The Experience of Thinking

How the fluency of mental processes influences cognition and behavior

Edited by
Christian Unkelbach and
Rainer Greifeneder

Halberstadt, J. & Winkielman, P. (2013). When good blends go bad: How fluency can explain when we like and dislike ambiguity. In C. Unkelbach & R. Greisfelder. The experience of thinking: How feelings from mental processes influence cognition and behavior (pp. 133 – 151). Psychology Press, NY.

Psychology Press
Taylor & Francis Group
LONDON AND NEW YORK

9 When good blends go bad

How fluency can explain when we like and dislike ambiguity

Jamin Halberstadt and
Piotr Winkielman

Abstract

We argue that some of the appeal of faces (and their owners) comes from the ease with which they are processed: their cognitive fluency. Individuals that are good examples of their group are processed more fluently, which produces positive affect that can generalize to evaluations. Conversely, when a person's individual or group membership is ambiguous, the difficulty in resolving that ambiguity produces negative affect that can compete with or override an otherwise beautiful appearance. We apply this analysis to the apparent paradox of morphed faces, which are robustly attractive despite their ambiguous identity, by introducing the notion of context-dependent fluency. We show that the same morphed face can be a good example of one group but a poor example of another, and by changing the salience of potential social categories, we can change the ease with which the target is classified, and in turn how beautiful it appears. We also extend our analysis to blends of racial groups and illustrate implications for stereotype theory, including the novel prediction that good examples may be preferred to poor examples of stereotyped groups, even when those groups are negative.

A man looks at Mona Lisa's face. At first, he really likes it. Why? One source of his reaction is "what" he sees. He appreciates the gentle smile, the smooth complexion, the inquisitive eyes. But another source of affect is "how" he sees, how easy (fluent) are the mental operations in processing Mona Lisa's face. She triggers a warm feeling because she is so "easy on the eyes," so smooth and familiar, as if da Vinci captured the very essence of the human condition. But now something happens! The man realizes there is something "off" in the face he sees. Is this really Mona Lisa? There seems to be a bit of da Vinci's own face in the picture. Is this really a woman? The face seems partially male. And what about that mysterious smile? Or is it a smirk? The man is now confused. He does not know what to make of the picture, even though he's trying hard to decipher it perceptually and conceptually. Perhaps he does not like Mona Lisa after all!

This chapter deals with such complexities of facial attractiveness. We argue that some of the appeal of faces comes from the ease with which they are classified. Stimuli that are good examples of their type, like Mona Lisa is of human faces, are processed more easily, and the affect associated with this fluent processing generalizes to evaluations of the stimuli themselves. Conversely, when the identity or group membership of a target is ambiguous, the difficulty in resolving that ambiguity produces negative affect that can compete with or override an otherwise beautiful appearance.

The tension between beautiful content and effortful process is starkest in the case of facial blends, which are robustly attractive despite their "objective" ambiguity. We resolve this apparent paradox by introducing the notion of context-dependent fluency, in which ambiguous faces are disliked, but whether a face is ambiguous depends on the social category of which it is seen to be a member. We show that the same blended face can be a good example of one group but a poor example of another, and by changing the salience of potential social categories, we can change the ease with which the target is classified, and in turn how beautiful it appears. In what follows we apply this approach first to understanding the attractiveness of blends of individuals, then extend our analysis to blends of racial groups. Finally, we illustrate implications for stereotype theory, and show how a dynamic, context-specific theory of fluency makes surprising predictions when stereotype content and processing ease conflict.

Beauty-in-averageness

Galton (1878) first noted that photographically blending faces (in some cases, faces of criminals) increases their attractiveness. This phenomenon, replicated many times in the modern psychological literature, reflects the tendency for category exemplars to be more attractive as a function of their proximity to the category's central tendency (Langlois & Roggman, 1990). Halberstadt (2006) documented the generality of the "beauty-in-averageness" effect, reporting a correlation between judged typicality and (independently) judged attractiveness in a wide variety of natural and artificial categories, from random dot patterns ($r = .37$) to birds ($r = .50$) to wristwatches ($r = .65$). In nearly all cases, category exemplars were liked better as a function of their similarity to the category prototype.

Empirically, morphing appears to produce faces that are more "face-like," more similar to the perceiver's image of the typical face, than the images used to create them. Busey (1998), for example, found that morphs are located closer to the center of multi-dimensional face space than their parents (as well as the emergence of other features such as youth and adiposity). Convergently, Rhodes and Tremewan (1996) found that faces that were shifted toward the center of face space (by distorting them in the direction of a blend of all faces in the study) were more attractive than faces shifted away from the center (Halberstadt & Rhodes, 2003, reported the analogous effect using line drawings of birds). Rhodes et al. (2003) found that the center of face space itself can be shifted (via exposure to

systematically distorted exemplars), with predictable effects on attractiveness. In all, these data suggest that blended faces are more similar to the perceivers' face prototype, which is itself a composite of the faces to which they have, over a lifetime, been exposed (cf. Potter & Corneille, 2008).

Just why prototypicality is attractive has been debated. In the case of faces, evolutionary psychologists have naturally invoked sexual selection, for example arguing that typicality is a proxy for various fitness-enhancing traits, such as heterozygosity or overall health (Symons, 1979; Thornhill & Gangestad, 1993). However, the demonstration of prototypicality effects in reproductively-irrelevant stimulus categories has led researchers to explore more general, proximate cognitive mechanisms.

