Generalization of the Features of Emotional Faces

Diane Berg

Utrecht University

Abstract

Emotional faces contain important social information and draw attention automatically. Some emotional expressions draw attention more rapidly than others; however, various studies disagree as to which emotional expression consistently draws attention the fastest, with both positive (e.g. happy) and negative (e.g. angry) emotional expressions showing faster reaction times depending on the study and task. This brings into question whether differences in reaction time to emotional faces are due to valence alone or factors such as low level image features (e.g. contrast and orientation). Additionally, if these low level features, particularly spatial frequencies, are involved in the rapid processing of emotional faces, then non-face objects with similar spatial frequency content would have similar reaction time effects. In this study, we examined the role of spatial frequency content in access to awareness of images of emotional faces. We used car images to test for generalizability based on low level features. Using the spatial frequency content from angry, happy, and neutral faces, we used machine learning to classify car images, both frontal and side views, as happy or angry based on their spatial frequency content. Using breaking continuous flash suppression (b-CFS) and a forced choice task, we measured reaction time to access of awareness as well as participants' subjective rating of images of emotional faces and classified "emotional" car sides and fronts. No significant differences were found between either image type or emotion in b-CFS, and notably, faces did not reach access to awareness faster than car images. In the rating task, however, human faces were rated as expected (e.g. happy faces as happy) even though car images were rated neutrally.

Table of Contents

Introduction

Emotional Expressions

There is no small amount of research devoted to understanding how and why faces are so important. It is clear that there are many ways that the face could play a major part in communication. In general, people are thought to be drawn towards faces automatically (Johnson, 2005; Palermo & Rhodes, 2007). This preference of faces over non-face objects is so fundamental that it can found as early as infancy (Valenza, Simion, Cassia, & Umiltà, 1996), and faces are preferred even when people are highly familiar with competing non-face objects (Stein, Reeder, & Peelen, 2015). Faces are a rich source of social information, and orienting attention towards faces can help us identify each other, regulate interpersonal interactions, and communicate emotional signals (Sorce, Emde, Campos, & Klinnert, 1985; Yik & Russell, 1999; Palermo & Rhodes, 2007).

Results from past research show that emotional faces draw attention faster than neutral faces (Palermo & Rhodes, 2007). More than neutral faces, faces showing emotional expressions could portray information about an individual's intentions and internal state, as well as potential information about their surroundings (Ekman, 1992; Yik & Russell, 1999). For instance, a happy face could communicate the desire for social interaction, whereas an angry face could act as a warning or deterrent to others (Becker, Anderson, Mortensen, Neufeld, & Neel, 2011; Yik & Russell, 1999). In a study conducted by Hansen and Hansen (1988), subjects were faster at picking out angry expressions than happy expressions from a "crowd" in a visual search task. In other words, angry expressions "popped out." They called this the "face-in-the-crowd" effect, arguing that certain emotional expressions draw attention not only faster than neutral faces, but also faster than other emotional expressions. This face-in-the-crowd effect has been demonstrated across many visual search studies, wherein certain emotional faces were found faster among a group of neutral or competing emotional faces (Hansen & Hansen, 1988; Becker et al., 2011; Williams, Moss, Bradshaw, & Mattingley, 2011).

Some emotional expressions could draw attention faster due to the type of information they convey. This attentional bias towards certain emotion expressions over other emotion expressions is referred to as an emotional superiority effect. For instance, angry faces could draw attention faster than happy faces because anger is a possible indicator of threat (Hansen & Hansen, 1988; Putman, Mermans, & van Honk, 2004), resulting in an anger superiority effect. There are disagreements, however, about which emotional expressions consistently draw attention faster than others. Hansen and Hansen's original study found an anger superiority effect, but other studies have found other emotional superiority effects, such as for happy or fearful expressions (Savage, Lipp, Craig, Becker, & Horstmann, 2013, Williams et al., 2011, Putman et al., 2004, Yang, Zald, & Blake, 2007). For example, Becker et al (2011) found a happiness superiority effect across a number of experiments, and a study by Savage et al (2013) found both anger and happiness superiority effects in several visual search tasks.

This lack of agreement presents an issue regarding the notion that attention is drawn because of conveyed social information. If emotional expressions draw attention due to the social information they carry, then it would be expected that superiority effects would be consistent across studies. If, for instance, angry expressions draw attention faster than other emotional expressions directly due to the social information they convey as an indicator of possible threat, then angry expressions should remain more important relative to other emotion expressions regardless of the task. The meaning and social information contained in emotional expressions, and by extension the attention oriented towards them in response, should not change relative to each other across studies. In practice, however, this is not the case even within studies using the same type of task.

A number of possible reasons for these inconsistencies have been proposed. One of these potential causes is the type of tasks used to study emotional superiority effects. Visual search is common and was the task used by Hansen and Hansen (1988) in the original face-in-the-crowd study. However, visual search tasks may have a number of potential pitfalls. A paper by Hampton, Purcell, Bersine, Hansen, & Hansen (1989) examined the methods used by Hansen and Hansen, noting that the position of images within the visual search task may have affected the results. They argued that "pop-out" effects should not be dependant on crowd size or image position. Despite these criticisms, visual search tasks continued to be common. A later paper by Savage et al (2013) examined the visual search task further and found a number of additional potential problems, including issues with variations in the task and how participants search through images.

