

1 Previous motor actions outweigh sensory 2 information in sensorimotor learning

3
4 Barbara Feulner^{1,2*}, Danilo Postin^{1,3}, Caspar M. Schwiedrzik^{4,5,6}, Arezoo Pooresmaeili^{1,6*}

5 ¹ *Perception and Cognition Lab, European Neuroscience Institute Göttingen- A Joint Initiative of the*
6 *University Medical Center Göttingen and the Max-Planck-Society, Grisebachstrasse 5, 37077 Göttingen,*
7 *Germany*

8 ² *Bioengineering Department, Imperial College London, Exhibition Rd, South Kensington, London SW7*
9 *2BU, United Kingdom*

10 ³ *Department of Psychiatry, School of Medicine & Health Sciences, Carl von Ossietzky University of*
11 *Oldenburg, Hermann-Ehlers-Str. 7, 26160 Bad Zwischenahn, Germany*

12
13 ⁴ *Neural Circuits and Cognition Lab, European Neuroscience Institute Göttingen- A Joint Initiative of the*
14 *University Medical Center Göttingen and the Max-Planck-Society, Grisebachstrasse 5, 37077 Göttingen,*
15 *Germany*

16 ⁵ *Perception and Plasticity Group, German Primate Center – Leibniz Institute for Primate Research,*
17 *Kellnerweg 4, 37077 Göttingen, Germany*

18 ⁶ *Leibniz Science Campus Primate Cognition, Göttingen, Germany*

19

20 * Correspondences should be made to: b.feulner18@imperial.ac.uk or a.pooresmaeili@eni-g.de

21 **Short title: Learning of spatial priors through saccadic eye movements**

22 Number of pages: 30

23 Number of figures: 6

24 Conflict of interests: The authors declare no competing interests.

25 Acknowledgements: We thank Tomke Schoss for her help with the pilot study. This work was initially
26 supported by a seed fund grant from Leibniz Science Campus Primate Cognition, Göttingen, Germany to
27 CMS and AP and continued by an ERC Starting Grant (no: 716846) to AP. CMS was supported by an
28 Emmy Noether Grant from the German Research Foundation (SCHW1683/2-1).

29 Authors' contribution: BF and AP conceptualized the project. BF, CMS and AP designed the task. BF
30 implemented the experiment. DP collected the data. BF analyzed the data. BF and AP interpreted the results
31 and wrote the first draft of the manuscript. All authors revised the manuscript. AP and CMS acquired
32 funding.

33 **Abstract**

34 Humans can use their previous experience in form of statistical priors to improve decisions. It is
35 however unclear how such priors are learned and represented. Importantly, it has remained elusive
36 whether prior learning is independent of the sensorimotor system involved in the learning process
37 or not, as both modality-specific and modality-general learning have been reported in the past.
38 Here, we used a saccadic eye movement task to probe the learning and representation of a spatial
39 prior across a few trials. In this task, learning occurs in an unsupervised manner and through
40 encountering trial-by-trial visual hints drawn from a distribution centered on the target location.
41 Using a model-comparison approach, we found that participants' prior knowledge is largely
42 represented in the form of their previous motor actions, with minimal influence from the previously
43 seen visual hints. By using two different motor contexts for response (looking either at the
44 estimated target location, or exactly opposite to it), we could further compare whether prior
45 experience obtained in one motor context can be transferred to the other. Although learning curves
46 were highly similar, and participants seemed to use the same strategy for both response types, they
47 could not transfer their knowledge between contexts, as performance and confidence ratings
48 dropped to naïve levels after a switch of the required response. Together, our results suggest that
49 humans preferably use the internal representations of their previous motor actions, rather than past
50 incoming sensory information, to form statistical sensorimotor priors on the timescale of a few
51 trials.

52

53 *Keywords: learning, sensorimotor, prior, probabilistic, saccade*

54

55

56

57

58

59

60 Introduction

61 We often have to make decisions based on sparse and uncertain sensory information. Previous
62 research has shown that in these cases humans use Bayesian inference where the current sensory
63 information (likelihood) and the previously acquired knowledge (priors) are integrated, each
64 weighted by their respective uncertainty [1,2]. While the majority of previous studies have
65 examined whether the perceptual and sensorimotor decisions follow the rules of a Bayesian
66 framework [1,3], less emphasis has been placed on understanding how likelihoods and especially
67 statistical priors are learned and represented in the first place. A number of elegant recent studies
68 have tried to bridge this gap by investigating how people learn likelihoods [4] and priors [5] to
69 perform Bayesian computations. Interestingly, the timescale of the two types of learning varied
70 vastly, with fast learning of likelihood but slow learning of prior distributions. It remains unknown
71 why such an asymmetry should exist, as theoretically both types of learning are equivalent. It has
72 been hypothesized that learning about the likelihood versus learning about the prior involves
73 different neural mechanisms, potentially hinting to the fact that their respective distributions might
74 be represented in different regions of the brain [6].

75 Learning of statistical priors is itself not a uniform process as it shows dependencies on the specific
76 context where the learning occurs. In Bayesian framework, priors are a form of abstract knowledge
77 [7–9], which can be generalized across different contexts. However, previous findings regarding
78 the generalization of statistical priors have been mixed. Some studies have shown that statistical
79 learning of priors is very narrow-band and context/modality-specific in perceptual [10] and
80 sensorimotor domains [11,12], thus preventing learned information to transfer to different
81 contexts/modalities. Other studies, on the other hand, provided evidence for generalization [4,13],
82 although generalization, in some instances, seemed to occur differently for different parameters of
83 a statistical distribution; e.g. the mean and the variance of a distribution [14]. The finding that
84 some aspects of learning could generalize, while others could not, was confirmed by a recent study
85 showing that, for instance, in Bayesian time estimation, priors can be generalized across stimuli,
86 but not motor actions [15].

87 Therefore, despite an increasing number of studies testing generalization and transfer, the exact
88 rules determining generalizability remain unclear. One potential reason for these seemingly
89 contradictory results is a lack of a formal definition of *what* is learned. It has been argued that
90 when learning is not generalized, a *policy* (i.e. a specific rule for action) rather than *knowledge* (i.e.

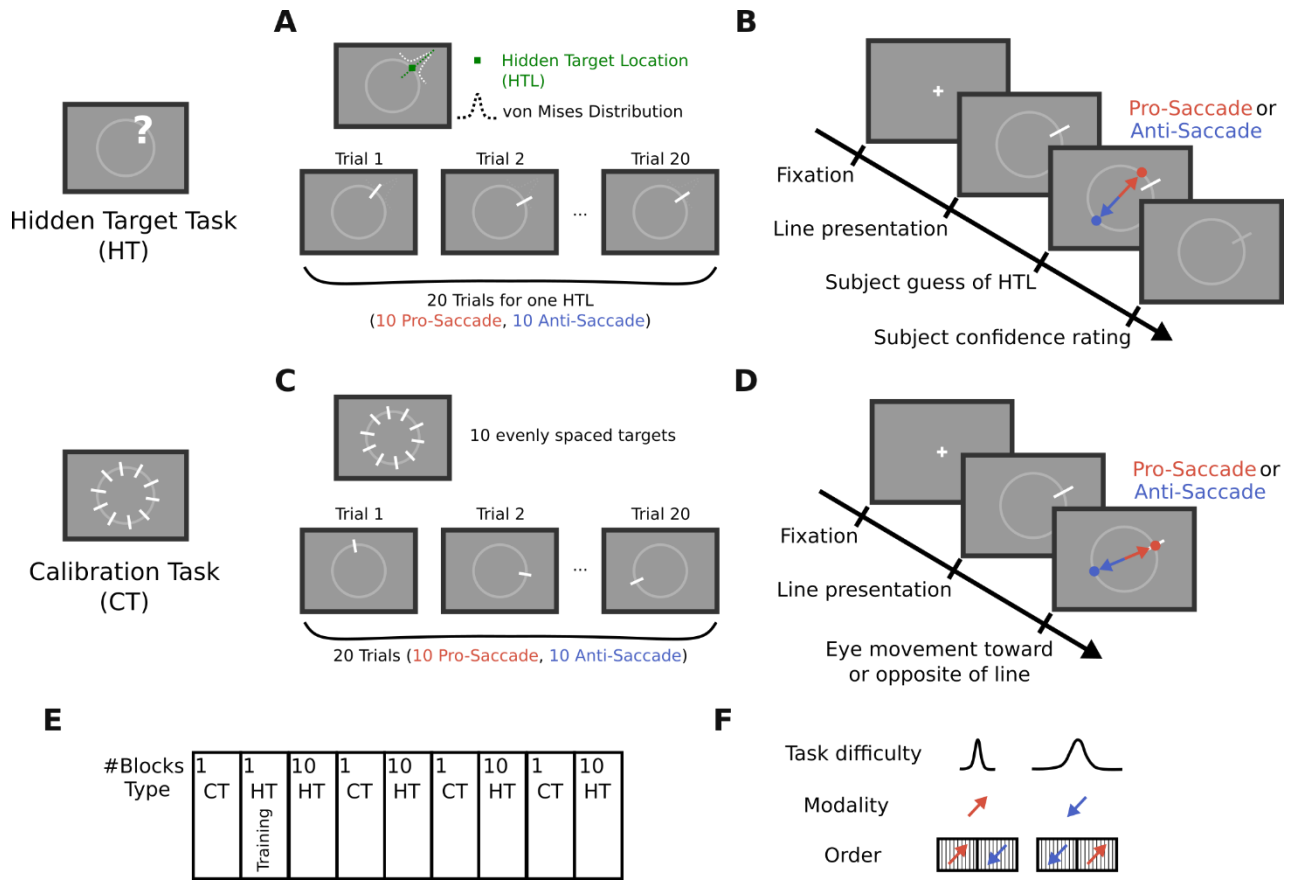
91 abstract and context-independent information) is acquired through learning [16]. However, it is
92 not clear what features of the learning dynamics determine whether a *policy* or *knowledge* is
93 acquired during encounters with the learned information.

94 To investigate learning and generalizability of a prior distribution, we employed a saccadic eye
95 movement task similar to the design of a previous study [17], where participants had to learn to
96 locate a hidden target. The location of the hidden target corresponded to the mean of a circular
97 normal distribution. In each trial, a visual hint sampled from the underlying distribution was shown
98 and participants indicated their current estimate with a saccadic eye movement, looking either
99 towards (pro-saccade) or to the exact opposite direction of the target (anti-saccade). To
100 successfully estimate the hidden target location, participants had to combine information across
101 multiple trials. This design allowed us to investigate whether participants formed their prior
102 knowledge by combining the visual information or by combining the previous motor actions across
103 time, under different saccadic response contexts. Our results from two experiments indicate that
104 sensorimotor learning of a spatial prior in both response contexts is largely guided by previous
105 motor plans, rather than by previous sensory input in form of visual hints. Despite the high degree
106 of similarity of pro- and anti-saccades in their learning behavior, suggesting a motor-independent
107 learning algorithm, the learned prior in one context did not generalize to the other. We propose
108 that the lack of transfer between the two contexts is a natural consequence of their shared learning
109 algorithm in which previous motor actions outweigh sensory information.

110 **Results**

111 To investigate the dynamics of sensorimotor learning of a spatial prior and its dependence on the
112 response modality, we designed an experiment where participants had to find a hidden target and
113 indicate their guess by either a pro- or an anti-saccade. Participants learned the location of each
114 hidden target within twenty trials, of which a block of ten trials required pro-saccades and the other
115 block of ten trials required anti-saccades as the response modality (**Fig.1 A**). This design allowed
116 us to probe whether the learning dynamics shows dependencies on the response modality, hence
117 being modality-specific. We also used another task, referred to as the calibration task, to estimate
118 each participant's motoric noise during the visually driven execution of pro- or anti-saccades. In
119 contrast to the calibration task, in the hidden target task the main error source is the uncertainty
120 regarding the hidden target location. As the same motor system is used in both the hidden target

121 task as well as the calibration task, we assumed that the motor noise affecting participants’
 122 performance in both tasks is equal. As the motor noise in the calibration task is not time-dependent,
 123 we assumed that all time-dependent performance improvement during the hidden target task
 124 reflected statistical learning, defined as the reduction in the uncertainty regarding the location of
 125 the hidden target.