Fluency as an explanation of the beauty in averageness effect

We have shown that classification fluency—the ease with which a stimulus can be classified as a member of its group—is an important part of the story (Winkielman, Halberstadt, Fazendeiro & Catty, 2006). In recent years, much research has been devoted to the idea that evaluative reactions to stimuli are partially caused by the ease with which they can be processed. There are several possible reasons for this association. One is that the experience of fluent processing might be pleasant in itself, because it signals low cognitive conflict and tells the organism about the success (or good trajectory) of recognition or categorization operations (Winkielman, Schwarz, Fazendeiro, & Reber, 2003). Fluency might also be positive because it is probabilistically associated with positive stimulus features (e.g., familiarity or safety), thereby functioning as a cue to stimulus evaluation (Gigerenzer, 2007).

Many studies have manipulated processing ease, using a variety of techniques. One is a simple priming manipulation, where an object is preceded by a perceptually congruent or incongruent stimulus. For example, in one study (Reber, Winkielman & Schwarz, 1998, Study 1), participants saw pictures of everyday objects, such as a desk or a plane, whose fluency was enhanced or reduced by prior subliminal presentation of matching or mismatching visual contours. The data showed that pictures preceded by matched contours were recognized faster, indicating greater fluency, and that these fluent pictures were liked more than disfluent pictures, preceded by mismatched contours. Other manipulations, such as varying viewing conditions (via contrast or clarity), stimulus duration, symmetry, comprehensibility of stimulus names or descriptions, predictability, and even rhymes, have yielded predictable changes in stimulus appeal (Reber, Schwarz & Winkielman, 2004; Winkielman et al., 2003). Furthermore, Winkielman and Cacioppo (2001) have shown that high fluency is associated with stronger activity over the zygomaticus muscle region (the "smiling" area, indicative of positive affect), but not with stronger activity of the corrugator region (the "frowning" muscles, indicative of negative affect).

Most important for our thesis, the *prototypicality* of an exemplar is a robust predictor of its fluency. In their classic study of category learning, Posner and Keele (1968) found that dot patterns were classified more slowly to the extent they were

distorted from a prototype; the prototype itself was classified fastest, despite never having been presented in the study. Indeed, classification time has been often used as a measure of the distance of a stimulus from its central tendency (e.g., Dovidio, Evans, & Tyler, 1986; Rosch & Mervis, 1975; Smith, Shoben, & Rips, 1974).

The independent relations between typicality, fluency, and positive affect, raise the possibility that prototypes are liked because they are fluent. Winkielman et al. (2006) examined this idea in a series of three experiments. Participants first learned a category of random dot patterns (Experiment 1) or of common geometric patterns (Experiment 2) and then were presented with novel patterns varying across different levels of similarity to a prototype pattern. Participants classified these patterns into their respective categories as quickly as possible (a measure of fluency), and also rated the attractiveness of each. A close relationship between fluency, attractiveness, and the level of prototypicality was observed. Both fluency and attractiveness increased with prototypicality. Importantly, when fluency was statistically controlled, the relation between prototypicality and attractiveness dropped by half (though it remained significant). Finally, Experiment 3 showed that viewing prototypical, rather than non-prototypical patterns elicited significantly greater electromyographic activity, suggesting that viewing prototypes involves genuine affective reactions.

Importantly, these effects are not limited to low-level perceptual judgments but can occur even if the task is a high-level classification. For example, in a rich stimulus domain—abstract art—Halberstadt and Hooton (2008) found that typicality (i.e., how well a painting exemplified an artist's style) predicted its classification fluency (i.e., the time it took to associate a painting with its painter). Further, Halberstadt and Hooton found that the fluency of abstract paintings predicted their appeal: people preferred the paintings that were easiest to classify (which were, in turn, the best examples of the artist's work).

Can fluency explain the beauty-in-averageness effect in faces?

As described above, Winkielman et al. (2006) showed that dot patterns, as they approached the prototypes from which they were generated, were easier to classify, and consequently more beautiful. Similarly, perhaps faces, as they approach *their* prototypes (via morphing) get easier to classify and more beautiful. Can fluency explain the beauty-in-averageness effect in faces?

At first glance, the answer seems to be a surprising but resounding "no." Face blends are beautiful but indistinct, similar on all perceptual dimensions to the faces that constitute them, and therefore difficult to identify and discriminate from other faces. Such difficulty, according to fluency theory, should translate into *negative affect*, which should generalize to the faces themselves to make them *less* attractive. Nevertheless, we argue that an appreciation of the *category relativity* of both categorization and fluency can not only explain the appeal of such facial blends, but also provide a unified account of the relations among beauty, prototypicality, and classification.

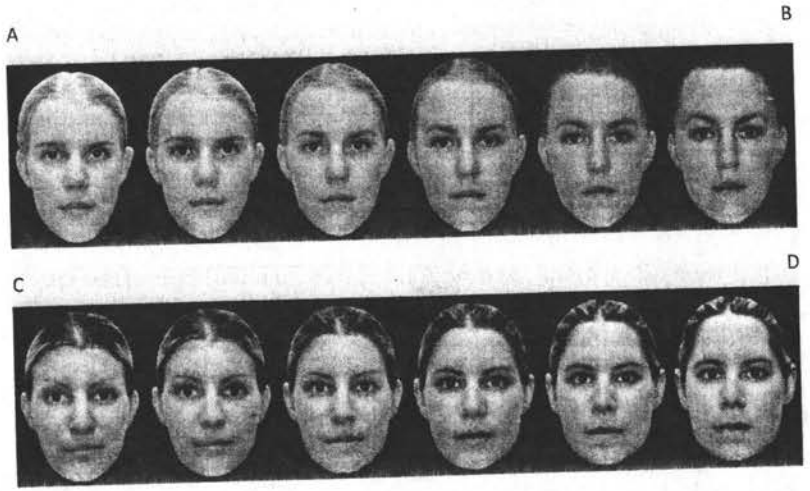

Figure 9.1 Examples of two morph continua. The four faces in the middle of each continuum are blends of the "parents" at the endpoints. Faces typically get increasingly attractive as they are increasingly blended.

