Emoji and Chernoff – A Fine Balancing Act or are we Biased?

Ricardo Colasanti[*] Swansea University
Rita Borgo[†] King's College London
Mark W. Jones[‡] Swansea University

ABSTRACT

We seek to answer the question on whether different geometrical attributes within a glyph can bias interpretation of data. We focus on a specific visual encoding, the Emoji, and evaluate its effectiveness at encoding multidimensional features. Given the anthropomorphic nature of the encoding we seek to quantify the amount of bias the encoding itself introduces, and use this to balance the Emoji glyph to remove that bias. We perform our analysis by comparing Emoji with Chernoff faces, of which they can be seen as direct descendant. Results shed light on how this new approach of feature-tuning in glyph design can influence overall effectiveness of novel multidimensional encodings.

Keywords: Emoji Glyphs, Glyph Bias, Balancing Glyphs, Chernoff, Multidimendional Data.

1 INTRODUCTION

Understanding multi-dimensional data involves tasks such as trend and sequence analysis, outlier detection, and gaining familiarity with a large set of data. A key aspect of data visualization is to be able to identify a suitable visual encoding for multidimensional data points. Glyph based visualization has emerged as a strong contender for data representation where multiple dimensions (up to 20 [15]) are mapped to distinct visual elements within the representation. Questions arise surrounding the human use of the visual encoding in terms of learnability and usability. Learnability relies on previous familiarity with the chosen token in a different context or that the token places low cognitive demands on the user. User studies provide evidence of good "learnability" if novel encodings provide comparable or better accuracy and speed to existing approaches.

We explore questions in this area through this research. We select a familiar visual encoding in the form of Emoji, of which usage in the visualization domain is, to our knowledge, novel. Emoji are highly stylized cartoon like faces, and are therefore strongly related to Chernoff faces [11] which have a long history in glyph visualization. We conducted a user study to test the usability of Emoji for tasks related to multi-dimensional data exploration, we explore performance when compared with conceptually similar layouts like Chernoff faces and standard representations like Star glyphs [10,29].

We also hypothesize that where multiple channels are encoded within a glyph, the individual visual elements comprising those channels can lead to visual bias, smile for example is visually highly salient and grabs attention automatically potentially influencing viewer's perception of non-happy/happy eyes in a face [8]. We think that it is vital to understand how a glyph can attach varying importance to each dimension in accordance to the weighting perceived through the corresponding visual channel. To demonstrate this idea, we introduce glyph balancing on Emoji. We seek to determine how visually effective each channel is within a glyph, and then adjust the size of that channel accordingly. We use as control the Star glyph which is trivially balanced, and Chernoff faces in which different features are probably more visually important than others. We aim to balance the Emoji, and test this balancing to see if it has any impact during usage.

[*]E-mail: r.l.colasanti@Swansea.ac.uk
[†]E-mail: r.borgo@kcl.ac.uk
[‡]E-mail: m.w.jones@swansea.ac.uk

Therefore our contributions are: (i) a proposal that glyph balancing is an important attribute of glyph design to remove visual dominance as a bias on the data, (ii) a methodology for a controlled user study to achieve glyph balancing, (iii) the use of this methodology to introduce a new balanced glyph – Emoji, (iv) the comparison of performance between the new glyph and existing Star and Chernoff faces glyphs.

2 RELATED WORK

In this section we present some of the related work and current approaches in the literature for glyph design guidelines and user studies, along with a discussion of the evolution of Emoji from representing emotion to other ideas.

Fuchs *et al*. [22] present an extensive overview and classification of user studies involving glyphs. They categorized tasks into visual search, similarity search and trend detection and identified study goals which were: comparing different glyph designs, variations and with reading data tables and text. We focus on similar tasks and comparing different glyph designs with an overall goal of identifying whether familiarity affects performance. We obtain accuracy and completion speed.

Fuchs *et al*. [21] present a user study comparing variations of Star glyphs with respect to accuracy and speed of picking similar stimuli amongst other experiments.