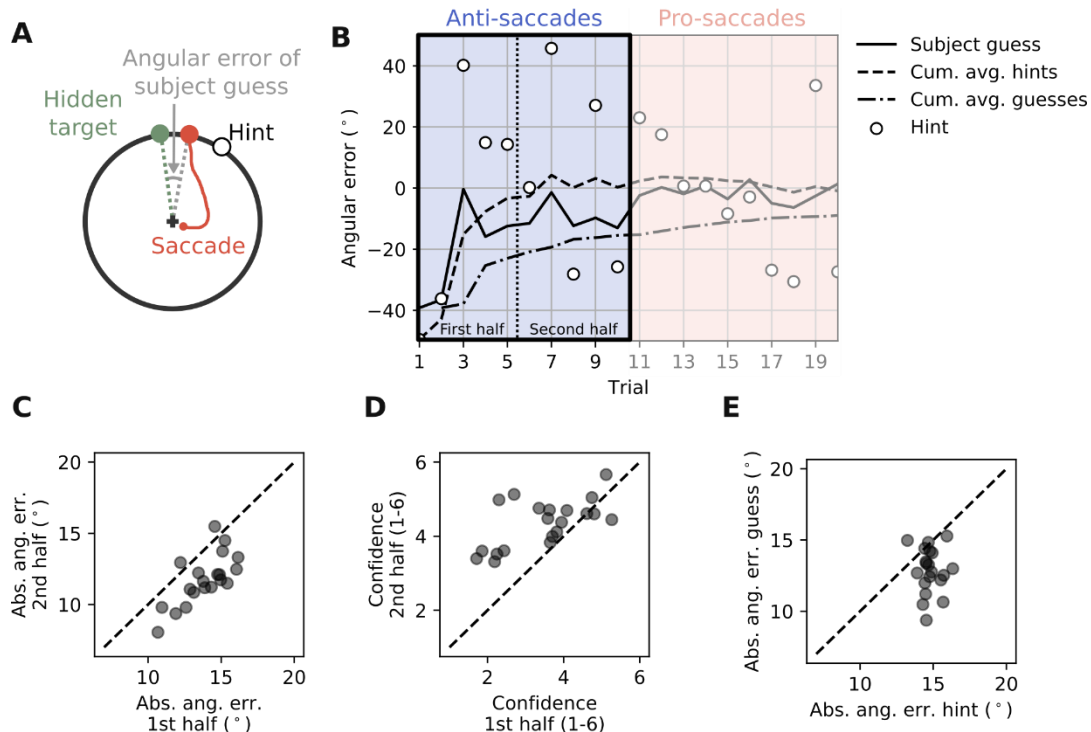
126

127 **Figure 1:** Experimental design of the hidden target task employed to study the statistical learning of a spatial prior
 128 in two different visuo-motor contexts. (A-B) Main task of the experiment. (A) Participants were told to estimate the
 129 location of a ‘hidden treasure’ on a ring by observing and combining information provided by the visual hints across
 130 trials. The hidden target location was defined as the mean of a von Mises distribution and the hints, presented at each
 131 trial, were samples drawn from this underlying distribution. Participants had twenty trials to estimate the location of
 132 the hidden target, after which a new hidden target had to be found. Participants used their gaze to indicate their
 133 responses. (B) Each trial started with a fixation period, after which the hint was presented, and participants had to
 134 indicate their guess about the location of the hidden target by either looking at it (pro-saccade response) or by looking
 135 exactly opposite to it (anti-saccade response). In half of the trials (i.e. consecutive 10 trials), participants had to use
 136 pro-saccades, and in the other half they used anti-saccades, with a randomized order across blocks. (C-D) Calibration
 137 task used to estimate the motoric error of each participant for pro- and anti-saccades. Participants had to directly look
 138 either at the lines (pro-saccade response) or exactly opposite to the lines (anti-saccade response). (E) Block-design of
 139 the experiment. (F) We compared learning across two levels of difficulty and two different response types. Task
 140 difficulty was varied by changing the standard deviation of the von Mises distribution. Finally, we tested whether
 141 knowledge could be transferred from one visuo-motor context to the other. For this we also varied the order of pro-
 142 saccade and anti-saccade responses across blocks.

143

144 **Participants successfully accumulate information and learn on a short time scale**

145 To establish that participants were in general able to learn on a short timescale, we initially focused
 146 on the first ten trials of each hidden target block (**Fig.2 B**). In this case, in all ten trials participants
 147 responded by using the same modality, either exclusively by pro-saccades, or exclusively by anti-
 148 saccades. To quantify performance, we calculated the absolute angular error between the saccade
 149 endpoint and the hidden target location (**Fig.2 A**). By comparing the average absolute angular error
 150 in the first five trials with the average absolute angular error in the last five trials, we found that
 151 most of the participants were able to improve their estimates of the target location (i.e. their
 152 guesses) during this short timescale (paired t-test: $t=7.25$, $p<0.0001$, $N=20$) (**Fig.2 C**). In line with
 153 performance, participants' confidence about the accuracy of their guesses increased during the ten
 154 trials (paired t-test: $t=-4.39$, $p=0.0003$, $N=20$) (**Fig.2 D**). Additionally, the absolute angular error
 155 of participants' guesses was lower than the absolute angular error of the visual hints, meaning that
 156 participants' guesses were closer to the center of the von Mises distribution compared to the
 157 presented visual hints (paired t-test: $t=4.92$, $p=0.0001$, $N=20$) (**Fig.2 E**). This shows that
 158 participants were able to combine information across trials and thereby improve their estimates of
 159 the target location, rather than just following the current visual hint. Hence, we can conclude that
 160 participants showed some form of statistical learning during the first ten trials of the hidden target
 161 task.



162

163 **Figure 2:** *Participants successfully learn the most probable location of the target on a short time scale.* (A) The
164 angular difference between the participants' guess and the true location of the hidden target was used to measure
165 learning. (B) Example block. To test learning, we compared the performance in trial 1-5 (first half) to the performance
166 in trial 6-10 (second half). (C) The absolute angular error in the second half is lower than in the first half (paired t-
167 test: $t=7.25$, $p<0.0001$, $N=20$). (D) Participants' confidence is higher in the second half than in the first half (paired t-
168 test: $t=-4.39$, $p=0.0003$, $N=20$). (E) The absolute angular error of participants is lower than the absolute angular error
169 of the visual hints, i.e., participants' guesses are closer to the center of the von Mises distribution compared to the
170 presented visual hints (paired t-test: $t=4.92$, $p=0.0001$, $N=20$).

171 Learning curves are stereotypic across response modalities

172 To work out whether statistical learning is modality-dependent or not, we contrasted the learning
173 performance in pro- versus anti-saccade trials in the hidden target task, as well as in the calibration
174 task. First, we quantified the mean and the standard deviation of the respective angular error
175 distributions, pooled across participants (**Fig.3 A&B and Table 1**). We did not find a significant
176 bias towards a specific direction, either clockwise or counter-clockwise (one-sample t-test; mean
177 different from zero; $N=20$; **Table 1**), for any of the distributions.

178

Task	Modality	Mean	SD	p-value
Calibration	Pro-saccade	-0.1°	6.3°	0.48
Calibration	Anti-saccade	0.7°	9.5°	0.15
Hidden target	Pro-saccade	1.1°	16.9°	0.11
Hidden target	Anti-saccade	0.1°	18.4°	0.99
Hidden target	Hints	0.1°	20.6°	0.60

179 **Table 1:** *Angular error distribution for calibration and the hidden target task.* Mean and standard deviation for
180 distributions shown in Fig. 3A&B. We tested whether the mean of either of the five distributions was significantly
181 different from zero (one-sample t-test; $N=20$). The corresponding p-values are shown in the last column.

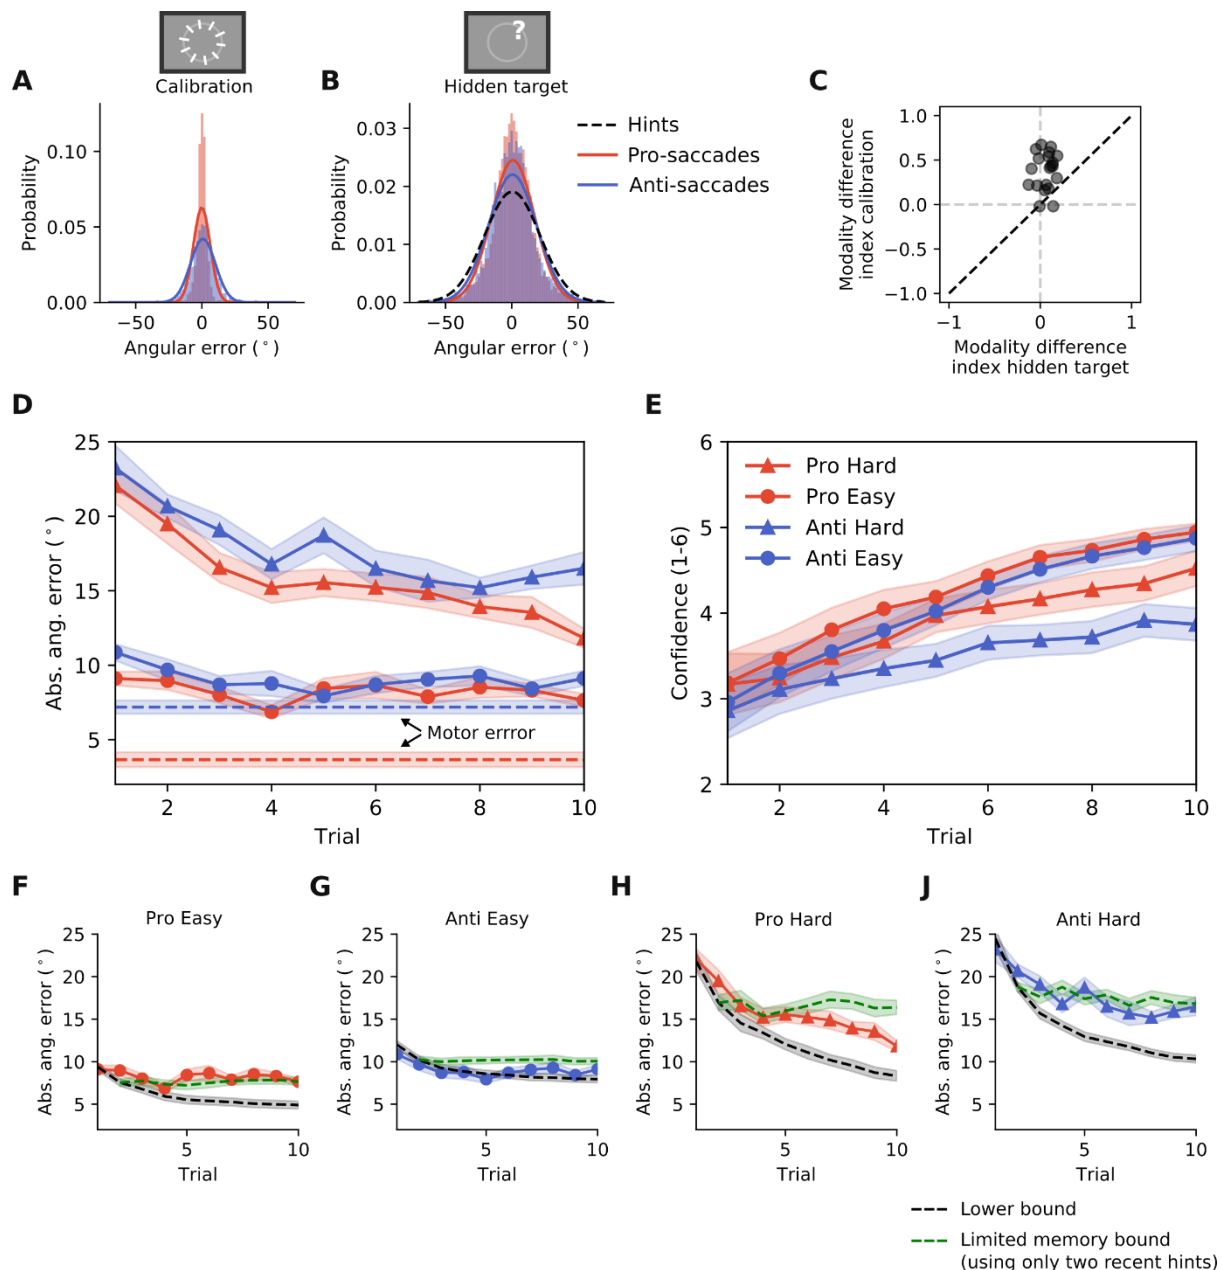
182 For both modalities, the standard deviation of the angular error was higher in the hidden target task
183 than in the calibration task (pro-saccades: paired t-test, $t=-11.8$, $p<0.0001$, $N=20$ / anti-saccades:
184 paired t-test, $t=-10.9$, $p<0.0001$, $N=20$), which is plausible as in the former task participants were
185 more uncertain of the target location. Furthermore, the standard deviations of pro- and anti-
186 saccades were only significantly different during the calibration task but not during the hidden
187 target task (calibration: paired t-test, $t=-3.8$, $p=0.0011$, $N=20$ / hidden target task: paired t-test, $t=-$
188 1.7 , $p=0.1108$, $N=20$). To quantify this difference further, we calculated a modality difference
189 index (described in Materials and Methods section), which was significantly higher in the

190 calibration task compared to the hidden target task (paired t-test, $t=-6.44$, $p<0.0001$, $N=20$,
191 **Fig.3 C**). These results provide the first evidence that statistical learning is modality-independent.

