Supplemental Data:

Rapid extraction of mean emotion and gender from sets of faces

Jason Haberman and David Whitney

Supplemental Figures

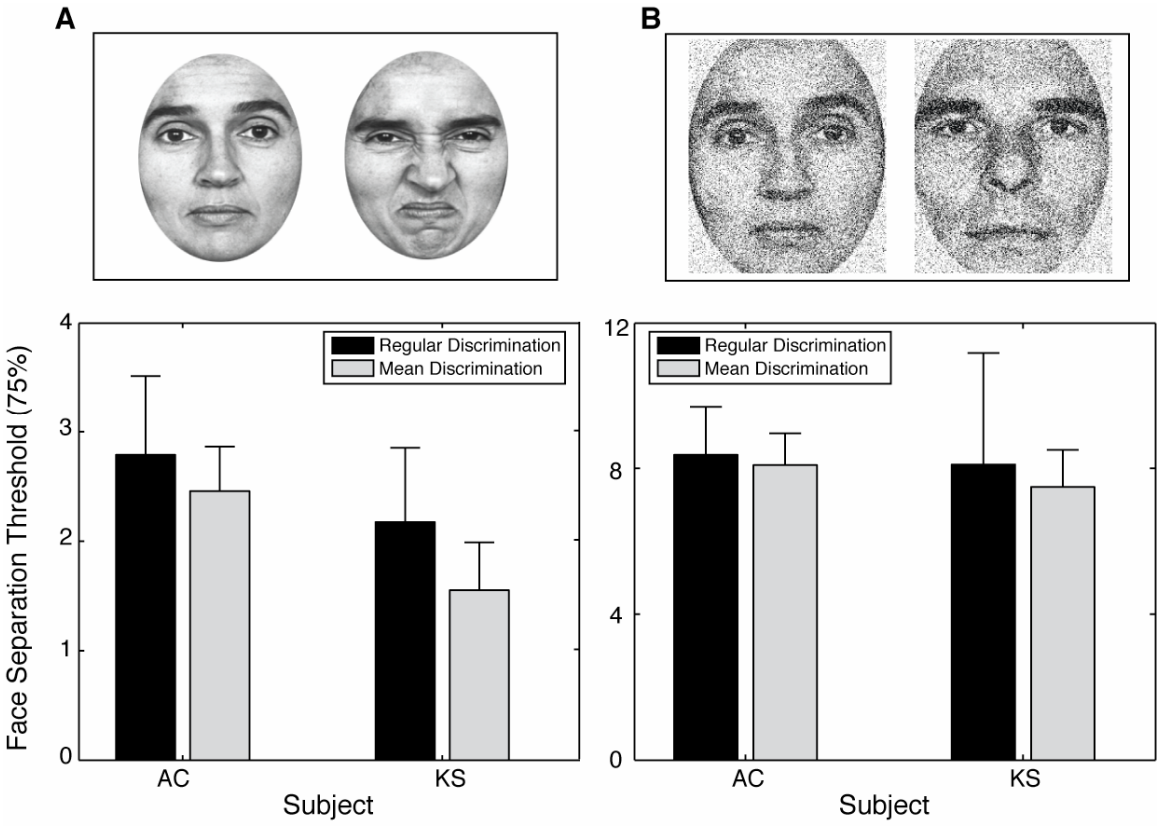

Figure S1. Stimuli and results for two follow-up experiments. We tested how well mean extraction would generalize to other face dimensions. For both neutral to disgusted morphs (**A**) and female to male morphs (**B**), participants were able to extract a mean

representation as precisely as they could discriminate any two faces, within the same dimension. In the gender experiment, 30% Gaussian noise was added to reduce the influence of low-level feature cues, such as contrast differences. Similarly precise mean discrimination was obtained with and without the added noise. The results in both experiments here parallel the finding for happy to sad morphs and suggests that statistical representation extends to other emotions and even other face dimensions.

