A GENERAL RECOGNITION THEORY STUDY OF RACE ADAPTATION

Leslie M. Blaha
Air Force Research Laboratory, Wright-Patterson AFB, Ohio, Leslie.Blaha@wpafb.af.mil

Noah H. Silbert
Center for Advanced Study of Language, University of Maryland

James T. Townsend
Department of Psychological and Brain Sciences, Indiana University, jtownsen@indiana.edu

Abstract

Studies of race aftereffects show that adaptation biases responses away from an adapting stimulus. However, it remains unclear if shifts in response frequencies result from changes in perceptual representations or in decisional mechanisms supporting race classification. General Recognition Theory (GRT) provides a modeling framework within which we investigated adaptation-induced changes on perceptual and decisional mechanisms. In a series of experiments, we replicated previous findings that adaptation shifts perceived features away from the adapting stimulus and showed similar effects for skin tone adaptation. GRT modeling was based on five complete-identification tasks. Baseline models derived from the no-adaptation condition found positive correlations between features and skin tone within and across stimuli. Adaptation-induced changes in representation from four adaptation conditions revealed shifts in perceptual representations away from adapting stimuli, variability in the within-stimulus correlations, and shifts in the decision bounds toward the adapting stimulus.

Adaptation Aftereffects

It is well known that extended exposure to a stimulus results in adaptation to that stimulus, which can produce aftereffects that alter the perception of subsequently presented stimuli. For example, visual adaptation to a distorted (e.g., extremely compressed or expanded) face influences the perceived distortion of an unaltered face (Burkhardt, et al, 2010). However, it is not clear if adaptation aftereffects are the result of changes in perceptual or decisional mechanisms. Adaptation aftereffects may result from a shift in the perceptual norm to which stimuli are compared or a change to the representational distances between percepts (e.g., Burkhardt, et al, 2010; Leopold, et al., 2001; Rhodes, et al., 2003). Alternatively, adaptation may influence decision-making by shifting a category boundary or gain control mechanism (e.g., Anderson & Wilson, 2005; Webster, et al., 2004). Of course, it is possible that adaptation has both perceptual and decisional influences.

The present work utilizes adaptation to cues to the racial categories of faces, which is known to produce aftereffects which alter the perceived race of faces (Webster, et al., 2004). The cues of interest here are facial features and skin tone cues to African-American and Caucasian race categories. We employ the General Recognition Theory framework in order to simultaneously investigate distinct perceptual and decisional adaptation aftereffects while also analyzing possible perceptual interactions between facial features and skin tone.

General Recognition Theory

General Recognition Theory (GRT) is a mathematical framework developed to model perceptual and decisional relationships between dimensions. GRT relies on two fundamental assumptions. First, it is assumed that perception is stochastic, such that the presentation of a stimulus produces a random perceptual effect, and so, over time, a perceptual distribution. Second, it is assumed that the perceptual space is exhaustively partitioned into response regions by decision bounds. The model predicts response probabilities as the multiple integrals of the perceptual distributions over the appropriate response regions.

In the simplest full GRT framework, the perceptual distributions and response regions correspond to the factorial combination of two levels on each of two dimensions, here facial features and skin tone (see below). Three kinds of relationships between dimensions are defined in GRT, two perceptual and one decisional. *Perceptual independence* is a within-stimulus (or local) notion of independence, and *perceptual separability* and *decisional separability* are between-stimuli (or global) notions of independence.

Perceptual independence holds within a given perceptual distribution if stochastic independence holds between the (random) perceptual effects on the two dimensions; it fails of stochastic independence fails. Perceptual separability holds if the marginal perceptual effects on one dimension are invariant across levels of another dimension. Decisional separability holds if decisions about one dimension do not depend on another dimension, or, equivalently, if a decision bound is parallel to the appropriate coordinate axis.

Race Adaptation in the GRT Framework

The present study had three main goals. First, we used the GRT framework to analyze interactions between facial features and skin tone in race perception. Second, we identified the possibly distinct perceptual and decisional effects of adaptation in race perception. Third, we investigated any differences in the effects of adaptation between Caucasian and African-American participant groups.

The study took place in three phases. In the first phase, stimuli were calibrated to each individual participant’s perceived race categories. In the second phase, standard unidimensional adaptation was studied, based on the individually calibrated stimuli, for each of the two dimensions under consideration (facial features and skin tone). In the third phase, interactions between facial features and skin tone and distinct perceptual and decisional adaptation effects were modeled using a hierarchical Gaussian GRT model.