192 To see whether there is a difference between pro- and anti-saccades depending on the difficulty of
193 the task, we next looked at the influence of both experimental factors (pro-/ anti-saccade response,
194 broad/narrow hint distribution) on participants' performance. To this end, we conducted a linear
195 mixed model analysis of the single trial data (i.e., the absolute angular error), with *task difficulty*
196 and *response type* as fixed effects and participant identity as random effect. Using this statistically
197 more powerful approach, we found a main effect of task difficulty ($t=20.6$, $p<0.0001$; contrast:
198 easy-hard= $-8.28^\circ\pm 0.27^\circ$ (SE), $t=-31.2$, $p<0.0001$), as well as response type ($t=2.0$, $p=0.0469$;
199 contrast: pro-anti= $-1.32^\circ\pm 0.27^\circ$ (SE), $t=-5.0$, $p<0.0001$). Additionally, an interaction effect was
200 found between the task difficulty and response type ($t=2.2$, $p=0.0299$), indicating that there is only
201 a difference between pro- and anti-saccade responses in the hard task condition (hard_pro-
202 hard_anti= $-1.90^\circ\pm 0.38^\circ$ (SE), $t=-5.035$, $p<0.0001$; easy_pro-easy_anti= $-0.74^\circ\pm 0.37^\circ$ (SE), $t=-$
203 1.988 , $p=0.1927$). The same pattern was found when we examined participants' confidence
204 ratings: participants were more confident in pro-saccade trials and in the easier task condition
205 (main effect *response*: $t=-4.2$, $p<0.0001$, contrast: pro-anti= 0.28 ± 0.03 , $t=10.3$, $p<0.0001$ / main
206 effect *difficulty*: $t=-9.25$, $p<0.0001$, contrast: easy-hard= 0.47 ± 0.03 , $t=17.4$, $p<0.0001$ / interaction
207 effect: $t=-4.4$, $p<0.0001$). These results show that there is a small influence of the used visuo-motor
208 modality on participants' performance and confidence in the hidden target task, if the hint
209 distribution is broad (hard task condition). As shown in **Fig.3 A-C**, this difference is much smaller
210 than the difference we observed in the calibration task.

211 To investigate whether the response type had an influence on the statistical learning itself, we
212 included the trial number as another fixed effect in our linear mixed model (as a continuous
213 variable). If the response type influenced participants' performance in a trial-independent manner,
214 it would act as a general offset. In contrast, if statistical learning were modality-dependent, we
215 would expect an interaction between the trial number and the response type (pro- versus anti-
216 saccades). Visual inspection of the learning curves for pro- and anti-saccade responses suggested
217 the former, and the main effect influencing the shape of the learning curves seemed to be the task
218 difficulty (**Fig.3 D-E**). For participants' performance, measured as the absolute angular error, we
219 found a significant main effect of difficulty ($t=21.4$, $p<0.0001$) and a significant interaction
220 between trial number and difficulty ($t=-7.807$, $p<0.0001$), but no significant interaction between

221 trial number and response type ($t=0.650$, $p=0.516$). Further analysis of the interaction effect
 222 showed that only for blocks with high task difficulty there is a significant effect of trial number
 223 (hard: slope= -0.84 ± 0.08 , $t=-10.57$, $p<0.0001$ / easy: slope= -0.12 ± 0.08 , $t=-1.50$, $p=0.13$). This
 224 result was further supported by the analysis of the confidence ratings, as participants' confidence
 225 time course (**Fig. S1**) was also mainly dominated by the influence of task difficulty (interaction
 226 effect *trial x difficulty*: $t=-8.858$, $p<0.0001$) and not by the response type (interaction effect *trial x*
 227 *response*: $t=-1.006$, $p=0.315$). Together, these results show that both performance and confidence
 228 increase in a modality-independent manner during learning of a spatial prior.



229

230 **Figure 3:** *Learning curves are stereotypic across response modalities.* (A) The distribution of the angular error for
231 pro- and anti-saccade response in the calibration task. (B) The distribution of the angular error for pro- and anti-
232 saccade response in the hidden target task. (C) Comparison of pro-/anti-saccade performance difference between the
233 hidden target task and the calibration task. (D) Time course of the absolute angular error for each of the four different
234 conditions (two response types x two difficulties). Here and in the following panels, except stated otherwise, shaded
235 areas represent the standard error of the mean (N=20) (E) Time course of the confidence ratings for each of the four
236 different conditions. (F-J) Participants' learning curves compared to the lower bound and the limited memory bound.
237 The lower bound is given by taking the cumulative average of all hints presented so far and adding the error due to
238 the noise in the motor plans, estimated from the calibration task. The limited memory is similarly calculated, except
239 that only the latest two hints are used.

240 **Performance is suboptimal**

241 Since at each point in time, participants have only seen a limited number of samples from the
242 underlying distribution of the target location (in form of visual hints), it is theoretically impossible
243 to correctly estimate the distribution's mean, i.e., the hidden target location in our task. In principle,
244 only for an indefinite number of samples, the sample mean equals the population mean. For the
245 hidden target task, this means that participants could theoretically only reach zero absolute angular
246 error if an infinite number of trials were seen. Taking this into account, we can compare
247 participants' performance to the time course of the theoretical statistical uncertainty due to seeing
248 a limited number of samples (i.e., sampling error, see also Materials and Methods section). This
249 provides a lower bound on participants' uncertainty, indicated by the variance of angular errors
250 σ^2 :

$$251 \quad 1) \quad \sigma^2 = \sigma_{motor}^2 + \sigma_{sample}^2 = \sigma_{motor}^2 + \frac{\sigma_{dist}^2}{n}$$

252 where n is the number of visual hints seen so far, σ_{dist}^2 is the variance of the von Mises distribution
253 and σ_{motor}^2 is the motoric noise measured in the calibration task. We can then calculate the
254 theoretical optimal performance in terms of the mean absolute angular error from this optimal
255 variance estimate (see Eq 2 in Materials and Methods section). Comparing participants'
256 performance to this lower bound shows that for the first trial in each hidden target block,
257 participants perform as well as they possibly can (**Fig.3 F-J**), independent of the task difficulty or
258 the used response modality – pro- or anti-saccades. However, as soon as they need to combine
259 multiple samples to estimate the hidden target location, they perform suboptimal (**Fig.3 F-J**, black
260 dashed line is optimal). We hypothesized that the reason for this suboptimal behavior is limited
261 memory, as to perform optimally would mean to remember every single hint which has been
262 presented so far. So instead of calculating the cumulative mean of all visual hints presented so far
263 (which results in optimal performance), we calculated the performance estimate for an alternative

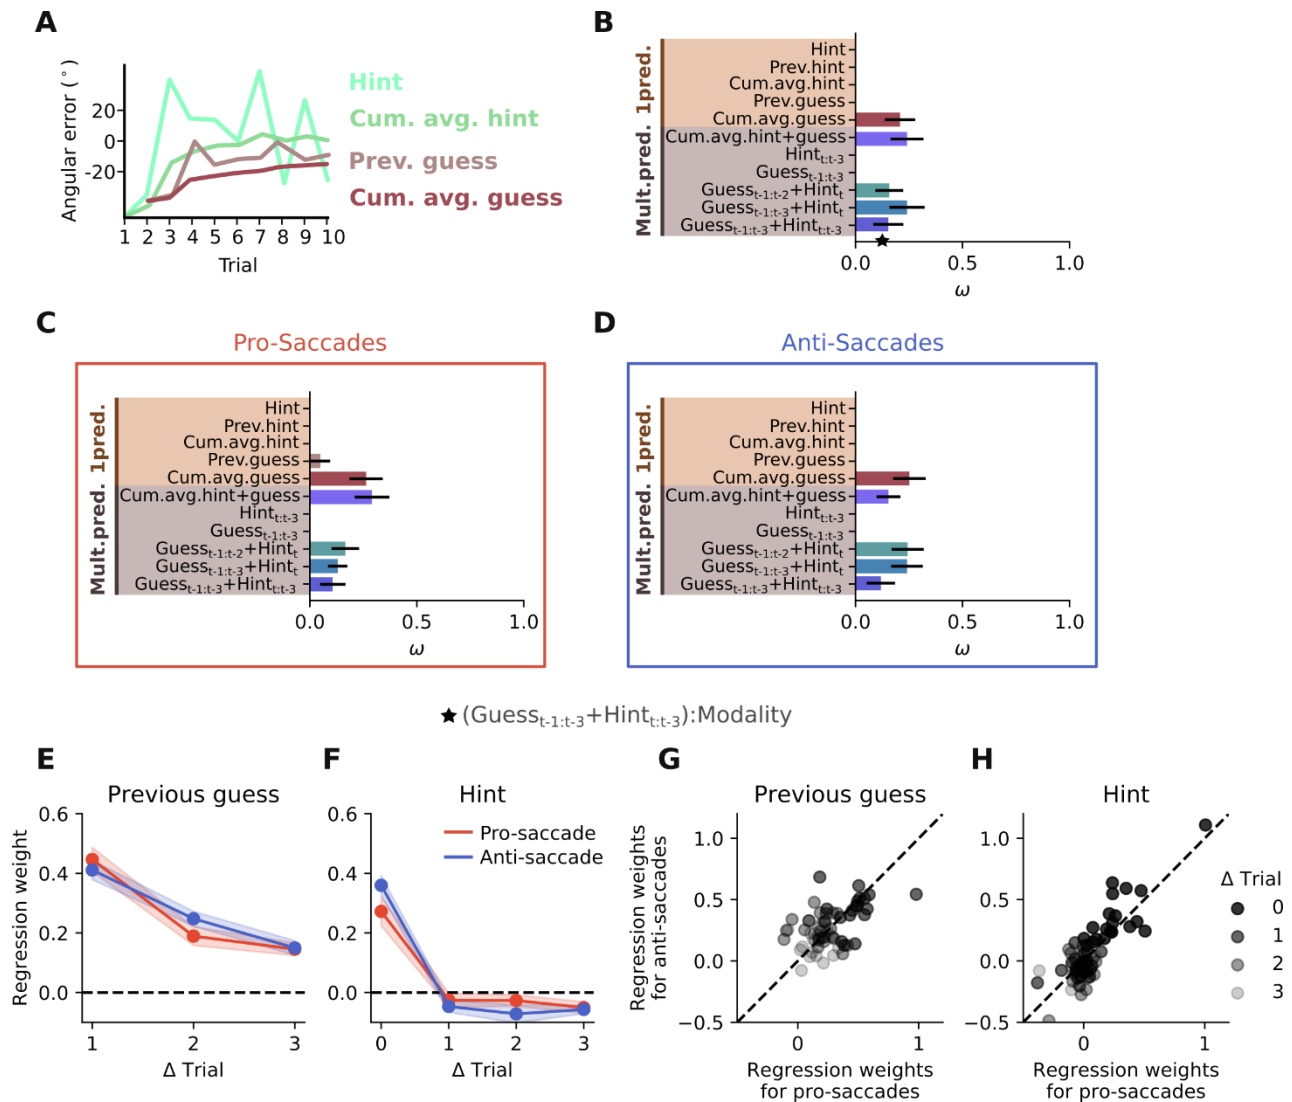
264 strategy where subjects only average the latest two visual hints (**Fig.3 F-J** green dashed line). This
265 limited memory estimate could describe participants' performance qualitatively better, as it
266 captured the shallower learning curves for both response types. In summary, the analysis of
267 different aspects of performance did not show large differences in statistical learning between pro-
268 and anti-saccades.

269 **Strategy is modality-independent**

270 So far, we have seen that statistical learning in the pro- and in the anti-saccade context is similar
271 in terms of the absolute performance and the shape of the learning curve. Next, we wanted to test
272 whether participants use different strategies in the pro- versus in the anti-saccade blocks. The first
273 candidate strategy would be to only look at the visual hints in each trial, which would mean that
274 no learning is happening, and that participants' behavior is only visually-driven. Instead, the
275 optimal strategy, which would result in the lower bound we calculated before, would be to
276 calculate the cumulative average of every hint seen so far in each trial. Besides taking into account
277 the hints, visually presented to the participants, we can also imagine that the behavior is driven by
278 an internal state that promotes looking close to where one has been looking before, i.e., to follow
279 previous guesses (**Fig.4 A**). To test these different hypotheses, we fitted several single predictor
280 linear regression models, for each proposed strategy. Through a model comparison (see Materials
281 and Methods for the details) we found that the best single predictor model is the cumulative
282 average of previous guesses (**Fig.4 B**). Splitting the data into pro- and anti-saccade blocks and
283 repeating the analysis showed that the best single predictor model does not depend on the used
284 response type (**Fig.4 C-D**). Thus, this provides another piece of evidence that the statistical
285 learning happening in our task is independent of the response type.

286 In a second step, we also tested several multi-predictor models to get a more detailed view on
287 participants' strategy. We tested multiple combinations of past guesses, and past and present visual
288 hints. Models based on the previous three guesses or the current and previous three visual hints
289 did not perform as well as the best single predictor model. In contrast, models which relied on
290 combinations of external and internal information, i.e., the visual hints and previous guesses,
291 performed as well as the best single predictor model, even if they only looked three time steps in
292 the past (**Fig.4 B**). Again, splitting the data into pro- and anti-saccade blocks did not affect the
293 main trend (**Fig.4 C-D**).

294 Lastly, to find the exact weighting participants put on their previous guesses and the current and
295 previous hints, depending on the used response type, we fitted a model with all three previous
296 guesses, the current and three previous hints, and the response type as predictors. We found that
297 participants used external information only from the current trial, ignoring the hints from previous
298 trials (**Fig.4 F**). Instead, to combine information across trials they relied on their internal
299 estimations from the past (**Fig.4 E**). We obtained similar results when models that only included
300 either the past guesses or the past visual hints were tested (**Fig. S5**). Including the response type in
301 the model allowed to estimate separate regression weights for pro- and anti-saccades. Again, we
302 didn't find any significant difference in the weighting participants put on their own guesses versus
303 given hints, depending on whether they use pro- or anti-saccades for response (**Fig.4 E-H**; paired
304 t-test: n.s.). In summary, we found that participants used the same strategy, regardless of the
305 response type, to solve the task, which provides further evidence for the modality-independent
306 learning hypothesis.