This account begins with the critical assumption, easy to overlook, that both prototypicality and fluency are products not of the stimulus, but of the cognitive system. A stimulus is only typical with respect to a particular category, and the same stimulus can be a good example of one category and a poor example of another, or both a good and bad example of the same category in different contexts (Barsalou, 1983). Changes in context activate different features of a stimulus (e.g., a newspaper in the context of a library versus a house fire), as evidenced by property verification times (Barsalou, 1982), and individuals themselves can shift their cognitive perspective to produce different—but reliable and accurate—category structures from different points of view (Barsalou & Sewell, 1984). In the most dramatic example of category relativity, individuals can derive reliable graded structures of entirely new ("ad hoc") categories in relation to specific and momentary goals (Barsalou, 1982). In such cases, we propose that the prototypicality of a stimulus, the ease with which it can be categorized, and therefore the fluency-derived positive affect associated with it, can change quickly in response to environmental input or idiosyncratic goals.

As an initial examination of this hypothesis with respect to faces, we examined the fluency of the same morphed faces in two different contexts, one in which the faces should be more fluent, and one in which they should be *less* fluent, as they are increasingly blended. The critical stimuli were 26 Caucasian faces—two "parents," which we called A and B, and 24 equally-spaced morphs between them (i.e., 96 percent A/4 percent B, 92 percent A/8 percent B, etc.; see Figure 9.1 for examples). In one condition, participants were shown A and B,

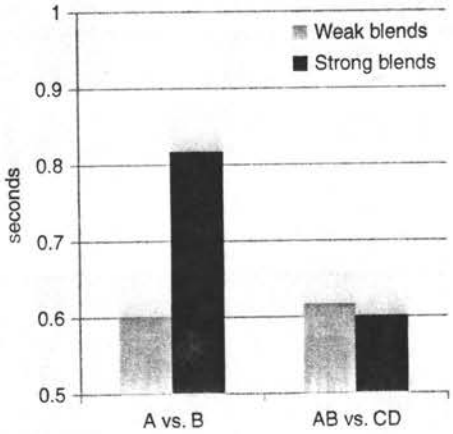

Figure 9.2 Classification time for weakly blended (at least 80% similar to one parent) and strongly blended morphs of A and B faces as a function of classification task (see text for details). Examples of AB and CD blends appear in Figure 9.1.

whom they were told represented members of two different groups. They were then shown the 26 A–B morphs and asked to judge whether each was an A group member or a B group member as quickly as possible. In a second condition, participants were again shown A and B but told they represented two members of the *same* group. They were then shown the 26 A–B morphs randomly intermixed with a different set of morphs—26 blends of two different faces, C and D—and asked to judge whether each was an A–B group member or a C–D group member as quickly as possible.

This paradigm provides an elegant test of the context-dependence of face fluency, because A–B morphs' classification fluency should be precisely reversed in the two experimental conditions. That is, when A and B are presented as members of different groups, blends of A and B should be relatively difficult to classify, because they are similar to both. However, when A and B are presented as members of the same group, blends should be relatively *easy* to classify because their similarity to both makes them highly prototypical A–B group members. In other words, an A–B blend is a poor exemplar of either A or B as individuals but a good exemplar of a group consisting of A and B members. Because good exemplars are fluent, blends should be classified slowly in the first case but quickly in the second. Indeed, as seen in Figure 9.2, more strongly blended faces (less than 80 percent similar to either parent) were easier to classify as A–B group members, but particularly difficult to classify as As and Bs, relative to weakly blended faces.

These data provide the first evidence that the fluency of facial composites is context-dependent (in this case on the perceiver's goal-derived category structure). If faces derive some of their attractiveness from their fluency, then this context dependence should have implications for the appeal of the A–B blends. In particular, their normally robust attractiveness should be reduced or reversed when they are classified in terms of their constituent faces, rather than in terms of a higher-order group containing both of them.

As a test of the attractiveness hypothesis, we solicited attractiveness judgments of the morphed faces described above, this time using the structure of the judgment task itself to situate the stimuli in different categorical contexts. Participants were again shown two faces and were told they represented members of two different groups, A and B, which was associated with different "family names." They were then presented with 26 A–B morphs of the faces (the A and B "parent" faces and 24 equally spaced blends between them) and then asked, in a control condition, to rate their attractiveness on a single sliding scale anchored at 1 ("very unattractive") and 9 ("very attractive"). In an experimental condition, participants were presented with *two* such scales, each labeled with a family name, and asked to make their rating on whichever scale was appropriate for a given face. Thus, in the experimental condition, but not the control condition, judgments require implicit classification of a composite face in terms of its group membership immediately prior to rating it, in order to use the correct scale. We expected that implicitly forcing participants to classify the ambiguous faces would create a feeling of disfluency, rendering blended faces relatively unattractive compared to the control condition.