Methods

General Methods for All Tasks

18 participants completed all three phases, 9 Caucasian and 9 African-American. 6 were male, 12 female, and 4 of the 18 were left-handed. All had normal or corrected-to-normal vision. In a given session, the initial adaptation exposure in all phases lasted for six minutes. Trials consisted of 300 ms of blank screen, a 250 ms presentation of a test face, a response window, and a "top-up" adaptation exposure lasting for six seconds. The stimulus sets consisted of faces produced by morphing a female Caucasian and a female African-American face (for the feature dimension) and by manipulating the grayscale brightness (for the skin tone dimension).

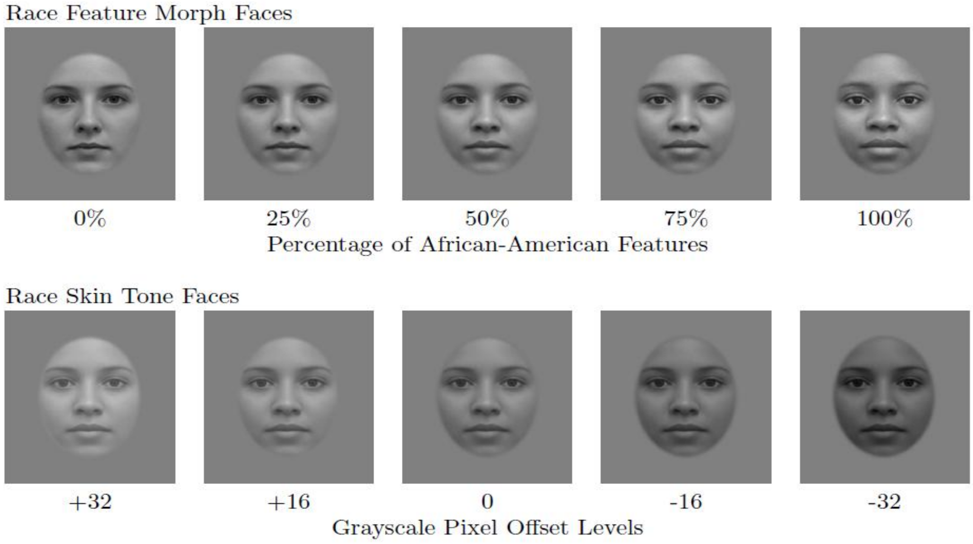

- Figure 1. Example Stimulus Faces

Phase 1: Stimulus calibration

Stimuli were calibrated in order to find threshold values on each dimension to construct the GRT stimuli for Phase 3. First, the points of subjective equality (PSE), or 50% threshold, on each dimension were estimated. Two-choice classification responses to a stimulus series along each dimension, like those depicted in Figure 1, were collected, and thresholds were estimated by fitting a Normal CDF to the observed proportions of African-American responses. Secondly, the 80% discrimination just noticeable difference (JND) was measured separately for each dimension using a same-different task. JND values were estimated by the fit of a corrected-for-guessing Weibull distribution to the percent of correct responses.

Phase 2: Unidimensional Adaptation

Unidimensional adaptation tasks were conducted on each dimension. The outcome measure in each case was the proportion of African-American responses. Participants completed baseline tasks with no adaptation as well as two adaptation tasks on each dimension wherein each of the endpoint faces on each dimension served as adaptors. Response rates reflected shifts away from the adapting dimension, such that adaptation to Caucasian (African-American) skin tone/features increased (decreased) the rate of responding African-American to each face.

Table 1 shows the results of a linear mixed effects model analysis of the unidimensional adaptation data. Robust effects of adaptation condition were observed for each dimension such that the two adapting images influenced classification responses in the expected, opposite directions. A statistically significant interaction between participant group and adaptation condition was also observed, but only with respect to facial features.

GRT and Adaptation

For each participant, a two-by-two, features-by-skin tone stimulus set was constructed such that the face space was centered on the PSE values for each dimension and the four test faces were each one JND apart on the individual dimensions. In the GRT task, trials proceeded as in the unidimensional adaptation task with respect to "top-up" adaptations, though in this task both classification responses on both skin tone and facial features were required on each trial.

- Table 1: Unidimensional Adaptation Results

	Adapt Features	Adapt Skin Tone
Group (G)	F(1,16) = 1.05, p = 0.32	F(1,16) = 0.64, p = 0.44
Adaptation Condition (A)	**F(2,32) = 141.2, p < 0.001**	**F(2,32) = 22.2, p < 0.001**
G x A	**F(2,32) = 5.3, p = 0.01**	F(2,32) = 1.8, p = 0.18

Identification-confusion data were analyzed with a newly developed hierarchical Gaussian GRT model (Silbert, 2011) fit to each participant's data across the baseline and four adaptation conditions. This model provides, on the one hand, a clear and concise picture of any perceptual interactions between feature and skin tone dimensions, and, on the other hand, distinct perceptual (*d'*) and decisional (*c*) measures of adaptation effects.