307

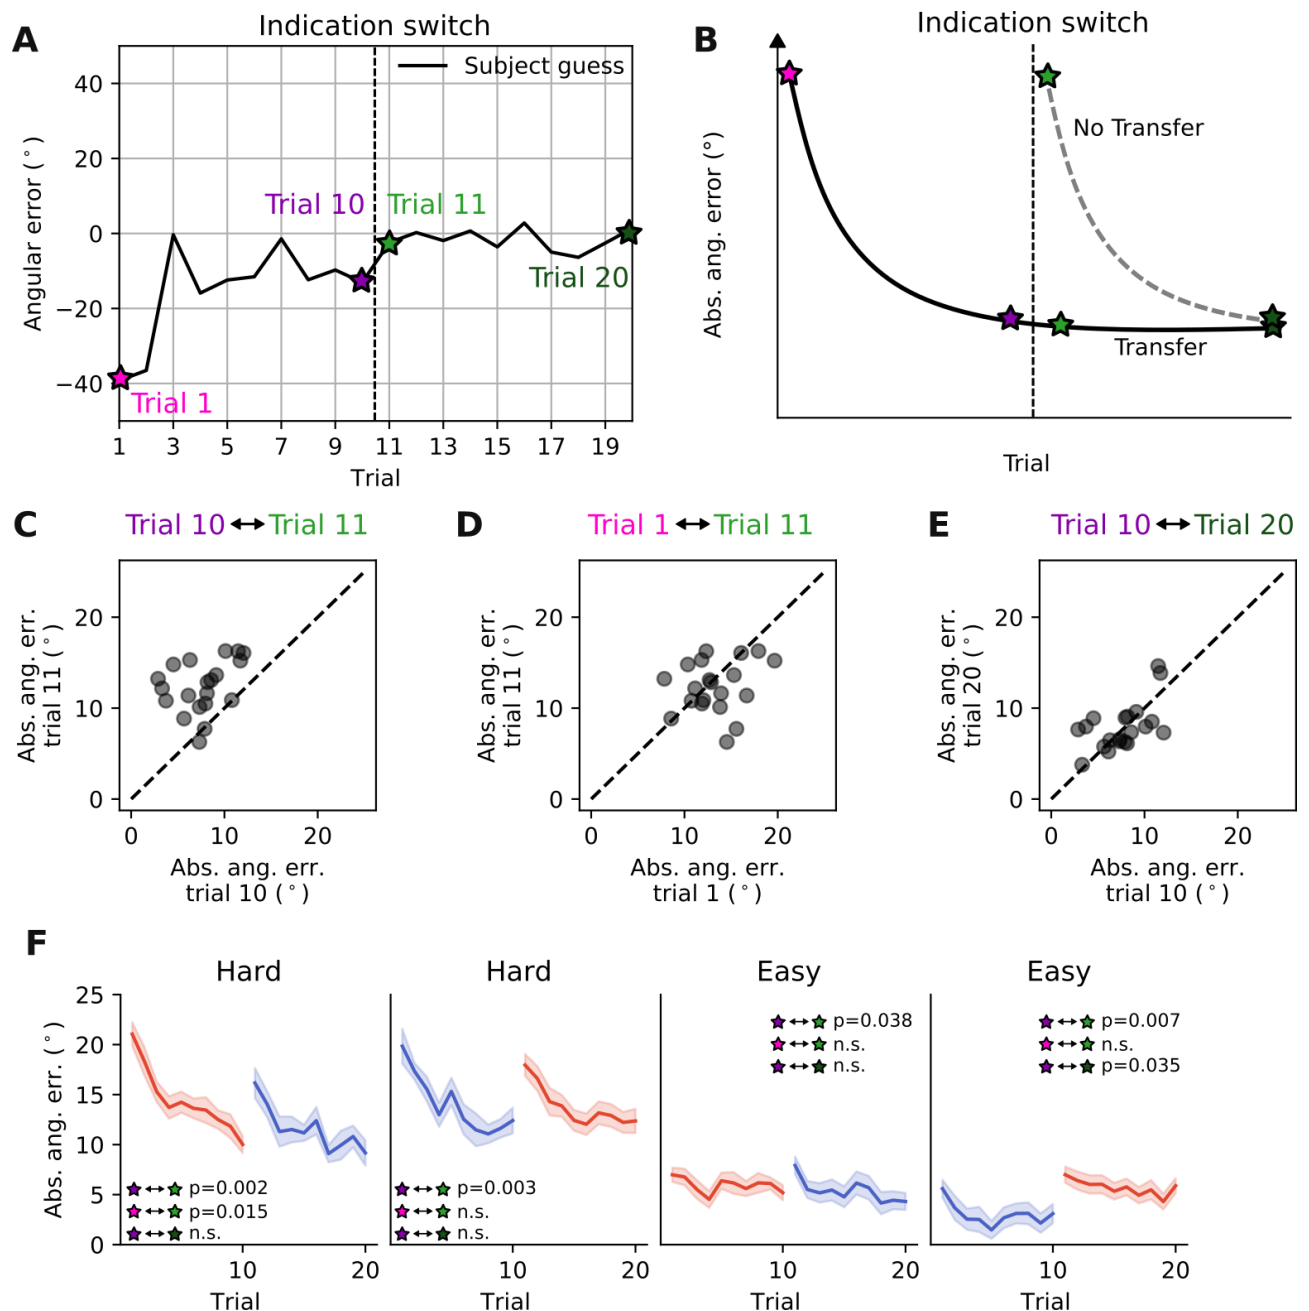
308 **Figure 4: Strategy is modality-independent.** (A) Different predictors used to explain the participants' single trial
 309 estimates. (B) Model comparison between various single and multiple predictor models. Shown are the weights ω ,
 310 which represent the probability that a model is the best among the ones considered. (C-D) Same as (B) but performed
 311 on two different data sets, one consisting only of pro-saccade response trials, the other consisting only of anti-saccade
 312 response trials. (E-H) Regression weights for a model including participants' last three guesses and the current and
 313 last three visual hints. (E) Regression weights put on the last three guesses. (F) Regression weights put on the current,
 314 as well as the last three hints. (G-H) Participants put similar weight on guesses and hints in pro- and anti-saccade
 315 response trials (paired t-test; n.s.).

316

317 Drop in performance after response switch

318 Despite the similarity in performance and observed strategy, hinting at a general algorithm used
 319 for statistical learning in our task, it is still possible that learning happens for each visuo-motor

320 modality in a very specific manner and that both just look similar in terms of performance and
321 strategy. In this case, it would not be possible to generalize between modalities. To test this, we
322 included trial eleven until twenty where participants had to continue looking for the same hidden
323 target (and were also instructed that all twenty trials belong to one hidden target), but had to use
324 the other visuo-motor modality than in trial 1 until 10 (**Fig.5 A**). We considered two alternative
325 hypotheses. The first states that participants learn in a modality-independent fashion and store the
326 acquired knowledge in an abstract form, which allows a complete transfer of previous experience
327 to a new response modality after a switch (**Fig.5 B**, black solid line). The second hypothesis would
328 be that participants perform trial eleven until twenty as if there was no previous experience, which
329 would suggest a modality-specific implementation of the learned knowledge (**Fig.5 B**, gray dashed
330 line). To test which of these hypotheses is true, we compared the performance in trial 10 (last trial
331 before response switch) to the performance in trial 11 (first trial after response switch). We found
332 that the performance is significantly worse in trial 11 than trial 10 (paired t-test, $t=-6.5$, $p<0.0001$,
333 **Fig.5 C**). The performance decrease in trial 11 was consistent across all 4 conditions (**Fig.5 F**).
334 This shows that there is no direct transfer of knowledge across response types. After ruling out the
335 modality-general implementation, we wanted to test if there is any measurable interaction between
336 the first ten trials and the second ten trials. For this, we compared the starting points of each
337 learning curve – trial 1 and trial 11 (**Fig.5 D**), as well as the end points of each learning curve –
338 trial 10 and trial 20 (**Fig.5 E**). There was no significant difference between these, suggesting that
339 the learning curves are identical. These results were corroborated by analyzing confidence ratings,
340 as we also observed a drop in confidence at trial 11 compared to trial 10, in line with our results
341 on performance (paired t-test trial 10 - trial 11, $t=3.7$, $p=0.0016$, **Fig. S1&S2**)



342
 343 **Figure 5: Drop in performance after response switch.** (A) To test the knowledge transfer hypothesis, we analyzed all
 344 trials within a block, encompassing trials before and after the switch. We specifically focused on the difference
 345 between trial 10 (before response switch) and trial 11 (after response switch). (B) If knowledge is transferred, we
 346 expect similar performance in trial 11 as in trial 10. In contrast, if no knowledge is transferred, we expect trial 11 to
 347 show similar performance to trial 1. (C) Performance in trial 11 is worse than in trial 10 (paired t-test, $t=-6.5$,
 348 $p<0.0001$). (D) Performance in trial 11 is similar to performance in trial 1 (paired t-test; n.s.; $N=20$). (E) Performance
 349 in trial 10 is similar to performance in trial 20 (paired t-test; n.s.; $N=20$). (F) Performance time course for different
 350 difficulty levels and pro-/anti-saccade orders. To illustrate the difference due to statistical learning only, we subtracted
 351 the motor error estimated from the calibration task. Insets show the result for the same statistical test as in (C-E), but
 352 performed separately on the data of each condition.

