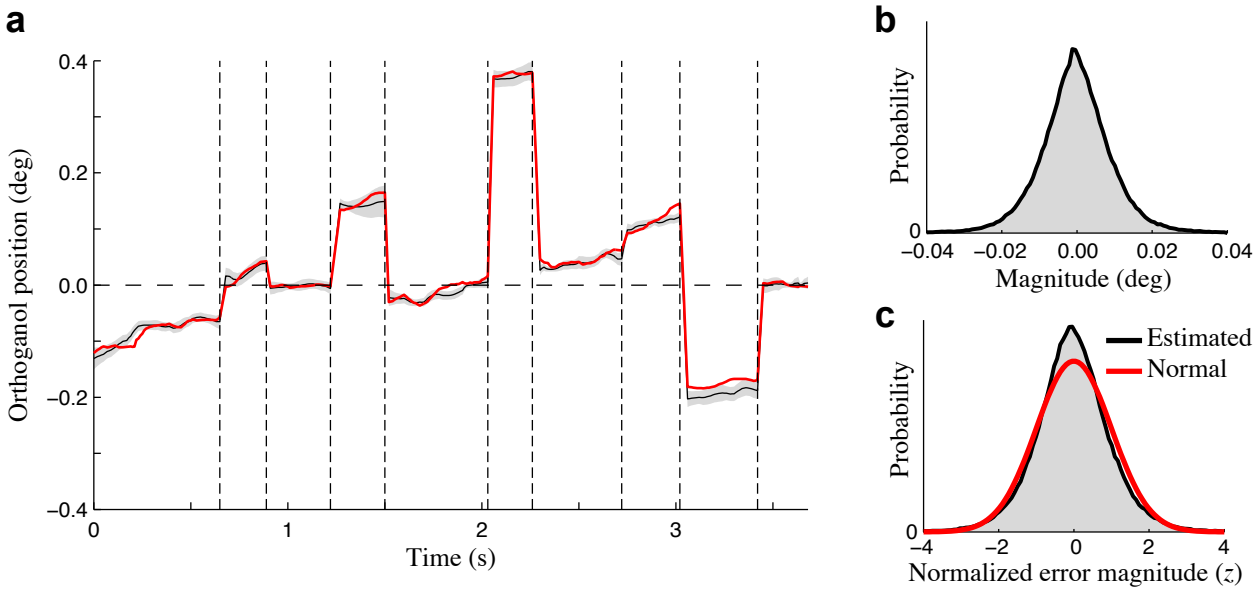

Supplementary Figure 1: Validation of eye position uncertainty estimates
(a) To test the validity of our eye position uncertainty estimates, we simulated the activity of a population of neurons matched to an example recording form the array (as in Fig. 6). The inferred eye positions (black) were in close agreement with the 'ground truth' eye position (red) used for the simulation. Note that for the example trace the differences stay largely within the estimated uncertainty intervals. **(b)** Errors in the estimated eye position were tightly distributed around zero (SD = 0.0082°, or 0.49 arc min). **(c)** To determine if the estimated eye position uncertainties were consistent with the observed error distribution we computed the uncertainty normalized errors (z-scored). The distribution of normalized errors was similar to the expected distribution under the assumption of Gaussian errors (red), although with somewhat heavier tails reflecting the fact that the estimated eye position distributions were often non-Gaussian. This shows that the uncertainty estimates of our model accurately characterize the true likelihood of different eye positions, in the case when the assumptions of the model are satisfied (with simulated data).