Indeed, in the control condition, when faces were simply judged on a single attractiveness scale, morphs were rated higher than their parents, creating an "inverted U" relationship between morph level and attractiveness, a classic beauty-in-averageness effect. In the experimental condition, however, in which each judgment was preceded by classification, this quadratic relationship flattened out (see Figure 9.3, left panel), presumably because the difficulty associated with classifying the faces depressed their otherwise strong appeal. This account is supported by speeded classification data, generated by same participants, that confirmed that increasingly blended faces were increasingly difficult to categorize, and that categorization speed is associated with the extent to which the quadratic attractiveness function is flattened in the experimental condition.

It could be argued, however, that these data represent not a genuine fluency-driven change in the appeal of blended faces, but rather a change in standards against which the faces were judged. To verify that changes in the attractiveness of blends were driven by the difficulty of classifying them, we conducted a replication which required all participants to judge face morphs on family-specific scales, but which varied the cognitive difficulty involved in doing so by informing half of the participants of the family to which a stimulus belonged (and therefore which of the two scales should be used to rate it). Thus, the rating conditions were equated in terms of their use categories and multiple scales

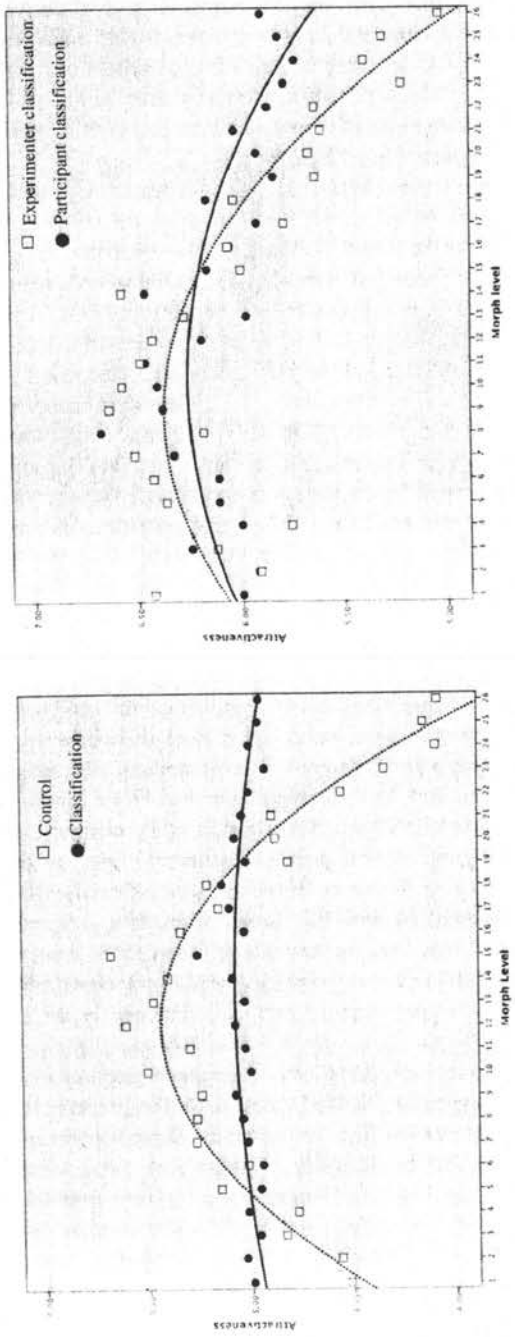

Figure 9.3 Attractiveness as a function of morph level (in 4% blend increments) and classification task (see text for details). In both studies, classification difficulty was associated with a decrease in the beauty-in-averageness effect.

(and in turn the standards for those scales); they varied only in whether the participants were required to classify the faces for themselves. As shown in the right side of Figure 9.3, the attractiveness of morphs relative to their parents (the beauty-in-averageness effect) was still reduced when judgments required implicit identification of the group to which the faces belonged, than when the family membership was provided, suggesting that the cognitive effort involved in classifying the faces played a role in their attractiveness.

In sum, these studies show that the relative attractiveness of blended faces depends critically on the effort required to classify them, which in turn depends on what they are being classified *as*. A face blend represents both a distortion of the faces that compose it, and at the same time a good example of the superordinate category (e.g., "faces") that includes it. Thus, by implicitly inducing participants to use the more specific categories, we dramatically altered the prototypicality of the blends, making them more difficult to process and consequently lowering (though not entirely reversing) their attractiveness. Most tellingly, blends were only less attractive when participants actively classified them: it was not enough for a blend merely to be an atypical category member; participants needed to experience the disfluency for themselves for typicality to have an effect on attractiveness.

Fluency and cross-race faces

Our model of context-dependent fluency can also be applied to "natural" morphs—bi- or multi-racial faces which, like laboratory blends, are notably attractive (Rhodes et al., 2005), as well as self-evidently difficult to classify. Indeed the inscrutability and mystique of cross-race faces has sometimes been cited as the source of their beauty (e.g., Beech, 2001), which if true would seem to falsify a fluency account of their attractiveness.

We believe the seemingly paradoxical, disfluent beauty of cross-race faces can be resolved by assuming that they are both typical and atypical of different categories under different conditions. That is, if cross-race faces are "faces" when judged for attractiveness, but race-specific exemplars when classified by race, fluency could explain their attractiveness in both cases. If so, then we should be able to reduce, or even reverse, the attractiveness of cross-race faces by encouraging race-specific processing prior to judgment. Barack Obama, for example, should be a fluent, and attractive, man, but a relatively disfluent and unattractive "Black man" (or "Caucasian man"), with category-relative fluency explaining his attractiveness in both cases.

