Supporting Information

For “Visual perception as retrospective Bayesian decoding from high- to low-level features”
by Stephanie Ding, Christopher J. Cueva, Misha Tsodyks, and Ning Qian

Variances and biases of absolute distributions were larger in the 2-line condition than in the 1-line conditions

Fig. 2 of the main text showed a naïve subject’s 1-line and 2-line absolute distributions. Fig. S1 compares these two types of absolute distributions in terms of their variances and biases across all subjects. The 2-line absolute distributions had larger variances (left panel) and bias magnitudes (right panel) than the 1-line absolute distributions. We used the bias magnitude instead of signed bias because, for example, a bias of -4^o is worse than a bias of +0.1^o. The bias magnitude was calculated by taking the absolute value of the standard bias of a distribution. In each panel, the results for the 50^o and 53^o stimulus orientations are represented by open dots and crosses, respectively. The mean SDs of the 2-line and 1-line absolute distributions are 4.5^o and 3.7^o, respectively, and the difference is significant (two-tailed Wilcoxon signed rank test, $p = 0.016$). The mean bias magnitudes were 3.6^o and 1.7 ^{o}, respectively and the difference is also significant (two-tailed Wilcoxon signed rank test, $p = 0.0024$).

The subject of Fig. 2 in the main text showed negative biases in all four absolute distributions (i.e., the mean reported orientation of each line was smaller than the true value). Other subjects showed positive biases or different bias signs in different conditions. Girshick et al (1) demonstrated that compared with low-noise stimuli, high-noise stimuli are biased toward the cardinal axes. Since we only used low-noise stimuli (single lines with a contrast near 1) in our experiment, the two studies are not directly comparable. One possible explanation of seemingly arbitrary biases in our study is that orientation representations distorted during the relatively long delays in noisy working memory producing random biases with either sign. As Girshick et al pointed out, lower-noise stimuli must have smaller biases toward the cardinals (the peak locations of natural-image orientation distributions), and for our very low-noise stimuli, this weak bias toward the cardinals must be overwhelmed by the memory-based random biases.

Joint and relative distributions of orientation judgments cannot be explained by the absolute-to-relative assumption

Fig. 2 of the text shows one subject’s joint and relative distributions of the 2-line condition, and the predicted distributions according to the absolute-to-relative assumption applied to the 1-line absolute distributions. The results of the other 11 subjects are plotted in Fig. S2. Different rows are for different subjects, and the left and right columns are for the joint and relative distributions, respectively. The gray dots and histograms show the observed joint and relative distributions, whereas the blue dots and histograms are the corresponding predictions of the absolute-to-relative assumption. All subjects showed similar patterns including significant inter-report correlation (mean Pearson correlation coefficient 0.56 ± 0.04; all 12 p’s < 0.025) which cannot be explained by the absolute-to-relative assumption. The difference between the observed relative distribution and that predicted by the absolute-to-relative assumption was significant for each subject (two-sided Kolmogorov-Smirnov test, all 12 p’s $< 4.7 \times 10^{-4}$).

The retrospective Bayesian decoding theory explained the data of individual subjects

Fig. 6 of the text shows that the retrospective Bayesian decoding theory explained one subject's joint and relative distributions of the 2-line condition. The results of the other 11 subjects are plotted in Fig. S3. Different rows are for different subjects, and the left and right columns are for the joint and relative distributions, respectively. The gray dots and histograms show the observed joint and relative distributions, whereas the green dots and histograms are the corresponding simulations of the retrospective Bayesian decoding theory. The blue dots and histograms indicate simulated memory representations before the application of the Bayesian prior (see text). The retrospective Bayesian decoding theory explained the data well for 10 of the total 12 subjects. For the other two subjects (the second and tenth rows of Fig. S3), the simulated angular differences were smaller than the observed. This discrepancy can be eliminated by introducing a free parameter (see text).

Control conditions and analyses

One might argue that our reporting method, which involved rotation of the marker dots, introduced large variabilities into the measured orientation distributions. In other words, subjects could not place the marker dots at intended orientations because of limitations in their fine motor control or the experimental setup or both. We included two control conditions to exclude this possibility. They were identical to the 1-line test conditions except that the stimulus line in each trial stayed on until the subjects reported the perceived orientation. Since visual feedback was available before reporting, these control conditions measured subjects' ability to place the marker dots at the intended orientation. The absolute distributions from these control conditions have little variances or biases. The mean SD and bias of all 12 subjects were only 0.35^o and 0.15^o, respectively. Moreover, no subject had overlap between the 50^o and 53^o distributions. Therefore, for our setup, variabilities from fine motor control and experimental setup were negligible and subjects were able to place the marker dots at intended orientations.

One might also argue that in the 2-line condition, the subjects had to rotate the marker dots by a large angle to report the first orientation, but by a small angle to report the second orientation relative to the first, generating more motor noise for the first absolute judgment than for the relative judgment. This was unlikely because as shown above, the control conditions excluded the possibility that motor noise prevented subjects from reporting their intended orientations. To further exclude this possibility, when two subjects (one naïve) did the experiment using bite bars (see Methods), we modified the 2-line condition such that the initial marker-dot orientation was on either side of the first line with an angle equal to half of each subject's mean reported relative orientation in the original 2-line condition. Their rotation amplitudes for reporting the first line were 1.1^o and 3.8^o *smaller* than those for the second line. We analyzed data in the exactly same way as before, and found that for each subject, the relative distribution still significantly differed from that predicted by the absolute-to-relative assumption (two-sided Kolmogorov-Smirnov test, p = 4.9×10^{-13} and 1.4×10^{-15} for the two subjects). This ruled out the rotation-amplitude based explanation.

Possible learning effects and inter-trial correlations

We concluded in the main text that applying ordinal relationship to constrain absolute decoding not only is justified by our working-memory considerations but also accounts for the data. If we used the actual 3° angular difference, instead of the ordinal relationship, to produce the prior in the model, then the predicted angles would be close to 3°, much smaller than the observed angles in the 2-line condition (Fig. 4b). One might argue that over a block of the 2-line trials, subjects could gradually learn the true angular difference of 3° and thus show a trend toward 3° for later trials in the block. We tested this possibility by dividing each subject's 2-line trials into the earlier and later halves. Fig. S4, left panel, plots the mean angular difference for the later half against that for the earlier half, showing no significant difference (two-tailed Wilcoxon signed rank test, $p = 0.27$). There was actually an opposite (but insignificant) trend, with a little larger angular difference in the later half than the earlier half. Similarly, the SDs of the angular difference in the early and later half trials did not show a significant difference (Fig. S4, right panel; $p = 0.97$). For both the 1-line and 2-line conditions, we also checked the possibility that subjects gradually learned the true absolute orientation of each line, which would predict smaller biases and variances of the absolute distributions in the later half trials compared with the early half. We again failed to find such trends (Fig. S5, all p's > 0.05). The only comparison that almost reached significance is the SD of the 53° line of the 2-line condition ($p = 0.052$), but again the trend is in the opposite direction of the learning prediction: the SD of the later half trials was a little larger than that of the early half. Note that we tested the same hypothesis family (differences between the two halves of trials) multiple times, and did not find a significant result even without the Bonferroni correction. (Such multiple comparisons did not occur in other tests of this paper.) We conclude that the subjects did not show significant learning of the true angular difference in the 2-line condition or the absolute orientations in the 1-line and 2-line conditions.