These problems could have origins in the design of the task itself. In visual search tasks, the participant finds a target face among distractor images. One of the potential drawbacks of this design is that, by having to look through multiple images, there are possibilities for search strategies that could skew the results (Savage et al., 2013; Hampton et al., 1989; Purcell, Stewart,

& Skov, 1996; Becker et al., 2011). The face-in-the-crowd effect assumes that an expression will "pop out" and draw attention faster, but different aspects of the visual search task, such as the array size and whether the target is fixed (one target expression, such as angry, across trials) or variable (target expression can change between trials), could influence how the participant searches through the distractor images, and by extension, the time it takes to find the target expression (Savage et al., 2013). Additionally, the assumption that an expression will "pop out" due to its emotional meaning may not hold for all emotional expressions, as Williams et al (2011) suggest that while angry faces indicate threat and draw attention, fearful faces may indicate threat in the surrounding environment and actually divert attention away (Williams et al., 2011). Thus, the "crowd" of multiple distractor images in a visual task may not show an emotional superiority effect that might otherwise be found for expressions such as fearful expressions. For this reason, it may be difficult to determine whether visual search tasks can adequately test for an emotional superiority effect that is inherent to an emotional expression, rather than a potential interplay or relationship between that expression with its surroundings. When testing whether an emotional expression consistently draws attention over other expressions, this potential for interplay could make it difficult to determine if emotional superiority effects found in visual search are due to the social information conveyed through the expression itself, as emotional expressions that draw attention faster may only do so in the context of the visual search environment, but may not demonstrate the same effect in other tasks or situations.

In addition to visual search, another task used in studying emotional superiority effects is breaking continuous flash suppression (b-CFS). When people perceive their environment, not everything can be given equal levels of awareness due to the limits of sensory processing resources. In order to orient attention and give preference to more important stimuli, some level of preconscious processing must take place. B-CFS measures which stimuli gain access to awareness faster, and thus may be given preference in preconscious processing (Stein & Sterzer, 2011; Yang et al., 2007; Alpers & Gerdes, 2007; Hedger, Adams, & Garner, 2015). B-CFS was developed from binocular rivalry, and as such, b-CFS works by presenting a different stimuli into each eye. The participant perceives only one of these stimuli at a time, resulting in one image being suppressed (Tsuchiya & Kock, 2005). In b-CFS, one eye is shown an image while the other is shown a high-contrast changing mask (Tsuchiya & Kock, 2005; Gayet, Van der Stigchel, & Paffen, 2014). This gives more control over the duration and strength of interocular suppression, which makes b-CFS suitable for studying preconscious processing (Tsuchiya & Kock, 2005; Yang et al., 2007). The dynamic mask suppresses the static image, so at first, the participant only perceives the mask. Over time, the static image breaks interocular suppression and reaches awareness (Tsuchiya & Kock, 2005; Gayet et al., 2014). The stimuli that breaks interocular suppression and reaches access awareness faster may receive prioritized preconscious processing.

Past research using b-CFS shows that stimuli that are emotionally relevant may reach access to awareness faster, and thus may be among the stimuli that receive prioritized preconscious processing. A study by Alpers and Gerdes (2007) using binocular rivalry found that emotional faces predominated over neutral expressions. This preference for emotional faces held in a number of studies using b-CFS; however, like visual search, there are inconsistencies regarding the emotional superiority effects found. For instance, studies by Yang et al. (2007) and Hedger et al. (2015) found that fearful faces broke suppression faster, whereas Stein and Sterzer (2011) found an effect for positive emotions. Hedger et al. (2015) stated that rather than emotional relevance, they believe the faster access to awareness of fearful faces may be due to low level image features, such as effective contrast. If that is the case, then the images themselves, and not just the information conveyed through the emotional expressions portrayed, may be contributing to the superiority effects found.

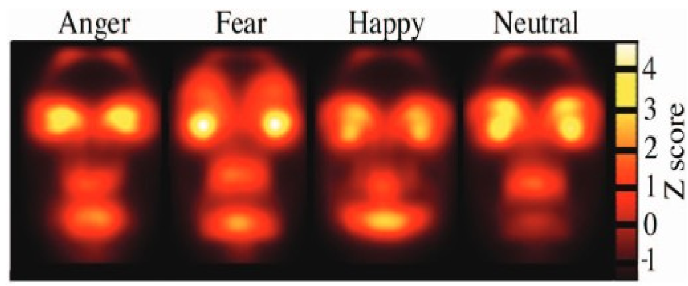

Figure 1. *From Hedger et al. (2015). Fearful faces have a sensory advantage in the competition for awareness. Variations in effective contrast in emotional faces.*

Thus, in addition to the task, the stimulus set chosen, and by extension the images used, may play also a role in affecting the type of superiority effect found. Savage et al. (2013) recreated past studies using two different, commonly used stimulus sets: the Ekman & Friesen (1976) database and the NimStim database (Tottenham et al, 2009). If certain emotional expressions draw attention due to their social information, it would be expected that images of the same emotional expressions in different stimulus sets would have comparable results because the images' emotional content, and thus conveyed social information, are the same. Savege et al. (2013) found, however, that the set used influenced the results. When using the NimStim database, they found an anger superiority effect. In contrast, when using the Ekman & Friesen database, they found a happiness superiority effect. Even within the same set, the effect could be influenced through using different subsets of images of the similar emotional expressions. With the NimStim database, they found the previously mentioned anger superiority effect while using the "open mouth happy" expression, but when they used the "exuberantly happy" expression, the results returned a happiness superiority effect. Savage et al. (2013) concluded that which emotional superiority effect a study finds may be largely due to which stimulus material was used.