To examine these ideas, we asked participants to rate the attractiveness of a series within-race and cross-race blends of Caucasian and Asian faces on a 1 to 7 scale (see Figure 9.4 for examples). We manipulated social category salience by asking some of the participants to classify the faces in terms of their race, as quickly and accurately as possible, immediately prior to rating them. Control participant made an unrelated prior judgment. As seen in Figure 9.5, cross-race blends were more attractive than within-race blends, as typically observed, but

Figure 9.4 Examples of within-Asian, between Asian-Caucasian, and within-Caucasian morphs. All images are blends of two faces.

their advantage was significantly attenuated when participants first identified their race. An analysis of race classification times revealed that cross-races faces were also classified particularly slowly relative to within-race faces, consistent with the hypothesis that at least part of the difference in attractiveness is due to their ambiguity, and consequent disfluency, with respect to their group identity.

In a second study we measured *nonverbal* affective responses to cross-race faces, using facial electromyography (EMG) to record electrical activity from the zygomaticus (the "smiling") muscle regions. Facial EMG is a well validated nonverbal indicator of incipient emotional response, particularly appropriate for

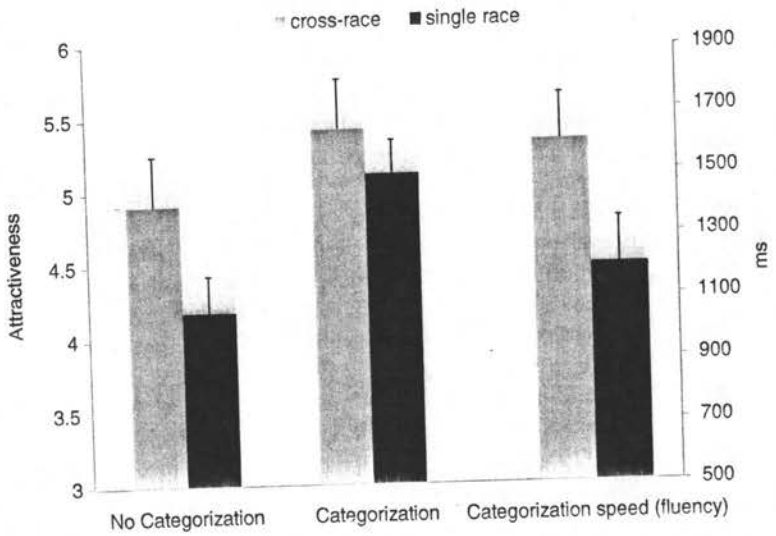

Figure 9.5 Attractiveness (left axis) and fluency (right axis) of within-race and between-race morphs as a function of experimental condition.

measuring the appeal of cross-race faces, which is fraught with self-presentational bias. Twelve participants were exposed to both within and between-race blends, which half of whom had previously classified by race. Analysis of baseline-adjusted zygomaticus activation revealed that participants smiled more to cross-race morphs in the control condition, but *less* to cross-race morphs (and more to within-race morphs) in the categorization condition.

Together, these studies show that racial classification reduces the appeal of cross-race faces. This is presumably due to the role cognitive fluency plays in the attractiveness of facial averages. With no need for racial classification, multi-race individuals are perceived fluently, which partly contributes to their appeal. When racial classification is required, however, the difficulty of performing such a task on multiracial individuals leads to a decline in their appeal, as reflected in participants' self-reports and their physiological responses.

Fluency and stereotyping

When the flexibility of perceivers' category use is taken into account, fluency provides a plausible account of why cross-race faces are generally attractive, and also predicts the conditions when they should be *un*attractive: when race categorization is required and/or salient. Interestingly, when considered more broadly, both positive *and* negative attitudes toward category blends (relative to unambiguous exemplars) appear at odds with intuitive stereotyping theory (formal theory being mostly silent on the issue), and illustrate the need for the incorporation of processing dynamics into any complete account of affect and social categorization.

Stereotypes are summary judgments of the perceived central tendency of a social group: category prototypes. As such they should form the basis of a graded group structure, such that individuals are better or worse examples of the group depending on their similarity to the stereotype (Cantor & Mischel, 1979). For example, Arnold Schwarzenegger is a prototypical "action hero" but a very atypical "politician;" Ronald Reagan is the opposite.

The question of how the affective component of stereotypes is generalized to group members as a function of their typicality has, oddly, received little attention (see Fiske, 1982, for an exception). Although many theories address the likelihood that a stereotype will be applied to a given individual (e.g., Devine, 1989; Fiske & Neuberg, 1990; Kunda & Spencer, 2003), these theories apparently assume that *if* the stereotype is applied, it is applied in full. Thus, if a perceiver hates politicians, then (all things being equal) she hates Ronald Reagan whenever he is seen as a politician (although a variety of individual and situational factors will influence whether he is indeed classified as such). Intuitively, however, the affect associated with a member of a stereotyped group should apply to group members in proportion to the member's fit to the stereotype: when both are being judged as politicians, Reagan should be judged more negatively than Schwarzenegger by people who hate politicians, and more positively by people who love them. When both are being judged as action heroes,

Schwarzenegger should be judged more positively than Reagan by people who love action heroes, and more negatively by those who hate them.