Finally, we examined inter-trial correlations as possible shorter-term learning effects by plotting a quantity for trial *n* against that for trial *n-1*. If, for example, a larger value in trial *n-1* more likely led to a smaller value in trial *n*, then a negative correlation would occur. The quantity we examined was the reported orientation for each of the 1-line conditions, and the angular difference for the 2-line condition. This analysis was done for each of the 12 subjects separately, with three tests (Pearson correlation) per subject. If the significance level was set to (0.05/3) according to the Bonferroni correction, then only one subject in one test (the 53° 1-line condition) showed a significant, positive correlation. Without the Bonferroni correction, two subjects showed significant, positive correlations for the 50°1-line condition, three subjects showed significant, positive correlations for the 53°1-line condition, and two subjects showed significant correlations (one positive, one negative) for the 2-line condition. The plots for the naïve subject in Fig. 2 of the main text are shown in Fig. S6 as an example. The open dots and crosses in the left panel are for the 50° and 53°1-line conditions, respectively. The right panel is for the 2-line condition where the small number of points away from the main cluster are produced by the four trials of incorrect ordinal discrimination. This subject was one of the two who showed a significant (without the Bonferroni correction), positive inter-trial correlation for the 50°1-line condition (open dots in the left panel of Fig. S6). We conclude that for each condition, most subjects did not show inter-trial correlations. There might be a weak trend of positive correlation but it did not produce consistent learning of stimulus statistics over a block of trials (Figs. S4 and S5).

The lack of learning is not surprising because we randomized the initial marker-dot position and did *not* provide any feedback to the subjects at any time.

Reference

1. Girshick AR, Landy MS, & Simoncelli EP (2011) Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nat Neurosci* 14(7):926-932.

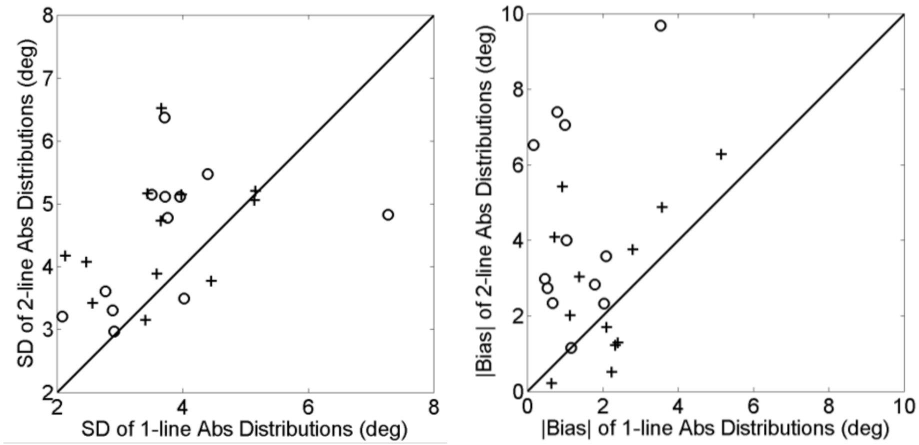

Fig. S1. Related to Fig. 2. Comparison of variances and biases between the 1-line and 2-line absolute distributions. Left panel: the SD of the 2-line absolute distribution vs. that of the 1-line absolute distribution for each subject and stimulus orientation. Right panel: the bias magnitude of the 2-line absolute distribution vs. that of the 1-line absolute distribution for each subject and stimulus orientation. In each panel, the results for the 50^o and 53^o stimulus orientations are represented by open dots and crosses, respectively.