Borgo *et al*. [7] link theoretical models of semiotics with glyph design. Design guidelines are presented, and current design guidelines informed by the literature are also reviewed. Glyph design guidelines are also suggested by Chung *et al*. [12], along with the concept of being able to visibly sort glyphs to aid visual analysis of rugby video data. Ward [41] surveys glyph placement algorithms. Lie *et al*. [33] provide design guidelines for mapping data onto glyph visual channels. The semantics of attributes should also be considered when mapping to glyph channels [7,33,38,41], particularly to avoid introducing bias [34] where some visual channels of the glyph will have more discriminating power compared to others. Maguire *et al*. [34] introduce a qualitative approach to channel balancing, whereas we are introducing a quantitative approach and demonstrate through the user study that the bias effect really exists and can be removed. Metaphoric glyphs or pictograms are shown to reduce cognitive load by exploiting familiarity [31]. Glyphs for pre-attentive visualization are discussed by Healey *et al*. [23].

Barbieri *et al*. [4] studied the semantics of Emoji by a user study on how people perceive the meaning and relationship between Emoji, and also using twitter data to derive Emoji semantics. This demonstrates connection between perception and meaning by Emoji which could inform how they could be used as surrogates for visualization. Cappallo *et al*. [9] present an image classifier that produces Emoji to capture the semantics of the scene rather than a textual description.

Kelly *et al*. [30] demonstrate that Emoji are being appropriated for tasks other than creating emotional context or intonation in messages. In this user study, they categorized other uses, such as maintaining a conversation, playful interaction or a shared context specific to that relationship. The study focused on use within relationships. Moschini [37] traces the history and usage of Emoji (specifically the tears of joy). It documents that Emoji can move beyond portrayal of emotions, to also convey ideas. We take this a step further by exploiting their use for the visual representation of data.

Channel	Attribute
0	Left Eye
1	Right Eye
2	Left Brow
3	Right Brow
4	Smile

(a) Emoji Glyph

Channel	Attribute
0	Eye Width
1	Brow Width
2	Eye Slant
3	Nose
4	Mouth

(b) Chernoff Face

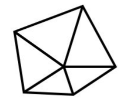

Channel	Attribute
0	Clockwise 15 min
1	Clockwise 27 min
2	Clockwise 39 min
3	Clockwise 51 min
4	Clockwise 03 min

(c) Star Glyph

Figure 1: Visual encodings design channel breakdown.

3 Visual Encodings

3.1 Face like Visual Representations

In order to represent multivariate data we have created a minimal face Emoji according to the Emoji design standard principles [14]. Our Emoji consists of a circular face with a background hexadecimal rgb color of #FFDD67. The Emoji has a "smile" of color #664E27. The Emoji has two eyes and two eyebrows. The eye consists of an outer circle in color White and an inner "pupil" circle of color #664E27. The eyebrows are arcs outside the area of the eye and are of color #664E27. The colors were taken from the Unicode®Emoji design guidelines [14], the same color pattern is also followed by Google® Emoji set [17]. We have limited our design to five attributes but further attributes could be mapped to Emoji features. The five mappable Emoji attributes are: depth of the smile, circumference of each outer eye, width of each eyebrow, as defined by the angle of the arc of the eyebrow. Fig. 1a shows a detailed breakdown by channel of our Emoji glyph design.

The creation of a multivariate face Emoji allows for direct comparison with two other multivariate glyphs: Chernoff face and Star.

