An Experimental Analysis of the Effectiveness of Features in Chernoff Faces

Christopher J. Morris, David S. Ebert, Penny Rheingans

University of Maryland Baltimore County¹

ABSTRACT

Chernoff faces have been proposed as a tool for scientific and information visualization. However, the effectiveness of this form of visualization is still open to speculation. Chernoff faces, it is suggested, make use of the human's apparently inherent ability to recognize faces and small changes in facial characteristics. Limited research has been conducted to assess how well Chernoff faces make use of this ability. So far, it is still unclear how humans recognize faces and whether or not a specific set of rules governs the process. A particular area of interest is whether or not certain features are pre-attentive. Furthermore, what effect a certain number of distracters (i.e. more faces) have on the attentiveness of various features is also of concern. This information could be used to maximize the effectiveness of Chernoff faces by providing an indication of which applications would be best served by the use of Chernoff faces. In order to address this issue, we have conducted a user study, which tested the effectiveness and pre-attentiveness of several features of Chernoff faces. Our user study indicated that the perception of eye size, a specific face, eyebrow slant, and the combination eyebrow slant with eye size is a serial process (not pre-attentive). Our study also indicated that for longer viewing times (two seconds), eye size and eyebrow slant were the most accurate features. These initial results indicate that Chernoff faces may not have a significant advantage over other iconic visualization techniques for multidimensional information visualization.

---

¹ICSEE Department, U. of Maryland Baltimore County, 1000 Hilltop Circle
Baltimore, MD 21250, {cmorris, ebert, rheingan}@csee.umbc.edu

INTRODUCTION

In scientific and information visualization, the goal is to find a balance between the volume of data presented and to display that data as effectively as possible. One novel approach that has been proposed is the use of Chernoff faces to represent scientific or information data [Ch73].

Chernoff faces were first proposed by Herman Chernoff [Ch73] in 1973, as a way to represent multivariate data in a manner that is easily discernible by the human viewer. The faces consist of two-dimensional line drawings that contain a variety of facial features. Five example Chernoff faces are shown in Figure 1. These facial features can be mapped to different dimensions in a multidimensional data set. The reasoning behind the use of Chernoff faces lies in the notion that humans are adept at face recognition and are able to notice reasonably acute changes in facial characteristics. Coupled with the ability to represent multivariate data, the

Figure 1: Five example Chernoff faces.

use of Chernoff faces is a provocative method to represent data. However, the question that still remains to be answered is the following: How effective are Chernoff faces in reality?

To address the aforementioned question, we decided to investigate the pre-attentiveness of the features that comprise the Chernoff face. If certain features are more pre-attentive than others, this must be taken into consideration when trying to design a robust and informative representation.

CHERNOFF FACES AND THE PSYCHOLOGY OF FACIAL RECOGNITION

Chernoff first developed the idea of using human facial characteristics as a means to visualize data [Ch73]. He generated faces that possessed up to eighteen distinct facial parameters. These parameters ranged from length of nose and size of eyes to radius to the corner of the face and eccentricity of the lower face. Chernoff used these facial representations to investigate two

proposed measures of efficiency. One experiment had an investigator group together similar faces. The second had an investigator look at a sequence of faces that corresponded to successive points in time, in order to evaluate at which points in time the faces' character exhibited a noticeable change. There have also been studies [Ma] that examine how people group faces.

The fundamental reason for using Chernoff faces is that humans can easily recognize distinct faces and notice changes in facial features, even small and subtle changes. To the contrary, however, it appears that the research concerning facial recognition is significantly far from being conclusive, and there are many questions that have yet to be answered. Faces are a specific class of stimuli that humans recognize [Da86]. However, the manner in which humans recognize faces does not contain a unique process that is not used in recognizing other objects. Therefore, other objects that are recognized with the same process would be just as effective for visualization as faces.

It was believed, at first, that there were some distinguishing characteristics governing facial recognition. Earlier experiments initially conveyed such an idea. For example, Yin [Yi69] believed that the novelty of facial recognition came from the human's apparent inability to recognize inverted faces. However, these findings were later deemed inconclusive [Da86] because there was no control over the familiarity or complexity of faces as compared with other stimuli.

One of the biggest problems in trying to define how people recognize faces is that various facial features have a quite different meaning, depending on who you ask. Within ethnicities, races, sexes, religions, and even social classes, faces and their features hold varying connotations. However, there is still a belief that certain faces or facial characteristics produce remarkably uniform impressions in perceivers (facial stereotypes) [A188]. The challenge comes in determining where the consistencies reside. As a result, psychologists are still uncertain about the determinants of facial stereotypes [A188].