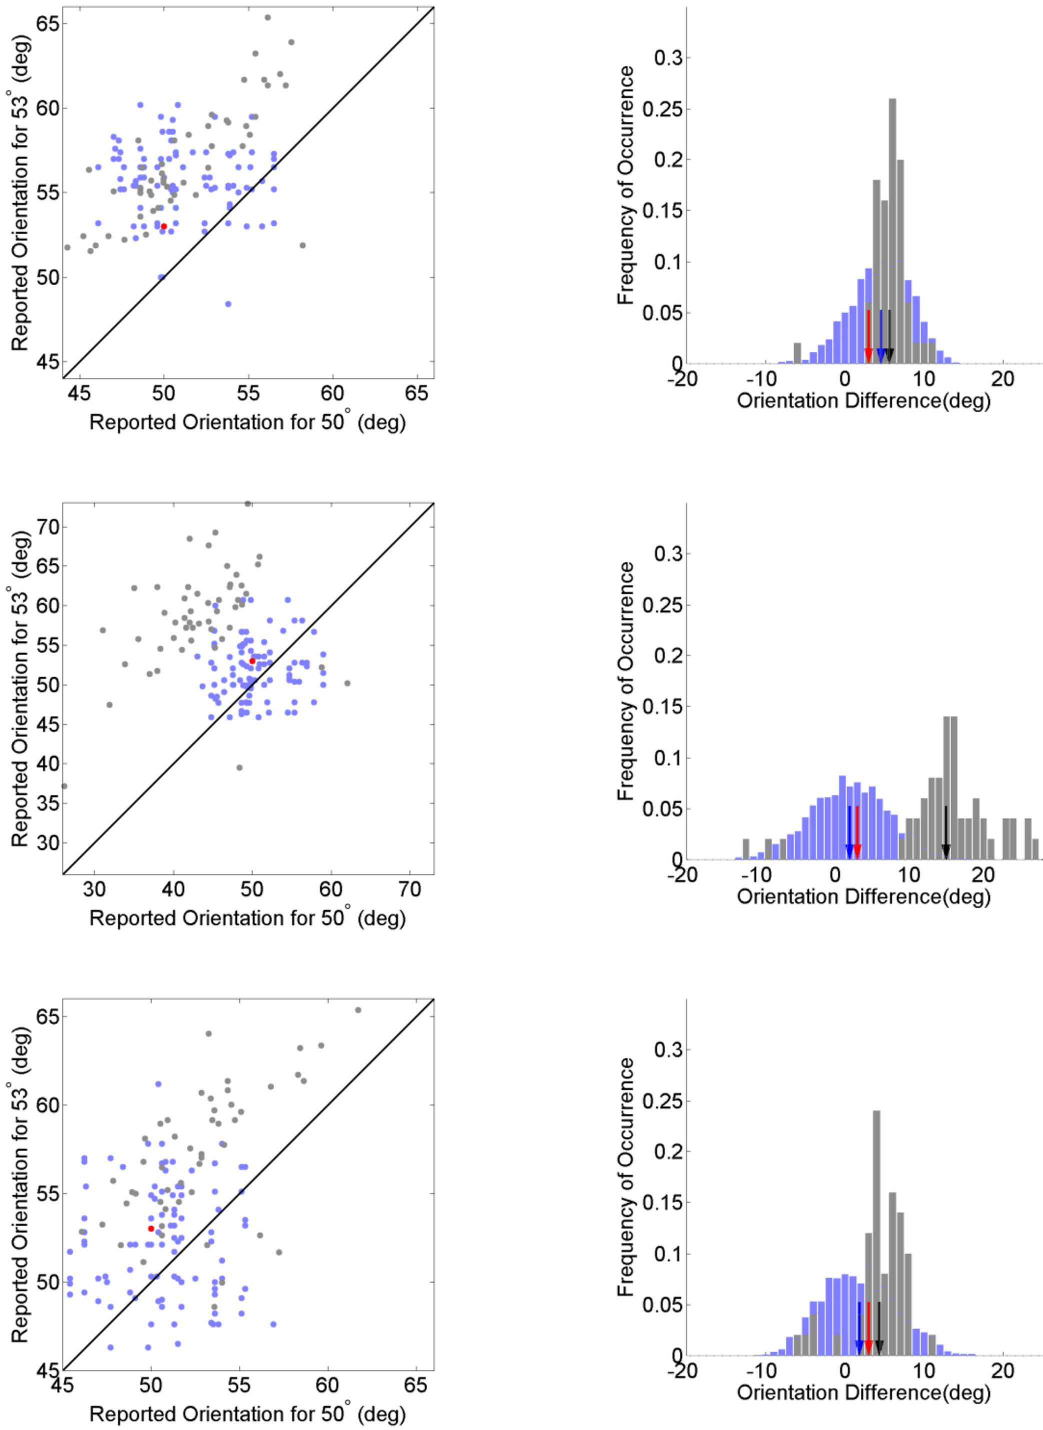

Fig. S2. Continue to next page.

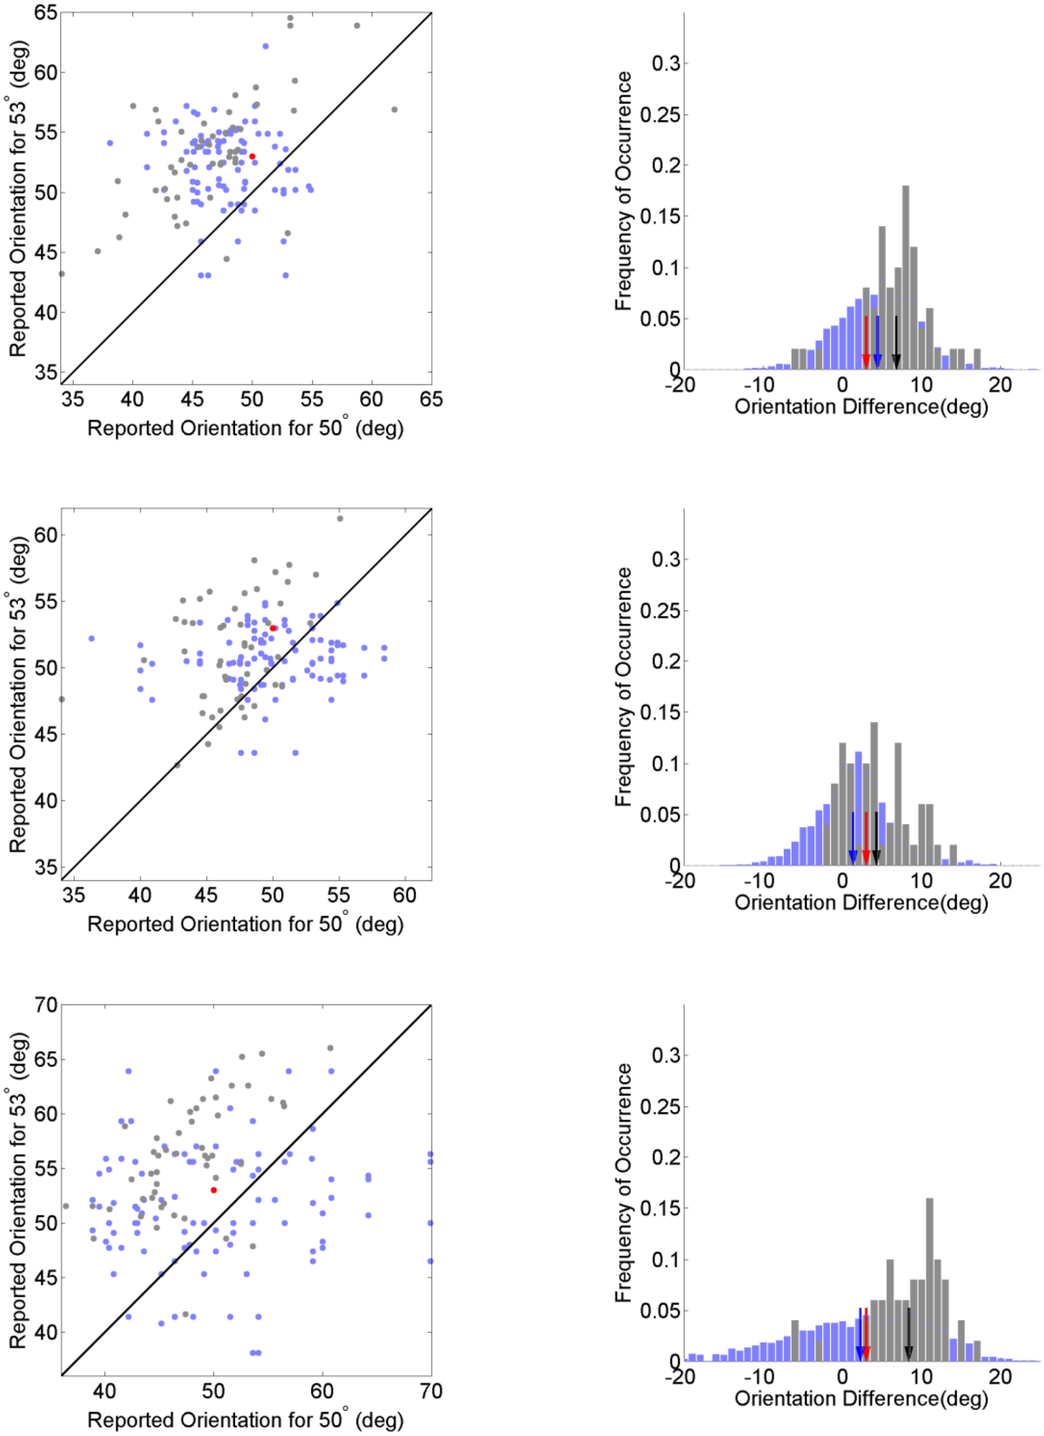

Fig. S2. Continue to next page.