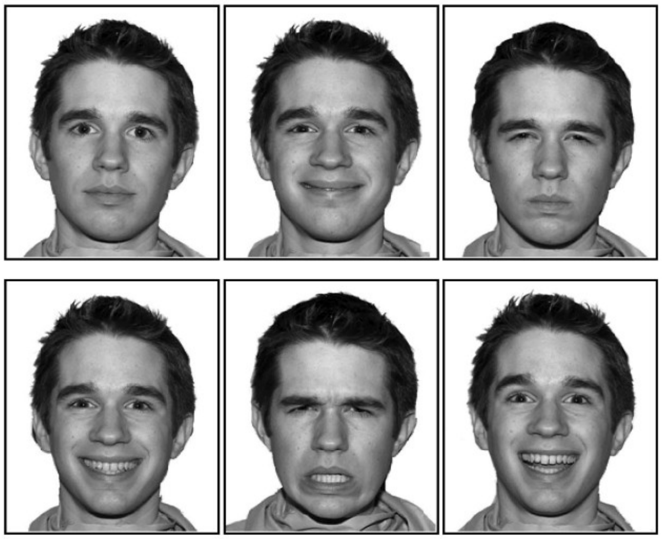

Figure 2. From Savage et al. (2013). In Search of the Emotional Face: Anger Versus Happiness Superiority in Visual Search. Neutral, happy, and angry expressions with variations: closed-mouth and open-mouth, and exuberantly happy in the lower right.

From this conclusion, Savage et al. (2013) proposed that low level image features could be a possible reason for this inconsistency in emotional superiority results when different stimulus sets are used. This position regarding the involvement of low level image features was shared to varying degrees with other studies (Purcell et al., 1996; Becker et al., 2011). Hedger et al. (2015) went so far as to suggest that we evolved certain expressions to take advantage of the preferences already in our visual system. Fearful faces, for instance, may have evolved to have wide eyes and mouth, and therefore higher effective contrast, in order to draw attention. Additionally, Becker et al. (2011) speculate that the happiness superiority effect they found was not due to the emotional content, but rather because happy expressions had evolved to be less ambiguous. Both arguments suggest that attention does not follow the emotional content of the expression, but instead the low level features in the image. Rather than emotional expressions drawing attention due to the social information they convey, emotional expressions evolved to take advantage of how attention was already being allocated based on low level features (Hedger et al., 2015; Becker et al., 2011; Horstmann & Bauland, 2006). Given this argument along with the previous observation that we would expect emotional superiority effects to be consistent if emotional superiority effects were driven by emotional content and social information, it seems likely that the inconsistencies between studies may be at least partially due to low level features.

Low Level Features

Every image can be decomposed into certain features, such as contrast, orientation, and spatial frequency. These low level features are among the first to be extracted by the visual system (Hubel & Wiesel, 1959) before higher level processing for meaning and content. If emotional superiority effects are not consistent based on emotional content, then it may be these low level features that draw attention.

There are a number of ways that low level features could affect attention. For instance, people show different sensitivities to different levels of contrast (Campbell & Robson, 1968). This contrast sensitivity has been defined in terms of the contrast sensitivity function, which defines

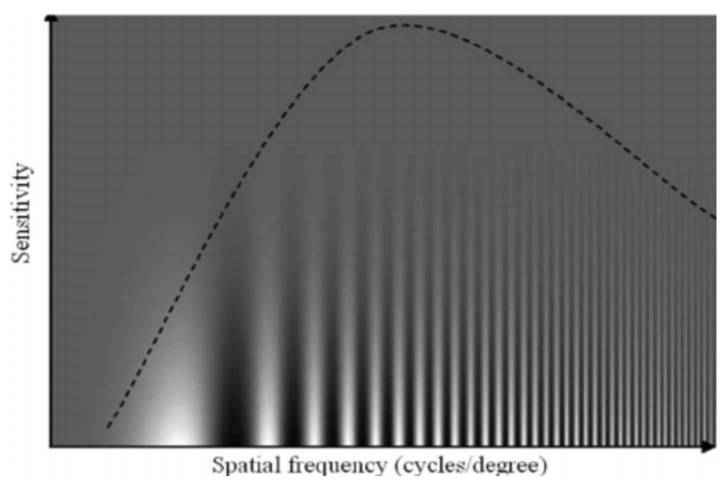

Figure 3. A visualization of the contrast sensitivity function.

a threshold for spatial frequency, or level of detail, and contrast in an image that people are sensitive to. Lower contrast and details that are too fine or too coarse lead to lower sensitivity, which makes them harder to perceive (Campbell & Robson, 1968). People tend to prefer certain orientations as well. The oblique effect describes the tendency for people to prefer horizontal and vertical orientations over oblique orientations (Li, Peterson, & Freeman, 2002; Campbell, Kulikowski, & Levinson, 1966). The features of contrast, spatial frequency, and orientation in an image therefore may affect how readily it draws attention.

In addition to this contrast sensitivity, spatial frequencies have also been found to affect processing of emotional expressions. When viewing an image, people process a broad spectrum of spatial frequencies. However, this spectrum can be broken up into more specific ranges, such as high or low spatial frequencies, which contain fine or coarse detail respectively. In research, these ranges have been used to study how people process emotional expressions. For example, Vuilleumier, Armony, Driver, and Dolan (2003) found that there are different neural pathways for processing different ranges of spatial frequencies. Low spatial frequencies were processed

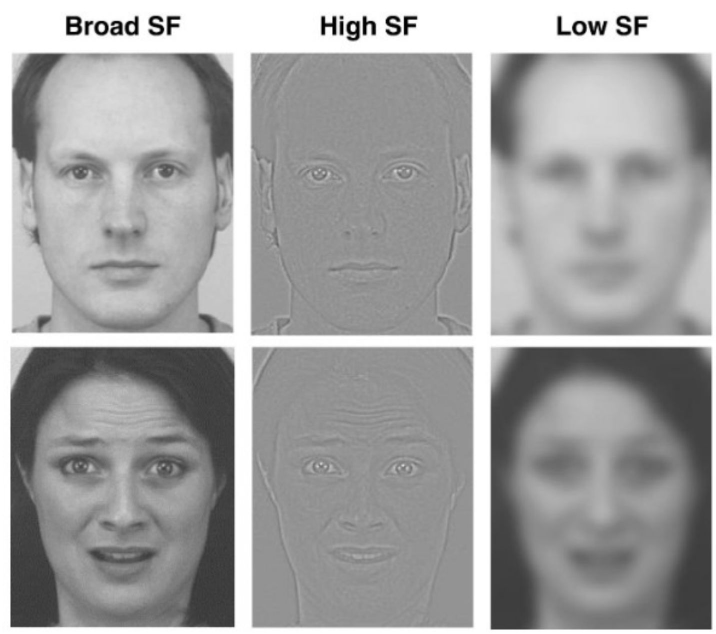

Figure 4. *From Vuilleumier et al. (2003), Distinct spatial frequency sensitivities for processing faces and emotional expressions. An example of how researchers break the broad spatial frequency spectrum into ranges.*