353

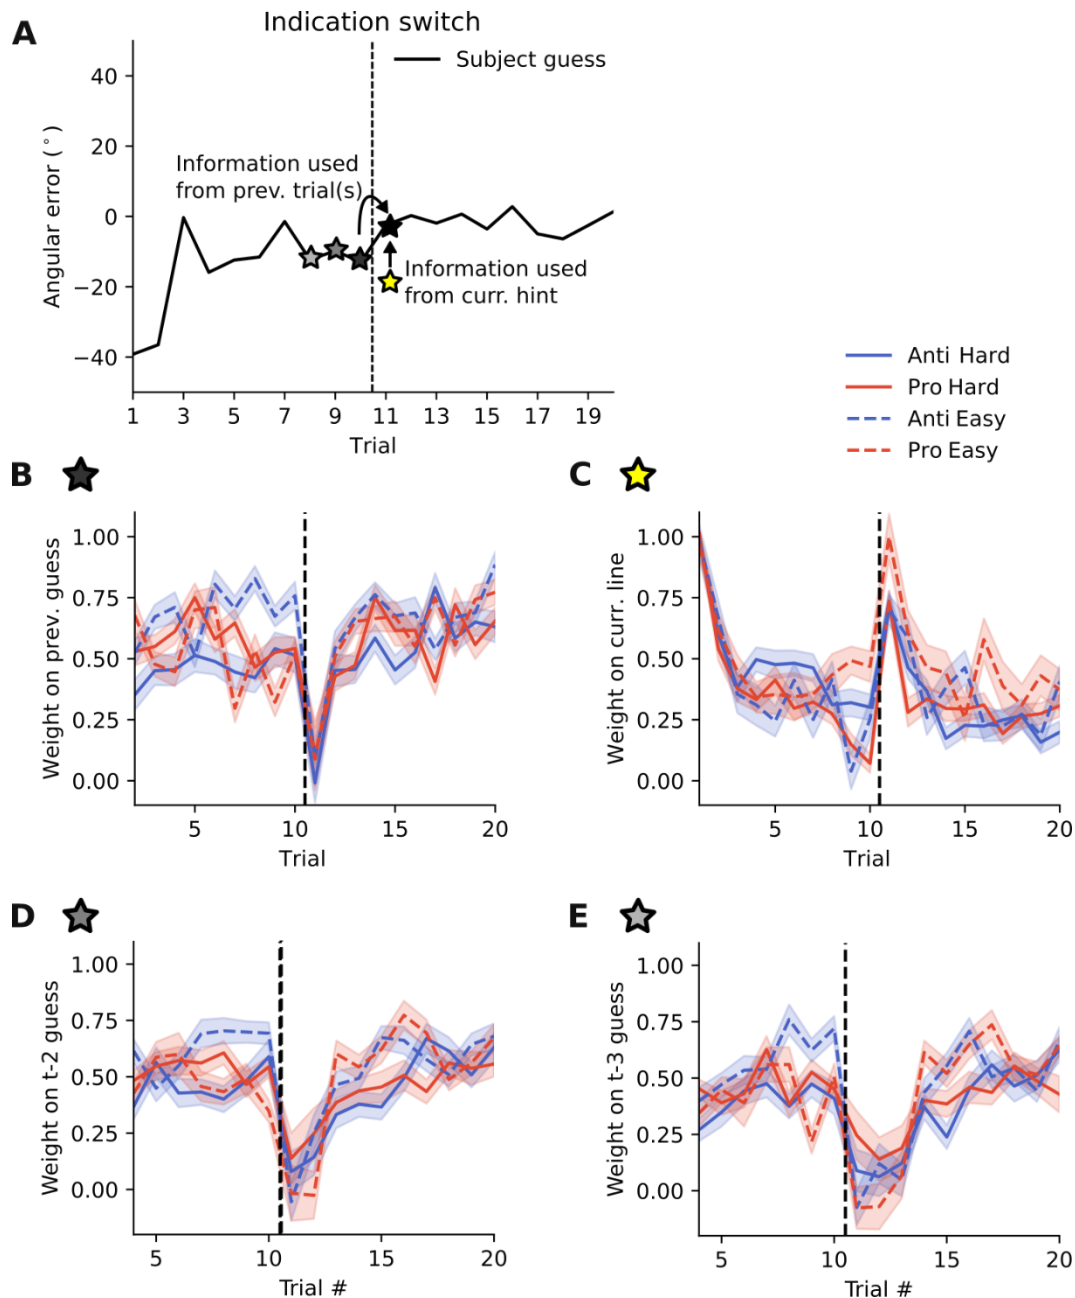
354

355 **No knowledge transfer between visuo-motor modalities**

356 As modeling showed that participants rely on their previous guesses rather than previously seen
357 visual hints (**Fig.4 E**), we wished to test whether this is also the case across the time that a response
358 switch occurs. For this, we calculated the regression weights on previous guesses in a time-
359 dependent manner (**Fig.6 A**). If knowledge is not transferred between different visuo-motor
360 modalities we would predict: 1) Participants highly rely on the visual hint in trial 11 and do not
361 use previous guesses to inform their decision. 2) They are also not able to use previous guesses
362 further in the past, if these guesses were obtained before the response switch. Our analysis showed
363 that indeed these were the case. We found that participants' estimates in trial 11 (after response
364 switch) are independent of their guess in trial 10 (before response switch). In contrast, at all other
365 time points they used previous guesses to inform their current decision (**Fig.6 B**). As participants
366 did not use their previous experience in trial 11, we expected that they instead rely highly on the
367 hint presented in trial 11. Regression analysis confirmed this hypothesis (**Fig.6 C**), although we
368 found that in three out of the four conditions the weight on the current hint in trial 11 is lower than
369 in trial 1. Only for the easy task difficulty level and the ordering of first anti-saccades and then
370 pro-saccades, the weight on the presented hint at trial 11 was as high as at the beginning of the
371 block (**Fig.6 C** red dashed line). Also inspired by the initial modeling results on participants'
372 strategies (**Fig.4**), we tested the influence of previous guesses further in the past (two trials (D)
373 and three trials (E)). Again, we found that there is no knowledge transfer across the response
374 switch. In summary, these results show that participants were unable to use their past estimates,
375 made with another response type, to inform their current guess.

376 One possible explanation for these results is that participants had difficulties in understanding that
377 the hidden target location had remained the same across first and second 10 trials (or between trial
378 10 to 11). To rule out this possibility, we performed a second experiment with an independent set
379 of participants (N=20) where we further emphasized that the target location was the same across
380 the two halves of the experiment. Furthermore, we asked the participants at the end of each block
381 whether they were aware that all twenty trials performed so far belonged to the same hidden target.
382 We obtained overall similar results in this experiment (**Fig. S3**), as learning performance, as well
383 as information transfer and confidence were disturbed after the response switch between trial 10
384 and 11, despite the fact that participants were explicitly instructed that all twenty trials belonged
385 to the same hidden target and that they were also aware of this rule (**Fig. S3 B**). This indicates that

386 the inability to transfer knowledge between response types is not because of a cognitive
 387 misunderstanding of the task.



388

389 **Figure 6:** No knowledge transfer between visuo-motor modalities. (A) To test if there is any knowledge transferred
 390 from the experience with one response type to the other, we regressed the current guess against previous
 391 guesses/current hint. (B) Participants' estimates at trial 11 (after response switch) are independent of the estimates at
 392 trials 10 (before response switch). In contrast, at every other time point, participants use previous experience to inform
 393 their current guess. (C) At trial 11, participants highly rely on the information coming from the current hint. (D-E)
 394 Besides transfer from one trial to the next (B), there is also no transfer from trials further in the past across the response
 395 switch.

396 **Discussion**

397 The aim of the current study was to characterize the dynamics of prior learning and its dependence
398 on the type of motor response used to report choices. We found that participants could learn a
399 sensorimotor prior within a few trials, with the learning time course being independent of the
400 response type (pro- or anti-saccades). By using a model-comparison approach, we further
401 demonstrated that participants relied more on their own guesses from previous trials compared to
402 visual hints provided in previous trials – again – independent of the response type. This suggests
403 that prior knowledge is represented in terms of previous motor actions and not incoming, external
404 information provided by the visual hints. To verify this hypothesis, we tested whether participants
405 could generalize their learned prior knowledge from one motor context to the other – a switch from
406 pro- to anti-saccades or vice versa. We found that switching the response type caused participants
407 to reset to naïve levels of performance, indicating that experience from one response type could
408 not be generalized to the other. This was the case even despite explicit instructions and
409 participants’ awareness that pro- and anti-saccade trials belonged to the same hidden target
410 location. Our results suggest that humans learn sensorimotor priors through monitoring their
411 previous motor decisions rather than external sensory hints. The dependence of learning on past
412 motor decisions discourages generalization of the learned knowledge to conditions where a
413 different visuo-motor mapping is needed.

414 Our findings suggest that prior knowledge is represented in a motor specific manner during early
415 learning, which is in line with previous studies reporting motor specific priors in different
416 paradigms [15,16]. We could furthermore identify one potential reason for why generalization is
417 not possible in such contexts, as we found that participants do not memorize the external
418 information from previous trials (visual hints in our case), but instead they memorize their own
419 actions in each trial (**Fig.4 E-F**). As our task required an estimation in every trial, indicated by
420 either a pro- or an anti-saccade, the memory of each trial’s decision was probably represented as
421 the motor action taken to indicate the guess. This could also explain why tasks which do not require
422 an explicit response via a motor action are more generalizable [15]. In these cases, the memory
423 from previous trials is potentially formed in a more abstract way, as participants only have to
424 ‘think’ about their decision, but not perform any specific action. The finding that prior knowledge
425 is built on internal decisions, compared to external cues, could therefore unify some previous
426 controversial findings about prior generalization.

427 Despite the suggested motor specific formation of prior knowledge, the algorithm to combine
428 previous experiences to inform the current decision seemed to be similar for both tested response
429 contexts (**Fig.4 E-H**). This suggests that there is a general procedure for how humans combine
430 previous experiences. However, whether prior knowledge can generalize or not depends on the
431 specific manner through which previous experience or decisions are stored in memory (e.g., in
432 terms of motor actions or more abstract decisions). In other words, although at an algorithmic level
433 learning is independent of the response modality, the learned information is stored with a format
434 that is specific for each modality. This explanation is in line with the previously proposed
435 dissociation between learning a *policy* versus *knowledge* [16], and further narrows the space of
436 testable predictions regarding the neural implementation of these different types of learning, as the
437 well-described neuronal machinery of pro- and anti-saccades [18,19] could allow a
438 characterization of how the two types of learning occur in the brain, for instance through using
439 neuroimaging techniques.

440 Our study is different from previous studies as it does not test prior learning in a condition where
441 there is also sensory uncertainty [5,17]. In these general task designs, participants are asked to
442 perform a task trial-by-trial and are not explicitly told to combine information from previous trials.
443 Prior learning in these cases is therefore implicit and potentially unconscious. Furthermore,
444 learning is mostly observed by analyzing how participants combine the noisy sensory information
445 in a given trial with the formed prior. It is therefore not directly possible to resolve which of the
446 two is learned, as observed changes in this combination could potentially come from a changed
447 likelihood distribution, from a changed prior distribution, or from changes in both distributions.
448 Because of these limitations, we designed our task such that the sensory information in each trial
449 was given by a clearly visible hint. We then explicitly asked the participants to combine the
450 information across trials to find the hidden target location. Compared to previous prior learning
451 studies [5], we could therefore directly look at the ability to learn statistical regularities over trials.
452 This more general statistical learning context has also been studied previously, for example
453 showing that learning can happen rapidly within a dozen trials if feedback is provided [20].
454 However, to our knowledge, pure statistical learning, without trial-by-trial feedback, together with
455 generalization has not been studied in this context.

456 The two different task contexts we investigated in this work are distinct to previous studies, as we
457 did not test generalization from one effector to another, such as performing a task with the right

458 hand and switching to the left [11,21], or switching from a motor to a perceptual task [16]. Instead,
459 our two contexts represent two distinct cue-action mappings, though performed with the same
460 modality (oculomotor system). By cue-action mapping we mean that participants had to indicate
461 their guess (the internal cue) with two different response types – pro- and anti-saccades (the
462 actions), dependent on the task context. Potentially, generalization could be easier between
463 modalities compared to generalization between different cue-action mappings. Given the specific
464 design of our task, we cannot differentiate whether participants form a motor independent spatial
465 prior, which is aligned to the given cue-action mapping, or whether they form their prior directly
466 at the motor level. In both cases generalization would fail, matching our experimental observation.
467 Potentially, participants learn a spatial representation in the pro-saccade context, where the correct
468 estimate lies close to their performed motor action endpoint. Then, in the anti-saccade context,
469 they don't follow this estimate and solely invert their motor plan, but instead they form a 'pro'
470 representation of the hidden target location in this new context, where again the performed motor
471 action endpoint is close to the acquired spatial estimate. In this interpretation, anti-saccades are
472 not really anti-saccades, but pro-saccades relative to the participants' estimates and only visual
473 information is inverted. What speaks for this interpretation is the fact that participants seemed to
474 be closer to the optimal learning performance in the anti-saccade condition (**Fig.3 F-G**), although
475 this was only the case when the visual hints were narrowly distribution.

476 Our setup allowed us to simultaneously evaluate participants' performance, as well as their
477 confidence in their given estimation. Interestingly, confidence also decreased after the response
478 type switch, suggesting that participants were aware that they cannot generalize between both
479 contexts. On the other hand, a control experiment, where we explicitly asked participants to
480 indicate after each block whether they were aware that both response type contexts belonged to
481 the same hidden target estimation, showed that that they knew that information could be combined
482 across both contexts (**Fig. S3 B**). Together, these results suggest that participants' inability to
483 transfer knowledge from one response context to the other was not due to conscious
484 misunderstanding, but more likely due to the specific mechanisms of how the prior is formed
485 unconsciously.

486 Although our results, especially the inability to generalize across different motor contexts, suggest
487 that statistical learning is implemented in a low-level, motor-dependent way (**Fig.5&6**), there
488 might be potentially alternative explanations for why we could not see generalization in our

489 experiment. In the following we will discuss these alternative interpretations. One potential
490 explanation for the lack of generalization could be an interference between the internal memory of
491 the target location, which is disturbed by the information slide, presented between trial 10 and 11
492 indicating response switch (**Supplementary Video 1**). Another possibility is that the short
493 timescale of only ten trials is not enough to form a general representation of the estimated target
494 location. Potentially, initial learning is motor-dependent, but over time this is transformed to a
495 motor-independent knowledge, which can then be transferred to other motor contexts. Another
496 possibility is that participants would need to train on our task for more than one session, such that
497 they can learn to adapt their way of forming the prior to something which allows for generalization
498 across response contexts. Finally, one major reason for a motor-context dependent learning could
499 be that we did not provide external feedback to guide the learning process (unsupervised learning).
500 Instead, participants had to rely solely on their internal feedback, potentially coming from the
501 motor system. Future studies will be needed to test these possible explanations.

502 Spatially directed movements such as saccadic eye movements and reaching are an integral part
503 of our daily activities and acquired skills (e.g., imagine a cellist or a tennis player). Both types of
504 movements are profoundly influenced by statistical priors [1,22]. Furthermore, saccadic eye
505 movements provide detailed sensory information about a scene and are tightly linked to the
506 allocation of attention, hence being instrumental for active vision [23,24]. Typically, the effect of
507 statistical priors on saccadic eye movements is investigated by using a fixed set of potential target
508 locations, where the probability of target appearing in some locations (hence the uncertainty of
509 that location) is more than the others. The typical finding of these studies [22,25] is that saccade
510 latencies became shorter with increasing prior probability of the corresponding target location.
511 One novel aspect of the current study is to test how spatial priors can be learned through eye
512 movements under more complex settings where the possible target locations are not fixed, and
513 probabilities could only be inferred through tracking noisy information over time. In comparison
514 to the saccadic eye movement, reaching movements have enjoyed a more rigorous characterization
515 of the learning dynamics [5,12,14,16]. Inspired by these studies, we characterized the dynamics of
516 statistical learning through eye movements, that are a more accessible motor plan to be tested in
517 the laboratory, and the learned information acquired through their execution can directly impact
518 the very way that the brain samples the sensory information [26].

519 **Materials and Methods**

520 In this study, we report the results of two experiments investigating how a spatial prior is learned
521 under different oculomotor response contexts. The second experiment was identical to the first and
522 served as a control for ensuring that participants were aware of how the location of the learned
523 prior varied across blocks of the experiment (see the description of the Experimental design). We
524 therefore, describe the methods common to both experiments and mention the difference where
525 they apply.