Results

Figure 2 shows a fitted face space from the baseline condition from a typical participant. The dimensions are defined as in the sample faces in each corner, so the lower-left quadrant corresponds to the Caucasian feature/Caucasian skin tone face, the lower-right quadrant to the African-American feature/Caucasian skin tone face, and so on. The typical pattern of results exhibited a mean-shift integrality such that the facial feature and skin tone cues to race identity interact synergistically. Faces with both Caucasian features and Caucasian skin tone and faces with African-American features and African-American skin tone are perceived highly accurately, whereas the faces with conflicting cues (e.g., Caucasian features and African-American skin tone, or vice versa) are frequently confused with one another and with the faces with consistent cues.

In order to probe the possibility that adaptation has distinct perceptual and decisional effects, we use the fitted models to calculate *d'* (sensitivity) and *c* (response bias) in the baseline and adaptation conditions. Note that each condition produces four *d'* measures (one

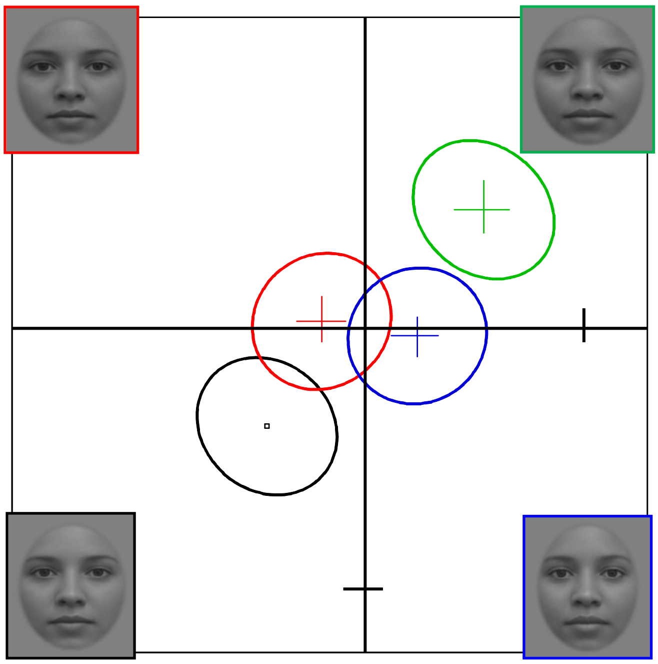

- Figure 2: Baseline GRT Face Space

for each level on each dimension) and two *c* measures. In order to test the supposition that adaptation influences perception and/or decision making, we take the differences between the baseline and adaptation condition for each of the *d'* and *c* measures.

Figure 3 shows changes in response bias (as measured by the criterion *c*) calculated as baseline *c* minus adaptation *c*. The left two panels show changes in decision criteria on the feature dimension, the right two panels on the skin tone dimension. The top two panels show changes in skin tone adaptation conditions, the bottom two panels in feature adaptation. Within each panel, the x-axis indicates the difference between the baseline and associated African-American adaptation condition, the y-axis between the baseline and the associated Caucasian adaptation condition. Zero values on each dimension indicate no adaptation-induced changes in *c*. Note that for the plots in Figure 3a, a positive value indicates a shift in criteria toward the Caucasian adaptor, and a negative value indicates a criterion shift toward the African-American adaptor. The data points show the median posterior modeled criteria values for the two adaptors on the dimension shown, with horizontal and vertical bars giving the statistical errors about the medians. Blue indicates Caucasian participants; red indicates African-American participants.

As can be seen in the top right and bottom left panels of Figure 3a, the fitted models exhibit the expected shifts in criteria along the adapted dimensions such that the criterion on the adapted dimension shifted toward the adapting images. For feature adaptation, when participants adapted to the African-American (Caucasian) feature face, *c* shifted right (left) along the x-axis. The same pattern was observed for Caucasian Participants on the skin tone dimension, such that the dark (light) adaptor shifted *c* toward itself. By way of contrast, the changes in tone criteria in feature adaptation and in feature criteria in tone adaptation are much more closely clustered near zero (top left and bottom right panels). Thus, adaptation induced minimal changes in the criterion on the non-adapted dimension.