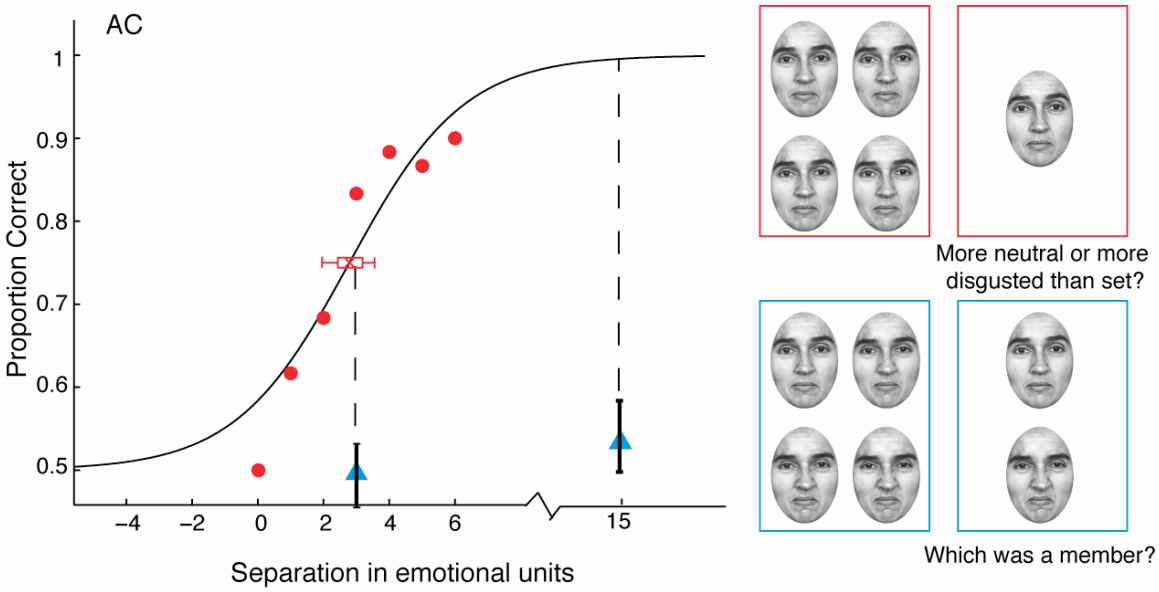

Figure S2. Neutral to disgusted discrimination data for two observers. The solid red circles indicate observers' ability to discriminate which of two faces was more neutral (example stimuli outlined in red; see Figure 1A for task details). To achieve 75% correct performance, subject AC required a separation of 2.2 emotional units. The triangles indicate performance on two conditions of the 2-AFC membership task (example stimuli outlined in blue; see Figure 1C for task details). In one condition, test faces were separated by three emotional units. In the other condition, test faces were separated by at

least 15 emotional units (right test panel, outlined in blue). In both conditions, the test faces were easily discriminable, as indicated by the vertical dashed lines. The same pattern of results was found for other subjects and other stimuli, including happy to sad morphs (Figure 1D). That is, discrimination of test faces separated by three units (Figure 1D) was significantly better than 2-AFC membership identification of the same faces. Observers were at chance in all versions of the 2-AFC experiment, indicating that they did not have a representation of the individual set constituents. This was true despite having a precise representation of the mean emotion of the set (Figure 1E and S1A).

