi:10.1038/nn2070

Levy, I., Schluppeck, D., Heeger, D. J., & Glimcher, P. W. (2007). Specificity of human cortical areas for reaches and saccades. *J Neurosci, 27*(17), 4687-4696. doi:10.1523/JNEUROSCI.0459-07.2007

Liu, B. H., Huberman, A. D., & Scanziani, M. (2016). Cortico-fugal output from visual cortex promotes plasticity of innate motor behaviour. *Nature, 538*(7625), 383-387. doi:10.1038/nature19818

Liu, C. C., & Watanabe, T. (2012). Accounting for speed-accuracy tradeoff in perceptual learning. *Vision Res, 61*, 107-114. doi:10.1016/j.visres.2011.09.007

MacInnes, W. J. (2017). Multiple Diffusion Models to Compare Saccadic and Manual Responses for Inhibition of Return. *Neural Computation, 29*(3), 804-824. doi:10.1162/NECO_a_00904

Manenti, G. L., Dizaji, A. S., & Schwiedrzik, C. M. (2023). Variability in training unlocks generalization in visual perceptual learning through invariant representations. *Curr Biol, 33*(5), 817-826 e813. doi:10.1016/j.cub.2023.01.011

McKean, J. W., & Schrader, R. M. (1984). A comparison of methods for studentizing the sample median. *Communications in Statistics - Simulation and Computation, 13*(6), 751-773. doi:10.1080/03610918408812413

Oldfield, R. C. (1971). The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia, 9*(1), 97-113. doi:10.1016/0028-3932(71)90067-4

Pakan, J. M., Francioni, V., & Rochefort, N. L. (2018). Action and learning shape the activity of neuronal circuits in the visual cortex. *Curr Opin Neurobiol, 52*, 88-97. doi:10.1016/j.conb.2018.04.020

Pedersen, M. L., Ironside, M., Amemori, K. I., McGrath, C. L., Kang, M. S., Graybiel, A. M., . . . Frank, M. J. (2021). Computational phenotyping of brain-behavior dynamics underlying approach-avoidance conflict in major depressive disorder. *PLoS Comput Biol, 17*(5), e1008955. doi:10.1371/journal.pcbi.1008955

Petrov, A. A., Dosher, B. A., & Lu, Z. L. (2005). The dynamics of perceptual learning: an incremental reweighting model. *Psychol Rev, 112*(4), 715-743. doi:10.1037/0033-295X.112.4.715

Petrov, A. A., Van Horn, N. M., & Ratcliff, R. (2011). Dissociable perceptual-learning mechanisms revealed by diffusion-model analysis. *Psychon Bull Rev, 18*(3), 490-497. doi:10.3758/s13423-011-0079-8

Pilacinski, A., & Lindner, A. (2019). Distinct contributions of human posterior parietal and dorsal premotor cortex to reach trajectory planning. *Sci Rep, 9*(1), 1962. doi:10.1038/s41598-019-39188-0

Poggio, T., Fahle, M., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science, 256*(5059), 1018-1021. doi:10.1126/science.1589770

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol Rev, 85*(2), 59-108. doi:10.1037/0033-295X.85.2.59

Ratcliff, R. (2014). Measuring psychometric functions with the diffusion model. *J Exp Psychol Hum Percept Perform, 40*(2), 870-888. doi:10.1037/a0034954

Ratcliff, R., & Childers, R. (2015). Individual Differences and Fitting Methods for the Two-Choice Diffusion Model of Decision Making. *Decision (Wash DC), 2*(4), 237-279. doi:10.1037/dec0000030

Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science, 9*(5), 347-356. doi:10.1111/1467-9280.00067

Rushworth, M. F., Nixon, P. D., & Passingham, R. E. (1997). Parietal cortex and movement. I. Movement selection and reaching. *Exp Brain Res, 117*(2), 292-310. doi:10.1007/s002210050224

Sagi, D. (2011). Perceptual learning in vision research. *Vision Res, 51*(13), 1552-1566. doi:10.1016/j.visres.2010.10.019

Schoups, A. A., Vogels, R., & Orban, G. A. (1995). Human perceptual learning in identifying the oblique orientation: retinotopy, orientation specificity and monocularity. *J Physiol, 483 (Pt 3)*, 797-810. doi:10.1113/jphysiol.1995.sp020623

Schwiedrzik, C. M., Singer, W., & Melloni, L. (2011). Subjective and objective learning effects dissociate in space and in time. *Proc Natl Acad Sci U S A, 108*(11), 4506-4511. doi:10.1073/pnas.1009147108

Series, P., Latham, P. E., & Pouget, A. (2004). Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nat Neurosci, 7*(10), 1129-1135. doi:10.1038/nn1321

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol, 86*(4), 1916-1936. doi:10.1152/jn.2001.86.4.1916

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci, 27*(3), 161-168. doi:10.1016/j.tins.2004.01.006

Snyder, L. H., Batista, A. P., & Andersen, R. A. (1997). Coding of intention in the posterior parietal cortex. *Nature, 386*(6621), 167-170. doi:10.1038/386167a0

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology), 64*(4), 583-639. doi:10.1111/1467-9868.00353

Székely, G. J., & Rizzo, M. L. (2013). Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference, 143*(8), 1249-1272. doi:10.1016/j.jspi.2013.03.018

Szpiro, S. F., Spering, M., & Carrasco, M. (2014). Perceptual learning modifies untrained pursuit eye movements. *J Vis, 14*(8), 8. doi:10.1167/14.8.8

Szul, M. J., Bompas, A., Sumner, P., & Zhang, J. (2020). The validity and consistency of continuous joystick response in perceptual decision-making. *Behav Res Methods, 52*(2), 681-693. doi:10.3758/s13428-019-01269-3

Szumska, I., van der Lubbe, R. H., Grzeczkowski, L., & Herzog, M. H. (2016). Does sensitivity in binary choice tasks depend on response modality? *Conscious Cogn, 43*, 57-65. doi:10.1016/j.concog.2016.05.005

Vogels, R. (2010). Mechanisms of visual perceptual learning in macaque visual cortex. *Top Cogn Sci, 2*(2), 239-250. doi:10.1111/j.1756-8765.2009.01051.x

Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Front Neuroinform, 7*, 14. doi:10.3389/fninf.2013.00014

Wilke, M., Turchi, J., Smith, K., Mishkin, M., & Leopold, D. A. (2010). Pulvinar inactivation disrupts selection of movement plans. *J Neurosci, 30*(25), 8650-8659. doi:10.1523/JNEUROSCI.0953-10.2010

Zhang, J. (2012). The effects of evidence bounds on decision-making: theoretical and empirical developments. *Front Psychol, 3*, 263. doi:10.3389/fpsyg.2012.00263

Zhang, J., & Rowe, J. B. (2014). Dissociable mechanisms of speed-accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model. *Front Neurosci, 8*, 69. doi:10.3389/fnins.2014.00069