This intuitive analysis, however, assumes a linear relationship between category fit and affect generalization, and ignores the fact that the process of fitting an individual to a category itself generates affect that may influence the perceiver's overall evaluation. In the case of facial prototypes, we have seen how the effort involved in classification by race can influence their attractiveness. In the case of social prototypes, classification fluency could analogously produce a counterintuitive, nonlinear relation between typicality and liking, such that atypical members of negative groups, for example, are judged less favorably than prototypical members. In other words, for the perceiver who hates politicians, the very typical politician might be liked *more* than the atypical one, because the latter derives negative affect both from the overall category evaluation and from the disfluency created by classifying an atypical politician. If politicians are negative, Reagan, the consummate politician, may ironically be favored over Schwarzenegger, the less typical one.

To test this curious prediction we examined a fundamental and automatic dimension of person classification—gender—creating ten sets of faces that varied continuously in their "femaleness" by blending pairs of male and female faces to eight different degrees (examples appear in Figure 9.6). Participants judged the attractiveness of these faces in one of two conditions, modeled on the face perception tasks described above. In the critical, experimental condition, prior to judging each face, participants were asked to classify it as male or female; in a control condition, participants made an unrelated judgment (pressing a key when a target face appeared on the screen) prior to each rating.

An analysis of attractiveness as a function of morph level and prior classification, illustrated in Figure 9.7, revealed that faces were overall more attractive to the extent they were feminized (consistent with the literature on facial attractiveness; Perrett et al., 1998). More important, however, an interaction between morph level and classification condition showed that, as predicted, when (and only when) faces were classified by their gender, more ambiguous faces were less attractive than relatively unambiguous ones. The pattern is particularly interesting in light of the general preference for feminized faces: even though men were less attractive than women, more "manly" men were *more* attractive than less manly men!

This nonlinear pattern cannot be accommodated by a simple model of stereotype application in which a blend is judged positively as a function of its fit with a positive group. Instead, it suggests that the ambiguity of the stimulus has an additive, depressing effect on their appeal, and the fact that the effect is only evident when the faces are categorized suggests that it is the disfluency created by the process of classification that accounts for the negativity.

An aesthetic preference for women over men, however, is not the same as a group attitude. To examine the stereotypes more directly, we conducted a conceptual replication using a weakly negative target group: homosexual men. Not only did this domain allow us to examine attitudes toward the targets (rather than

Figure 9.6 Examples of gender blends used.

their attractiveness), but the fact that the stereotype of homosexual men includes feminized features allowed us to vary social category fit via manipulation of facial appearance. Analogous to the study on gender classification described above, we predicted that although participants would express more positive attitudes toward heterosexual men than toward homosexual men, they would, remarkably, like "more gay" men (i.e., men with feminized features) more than "less gay" men.

Stimuli were images of 30 Caucasian men, each morphed 50 percent with a "supermale" or "superfemale" image (averages of male and female faces whose gender-specific features have been enhanced, created by Rhodes, Hickford, & Jeffery, 2000) to produce a masculinized and feminized version of each individual. Participants estimated how much they would like the feminized or masculinized version (within subjects) of each of the 30 men, who were all described as either homosexual or heterosexual (between subjects). Participants also completed

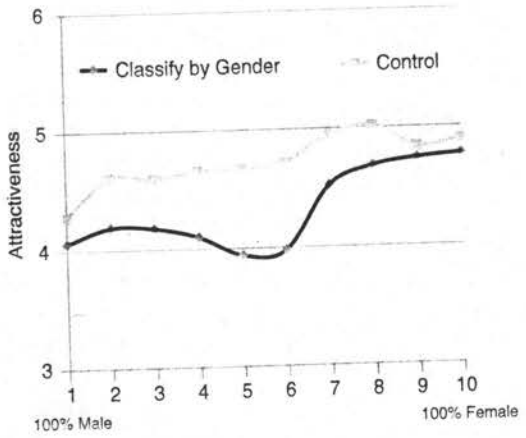

Figure 9.7 Attractiveness of morphed faces as a function of "femaleness" and classification condition.

Lamar and Kite's (1998) Components of Attitudes Toward Homosexuality scale, and were median-split in terms of their response to the item "Most gay men have identifiable feminine characteristics."

Analyses of liking judgments revealed that, although all participants preferred feminized over masculinized male images, the magnitude of the preference depended on whether the men were described as gay, and on participants' stereotypes about gay men's features. When the men were described as gay, participants who endorsed the stereotype that gay men have feminine features showed a greater preference for feminized men than participants who did not endorse the stereotype. The opposite was true when the men were described as heterosexual: participants who endorsed the stereotype showed a smaller preference for feminized men than participants who did not endorse the stereotype.

Importantly, participants did not differ in their *evaluation* of gay men as a group. For example, participants who thought gay men had feminine features did not think homosexuality was more or less "perverted" than participants who did not endorse this stereotype. Therefore, the data suggest that, as in the case of facial attractiveness, part of the evaluation of specific homosexual men was due to their similarity to the perceiver's stereotype of a homosexual man, regardless of whether the category itself was positive or negative. We propose that the affect associated with category fit is generated from fluent processing: Even when gay men are disliked as a group, "very gay" men (those who are good examples of the perceiver's category) will be easier to classify and paradoxically be evaluated more positively than "less gay" men.