526 Participants

527 In total, 41 participants were recruited for this study. 21 participants (age range 22-38 years,
528 $M=26.05$, $SD=3.80$, 10 females) took part in the first experiment. All participants had normal
529 ($N=9$) or corrected-to-normal vision ($N=12$). One participant was excluded from the data of the
530 first experiment as the post-experiment questionnaire indicated that the participant had
531 misunderstood the task. 20 participants (age range 23-38 years, $M=28.95$, $SD=5.07$, 11 females)
532 took part in the second experiment (with no exclusion) and all had normal ($N=16$) or corrected-to-
533 normal vision ($N=4$). Participants were recruited from the general population of the city of
534 Göttingen, Germany, using flyer and online advertisement and received cash financial
535 compensation for their participation. Participation was voluntary; all participants were informed
536 about the study procedure and gave written consent prior to the test session. The study was
537 approved by the local ethics committee of the “Universitätsmedizin Göttingen” (UMG), under the
538 proposal number 15/7/15.

539 Experimental setup

540 The stimuli were presented at the center of a calibrated ViewPixx/EEG monitor (VPixx
541 Technologies, QC Canada, dimension: 53 x 30 cm, refresh rate: 120 Hz) with a resolution of 1920
542 x 1080 pixels at a viewing distance of 60 cm. All experiments were scripted in MATLAB, using
543 Psychophysics toolbox [27]. Eye movements were measured using the EyeLink1000+ eye tracking
544 system (SR Research, Ontario, Canada) in a desktop mount configuration, recording the right eye,
545 with a sampling rate of 1000 Hz. A chin rest was used to stabilize the participant’s head. The
546 EyeLink camera was controlled by the EyeLink toolbox in MATLAB [28]. At the beginning of
547 each experiment, as well as after every 10 blocks of the *Hidden Target Task* (see the Experimental
548 design), the eye tracking system was calibrated using a 13-point standard EyeLink calibration
549 procedure. Calibration was repeated until an average error of maximum 0.5 visual degrees was
550 achieved and the error of all points was below 1. If the calibration accuracy dropped during the

551 experiment, e.g., due to the subjects' movement, the experimenter recalibrated the eye tracking
552 system again.

553 **Experimental design**

554 The experiment comprised two tasks: a calibration task (to estimate the motor error of pro- and
555 anti-saccades) and the main task, referred to as the 'hidden target task' (**Fig.1**). There were in total
556 four blocks of the calibration task (n=20 trials in each block) and forty blocks of the hidden target
557 task (n=20 trials in each block). Each experiment started with a block of the calibration task
558 followed by one training block for the hidden target task (n=10 trials in this block). The data from
559 this training phase was not analyzed. Thereafter, the experiment proceeded to the main task where
560 participants performed 10 blocks of the hidden target task followed by 1 block of the calibration
561 task. This sequence was repeated 4 times (**Fig.1 E**).

562 In a second experiment, we used the exact same experimental design but enforced the instruction
563 that the location of the hidden target remained the same across a block of 20 trials although the
564 response type changed halfway through (from anti- to pro-saccade or vice versa). For this, we
565 adapted the information slides shown during the experiment (**Supplementary Video 1&2**).
566 Furthermore, we asked the participants after every twenty trials of a hidden target whether they
567 were aware that the last twenty trials belonged to the same hidden target. Participants then had to
568 press a button to indicate their response, either yes or no.

569 **Hidden Target Task**

570 Participants were instructed to look for a 'hidden treasure' location on a ring with a radius of 7.5° ,
571 centered in the middle of the screen (**Fig.1 A**). The word 'hidden treasure' was used in our
572 instructions to the participants to make the task more realistic and engaging, however we will refer
573 to the task as the 'hidden target' task throughout. Each trial started with a fixation period, where
574 participants had to fixate for 0.5 s on the white cross (size = 0.1875° , color: white, displayed on a
575 half-grey background) in the middle of the screen (**Fig.1 B**). After that, a white line (length = 1.125° ,
576 color: white) was presented and participants had 3 s to estimate the hidden target location for this
577 trial and indicate their guess either by looking at it (pro-saccade) or by looking opposite of it (anti-
578 saccade) and fixate their estimated location for 0.5 s. Thereafter, participants rated their level of
579 confidence in their guess on an discrete scale from 1 to 6, where 1 means very uncertain and 6
580 means very certain about the target location. The confidence rating had to be done within 4
581 seconds.

582 Participants were told that they had twenty trials to guess the location of the hidden target, after
583 which a new hidden target had to be found. To estimate the hidden target location participants had
584 to closely monitor the location of a line that was presented in every trial and served as a visual
585 hint. The hidden target location was the mean of a von Mises distribution and each visual hint was
586 a sample drawn from this distribution (see below) [17]. Hence, by paying attention to the location
587 of the hint across trials, participants were able to infer the underlying distribution of the hidden
588 target location. Participants indicated their estimates by looking either at where they thought the
589 hidden target was located on the ring (pro-saccade), or at a location directly opposite to it (anti-
590 saccade). Ten consecutive trials of a block of twenty trials required pro-saccade responses, and the
591 other ten anti-saccade responses. The type of the required response (either pro- or anti-saccade)
592 was visually indicated by an instruction display presented every ten trials. Participants were
593 instructed to perform the same type of response for ten trials in a row, until the response type
594 changed.

595 In total, there were forty hidden target blocks. The target location of each block, which is the mean
596 of the von Mises distribution, was randomly drawn from a fixed set of twenty locations evenly
597 distributed on the circle. Thus, each location only appeared twice during the experiment. To
598 familiarize the participants with the connection between the hints and the hidden target location,
599 ten training trials were performed in the beginning of each experiment. After the ten training trials,
600 participants saw all the ten lines together on the screen, as well as the correct hidden target location.
601 Furthermore, after each hidden target block, participants saw where the actual hidden target
602 location was, but they did not receive feedback about their performance on a trial-by trial basis.
603 As such, in our experiments learning was unsupervised.

604 Four different experimental conditions, counter-balanced across blocks, were tested. Each block
605 was either *easy* or *hard*, controlled by adapting the concentration of the von Mises distribution,
606 and either ordered with first pro-saccade then anti-saccade, or first anti-saccade then pro-saccade
607 response (**Fig.1 F**). For the easy task condition the concentration of the von Mises distribution
608 (defined by κ which is a measure of dispersion, where $1/\kappa$ is equivalent to the variance σ^2_{dist} of the
609 distribution) was 30 ($\sigma_{\text{dist}} \sim 10^\circ$), for the hard task condition it was 5 ($\sigma_{\text{dist}} \sim 26^\circ$) and for the training
610 it was 80 ($\sigma_{\text{dist}} \sim 6^\circ$). As the concentrations of these distributions are relatively large, we could treat
611 the von Mises distribution as a normal distribution and use standard statistics.

612 Calibration Task

613 The aim of this task was to quantify the participant-specific motor error of the visually-driven pro-
614 and anti-saccades. Each block of this task consisted of twenty trials, from which the first ten were
615 pro-saccades and the last ten were anti-saccades. On each trial, one out of ten equally distributed
616 locations on the circle (same circle as in the hidden target task) were selected and a target line was
617 presented at that location (**Fig.1 C**). Participants were instructed to look either directly at the
618 displayed line (pro-saccade in the first ten trials), or directly opposite to where it appeared (anti-
619 saccades, second ten trials). Additionally, it was highlighted that this task is completely
620 independent of the hidden target task. Each trial consisted of an initial fixation phase, where
621 participants had to fixate on the white cross in the middle of the screen for 0.5 s. After that, a white
622 line appeared and participants had to either look at it or opposite of it (**Fig.1 D**). In the beginning
623 of a block of ten trials, participants received an instruction display indicating whether they had to
624 perform pro- or anti-saccades during the upcoming trials.

625 Successful response

626 For both tasks, a successful response was defined as follows. Participants had to move their gaze
627 from the central fixation point towards a peripheral location on the ring. As soon as they moved
628 away from the fixation and crossed a circular threshold of 5.375° a successful response was
629 possible. To complete the response, participants furthermore had to fixate on one specific point on
630 the screen, by holding their gaze for 0.5 s within an area with a radius of 1 .

631 Analysis

632 Data pre-processing

633 The recorded raw eye movement data was transformed to MATLAB files by using a MATLAB
634 library for eye movement analysis [29]. Participants' estimates in each trial were calculated offline
635 by averaging the eye movement data of the last 100 ms, out of the total 500 ms necessary for a
636 successful response. The main eye movement parameter that we analyzed was the angular distance
637 between the participant's estimate and the true hidden target location (**Fig.2 A**). Failed trials were
638 excluded from the analysis. A trial could fail in several ways. Firstly, there could have been a
639 disturbance with the eye tracking system or the participant's calibration so that the gaze position
640 was not correctly detected, which made it necessary to re-calibrate the eye tracker. Secondly, the
641 participant could have been too slow to indicate their guess in time (3 s). Thirdly, to exclude
642 erroneous pro- instead of anti-saccades and vice versa, we analyzed the distribution of angular
643 errors and found a bimodal distribution (**Fig. S4**). We set a threshold at 100° , which separated both

644 modes (**Fig. S4**). Thus, every trial with an absolute angular error bigger than 100° was categorized
645 as failed because of the wrong response type.

646 Data analysis

647 Statistical analyses were done using R and Python. To evaluate learning we used paired, two-sided
648 t-tests to compare several parameters within a participant. We did not use circular statistics as
649 subjects' responses were highly localized on the ring (**Fig.2 A-B & S4**). Learning was assessed
650 by measuring the decrease in the absolute angular error between the participant's estimate and the
651 true hidden target location across trials. Given the assumptions that the participant's estimates are
652 normally distributed and centered around zero, the mean absolute angular error can be related to
653 the variance of the Gaussian distribution by:

$$654 \quad 2) \text{ Mean}(|x|) = \sqrt{\frac{2}{\pi}} SD(x)$$

655 To compare the performance difference between pro- and anti-saccades, we calculated a 'modality
656 difference index' (**Fig.3 C**). For this, we first calculated the median of the absolute angular error
657 for pro-saccade and anti-saccade response trials. The modality difference index is then given by
658 the difference between the two medians, divided by their sum. We started our analysis by only
659 using the data from the first ten trials of each block (**Fig.1-4**). Only when we looked at the transfer
660 between modalities, we used all twenty trials of each block (**Fig.5-7**).

661 Theoretical bounds for learning performance

662 To evaluate participants' performance, we computed the theoretical lower bounds on the absolute
663 angular error they could potentially achieve by using all the information that was available to them.
664 Participants could use the previously seen visual hints and their memory of their previous motor
665 actions (referred to as the previous guess) to infer the most probable location of the target on each
666 trial. These sources of information are error-prone since on each trial participants had only seen a
667 limited number of visual hints (i.e., sampling error) and their previous responses contained motoric
668 noise. We assumed that these two sources of error are independent. The error due to the limited
669 number of visual hints was calculated as the cumulative mean of all hints seen so far, which
670 represents an optimal way of combining samples over time to estimate the mean of a distribution.
671 A constant motor error, measured for each participant during the calibration task, was used to
672 represent the motoric noise. The variance of the joint estimate, derived from combining visual
673 hints and motor actions, was then calculated by adding the variance of the two sources, based on

674 the assumption of their independence (cf. Eq.1). In addition, we also calculated a suboptimal,
675 ‘limited memory’ lower bound to account for the information loss across time. This model makes
676 use of the latest two visual hints, instead of all, to infer the statistical distribution of the target. To
677 understand how participants used these sources of information to make decisions on a trial-by-trial
678 basis, we employed a detailed model-comparison approach as described below.

679 **Modeling participants’ behavior**

680 We used linear regression models to analyze the single subject behavior. Four potential strategies
681 we wanted to test were: 1) looking directly at the visual hints in each trial, 2) estimating the
682 cumulative average of all visual hints so far, 3) looking at the same location as in the previous trial,
683 4) estimating the cumulative average of all previous guesses (as given by their saccadic responses)
684 (**Fig.4 A**). Strategy (2), calculating the cumulative average of all hints seen so far is directly linked
685 to the lower bound on performance described above. This is an optimal statistical strategy to
686 combine all observed hints and therefore produces the optimal performance. In principle, strategy
687 (2) is also equivalent to a Bayesian optimal strategy, as in our case each hint has objectively equal
688 certainty and should therefore be weighed equally.

689 We started with the four described single predictor models to see which of the four mentioned
690 strategies best describes participant's behavior. The dependent variable was the angular error of a
691 participant’s estimate in a certain trial. The independent variable was the angular error of the
692 estimate, given by one of the above-mentioned strategies. Only the data from trial 4 onward was
693 included to test all models on the same data, since we wanted to test the influence of up to three
694 trials in the past on the current trial. In a second stage, we tested linear regression models with
695 multiple independent variables. We included previous guesses from up to three time steps in the
696 past, as well as the visual hints from the current and up to three time steps in the past. Each of the
697 described regression models was fitted to the single subject data. To compare these models, we
698 calculated BIC, delta BIC, and Bayesian weights [30], to assess the likelihood of each model being
699 the best fit to the data (cf. Evaluating model performance). Since these values are normalized, they
700 can be used to determine the model that on average best fits the participants’ data.