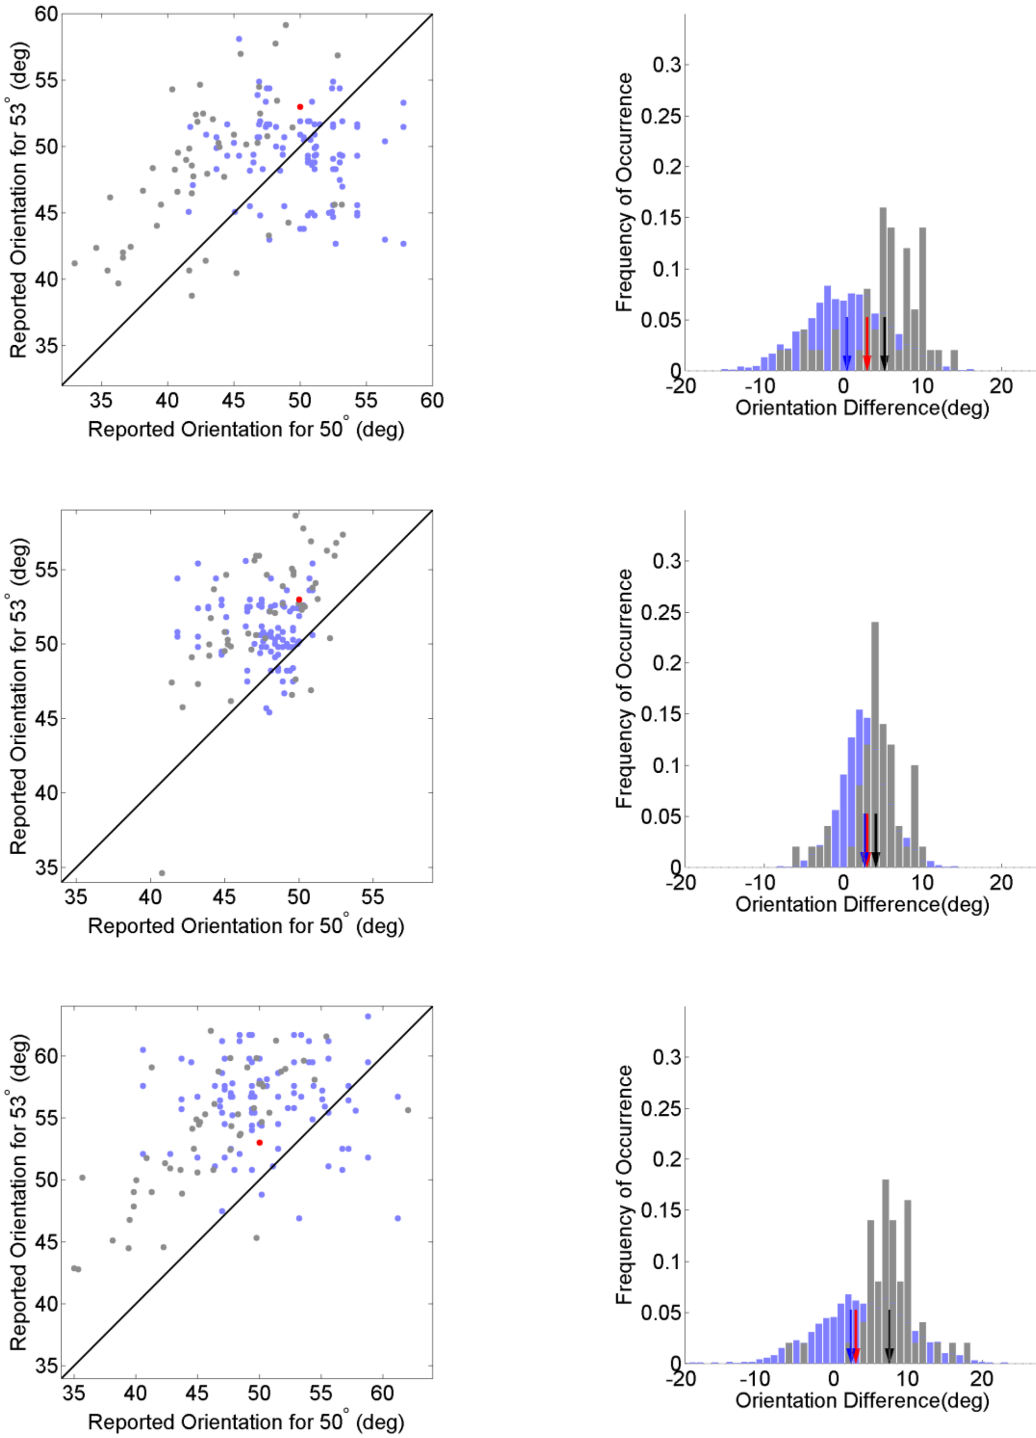

Fig. S2. Continue to next page.