It should be noted that confirmation of the mediating role of fluency still awaits confirmation from our ongoing research program, and that such confirmation may prove more complicated than it appears, depending on one's conception of fluency and its relation to positive affect. There are two general models of this relation: One assumes that fluency is a separate, temporally prior, and independent cause of positive affect. According to this "two step" model of fluency, some manipulation (e.g., blending) leads to a change in fluency, which participants subsequently explain by relating it to evaluative or other features of the stimulus (Jacoby, Kelley, & Dywan, 1989; Mandler, Nakamura, & Van Zandt, 1987). In this case, because fluency is clearly distinguishable from evaluative reactions, it can be statistically treated as such in causal and statistical models. However, a second general model (sometimes called "hedonic") assumes that fluency comes with an instant positive valence, either due to some intrinsic process (Winkielman et al., 2003) or to an automatic cue-learning history (Unkelbach & Greifeneder, Chapter 2, this volume). In this case—if positive affect is an integral and simultaneous component of fluency—statistical mediation becomes tricky, as by controlling for the mediator, one effectively controls for the DV itself. Future research may want to establish when fluency in attractiveness judgments functions more in line with the two-step or the intrinsic hedonic model, and what statistical models are most appropriate in which case.

Summary, implications, and conclusions

A recurring theme of this volume is that evaluations of the world are influenced not only by its content, but also by the ease with which that content is accessed and processed. In this chapter, we have shown how these two sources of affect combine to produce seemingly paradoxical effects on facial attractiveness and person perception. Using the principle of category-relative fluency, we showed how processing dynamics might play a role in the beauty-in-averageness effect, despite the apparent *ambiguity* of facial blends (i.e., the fact that they objectively resemble more than one group prototype). In fact, the ease with which a face can be classified and processed appears to depend on the category in which it is situated. In particular, blends, including bi-racial faces, can be subjectively seen as good examples of superordinate categories, such as "faces," making categorization fluent, which benefits attractiveness. However, they also can be subjectively seen as bad examples of the individuals or groups from which they were combined. So, making their componential structure salient decreases the blends' typicality, lowers the ease with which it can be categorized, and consequently hurts their aesthetic appeal. More generally, this research points out that 'objective' ambiguity (e.g., a stimulus created by averaging, morphing, blending, mixing, etc) matters less than how a stimulus is subjectively categorized. Superordinate categorization of a face (e.g., as a face, as a human, etc.) decreases its ambiguity by turning it into a good example of a higher-level category, while subordinate categorization (e.g., as a particular gender or member of a racial group) increases ambiguity by turning it into an, atypical, "bad" exemplar, with all the processing and preference costs this engenders.

Of course, some of these points have been made more or less explicitly before in the categorization literature. For example, classic approaches to categorization have long pointed out that the same two stimuli can be seen as belonging to just one or two separate categories. Further, the same stimulus can be perceived as typical or as atypical. What exactly happens depends on what dimensions are perceived as relevant, and whether the subjective representation of the relevant dimension is relatively fine-grained or relatively "coarse." These, in turn, depend on factors such as the relative range and frequency of stimuli in a category (Parducci, 1965), attention to the relevant dimension (Goldstone & Steyvers, 2001), and the current context and goals of the categorizer (Barsalou, 1982).

Analogously, work on intergroup relations has pointed out that individuals perceive others, and themselves, differently depending on the salience of their multiple social identities, and the level of social categorization (for a review, see Crisp & Hewstone, 2007). Thus, highlighting one group identity can make an individual seem typical, but highlighting another group identify can make the individual seem atypical. Similarly, highlighting superordinate categorization can make an individual seem like a "good" exemplar, but result in a loss of individual differentiation (Simon, Pantaleo, & Mummendey, 1995). Such superordinate categorization can be beneficial for the perception of common goals, common

fate and reduction of prejudice, but come at the well-known cost of 'we-ness' (Gaertner & Dovidio, 2000). Critically, as our chapter highlights, all these categorization shifts have consequences not only for the content, but also the subjective ease of processing, and ease-related affect. For example, our framework predicts that stereotype-based judgments will be a function not only of what stereotype is applied, but also how *easily* it can be applied, leading to the counterintuitive hypothesis that very typical members of negative groups may be preferred to atypical ones.

Finally, it is interesting to note a certain paradox about recent work on fluency: Fluency became a "hot" topic of research because it highlighted that social judgments could be based on low-level, non-analytic, "content-free" aspects of stimulus processing—a point that resonated with Zajonc's (1980) famous claim that "preferences need no inferences." But more recently, the focus of fluency research has turned toward its interaction with higher-order processes—categorization, expectation, learning, metacognitive theories, etc (see other chapters in this volume). In our opinion this is a welcome development. After all, social perception represents a dynamic interplay of cognitive content and process, producing rich, nonlinear relations between objective stimulus features, subjective representations, experiences, and judgments that we have only begun to explore.

Acknowledgment

We gratefully acknowledge the generous and insightful comments from the book editors and from Klaus Fiedler.

References

Barsalou, L.W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition, 10*, 82–93.

—— (1983). Ad hoc categories. *Memory & Cognition, 11*, 211–227.

—— (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A.C. Collins, S.E. Gathercole, & M.A. Conway (Eds.), *Theories of memory* (pp. 29–101). London: Lawrence Erlbaum Associates.

Barsalou, L.W., & Sewell, D.R. 1984. Constructing representations of categories from different points of view. Emory Cognition Project Technical Report .#2, Emory University.

Beech, H. (2001). Eurasian Invasion. Time Magazine, 23 April 2001.

Busey, T. (1998). Physical and psychological representations of faces: Evidence from morphing. *Psychological Science, 9*, 476–482.