The multivariate face Emoji is most similar to the well researched Chernoff face. The Chernoff face encodes data as a two dimensional facial caricature or cartoon. In a survey of glyph user trials by Fuchs [22] 25 of the 65 studies involved Chernoff faces. Within these studies the rendering of the Chernoff face varied considerably. For our study we are using the formulation described in the original paper by Chernoff [11]. The Chernoff face has 18 separate attributes. We have chosen five of the attributes that most closely correspond to our multivariate face Emoji: curvature of the smile; width of the eyes; length of the eyebrow; length of the nose; angle of the eye and eyebrow. Fig. 1b shows a detailed breakdown by channel of our Chernoff face design. For the remaining 13 features we have taken the midpoint values between the minimum and maximum size of each feature. Some studies also use an extra color variable [2], but in our study we used Chernoff's original rendition colors of black for outline and features, and white for background and fill color. We have scaled the size of the Chernoff face so that the face width matches the face width of the multivariate face Emoji. To assess the overall suitability of an Emoji image to represent multidimensional data, we compare the multivariate face Emoji and the Chernoff face with a non-face glyph.

3.2 Non-face like Visual Representations

Chernoff's original hypothesis was that the documented ability of humans to recognize facial expressions would improve the representation of multidimensional data. This hypothesis has been extensively tested [22]. The Star glyph has figured as the most common glyph used for comparison. For our studies we have used a five element Star glyph. The Star glyph combines a five sided polygon with a five armed whisker plot. Each attribute of our data set is mapped to each of the lengths of the whisker plot. The ends of each arm are joined by a line to their nearest radial neighbor. Fig. 1c shows a detailed breakdown by channel of our Star glyph design.

The width and height of the Star are scaled to match the circumference of the multivariate face Emoji. We follow the same color scheme used in most of the studies surveyed by Fuchs [22] where the Star glyphs were visualized with a foreground fill of black, and background fill white. The Star glyph adheres to the Gestalt principle according to which we have a preference for simple shapes [7]. Unlike the multivariate face Emoji and the Chernoff face, the Star glyph can be said to be abstract.

4 Glyph Balancing

In image correction pre-balancing color channels means to view and adjust the composite one channel at a time to blend it better with the background. Similar to color correction in images we attempt to balance the channels within an Emoji glyph.

Faces are important and salient visual stimuli, being central in human social interaction, providing critical information about age, gender, emotional state, intention, and identity of other human beings. Human neuropsychology [24] provides evidence that human perception of faces is different from human perception of objects or complex scenes. Moreover the active control of eye fixation plays an important functional role in a variety of cognitive and perceptual tasks with fixation sites and fixation durations closely time-locked to ongoing perceptual and cognitive processes [25]. The human visual system reorients the fixation point around the viewed image an average of three times each second, this is achieved via saccadic eye movements, i.e. the movement of our focus from one part to another. The path our eyes take across an object, or a picture, is a pattern of saccades. In both primates' and humans' saccades pattern, when looking at a picture of a face, follows a specific distribution as shown in Fig. 2. Saccade patterns when gazing at face like representations compare to the saccade patterns when gazing at human faces. Following this principle we have hypothesized the smile (S) to potentially be the strongest visual cues while eyes (E) and brows (B) channels to be the weakest.

$$S >> E|B \tag{1}$$

We only considered those features which were common between our design and Henderson *et al*. [24]. To validate our hypothesis we have run a smaller study involving a single task in a format similar to Task 1 (see Section 5.1). Participants were selected from a pool of visualization experts (both academic and postgraduate students) at Swansea University. Participants were shown pairs of Emoji glyphs stimuli, members of a pair differed by only a single channel. Participants were required to select, as quickly as possible, the glyph they believed represented the highest value overall. All participants reached almost 100% accuracy however differences were found with respect to response time. Fig. 3 shows the overall response time distribution. Interestingly the pattern resembles results from Henderson et al. study [24] with smile being the fastest channel.

Let our data X be a set of k-dimensional vectors $X_i = x^i_1, x^i_2, \ldots, x^i_k$, where the data is processed ready for glyph representation (by normalizing each channel). To depict the data using a glyph, a function

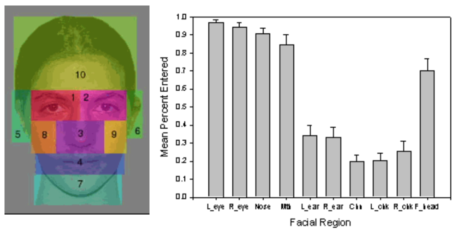

Figure 2: Mean percentage of fixation time per region (image courtesy of Henderson et al. [24]).