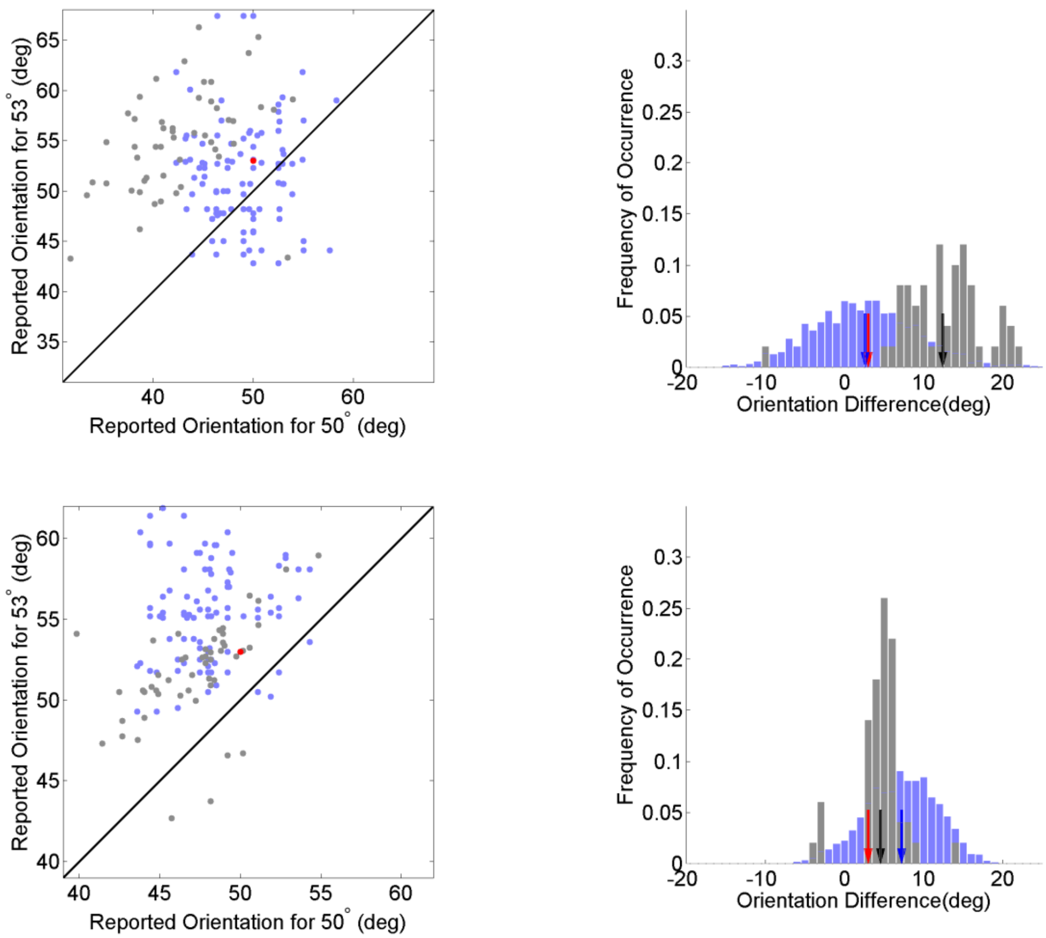

Fig. S2. Related to Fig. 3. Comparison between the absolute-to-relative assumption's prediction (light blue dots and histograms) and the data (gray dots and histograms) for the other 11 subjects. The two panels of each row are for one subject. The left and right columns show the joint and relative distributions, respectively. The format is identical to that of Fig. 3 in the text.

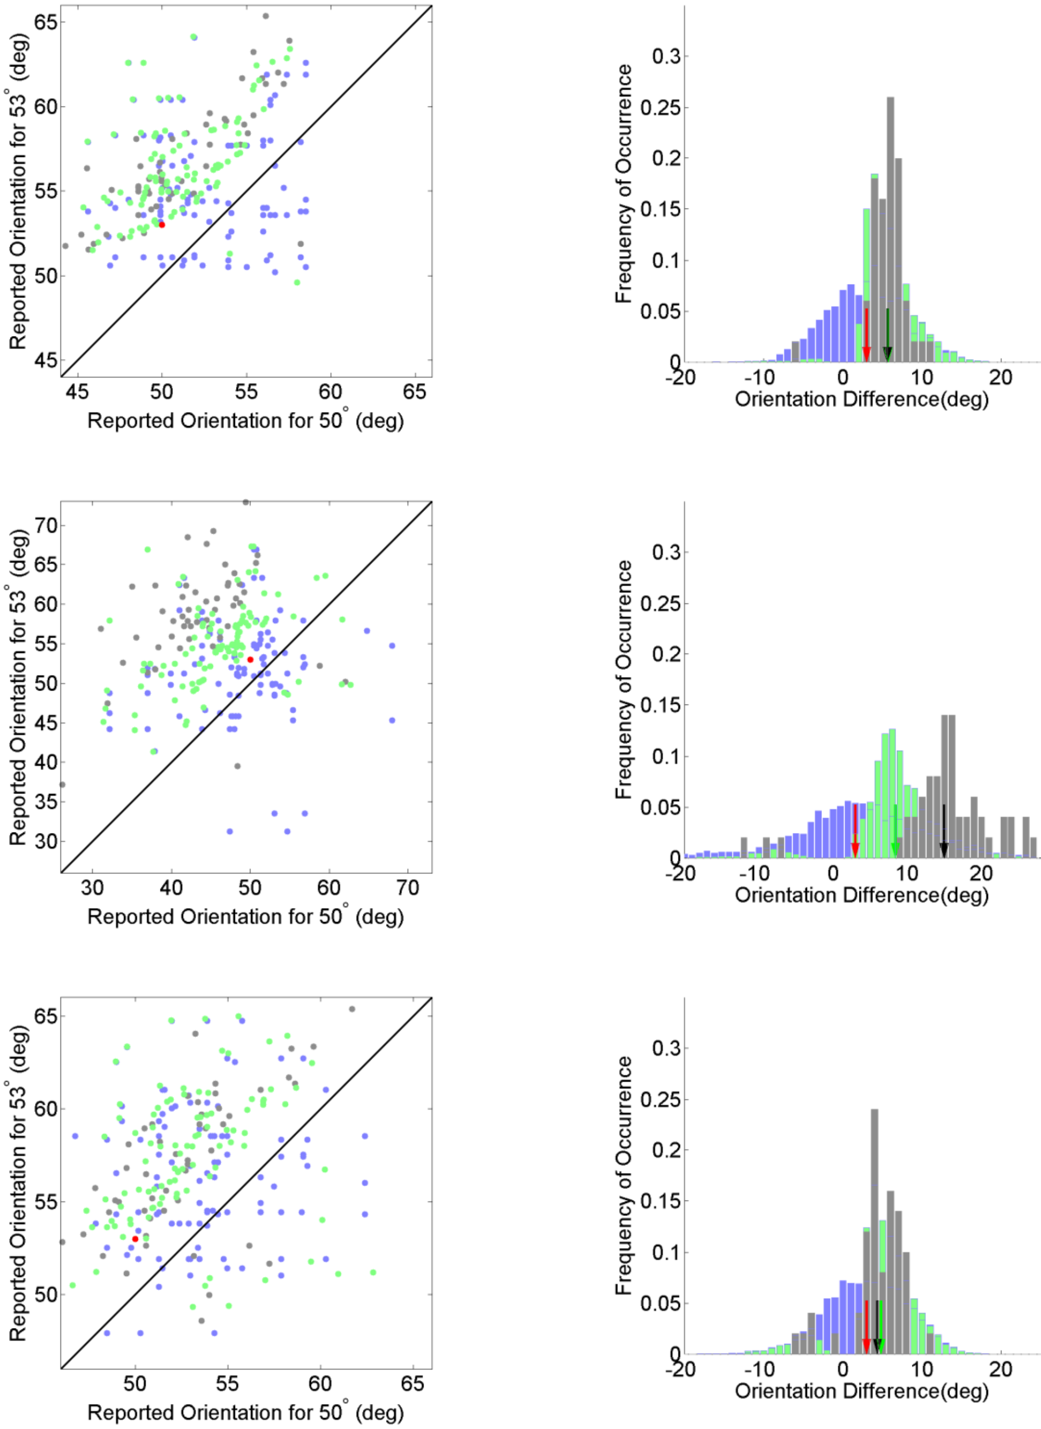

Fig. S3. Continue to next page.