701 **Evaluating model performance**

702 To compare the different linear regression models presented above we used a model comparison
703 evaluation based on Bayesian weights [30]. For this, we firstly calculated BIC values for each
704 model. Each BIC value was rescaled by calculating ΔBCI , which is calculating the difference to

705 the smallest BIC value in the group of models considered. This forces the best model to have
706 $\Delta\text{BCI}=0$ and the other models to have positive values. We then calculated Bayesian weights ω
707 with:

$$708 \quad 3) \quad \omega_i = \frac{e^{-0.5\Delta\text{BIC}_i}}{\sum_{n=1}^N (-0.5\Delta\text{BIC}_n)}$$

709 The Bayesian weights of all tested models in **Fig.4 B** sum up to 1 and define the probability of
710 being the best model, among the one tested.

711

712 References

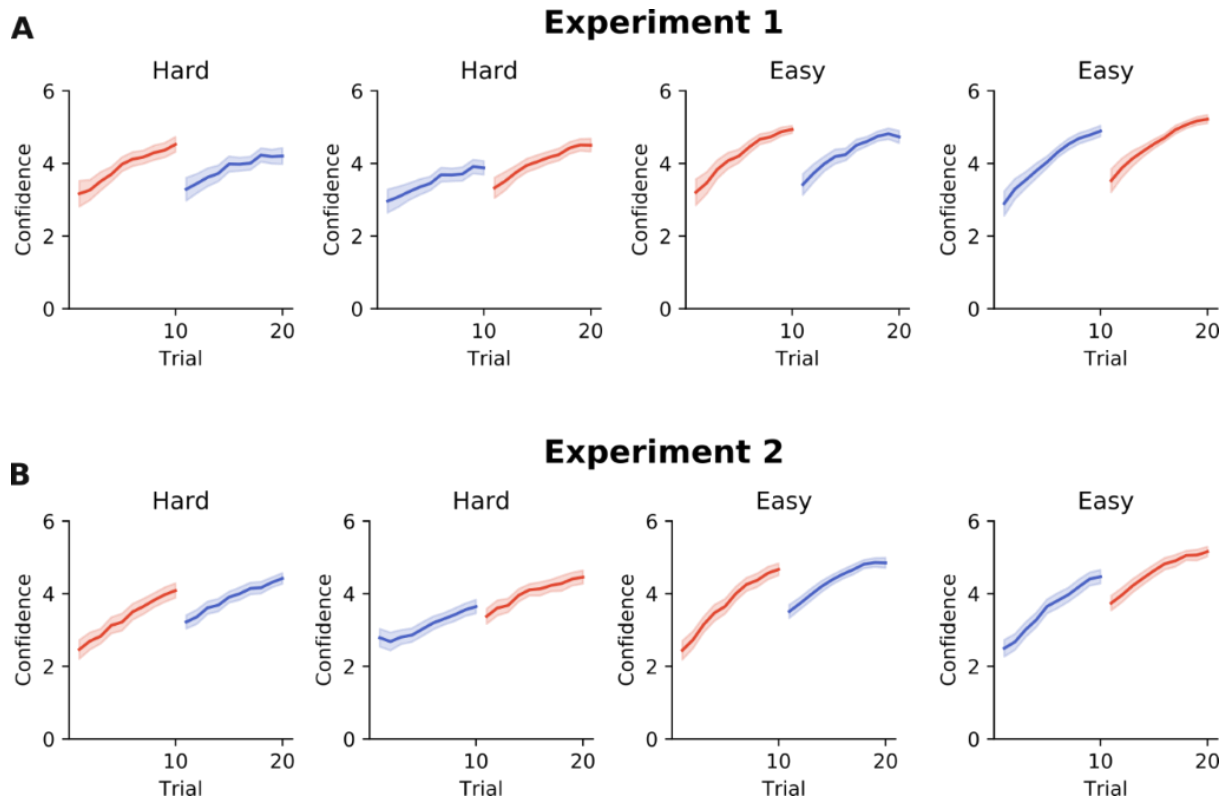
- 713 1. Körding KP, Wolpert DM. Bayesian integration in sensorimotor learning. *Nature*. 2004;427: 244–
714 247. doi:10.1038/nature02169
- 715 2. Wei K, Körding K. Uncertainty of feedback and state estimation determines the speed of motor
716 adaptation. *Front Comput Neurosci*. 2010;4: 1–9. doi:10.3389/fncom.2010.00011
- 717 3. Ernst MO, Banks MS. Humans integrate visual and haptic information in a statistically optimal
718 fashion. *Nature*. 2002;415: 429–433. doi:10.1038/415429a
- 719 4. Sato Y, Kording KP. How much to trust the senses: Likelihood learning. *J Vis*. 2014;14.
720 doi:10.1167/14.13.13
- 721 5. Berniker M, Voss M, Kording K. Learning Priors for Bayesian Computations in the Nervous
722 System. Brezina V, editor. *PLoS One*. 2010;5: e12686. doi:10.1371/journal.pone.0012686
- 723 6. Vilares I, Howard JD, Fernandes HL, Gottfried JA, Kording KP. Differential representations of prior
724 and likelihood uncertainty in the human brain. *Curr Biol*. 2012;22: 1641–1648.
725 doi:10.1016/j.cub.2012.07.010
- 726 7. Battaglia PW, Hamrick JB, Tenenbaum JB. Simulation as an engine of physical scene
727 understanding. *Proc Natl Acad Sci U S A*. 2013;110: 18327–18332. doi:10.1073/pnas.1306572110
- 728 8. Tenenbaum JB, Kemp C, Griffiths TL, Goodman ND. How to grow a mind: Statistics, structure, and
729 abstraction. *Science*. Science; 2011. pp. 1279–1285. doi:10.1126/science.1192788
- 730 9. Tenenbaum JB, Griffiths TL, Kemp C. Theory-based Bayesian models of inductive learning and
731 reasoning. *Trends Cogn Sci*. 2006;10: 309–318. doi:10.1016/j.tics.2006.05.009
- 732 10. Frost R, Armstrong BC, Siegelman N, Christiansen MH. Domain generality versus modality
733 specificity: The paradox of statistical learning. *Trends in Cognitive Sciences*. Elsevier Ltd; 2015. pp.
734 117–125. doi:10.1016/j.tics.2014.12.010
- 735 11. Hewitson CL, Sowman PF, Kaplan DM. Interlimb generalization of learned bayesian visuomotor
736 prior occurs in extrinsic coordinates. *eNeuro*. 2018;5. doi:10.1523/ENEURO.0183-18.2018
- 737 12. Yin C, Wang H, Wei K, Körding KP. Sensorimotor priors are effector dependent. *J Neurophysiol*.
738 2019;122: 389–397. doi:10.1152/jn.00228.2018

- 739 13. Kiryakova RK, Aston S, Beierholm UR, Nardini M. Bayesian transfer in a complex spatial
740 localization task. *J Vis.* 2020;20: 17–18. doi:10.1167/JOV.20.6.17
- 741 14. Fernandes HL, Stevenson IH, Vilares I, Kording KP. The generalization of prior uncertainty during
742 reaching. *J Neurosci.* 2014;34: 11470–11484. doi:10.1523/JNEUROSCI.3882-13.2014
- 743 15. Roach NW, McGraw P V., Whitaker DJ, Heron J. Generalization of prior information for rapid
744 Bayesian time estimation. *Proc Natl Acad Sci U S A.* 2017;114: 412–417.
745 doi:10.1073/pnas.1610706114
- 746 16. Chambers C, Fernandes H, Kording KP. Policies or knowledge: Priors differ between a perceptual
747 and sensorimotor task. *J Neurophysiol.* 2019;121: 2267–2275. doi:10.1152/jn.00035.2018
- 748 17. Dekleva BM, Ramkumar P, Wanda PA, Kording KP, Miller LE. Uncertainty leads to persistent
749 effects on reach representations in dorsal premotor cortex. *Elife.* 2016;5.
750 doi:10.7554/eLife.14316
- 751 18. Munoz DP, Everling S. Look away: The anti-saccade task and the voluntary control of eye
752 movement. *Nature Reviews Neuroscience.* Nature Publishing Group; 2004. pp. 218–228.
753 doi:10.1038/nrn1345
- 754 19. Ford KA, Goltz HC, Brown MRG, Everling S. Neural processes associated with antisaccade task
755 performance investigated with event-related fMRI. *J Neurophysiol.* 2005;94: 429–440.
756 doi:10.1152/jn.00471.2004
- 757 20. Chukoskie L, Snider J, Mozer MC, Krauzlis RJ, Sejnowski TJ. Learning where to look for a hidden
758 target. *Proc Natl Acad Sci U S A.* 2013;110: 10438–10445. doi:10.1073/pnas.1301216110
- 759 21. Kumar A, Panthi G, Divakar R, Mutha PK. Mechanistic determinants of effector-independent
760 motor memory encoding. *Proc Natl Acad Sci U S A.* 2020;117: 17338–17347.
761 doi:10.1073/pnas.2001179117
- 762 22. Brodersen KH, Penny WD, Harrison LM, Daunizeau J, Ruff CC, Duzel E, et al. Integrated Bayesian
763 models of learning and decision making for saccadic eye movements. *Neural Networks.* 2008;21:
764 1247–1260. doi:10.1016/j.neunet.2008.08.007
- 765 23. Munoz DP. Commentary: Saccadic eye movements: Overview of neural circuitry. *Progress in*
766 *Brain Research.* Elsevier; 2002. pp. 89–96. doi:10.1016/S0079-6123(02)40044-1
- 767 24. Wurtz RH. Brain mechanisms for active vision. *Daedalus.* 2015;144: 10–21.
768 doi:10.1162/DAED_a_00314
- 769 25. Carpenter RHS, Williams MLL. Neural computation of log likelihood in control of saccadic eye
770 movements. *Nature.* 1995;377: 59–62. doi:10.1038/377059a0
- 771 26. Parr T, Friston KJ. The active construction of the visual world. *Neuropsychologia.* Elsevier Ltd;
772 2017. pp. 92–101. doi:10.1016/j.neuropsychologia.2017.08.003
- 773 27. Brainard DH. The Psychophysics Toolbox. *Spat Vis.* 1997;10: 433–436.
774 doi:10.1163/156856897X00357
- 775 28. Cornelissen FW, Peters EM, Palmer J. The EyeLink Toolbox: Eye tracking with MATLAB and the
776 Psychophysics Toolbox. *Behav Res Methods, Instruments, Comput.* 2002;34: 613–617.
777 doi:10.3758/BF03195489
- 778 29. Manohar SG. Matlib: MATLAB tools for plotting, data analysis, eye tracking and experiment
779 design (Public). 2019. doi:10.17605/OSF.IO/VMABG

- 780 30. Burnham KP, Anderson DR. Multimodel Inference. *Sociol Methods Res.* 2004;33: 261–304.
781 doi:10.1177/0049124104268644
782
783
784

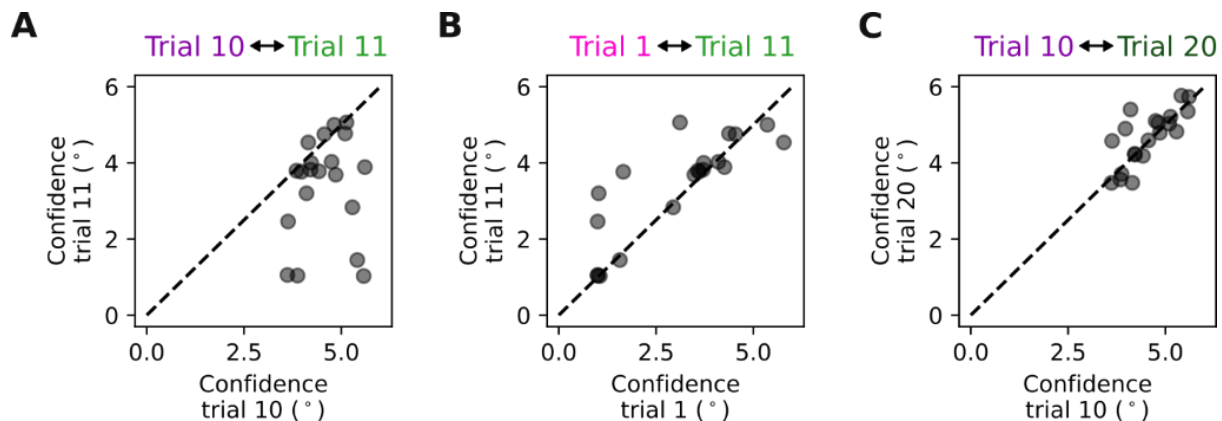
785 **Supplementary figures**

786



787
788
789
790
791
792

Figure S1: Confidence time course.