rapidly through the amygdala and aided in identification of emotional expressions such as fear, whereas high spatial frequencies were processed more through the fusiform cortex and were associated with facial recognition (Vuilleumier et al., 2003). Another study by de Jong, Engeland, and Kemner (2008) examined the connection between spatial frequencies and gaze shift cues in subjects with autism. Compared to the control group, emotional faces did not elicit gaze shifts in subjects with autism. Additionally, subjects with autism showed a preference for high spatial frequencies whereas control group subjects showed preference for low spatial frequencies. This indicates that a bias towards certain spatial frequency ranges could affect the ability to respond to emotional expressions and the social information they convey.

This evidence that low level features could draw attention and affect the processing of images of emotional faces may support the proposition that low level features could drive emotional superiority effects. Spatial frequency and orientation may give us insights to what features people are more sensitive to and are drawn to more readily. Because spatial frequency has also been connected to the processing of emotional expressions, it may be that these features, rather than emotional content, could be a major factor in emotional superiority effects and help explain the inconsistencies across past research.

Generalizability

As previously mentioned, low level features are the building blocks for images, and these low level features could contribute to emotional superiority effects. Because these features are in all images, they should not be unique to images of emotional faces. It is possible that other images share similar spatial frequencies to emotional faces, which would result in similar attentional effects. In other words, if emotional superiority effects are caused in part by low level features and spatial frequency in particular, then those effects should be generalizable to other images with similar spatial frequency content.

In order to test this, we used machine learning to classify cars using important features in emotional expressions (see Appendix). Cars were chosen because the front of cars are often thought to be similar to faces (Windhager et al., 2012; Kühn, Brick, Müller, & Gallinat, 2014), and therefore may have similar feature content. We used both the front and sides of cars to test whether this face-like configuration increases the likelihood that people react similarly to faces, or whether classified car sides have similar reactions despite being non-face-like.

Using machine learning to classify the images according to their similarity to spatial frequencies in images of emotional expressions could give us greater insight to the role of low level features than experimental design alone. In studies on spatial frequency, it is common to break the broad spectrum into ranges, like high and low (de Jong et al. 2008; Vuilleumier et al., 2003; Winston, Vuilleumier, & Dolan, 2003). This could potentially fail to capture interactions between spatial frequencies across multiple ranges, combinations, and strengths. However, identifying important subset combinations of spatial frequencies experimentally would be difficult and time-consuming as there would be too many possible combinations to test, and exploring the relationships and interactions would be highly impractical through experimental design alone. Using machine learning to identify the potential important features in the spatial frequencies of an emotional expression bypasses these limitations and provides a useful starting-point for experimentation.

Methods

Participants

12 participants (3 female) participated in this experiment, ranging in age from 21 to 28. Participants were recruited by word-of-mouth, social media, and flyers. Compensation was offered as credit through Utrecht University for participation. All participants had normal or corrected vision and had no history of epilepsy.

Apparatus and Stimuli

Stimuli were presented on two 27-inch ASUS pb278q monitors, with a resolution of 2560 x 1440 and a frame rate of 60 Hz, reflected via two mirrors, angled such that a separate image was reflected into each eye at a viewing distance of 45 centimeters. The stimuli were presented in the center of the screen within a 1/F noise frame, with a fixation dot present in the center of the frame. The remainder of the screen was gray. Before starting the experiment, the participant used the right and left arrow keys to adjust the frame on one screen to create stable binocular fusion.

Frontal gaze happy, angry, and neutral Caucasian faces from the Radboud face set (Langner, Dotsch, Bijlstra, Wigboldus, Hawk, & van Knippenberg, 2010) were used. Face images were cropped using the Viola Jones face detection algorithm (Viola & Jones, 2004) and resized to a 250x250 pixel square. Images of car fronts and car sides were selected from the CompCars dataset (Yang, Luo, Loy, & Tang, 2015) via classification by a combination of machine learning and models (see Appendix 1) and resized to a 250x205 pixel square. For each category, images of cars were classified as happy or angry based on their spatial frequency content (above 70% average accuracy). There were 39 images in each category: happy, angry, and neutral faces, happy and angry car fronts, and happy and angry car sides. All images were converted to grayscale and presented within the noise frame, which extended 50% larger than the stimuli size. A full-contrast black and white dynamic mask refreshing at a rate of 10Hz was prepared and presented in each trial in b-CFS.

Procedure

Task 1: Breaking Continuous Flash Suppression. The first task was a breaking continuous flash suppression (b-CFS) design. Each trial was led by a gray square and fixation point within the frame. The participant pressed the spacebar when they were ready to proceed. In each trial, a full-contrast black and white dynamic mask refreshing at a rate of 10 Hz was presented on one screen. On the opposite screen, an image was presented, beginning at 0% contrast and ramping up in contrast over 1 second until full contrast was achieved. There were seven conditions: faces with happy, angry, or neutral expressions, car fronts classified as happy or angry, and car sides classified as happy or angry, with 36 images per condition. A total of 546 trials were conducted, and trials were randomized and counterbalanced. Twenty-eight catch trials were included. Two types of catch trials were used: in one type, a mask was presented with no image, and the in the other type, an image was fused with the mask so that it was visible immediately. The participant was instructed to press the spacebar when any part of an image that was not the mask became visible.