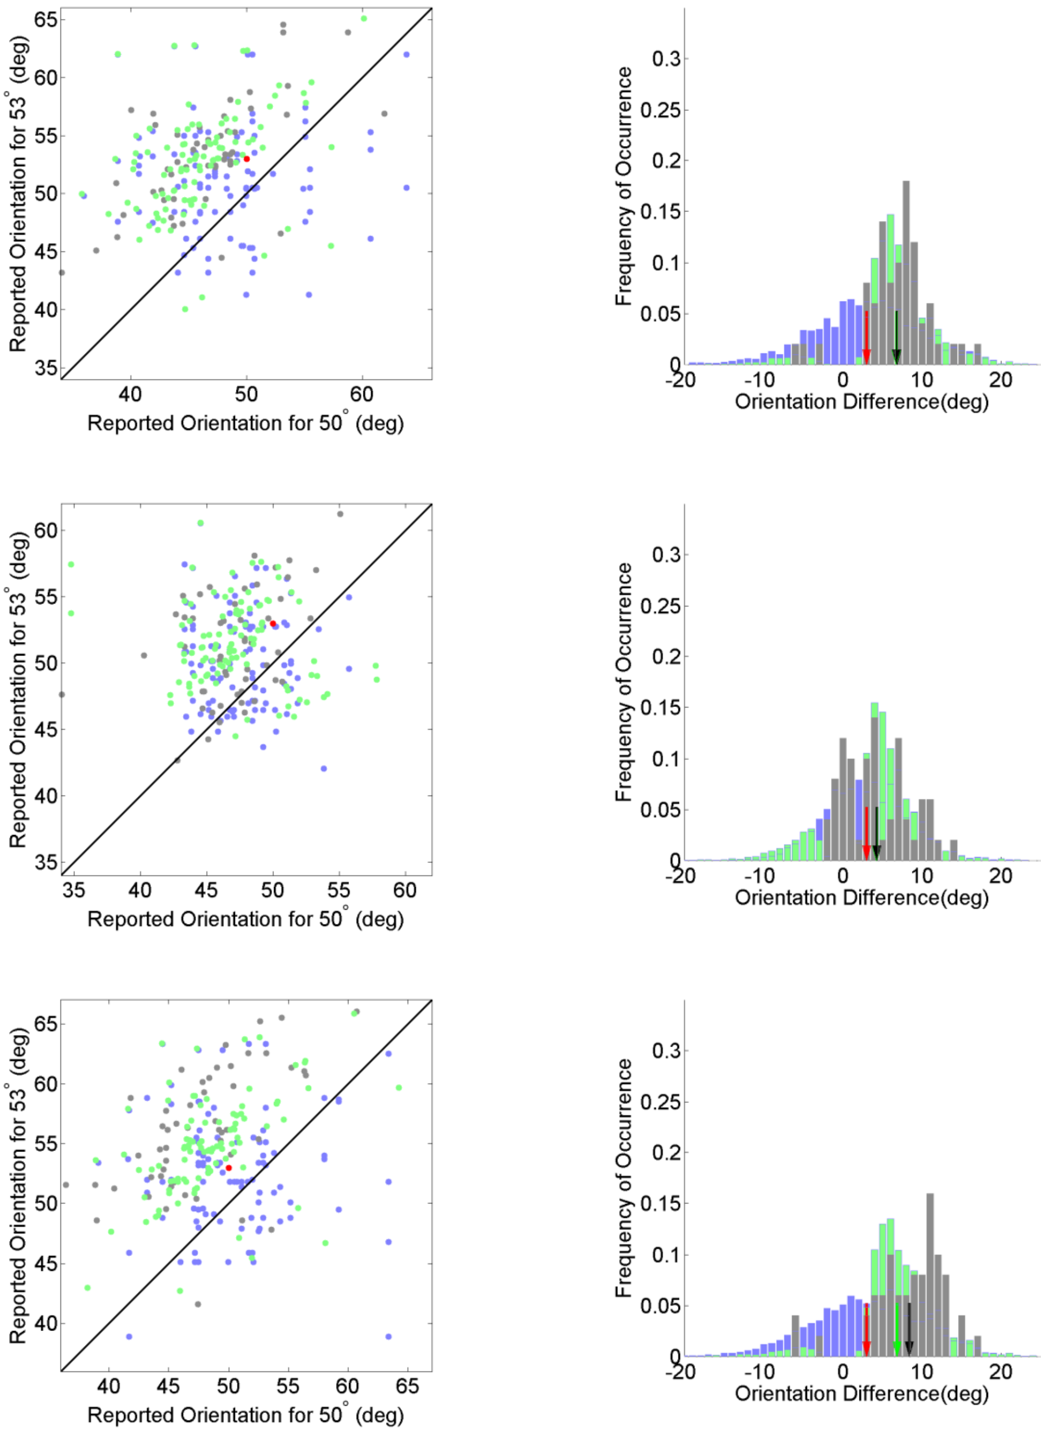

Fig. S3. Continue to next page.