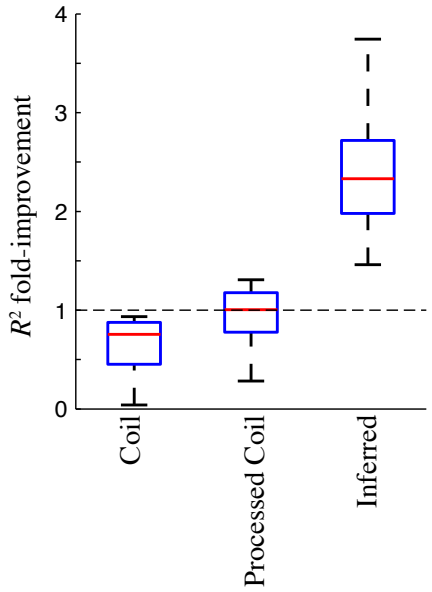

Supplementary Figure 2: Comparison of model fits using coil-based and inferred eye corrections
Comparison of R^2 values for all SUs (N = 38 SUs; 8 recording sessions) using different eye position signals to estimate the retinal stimulus. When using the coil-measured eye position signal we found that model R^2 values were significantly lower than assuming perfect fixation (median 0.76-fold difference, $p = 7.7\text{x}10^{-8}$; Wilcoxon signed rank test). We also computed a 'processed' coil signal by subtracting the median coil-measured eye position within each trial. This simple procedure eliminates a significant amount of erroneous slow drift, as illustrated by the consistent improvements in R^2 values compared to using the raw coil signal (median improvement 1.44-fold; $p = 7.7\text{x}10^{-8}$). Nevertheless, even with this preprocessing stage we did not observe significant improvements in R^2 compared with assuming perfect fixation (median improvement 1.01-fold; $p = 0.56$). In contrast, when using the inferred eye position signal, (cross-validated) R^2 improvements were observed in all SUs, ranging from 1.5-fold to 5.7-fold (median 2.3-fold; $p = 7.7\text{x}10^{-8}$).