TAKING ADVANTAGE OF PRE-ATTENTIVE PROCESSING

Pre-attentive processing is essential in grouping large sets of information, including visual stimuli. Understanding which features in a face are pre-attentive, in comparison to others, would have a very beneficial impact on how to effectively utilize Chernoff faces. According to Treisman, pre-attentive processing is visual processing that is apparently accomplished automatically and simultaneously for the entire visual field of view [Tr86].

Many studies have investigated which stimuli are pre-attentive. One common procedure is to measure the response time to find a target in a set of "distracters". If a stimulus is pre-attentive, the response time should be independent of the number and types of "distracters" presented with the stimulus. Another method is to display a group of elements, with one element different than the rest in some way, for a short period of time (commonly 250 milliseconds) and then determine whether or not the viewer was able to pick out the unique element.

Sophisticated image processing systems have been used to manipulate various facial features, and have shown how sensitive humans actually are to tiny variations in stimuli [Br88]. To some, small movement in the mouth or eyes can create the impression of an entirely new person [Mu98]. With all of this in mind, it appears that humans are able to detect changes in facial features and are able to quickly focus on certain features. The problem we chose to investigate is whether or not this holds true for the 2-D Chernoff faces and, if so, for which features.

EXPERIMENTAL METHODOLOGY

To test the pre-attentiveness and effectiveness of different features in Chernoff faces, we conducted a controlled experiment with twenty-four subjects. Display screens were created by a Java 1.1 application, which displayed Chernoff faces for a set amount of time. The experiments were run on a Sun Sparc 5.6 workstation.

This experiment employed a two-factor, within-subjects design. The two factors, which were systematically varied to create the test conditions, were target feature and number of faces in the display. Four levels of number of faces were tested (5, 10, 25, or 50). The order of presentation of the different number of faces was randomized.

The target feature factor took on four values:

- small eyes (Figure 1)
- a specific face (Figure 2)
- inwardly slanting eyebrows (Figure 3)
- combination of small eyes and inwardly slanting eyebrows. (Figure 4)

Figure 1: Small Eyes (the face on the left)

Figure 2: A Specific Face (the face on the left)

Figure 3: Inwardly Slanting Eyebrows (the face on the right)

Figure 4: Combination of Small Eyes and Inwardly Slanting Eyebrows (the face on the right)

Except for the head eccentricity and the pupil size², which remained constant, all of the other features on the face were randomized. As a result, there were a total of eight (seven for the coupling) other features that acted as “distracters”.

Initially, each subject was required to undergo a training session where two sets of faces were displayed for each target feature, one for a time of two seconds, another for a time of .4 seconds. .4 seconds was used, as opposed to the commonly accepted time for pre-attentiveness, .25 seconds[Tr86], in order to achieve a consistent display time within Java. Once a set of faces was displayed, the user was asked to identify whether they saw a face with a particular feature or not. Next, the user was able to click an “answer” button to view the set of faces with the correct face highlighted. If there was no correct face present, they were notified of this as well. Once the subjects completed the training session, they were given the actual experiment. Figure 5 details the display screen for the training session with five faces displayed. In this instance, the user was asked to look for a face with small eyes. Figure 6 is screenshot of the same display with the correct answer designated.

² The pupil size remained constant to eliminate the illusion of an eye being larger or smaller when, in actuality, only the pupil size changed.

Figure 5: Small Eyes Trial in Training Session

Figure 6: Small Eye Answer Screen in Training Session

For the experiment, the subjects were asked to undergo the same process as in the training session. However, for the experiment, they were not allowed to see the correct answers and, for each feature, there were two blocks of trial. Each block consisted of four displays of either five, ten, twenty-five, or fifty faces. One block consisted of the faces being displayed for two seconds and the other for .4 seconds. The block of trials for each target feature was presented in turn. The order of the feature blocks was determined by partial-counter balancing as was the order in which the user went through each display in a feature block.

For each trial, the user was asked to look for a face with a designated feature. Depending on

whether they saw the target face or not, they responded by circling YES or NO on a questionnaire. Whether or not there was a face with the target feature was randomized for each trial. The correct answers were recorded so they could be compared with the users' responses.

RESULTS

Subject accuracy in the short-time or "fast" trials was not significantly different from that which would result by chance if subjects simply guessed the answers, suggesting that perception of iconic face features is not a pre-attentive process. Table 1 contains the analysis of variance (ANOVA) for the fast trials. It is evident from the data that there was virtually no correlation between the number of faces and the number of correct responses. Furthermore, it can be seen that the averages mostly lie within a range that can be attributed to chance (45 - 55%). Accuracy in the long-time or "slow" trials was substantially higher and significantly greater than chance, further supporting the idea that face feature perception is a serial process. An ANOVA for the slow trials is provided in Table 2.

