Figure S7 - Convergence of the EM algorithm on held-out test data.

10-fold cross-validation tests show a significant improvement in the average log-likelihood per interaction (in bits), along 14 EM iterations. We show the improvement of the recognition preferences for different starting points. These include the prior knowledge parameters from Mandel-Gutfreund et al. (2001) and random starting points.