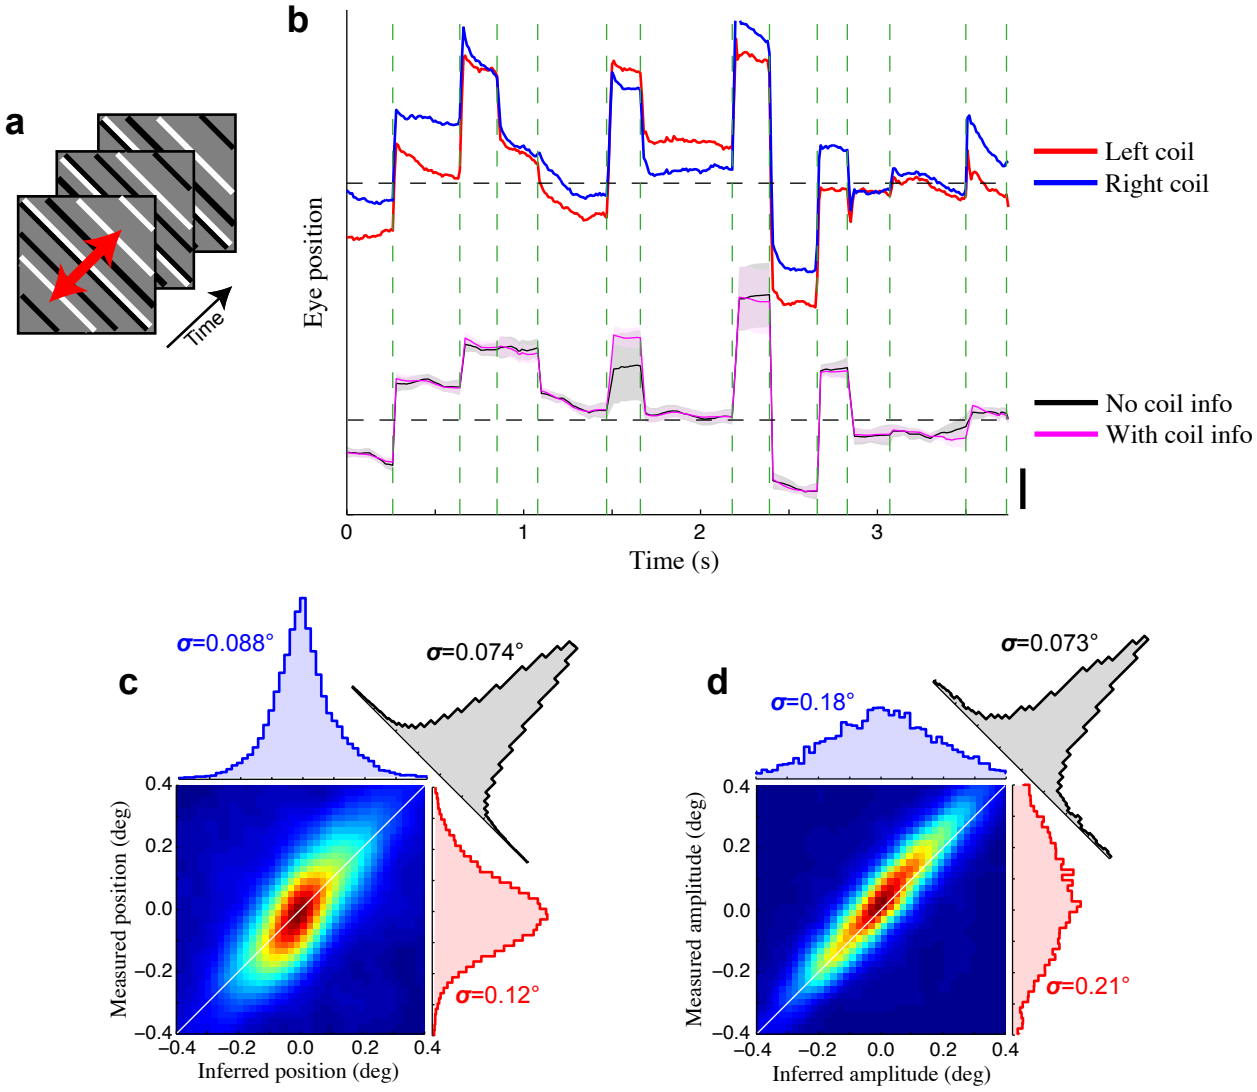

Supplementary Figure 3: Eye position inference using a linear electrode array
(a) For an example recording with the linear array, we used a stimulus consisting of random bars oriented at 135 deg, matched to the preferred orientations of the recorded units. We then inferred the orthogonal component of the eye position (red arrow). **(b)** Estimated eye positions were largely similar with (magenta) and without (black) using information from the eye coil signals. Information from the eye coils helped to refine the estimated eye position when evidence from neural activity was 'inconclusive'. The measured signals with the left (red) and right (blue) eye coils are shown above for comparison. Vertical scale bar is 0.1°. **(c)** Joint distribution of inferred and coil-measured (average of the two coils) eye positions. Inferred and coil-measured eye positions were strongly correlated (correlation = 0.73), though as with the planar array recordings, substantial differences remained. Color depicts the square root of the joint probability. **(d)** Inferred and coil-measured microsaccade amplitudes showed close agreement (correlation = 0.78), similar to that observed with the array recordings.