Task 2: Human Rating. The second task required participants to rate the images as "happy" or "angry" via a sliding scale. Along with the image frame, a slider scale was presented on screen.

One image was presented at a time, and the participant was instructed to rate the image with the question, "Is this more happy or angry?" using the slider scale provided. Each image was presented once, and the responses were collected as percentages between -200 to 200 from happy to angry.

Results

B-CFS. The median reaction time for each condition per participant was taken and averaged across all participants (see figure 5). Using a repeated measures analysis of variance (ANOVA) with factors image type (face, car front, car side) and emotion (happy, angry) revealed no significant difference in either emotion, $F(1, 33) = .21$, $p = .653$, or image type, $F(2, 33) = 3.03$, $p = .062$. Because the cars were classified by a rank order (see Appendix), and some cars were thus "more like" the spatial frequency content of emotional faces, a correlation analysis was conducted between reaction time and car image within each condition in order to determine if there was a relationship between rank order and reaction time. There was no significant correlation between car image and reaction time in any of the seven

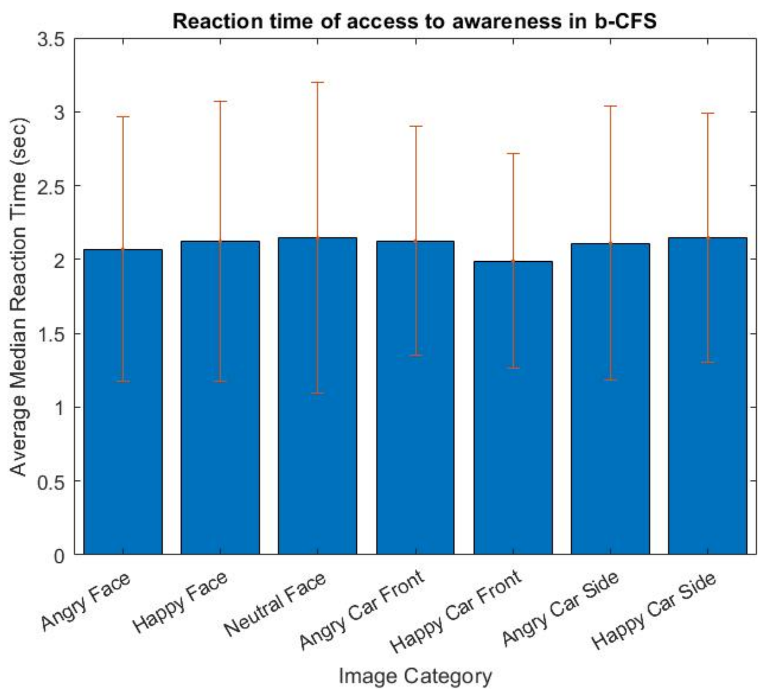

Figure 5. Average reaction times and standard deviation in b-CFS.

conditions.

Human Rating. For each category, the average rating was calculated (see figure 6). A repeated measures analysis of variance (ANOVA) revealed a significant difference between image type, $F(2, 33) = 140.33$, $p < .001$, and emotion, $F(1, 33) = 147.13$, $p < .001$. Between car images, however, there were no significant differences for emotion, $F(1, 22) = .220$, $p = .644$. Another correlation analysis was conducted between the ranked order classified car images and their average ratings to see if car images ranked higher, and therefore classified "more like" faces, were rated differently from lower ranked car images. No significant difference was found across any of the seven conditions between higher-ranked car images and lower-ranked car images for rating.

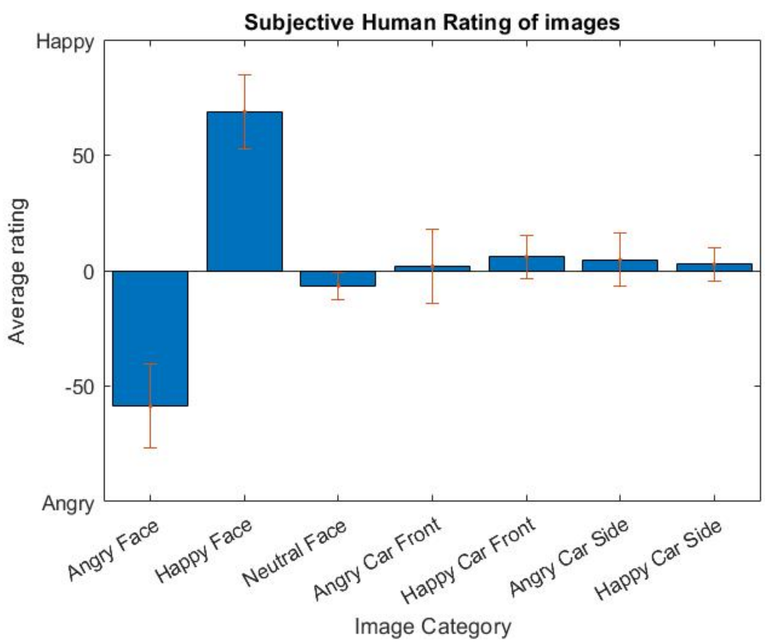

Figure 6. Subjective human ratings of images. Negative values indicate a higher anger rating, while positive values indicate a higher happy rating.