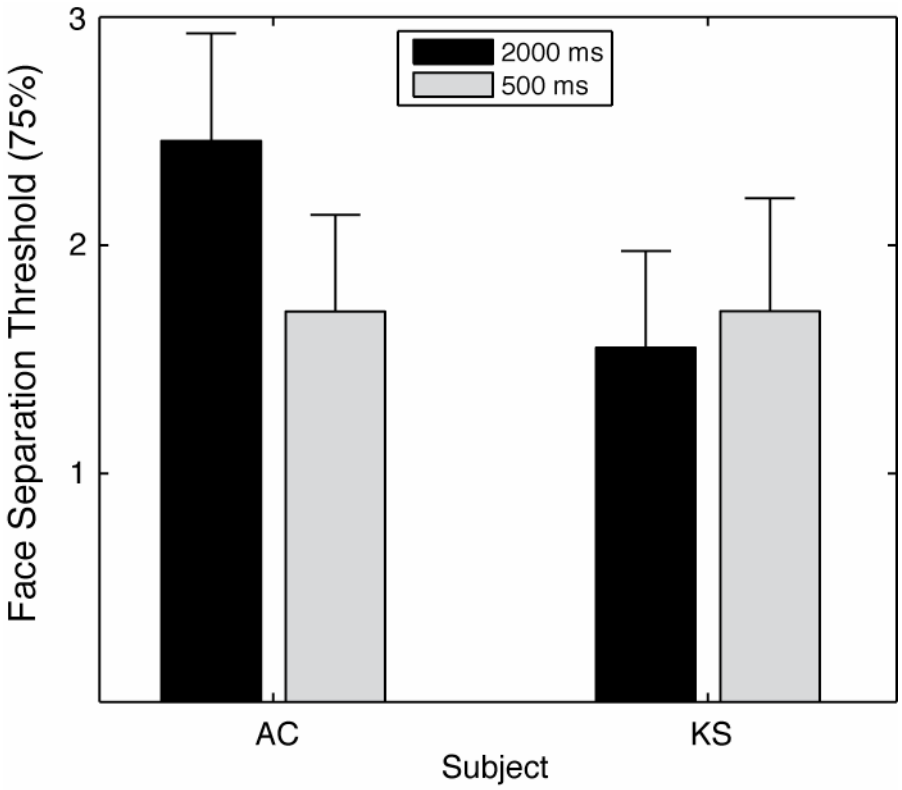

Figure S3. Mean discrimination performance for neutral to disgusted faces as a function of stimulus duration. Observers viewing a set containing up to 16 faces for only 500 ms

were able to extract a precise representation of the mean emotion. Performance at 500 ms did not differ significantly relative to performance at 2000 ms ($P > 0.05$, based on 5000 Monte Carlo simlulations [11]). Error bars indicate 95% confidence intervals.

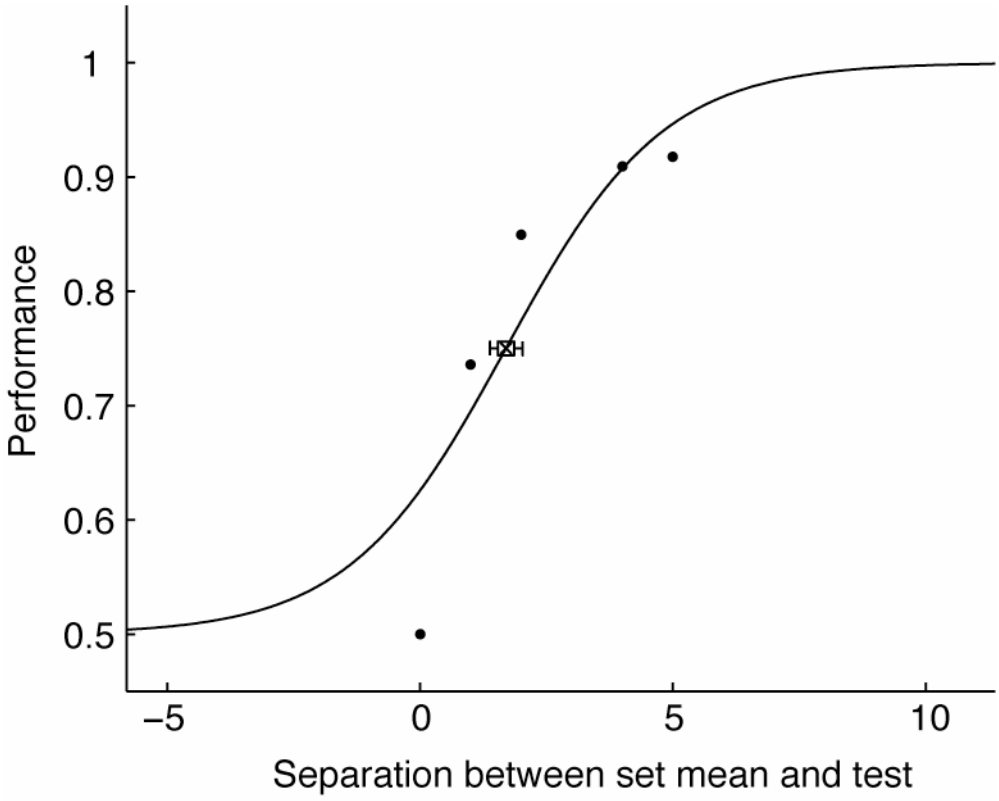

Figure S4. Performance collapsed across 11 naïve observers on mean discrimination. As in experiment 3, observers were asked whether the test face was happier or sadder than the mean emotion of the preceding set. Observers saw each possible set mean only once (a total of 32 trials per observer). The mean of each set was novel on every trial (no repeated means), and the faces were novel to each observer. Despite this limited exposure, observers were able to extract a precise representation of the mean emotion of each set, suggesting that this is not a statistical learning or prototype effect. We modeled

what the mean discrimination data would look like if observers had formed a prototype of the entire range of morphs, and found that, regardless of how well the prototype was coded, mean discrimination performance would have been significantly poorer than found here.