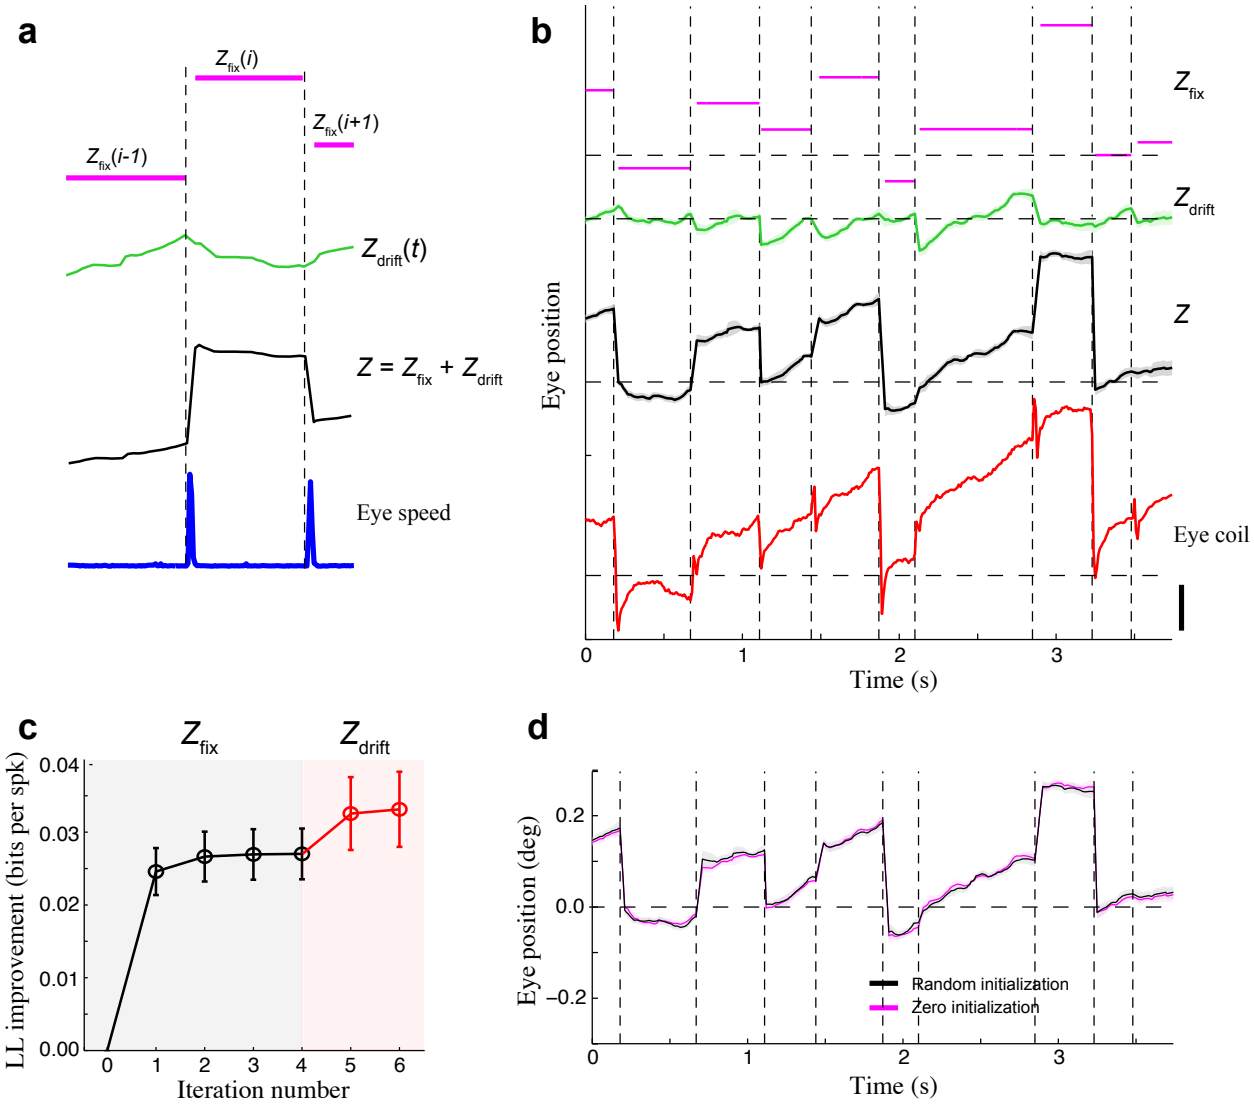

Supplementary Figure 4: Illustration of two-stage HMM estimation procedure
(a) Schematic illustration of the two-stage eye position estimation procedure. First, the average eye position within each fixation $Z_{fix}(t)$ is estimated (magenta). We then estimate a sequence of corrections $Z_{drift}(t)$ (green) capturing the within-fixation drift. The final estimate of eye position $Z(t)$ (black) is given by the sum of these two components. The instantaneous eye speed (as measured by the eye coils; blue) is used to detect saccade/microsaccade times, and hence to divide the data into fixations. **(b)** Estimates of the components described in (a) for an example trial. The simultaneous eye-coil measurement of eye position is shown in red. Vertical scale bar is 0.1°. **(c)** The log-likelihood of the stimulus model parameters **Θ** converged after a few iterations when first estimating Z_{fix}, as well as when subsequently estimating Z_{drift}. Here we plot the *LL* relative to its initial value (before corrections) in units of bits per spike. Error bars show the region mean ± SD across $N = 8$ recordings. **(d)** Estimated eye positions were largely independent of their initialization. We typically initialized Z_{fix} to zero (assuming perfect fixation; magenta); however we found nearly identical results when initializing Z_{fix} randomly (black). In the latter case, we initialized Z_{fix} to be Gaussian distributed (SD = 0.12°), but constant within each fixation.