An additional correlation analysis was conducted between rating and b-CFS reaction time per image across each category. A significant correlation was found for neutral faces, $r = .39$, $p < .05$, but no significant correlation was found for any of the emotional expression conditions in

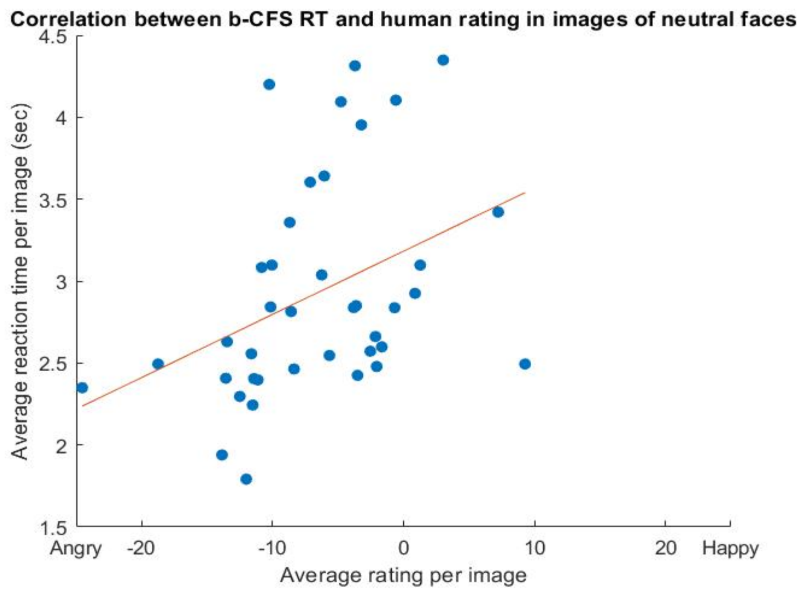

Figure 7. Correlation between human ratings and b-CFS reaction times for images of neutral faces.

either faces or classified cars.

Discussion

Our results showed that there was no difference between either image category (face, car front, or car side) or emotion in the b-CFS task. In contrast, in the rating task, human faces were rated as expected (e.g. happy faces were rated as happy), but images of cars were rated neutrally. The lack of an emotional superiority effect in either faces or car images in b-CFS makes it difficult to state the role of spatial frequencies in emotional expressions and how they affect access to awareness. As such, the inconsistencies in emotional superiority effects in past research remain open to further study.

Perhaps one of the most interesting aspects of our results is that, more than not demonstrating an emotional superiority effect, they were inconsistent with previous research regarding the tendency for emotional faces and faces in general to more readily draw attention than non-face objects. Our b-CFS results showed that not only did emotional faces not show shorter reaction time in access to awareness to neutral faces, reaction times to faces in general were not significantly shorter than to car images.

It has been well-established that faces draw attention in past research. In contrast, our results showed that faces did not reach access to awareness more readily, and thus did not necessarily show a prioritization in preattentive processing over non-face objects. Our results would either contest the notion that faces themselves are prioritized over non-face objects, or they would suggest that something in the stimuli we used prevented this effect from taking place. Given past research, it seems likely that faces do draw attention and have preference in preattentive processing. Studies such as Stein et al. (2015) suggest that objects such as cars may reach awareness faster when the participant has some expertise with the object; however, even with expertise, faces still had shorter suppression times in their study. Our study did not involve car experts and thus should not benefit from this addition level of expertise, yet car images did not have significantly longer suppression times than faces. Thus, we must consider whether the stimuli particular to this study may have contributed to this inconsistency with past research. It might be possible that by classifying the images of cars according to the spatial frequency content of emotional faces, the car images had properties similar enough to the face images to have similar access to awareness reaction times. Further studies may need to be conducted to see whether classified images of non-face objects have an advantage over non-classified images in b-CFS.

Our results follow the inconsistencies in emotional superiority effects in past research. We were not able to recreate an emotional superiority effect for either happy or angry expressions in faces using the Radboud face set. Given that the car images were classifying using the Radboud face set, the lack of emotional superiority effect in classified car images follows expectation, as their classification was based off of the spatial frequency content from the images in the Radboud face set. We would thus expect the resulting emotional superiority effect to be the same for classified cars as for faces. Whether the classified images would show similar superiority effects if those effects were to be found in face images remains a possibility, and the role of spatial frequency content requires more study.

In addition to lacking an emotional superiority effect, there was a significant correlation between the human ratings and reaction time in b-CFS only for images of neutral faces. None of the correlations between images of emotional faces or classified cars were significant between the tasks. This could indicate that neutral faces may be the only group that is rated by people in a way that follows how rapidly they reach access to awareness. The way people rate images of emotional faces and images of cars seem to be different from the way that these images are preattentively perceived in the b-CFS task.

In the human rating task, people were clearly able to identify happy and angry faces correctly. In comparison, car images tended to be rated as mostly neutral. While classified images of cars showed no difference from images of faces in access to awareness in the b-CFS task, there was a clear difference between these conditions in the human rating. This indicates that the identification and classification of emotional expressions in single static images do not necessarily follow the patterns of preferences in access to awareness and preattentive processing. Given the results from studies like Vuilleumier et al. (2003) and de Jong et al. (2008), the spatial frequency content in an image can affect how that image and the emotional expression portrayed in it is processed. Both studies split the broad spectrum of spatial frequencies into high and low ranges and found that low ranges were more associated with identification of emotional expressions. In the autism study by Jong et al. (2008), the autism group showed impaired gaze shifts along with a preference for high spatial frequencies over low spatial frequencies, potentially indicating that preference for certain spatial frequencies could directly impact the ability to respond to or identify emotional expressions. Given this understanding that spatial frequencies could play a role, it may be that in b-CFS, the classification of car images via spatial frequency content could have affected the speed of access to awareness for the car images, thus resulting in no difference between face images and classified car images. Regardless of how the participant judged the emotional content in the rating task, this lack of significant difference between image types in access to awareness could be related to spatial frequency content, and thus preferences in preattentive processing could still be related to low level features and spatial frequency content. In contrast, consciously perceiving the broad spectrum images in the rating task could allow more time to process the full spectrum of spatial frequencies as a whole and make higher-level judgments on the emotional content of the image. Although it is still unclear exactly how spatial frequency content affects access to awareness, the identification of emotional expressions in conscious awareness can still occur accurately, indicating that more information in the image is used in the process of conscious and accurate identification of emotional content.