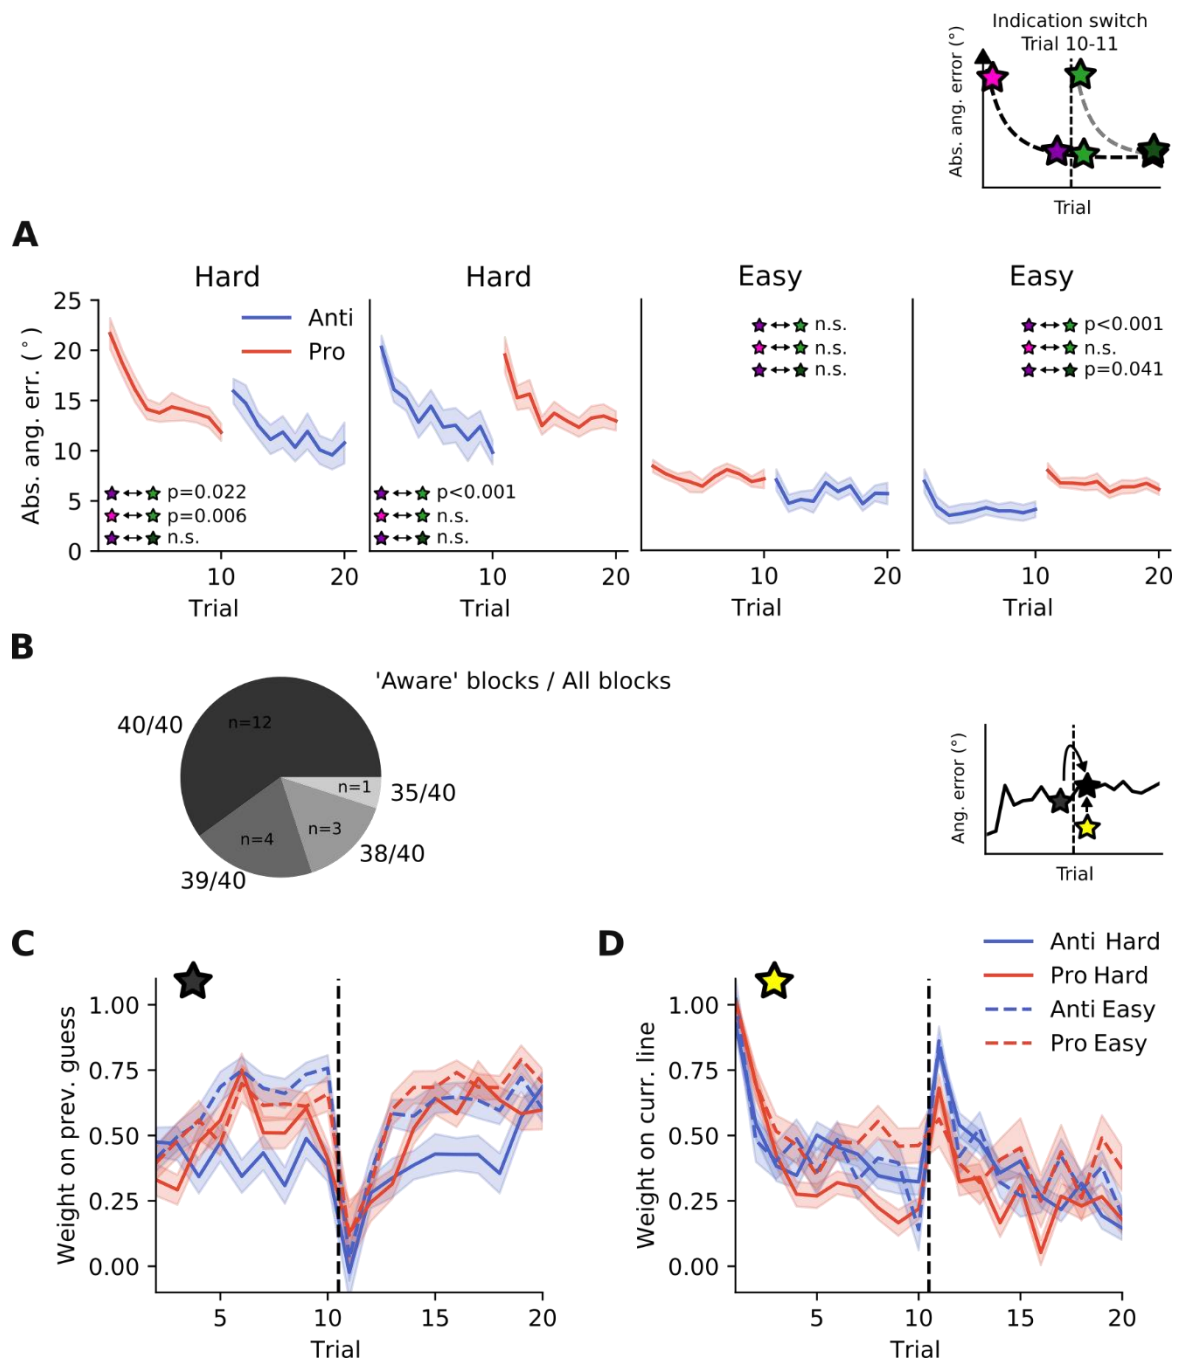


793
794

Figure S2: Confidence drops after response switch.

795 Figure S1 and S2 are supplementary to Fig.5 in the main text. The confidence ratings demonstrated the same results
796 as observed by analysing the absolute angular errors of the eye movements, as there was a decrement in confidence
797 between trials 10 and 11 as shown in S1-A) for Experiment 1, as well as in S1-B) for Experiment 2, for all experimental
798 conditions. The drop in confidence between trial 10 and 11 was significant (S2-A). The confidence of trial 11 was not
799 different from trial 1 (S2-B), and the last trial of a block of 20 trials was not significantly better than trial 10 (S2-C).
800 These results support the observation that after a switch in response modality, learning starts from a naïve level.

801



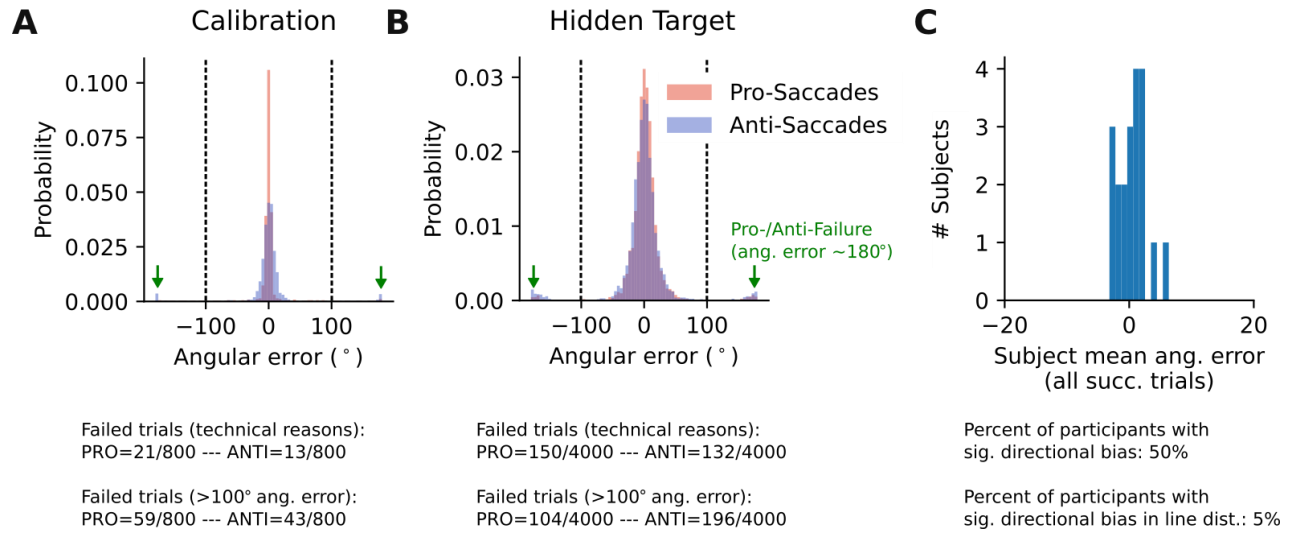
802

803 **Figure S3:** *Second experiment with reinforced instructions shows similar results.*

804 This figure is supplementary to Fig.5 and Fig.6 in the main text. These results demonstrate that we obtained similar
 805 results, when participants were explicitly instructed that the location of the hidden target remains the same after a
 806 switch. Additionally, participants had to report whether they were aware of this rule, thus reinforcing the instructions.
 807 A) Also in this experiment, performance dropped to naïve levels after a switch in response type. B) The majority of
 808 participants reported to be aware of the rule. C) The weighting of previous guesses dropped between trial 10 and 11
 809 (when the switch occurred) and D) instead more weight was put on visual hints.

810

811
812



813
814
815
816

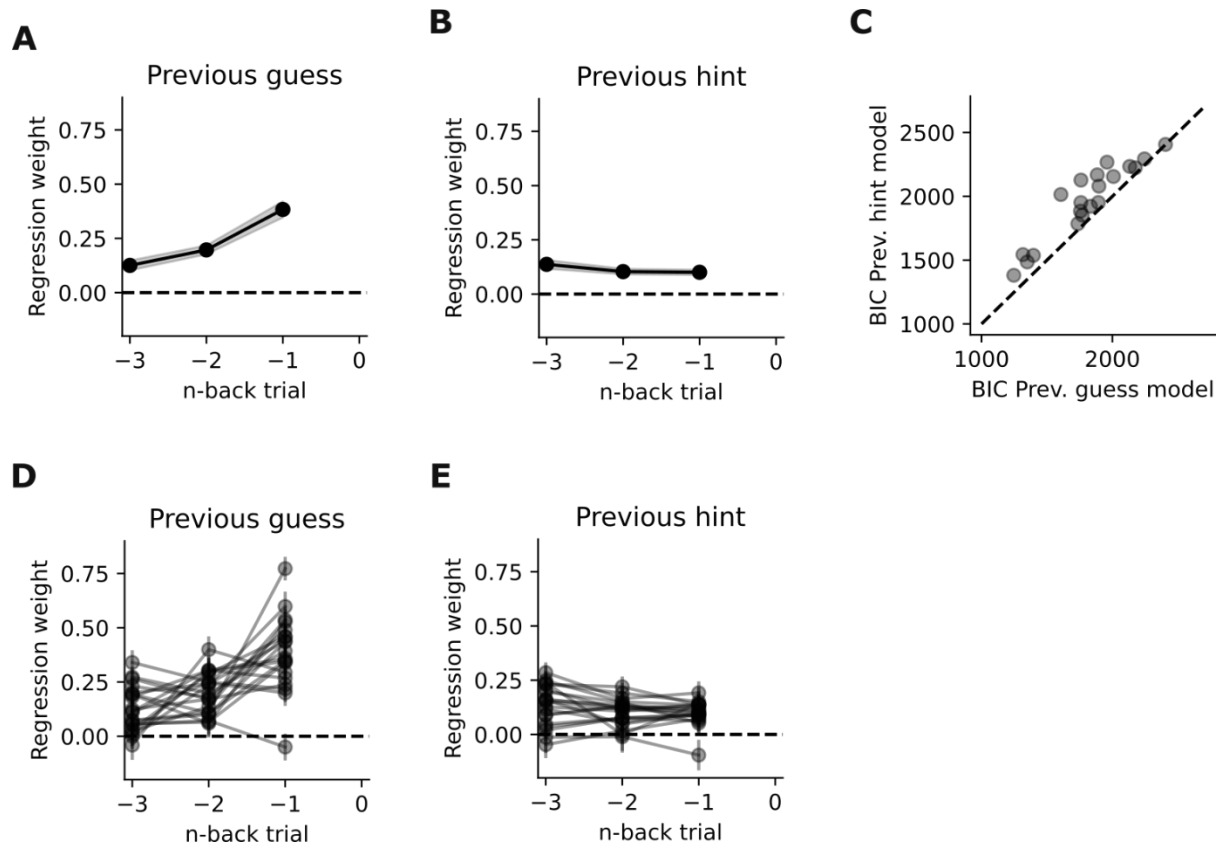
Figure S4: *Threshold for failed pro-/anti-saccades and subject-wise directional bias.*

817 This figure is supplementary to Fig.3 in the main text. It demonstrates the rationale for excluding erroneous pro- and
818 anti-saccades. A) Data of Calibration Task. B) Data of Hidden Target task. In both panels, subjects' estimates are
819 shown without excluding 'wrong' saccades due to pro-/anti-saccade error. As shown in A and B, there is a bimodal
820 distribution of angular errors, especially in the hidden target task. The peaks at $\pm 180^\circ$ error represents the failed pro-
821 /anti-saccades, meaning trials where pro-saccades should have been made but the subject responded with anti-saccades
822 and vice versa. We set a threshold at 100 degree to distinguish correctly aimed pro- or anti-saccades. C) Distribution
823 of the mean angular error across participants indicates that there was no directional (CCW or CW) bias.

824
825

826

827



828

829 **Figure S6:** *Memory traces for previous guesses and hints.*

830 This figure is supplementary to Fig.4 in the main text. Here, we quantified how much weight is put on previous hints
831 or guesses, tested in separate models each including either only the guesses or only the visual hints. A) Regression
832 weights of the model including previous guesses as a predictor B) Same as A for a separate model including visual
833 hints as a predictor 3) comparison of the model shown in A against the model shown in B. It can be seen that the
834 previous guess models consistently outperformed the visual information/hints models, confirming the finding that
835 subjects' behavior is better predicted by previous actions than external visual information D&E) Same as A&B, but
836 showing the single subject regression weights instead of the average weights.