	SS	df	MS	F
S	3.9583333	23	0.1721014	
F	0.8958333	3	0.2986111	1.0312826
N	1.1041667	3	0.3680556	1.5608195
FXN	0.9166667	9	0.1018519	0.4038308
FXS	19.979167	69	0.2895531	
NXS	16.270833	69	0.2358092	
FXNXS	52.208333	207	0.2522142	
Total	95.333333	383		

Table 1: ANOVA for "FAST" Trials

	SS	df	MS	F
S	5.9348958	23	0.2580389	
F	1.4661458	3	0.4887153	4.367285
N	6.0703125	3	2.0234375	12.018157
FXN	2.9609375	9	0.3289931	1.4458451
FXS	7.7213542	69	0.1119037	
NXS	11.617188	69	0.168365	
FXNXS	47.101563	207	0.2275438	
Total	82.872396	383		

Table 2: ANOVA for "SLOW" Trials

S = subjects
F = feature factor
N = number of faces factor
SS = sum of squared deviations
df = degrees of freedom
MS = mean squares
F = (MS of main factor/MS of interaction with subject)
(i.e. $MS_F = MS_F / MS_{FXS}$)

The ANOVA for the fast trials showed not significant effect for any of the factors. However, the ANOVA for the slow trials showed a significant effect for feature ($p < 0.01$), number of faces ($p < 0.001$). The factor p is the probability that the results were due to chance. In this instance, the p factor is low enough to assume that the results were strongly affected by the two factors.

Examination of the trial averages reveal that the eye size and eyebrow slant features resulted in the greatest accuracy (on average 16 percent more than with small eyes or 11 percent more than the specific face). The relatively low scores for the specific face come as somewhat of a surprise, however, they can probably be attributed to the amount of information, within the face, that the user felt he/she had to process before making a decision. Many subjects commented that they felt they had to look for many more features than just one or two to reach an accurate conclusion. Not surprisingly, subjects were more accurate in trials with fewer faces. The accuracy in long-time trials for the cases of 50 faces was not significantly greater than chance. This is further evidence for the serial nature of face feature perception.

CONCLUSIONS

The results of our user study indicate that Chernoff face feature perception is a serial process and is not pre-attentive. This suggests that the use of Chernoff faces for information visualization does not take advantage of human pre-attentive visual processing. Our user study also indicated that, of the face features studied, eye size and eyebrow slant produced the most accurate results. However, even when the user was given two seconds to view the faces, a display containing 50 faces did not produce results significantly more accurate than chance. This initial study indicates that Chernoff faces may not have significant advantage over other multivariate iconic visualization techniques. We plan to conduct further studies to validate our initial findings and explore other aspects of face features for information visualization.

REFERENCES

[Al88] Alley, T. R., "Physiognomy and Social Perception," *Social and Applied Aspects of Perceiving Faces*, pp. 167-185, 1988.

[Br88] Bruce, Vicki, *Recognizing Faces*, pp. 42-47, 1988.

[Ch73] Chernoff, Herman, "The Use of Faces to Represent Points in k-Dimensional Space Graphically", *Journal of the American Statistical Association*. June 1973, Vol. 68., No. 342, pp. 361-368.

[Da86] Davidoff, J.B., "The Specificity of Face Perception: Evidence from Psychological Investigations," *The Neuropsychology of Face Perception and Facial Expression*, edited by Raymond Bruyer, pp. 147-162, 1986.

[Ma] Marchette, David, "An Investigation of Chernoff Faces for High Dimensional Data Exploration," Computational Sciences and Informatics 803, George Mason University, www.farside.nswc.navy.mil/CSI803/Dave/chem.html

[Mu98] Muller, W., and Alexa, A., "Using Morphing for Information Visualization," *CIKM 98 Workshop on New Paradigms in Information Visualization and Manipulation (NPIV 98)*.

[Tr86] Treisman, A., "Features and Objects in Visual Processing," *Scientific American*, Vol. 255, Number 2, Nov. 1986, pp. 114-125

[Yi69] Yin, R. K., "Looking at upside-down faces," *Journal of Experimental Psychology*, Vol. 81, pp. 141-145, 1969.

[Ki96] King, W. and Jun Ohya. "The Representation of Agents: Anthropomorphism, Agency, and Intelligence". CHI96. April 13-18. pp. 289 - 290. 1996.