Along with not following the trends of faces between rating and b-CFS, the ranking of car images in classification did not affect either reaction time in b-CFS or human rating. When the car images were classified, some were "more like" faces in that their spatial frequency content matched more closely than others, and the top images in each category were chosen. All images were still above a certain threshold for classification, so none were below a standard for likeness to spatial frequency content of faces, but some were more highly ranked than others. This did not have any bearing over the results per each image, however, and there were no differences found over the ranking of the images. This could mean that the differences between spatial frequency content above the chosen classification threshold was minimal in terms of the tasks. It is also possible that the spatial frequency content did not have as strong a bearing on the results and thus the differences between high-ranked and low-ranked were negligible. Further study in the overall impact of spatial frequency content as discussed earlier may reveal more about how something like rank order may affect results.

The role of classification in this area of study remains open to exploration. The use of machine learning and artificial intelligence could reveal more interesting trends regarding how the visual system handles low level image features and allocates attention. By classifying images according to low level image features like contrast and spatial frequency, we can better understanding how the visual system rapidly processes inputs and automatically orients attention in the environment. While this can be done in part through experimentation alone through techniques like breaking the spatial frequency spectrum into ranges, such as in the experiments with Vuilleumier et al. (2003) and de Jong et al. (2008), these techniques may only capture a small portion of the way the visual system handles spatial frequency and other low level features. By using machine learning, we can explore, analyze, and classify via a much wider variety of ranges, as well as the potential interactions and combinations between ranges. Studies on low level features and their role in preattentive processing of images may benefit a great deal from continued use in machine learning and artificial intelligence in the future.

References

Alpers, G. W., & Gerdes, A. B. M. (2007). Here is looking at you: Emotional faces predominate in binocular rivalry. *Emotion*. Advance online publication.

Becker, D. V., Anderson, U. S., Mortensen, C. R., Neufeld, S. L., & Neel, R. (2011). The face in the crowd effect unconfounded: Happy faces, not angry faces, are more efficiently detected in single- and multiple-target visual search. *Journal of Experimental Psychology: General.* Advance online publication.

Campbell, F. W., Kulikowski, J. J., & Levinson, J. (1966). The effect of orientation on the visual resolution of gratings. *Journal of Physiology, 187.* 427-436.

Campbell, F. W., & Robson, J. G. (1968). Application of fourier analysis to the visibility of gratings. *The Journal of Physiology, 197*, 551-566.

de Jong, M. C., Engeland, H., & Kemner, C. (2008). Attentional effects of gaze shifts and spatial frequency, but not in autism. *Journal of the American Academy of Child and Adolescent Psychiatry, 47*(4), 443-454.

Ekman, P. (1992). An argument for basic emotions. *Cognition and Emotion, 6*(3-4), 169-200.

Ekman, P. & Friesen, P. W. (1976). *Pictures of Facial Affect*. Consulting psychologists Press, 1976.

Gayet, S., Van der Stigchel, S., & Paffen, C. L. E. (2014). Breaking continuous flash suppression: Competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology.*

Hampton, C., Purcell, D. G., Bersine, L., Hansen, C. H., & Hansen, R. D. (1989). Probing "pop-out": Another look at the face-in-the-crowd effect. *Bulletin of the Psychonomic Society, 27*(6), 563-566.

Hansen, C. H., & Hansen, R. D. (1988). Finding the face in the crowd: An anger superiority effect. *Journal of Personality and Social Psychology, 54*(6), 917-924.

Hedger, N., Adams, W., & Garner, M. (2015). Fearful faces have a sensory advantage in the competition for awareness. *Journal of Experimental Psychology: Human Perception and Performance, 41*(6), 1748-57.

Horstmann, G., & Bauland, A. (2006). Search asymmetries with real faces: Testing the anger-superiority effect. *Emotion, 6*(2), 193-207.

Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology, 148*(3), 574-591.

Johnson, M. (2005). Subcortical face processing. *Nature Reviews Neuroscience, 6,* 766–774.

Kühn, S., Brick, T. R., Müller, B. C. N., & Gallinat, J. (2014). Is this car looking at you? How anthropomorphism predicts fusiform face area activation when seeing cars. *PLoS ONE, 9*(12).

Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H. J., Hawk, S. T., & van Knippenberg, A. (2010). Presentation and validation of the Radboud Faces Database. *Cognition & Emotion, 24*(8), 1377-1388.

Li, B., Peterson, M. R., & Freeman, R. D. (2002). Oblique effect: A neural basis in the visual cortex. *Journal of Neurophysiology, 90*, 204-217.

Palermo, R., & Rhodes, G. (2007). Are you always on my mind? A review of how face perception and attention interact. *Neuropsychologia, 45*(1), 75-92.

Purcell, D. G., Stewart, A. L., & Skov, R. B. (1996). It takes a confounded face to pop out of a crowd. *Perception, 25*, 1091-1108.

Putman, P., Mermans, E., & van Honk, J. (2004). Emotional stroop performance for masked angry faces: It's BAS, not BIS. *Emotion, 4*. 305-311.

Savage, R. A., Lipp, O. V., Craig, B. M., Becker, S. I., & Horstmann, G. (2013). In Search of the Emotional Face: Anger Versus Happiness Superiority in Visual Search. *Emotion*. Advance online publication.

Sorce, J. F., Emde, R. N., Campos, J., & Klinnert, M. D. (1985). Maternal emotional signaling: Its effect on the visual cliff behavior of 1-year-olds. *Developmental Psychology, 21*(1), 195-200.

Stein, T., Reeder, R. R., & Peelen, M. V. (2015). Privileged Access to Awareness for Faces and Objects of Expertise. *Journal of Experimental Psychology: Human Perception and Performance*. Advance online publication.

Stein, T., & Sterzer, P. (2011). Not just another face in the crowd: Detecting emotional schematic faces during continuous flash suppression. *Emotion*. Advance online publication.

Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., … Nelson, C. (2009). The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research, 168*(3), 242–249.

Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages.

Nature Neuroscience, 8, 1096-1101.

Valenza, E., Simion, F., Cassia, V. M., & Umiltà, C. (1996). Face preference at birth. *Journal of Experimental Psychology: Human Perception and Performance, 22*(4), 892-903.

Viola, P., & Jones, M. J. (2004). Robust Real-Time Face Detection. *International Journal of Computer Vision, 57*(2), 137-154.

Vuilleumier, P., Armony, J. L., Driver, J., & Dolan, R. J. (2003). Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature Neuroscience, 6*(6), 624-631.

Williams, M., Moss, S., Bradshaw, J., & Mattingley, J. (2011). Look at me, I'm smiling: Visual search for threatening and nonthreatening facial expressions. *Visual Cognition, 12*(1). 29-50.

Windhager, S., Bookstein, F. L, Grammer, K., Oberzaucher, E., Said, H., Slice, D. E., Thorstensen, T., & Schaefer, K. (2012). "Cars have their own faces": Cross-cultural ratings of car shapes in biological (stereotypical) terms. *Evolution & Human Behavior, 33*(2), 109-120.

Winston, J. S., Vuilleumier, P., & Dolan, R. J. (2003). Effects of low-spatial frequency components of fearful faces on fusiform cortex activity. *Current Biology, 13*(20), 1824-1829.

Yang, L., Luo, P., Loy, C. C., & Tang, X. (2015). A Large-Scale Car Dataset for Fine-Grained Categorization and Verification. *In Computer Vision and Pattern Recognition* (CVPR).

Yang, E., Zald, D. H., Blake, R. (2007). Fearful expressions gain preferential access to awareness during continuous flash suppression. *Emotion, 7*(4), 882-886.

Yik, M. S. M, & Russell, J. (1999). Interpretation of faces: A cross-cultural study of a prediction from Fridlund's theory. *Cognition and Emotion, 13*(1), 93-104.

Appendix

Machine Learning

For this study, we employed machine learning to classify non-face images based on spatial frequency content extracted from emotional facial expressions. To accomplish this, we used two steps using different algorithms and models: 1) find the spatial frequencies that can be used to classify emotional faces most accurately and 2) use these features to classify images of cars. This appendix is a brief overview to the main points of that process.

To extract the features, we used a multi-model support vector machine (SVM). First, images of emotional faces (angry, happy, and neutral) were loaded, and a Fourier transform was used to find each image's frequency domain. As these domains contain values on a continuous spectrum, they were broken into 128 discrete sections, the number chosen based on processing resources. These sections are referred to as features and represent a portion of similar spatial frequencies and orientations.

To find the key features for classifying emotional faces, we used a multiple models based on support vector machines (SVM). Model performances were calculated based on the exclusion and inclusion of each feature. To start, features were removed, and if model performance in classifying emotional faces became worse without a certain feature, then it was considered likely that this feature should be included. If the model performance improved or remained the same after a feature was excluded, then that feature was less likely to be important in classification. After this exclusion protocol, starting combinations of the chosen features were compared. Features were then tested for inclusion, and features that improved model performance when added were included. The final result was a minimal model, and its performance underwent boot-strapped and counter-balanced testing. The resulting model's performance averaged above 70% accurate in classifying images of the desired emotional faces.

To classify the cars based on these features of emotional faces, the output of the features in the previous final model were ranked, and the top twenty features for each emotion were used. These features were weighted such that the top-ranked feature had the greatest impact on classification, while the lowest-ranked feature in the top-twenty had the least impact comparatively. For each of these top-twenty features, the frequency strength for that feature was averaged across the frequency domains of all of the face images in each category (angry and happy), resulting in a single average value of strength for each important feature for angry or happy expressions. Over 500 car images were collected, and a Fourier transform was conducted on each image in order to extract the features. The same top-twenty features for every car image were individually compared to the averages from the features of face images (see figure 7). For each feature, a score and a classification value was assigned based on the comparison between the

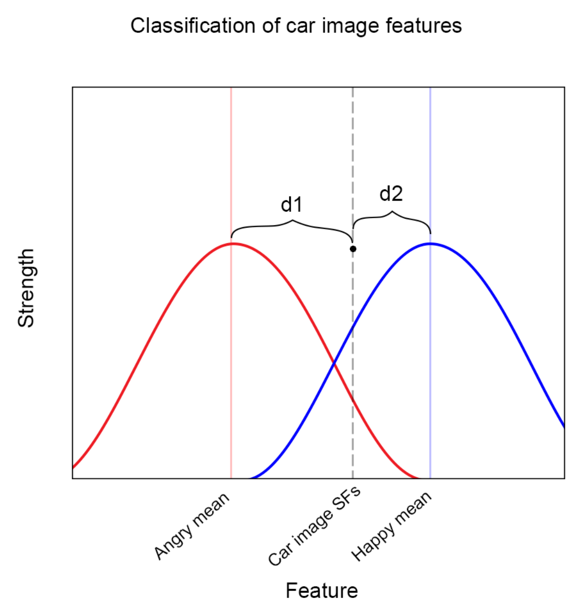

expression means. The classification was binary according to which expression mean the feature was more like, and the score was calculated based on the distance from the mean, or how “close” the feature was to the expression mean. After every feature was assigned a score and classification value, a combined score and classification was calculated based on the weights of the feature rankings. The combination of these weighted features resulted in a score of how strongly like one emotion expression the image was compared to the opposite, and the car image was given its final classification according to that score. After all the images were given a score and classification, the top 36 images in each of the four image categories (car front or car side per emotion) were chosen. As an additional precaution, we used a cut-off threshold score of 80% similar to the classified emotional expression, but all of the top images were well above this threshold.