Cantor, N., & Mischel, M. (1979). Prototypes in person perception. *Advances in Experimental Social Psychology, 12*, 3–52.

Crisp, R. J., & Hewstone, M. (2007). Multiple social categorization. In M. P. Zanna (ed.) *Advances in experimental social psychology* (vol. 39, pp. 163–254). Orlando, FL: Academic Press

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*, 5–18.

Dovidio, J. F., Evans, N., & Tyler, R. B. (1986). Racial stereotypes: The contents of their cognitive representations. *Journal of Experimental Social Psychology, 22*, 22–37.

Fiske, S. T. (1982). Schema-triggered affect: Applications to social perception. In M. S. Clark & S. T. Fiske (Eds.), *Affect and cognition: The 17th Annual Carnegie Symposium on Cognition* (pp. 55–78). Hillsdale, NJ: Erlbaum.

Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. *Advances in Experimental Social Psychology, 23*, 1–74.

Galton, F. (1878). Composite portraits. *Journal of the Anthropological Institute of Great Britain & Ireland, 8*, 132–144.

Gaertner, S. L., & Dovidio, J. F. (2000). *Reducing intergroup bias: The Common Ingroup Identity Model*. Philadelphia, PA: Psychology Press.

Goldstone, R. L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General, 130*, 116–139.

Gigerenzer, G. (2007). *Gut feelings: The intelligence of the unconscious*. New York: Viking Press.

Halberstadt, J. B. (2006). The generality and ultimate origins of the attractiveness of prototypes. *Personality & Social Psychology Review, 10*, 166–183.

Halberstadt, J., & Hooton, K. (2008). The affect disruption hypothesis: The effect of analytic thought on the fluency and appeal of art. *Cognition and Emotion, 22*, 964–976.

Halberstadt J., & Rhodes G. (2000). The attractiveness of nonface averages: Implications for an evolutionary explanation of the attractiveness of average faces. *Psychological Science, 4*, 285–289.

—— (2003). It's not just average faces that are attractive: Computer-manipulated averageness makes birds, fish, and automobiles attractive. *Psychonomic Bulletin and Review, 10*, 149–156.

Jacoby, L. L., Kelley C. M., & Dywan J. (1989). Memory attributions. In H. L. Roediger & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving*. (pp. 391–422). Hillsdale, NJ: Erlbaum.

Kunda, Z., & Spencer, S. (2003). When do stereotypes come to mind and when do they color judgment? A goal-based theoretical framework for stereotype activation and application. *Psychological Bulletin, 129*, 522–544.

LaMar, L., and Kite, M. E. (1998). Sex differences in attitudes toward gay men and lesbians: A multi-dimensional perspective. *The Journal of Sex Research, 35*, 189–196.

Langlois, J. H., & Roggman, L.A. (1990). Attractive faces are only average. *Psychological Science, 1*, 115–121.

Mandler, G., Nakamura, Y., & Van Zandt, B. J. (1987). Nonspecific effects of exposure on stimuli that cannot be recognized. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 646–648.

Parducci, A. (1965). Category judgments: A range-frequency model. *Psychological Review, 72*, 407–418.

Perrett, D. I., Lee, K. J., Penton-Voak, I., Rowland, D., Yoshikawa, S., Burt, D. M., Henzi, S. P., Castles, D. L., & Akamatsu, S. (1998). Effects of sexual dimorphism on facial attractiveness. *Nature, 394*, 884–887.

Potter, T., & Corneille, O. (2008). Locating attractiveness in the face space: Faces are more attractive when closer to their group prototype. *Psychonomic Bulletin & Review, 15*, 615–622.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77*, 353–363.

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review*, *8*, 364–382.

Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science*, *9*, 45–48.

Rhodes, G., Hickford, C., Jeffery, L. (2000). Sex-typicality and attractiveness: Are supermale and superfemale faces super-attractive? *British Journal of Psychology*, *91*, 125–140.

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

Rhodes, G., Lee, K., Palermo, R., Weiss, M., Yoshikawa, S., Clissa, P., Williams, T., Peters, M., Winkler, C., & Jeffery, L. (2005). Attractiveness of own-race, other-race, and mixed-race faces. *Perception*, *34*, 319–340.

Rhodes, G., & Tremewan, T. (1996). Averageness, exaggeration, and facial attractiveness. *Psychological Science*, *7*, 105–110.

Rosch, R., & Mervis, C. B. (1975) Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573–605.

Simon, B., Pantaleo, G., & Mummendey, A. (1995). Unique individual or interchangeable group member? The accentuation of intragroup differences versus similarities as an indicator of the individual self versus the collective self. *Journal of Personality and Social Psychology*, *69*, 106–119.

Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, *81*, 214–241.

Symons, D. (1979). *Evolution of human sexuality*. New York: Oxford University Press.

Thornhill, R., & Gangestad, S.W. (1993). Human facial beauty: Averageness, symmetry, and parasite resistance. *Human Nature*, *4*, 237–269.

Winkielman, P., & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face: Psychophysiological evidence that processing facilitation leads to positive affect. *Journal of Personality and Social Psychology*, *81*, 989–1000.

Winkielman, P., Halberstadt, J., Fazendeiro, T., & Catty, S. (2006). Prototypes are attractive because they are easy on the mind. *Psychological Science*, *17*, 799–806.

Winkielman, P., Schwarz, N., Fazendeiro, T., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion*. (pp. 189–217). Mahwah, NJ: Lawrence Erlbaum.

Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, *35*, 117–123.