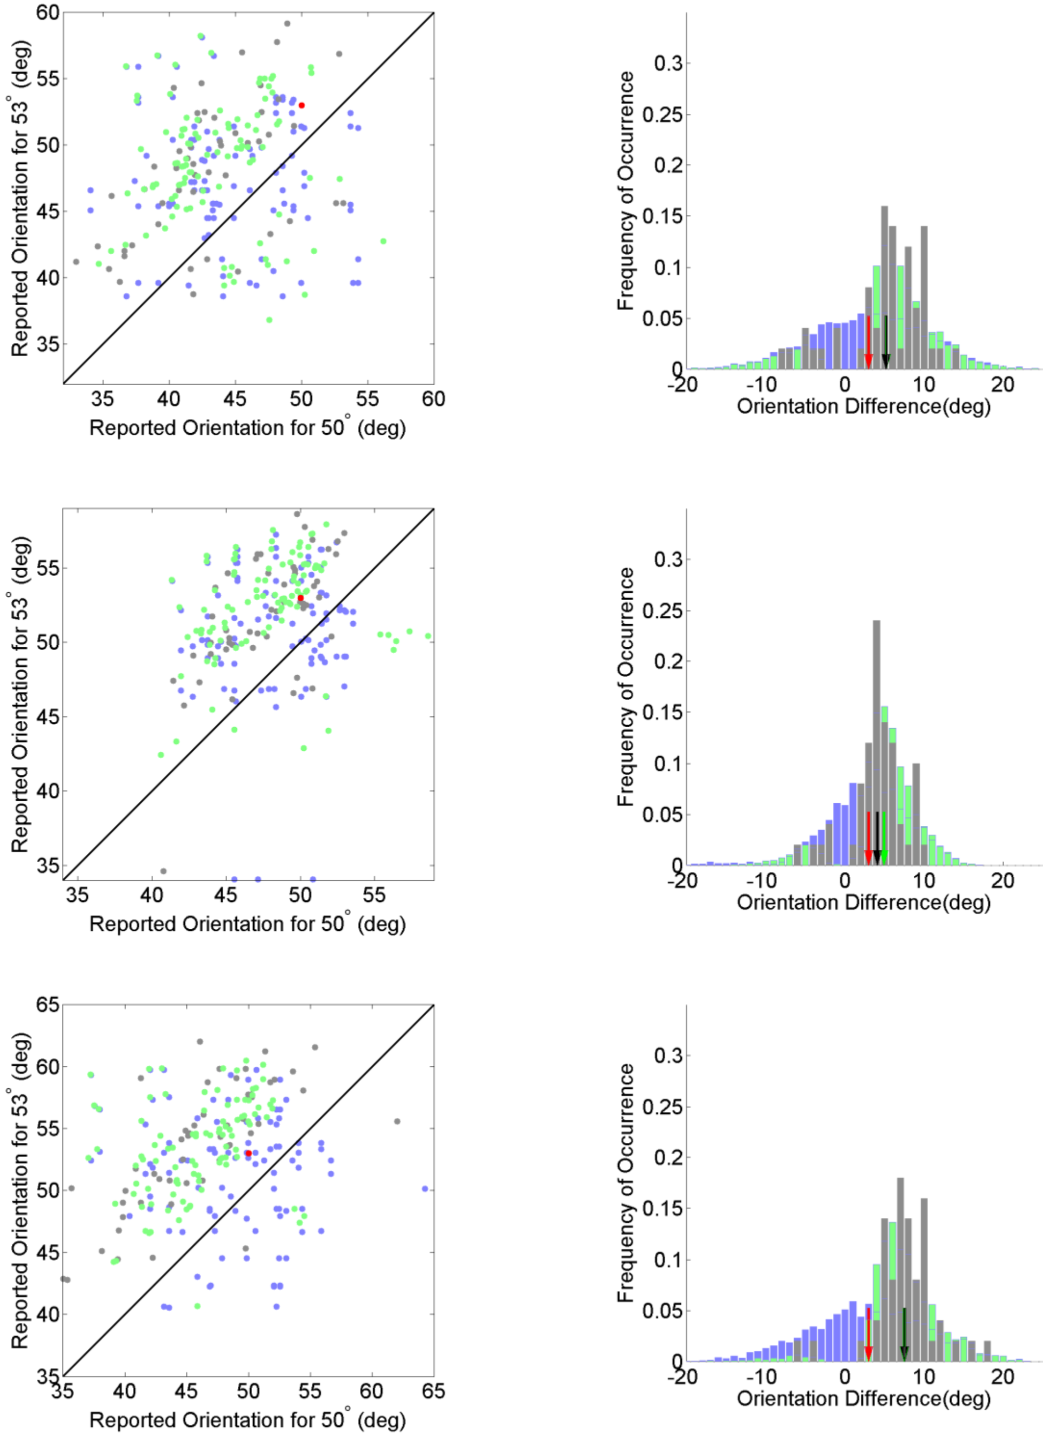

Fig. S3. Continue to next page.

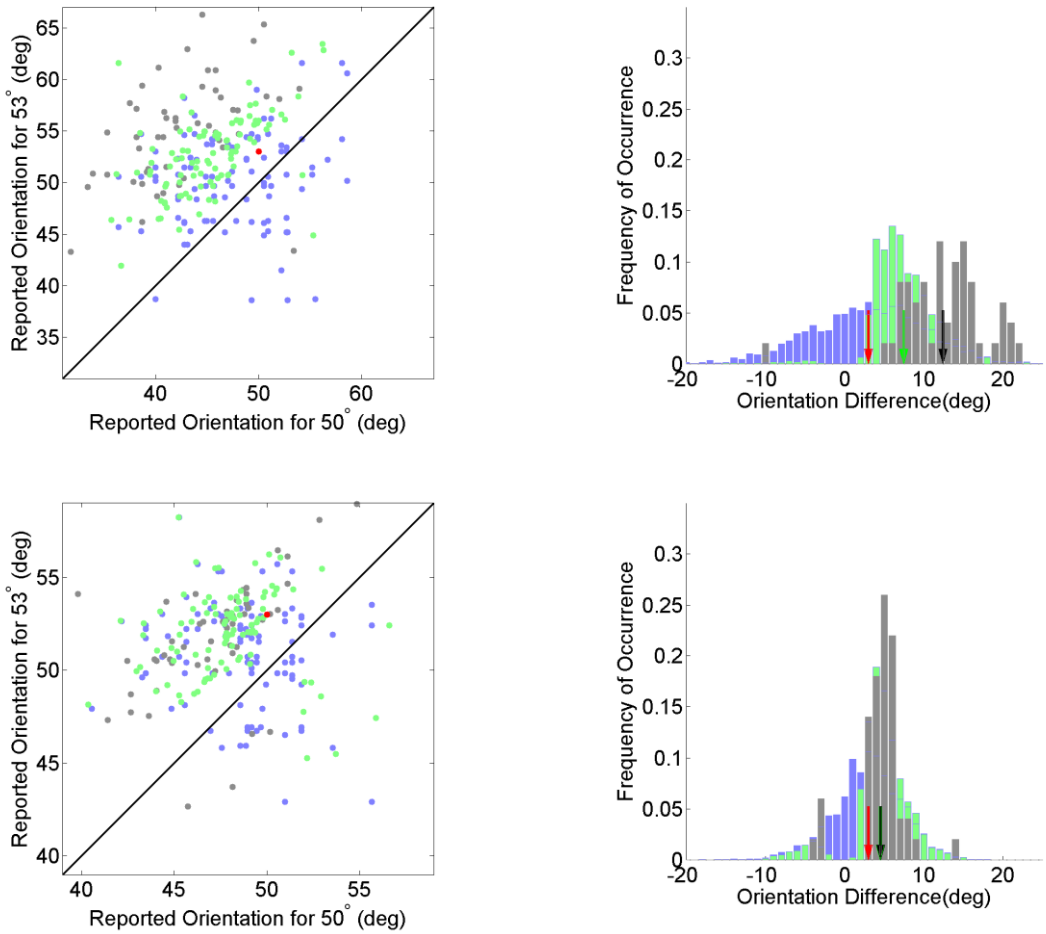

Fig. S3. Related to Fig. 6. Comparison between the retrospective Baysian decoding (green dots and histograms) and the data (gray dots and histograms) for the other 11 subjects. The two panels of each raw are for one subject. The left and right columns show the joint and relative distributions, respectively. The format is identical to that of Fig. 6 in the text. For some relative distributions, the green arrows are covered by the black arrows.