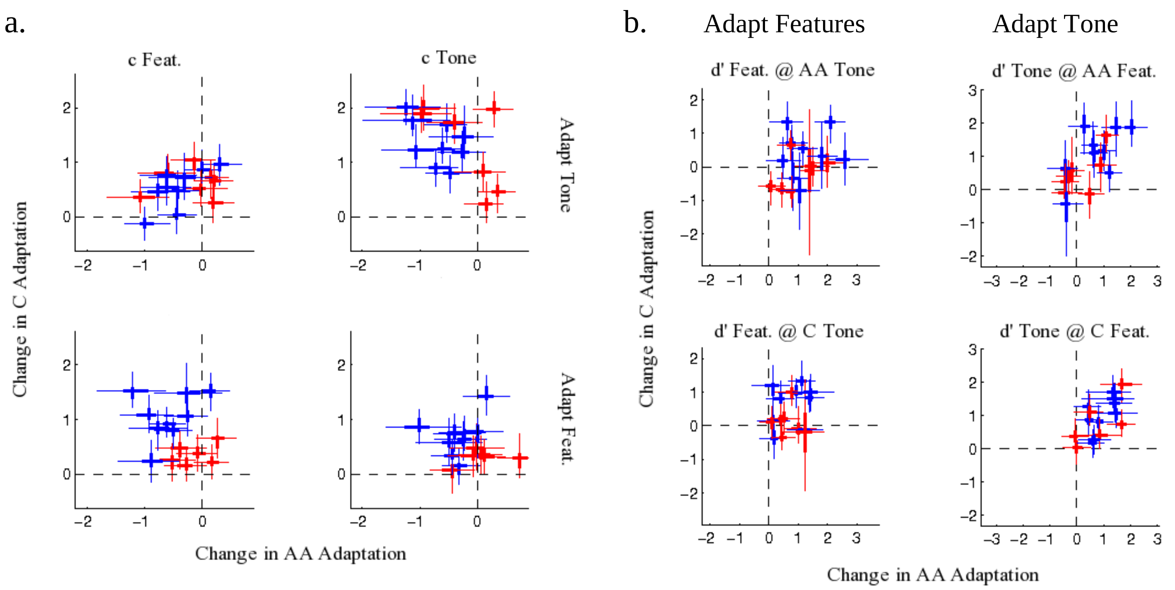

Figure 3: Criterion Shifts (a) and *d'* Shifts (b) Resulting from Adaptation

Figure 3b שגיאה! מקור ההפניה לא נמצא.

Discussion and Conclusion

Adaptation aftereffects may be based on changes in perceptual processes, decisional mechanisms, or both. Face race categories provide a complex, multidimensional space for modeling these aftereffects. The GRT framework provides a powerful mathematical tool for analyzing the relationships between multiple cues to racial categories and the locus of race adaptation effects.

In a three phase experiment, Caucasian and African American participants showed typical unidimensional adaptation effects on facial feature and skin tone dimensions, whereby adaptation biased responses about subsequently presented faces such that adaptation to Caucasian (African-American) resulted in more faces being classified as African American (Caucasian). It is not clear from unidimensional data if this change in responses is due to a change in decisional criteria or a change in perceived race similarity. In a novel two-dimensional adaptation experiment, and using a newly developed hierarchical Gaussian GRT model, we found (a) perceptual interactions between facial features and skin tone in the form of mean shift integrality and (b) consistent decisional *and* perceptual race adaptation effects. Adaptation resulted in shifts of criteria toward the adapting images and decrements in the perceived distance (increased similarity) between faces. We also found that both the decisional and perceptual adaptation effects tended to be larger and more consistent for the Caucasian participants than for the African-American participants. Our results suggest, therefore, that for complex face spaces, face adaptation aftereffects are the result of changes in both perceptual and decisional mechanisms.

References

Anderson, N. D. & Wilson, H. R. (2005). The nature of synthetic face adaptation. Vision Research, 45, 1815-1828.

Burkhardt, A., Blaha, L. M., Jurs, B. S., Rhodes, G., Jeffrey, L., Wyatte, D., Delong, J., & Busey, T. A. (2010). Adaptation modulates the electrophysiological substrates of perceived facial distortion: support for opponent coding. Neuropsychologia, 48(13), 3743-3756.

Leopold, D. A., O’Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. Nature Neuroscience, 4(1), 89-94.

Rhodes, G., Jeffery, L., Watson, T. L., Clifford, C. W., & Nakayama, K. (2003). Fitting the mind to the world: face adaptation and attractiveness aftereffects. Psychological Science, 14(6), 558-566.

Silbert, N. H. (2011). Two hierarchical Bayesian General Recognition Theory models. Talk presented at the 44th annual meeting of the Society for Mathematical Psychology.

Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural face categories. Nature, 428, 557-561.