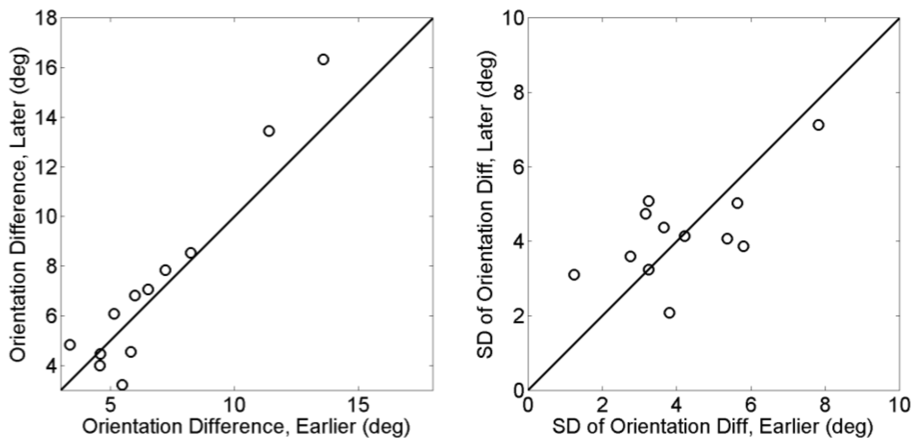

Fig. S4. Comparison of orientation difference (left panel) or its SD (right panel) between the earlier and later halves of the trials in the 2-line condition for all 12 subjects.

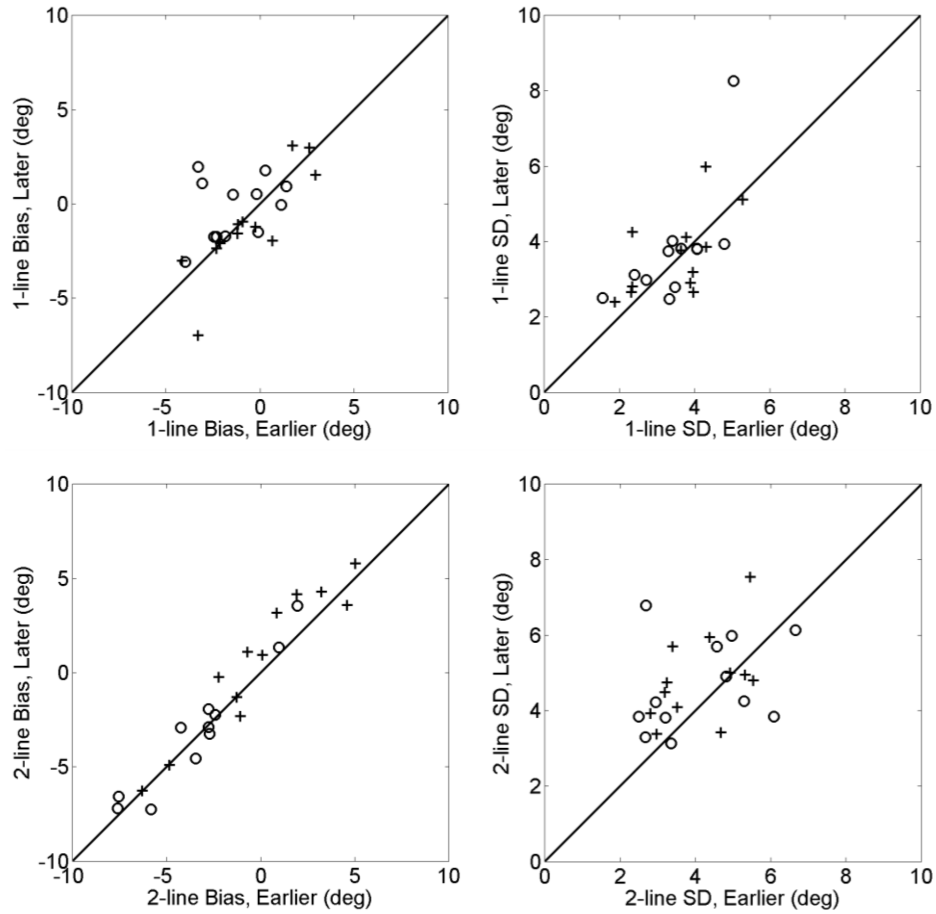

Fig. S5. Comparison of orientation bias (left column) or SD (right column) between the earlier and later halves of the trials in the 1-line conditions (top row) and the 2-line condition (bottom row) for all 12 subjects. In each panel, the results for the 50^{o} and 53^{o} stimulus orientations are represented by open dots and crosses, respectively.

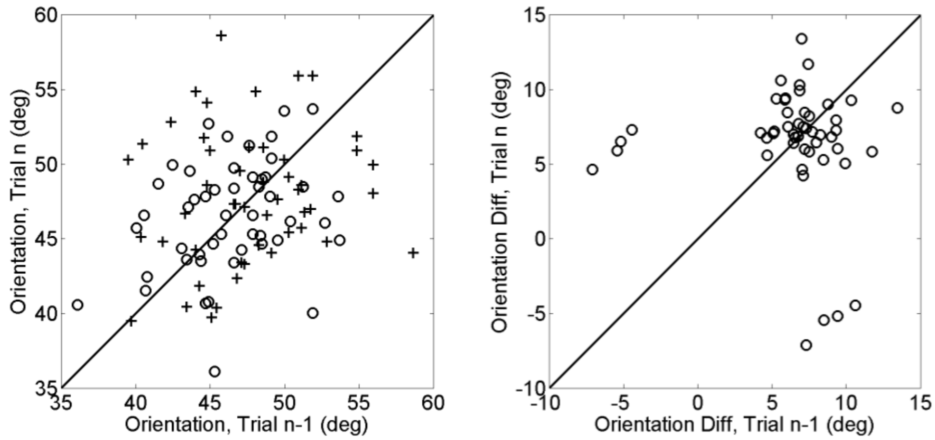

Fig. S6. Trial *n* plotted against trial *n-1* for the reported orientations in the 2-line conditions (left panel) or the reported orientation difference in the 2-line condition (right panel) for the naïve subject of Fig. 2 in the main text. In the left panel, the results for the 50^{o} and 53^{o} 1-line conditions are represented by open dots and crosses, respectively.