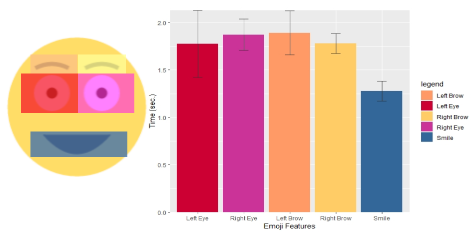

Figure 3: Mean response time per channel.

renders each attribute n, of a glyph, at size S_n, and for a specific glyph instance, G_i, each channel n is scaled by $S_n \times x_n^i$. Usually S_n is transformed such that it has a minimum value when $x_n^i = 0$, e.g., in the case of the star glyph, there is a minimum size for each axis so that it does not collapse to a point.

Glyph balancing results in a vector of global weights $\omega_1, \omega_2, ..., \omega_k$. The size of each channel is balanced by applying $S_n \times \omega_n \times x_n^i$ to its rendition.

The weights are discovered through the response times on the chosen task such that:

$$\omega_i = 1 + log(\frac{\overline{T}_i}{\overline{T}_k}) \tag{2}$$

where $\overline{T}_i$ is the mean response time from channel i and $\overline{T}_k$ is the mean response time over all remaining channels. Where there is a significant difference, we use the weight ω_i to rescale the respective glyph channel. This process is employed iteratively until there is no significant difference in response times, each time selecting the channel with the largest difference.

5 Experimental Overview

To evaluate our hypothesis over effects of glyph balancing over data interpretation we have devised a $3+1$ tasks empirical study. In the design of our study we have considered both elementary and synoptic tasks but opted for the latter. With respect to glyph-based designs in [22] elementary tasks are identified as those tasks where the user is concentrating on a specific, single visual attribute of a glyph and relating it to an attribute of the data set. Synoptic tasks are instead identified as those tasks that involve the user interacting with the glyph as a whole. The user is observing all of the visual attributes of the glyph rather than a specific attribute. Synoptic tasks normally involve two or more glyphs. We have chosen to investigate only synoptic tasks since our aim is to assess the relative performance of the three glyph types as a whole. Examples of synoptic tasks are visual search (outlier detection in our study), pattern detection and similarity estimation (parameters estimation in our study). It is these three tasks that inform our experimental design to assess the suitability of an Emoji image to represent multidimensional data.

Fig. 13, Fig. 14 and Fig. 15 in supplemental material show screenshots of how each task was implemented in our study. The order of the tasks within the study is fixed but the order of trials within each task is randomized, with the only constraint that all participants experience the same trial sequence within a task. For each task in our study we use a random distribution to generate artificial datasets. We opted to use artificially generated data to ensure that, for each task and trial, all glyphs within a display are separated by an amount equal or greater to the minimum channel just noticeable difference (JND, Section 5.5) to guarantee discriminability. More task specific design decisions such as sequence lengths for Task 3, and grid size for Task 2 and Task 3, were guided by results collected from a series of pilot studies, described in detail in Section 6.1, used to finalise the experimental design in terms of difficulty and feasibility.

5.1 Task 1: Parameters Estimation

The first of our tasks is parameters estimation. The participant is presented with a pair of distinct glyphs and required to select the one representing the highest value overall, as shown in Fig. 13. The source data is a 5-vector of numbers between 0 and 1. Four of the data channels are identical, with the remaining channel being the one the participant should identify to determine which glyph is highest. The task forces the participant to examine each channel both individually and as part of a collective, and to perform interpolation across the entire glyph to approximate the overall value. We generate a total of 90 trials corresponding to 6 repetitions per each channel per glyph.

5.2 Task 2: Outlier Detection

The second of our synoptic tasks is outlier detection. The user is required to find a target glyph within a square grid of glyphs of the same kind, as shown in Fig. 14. Within the grid both target and distractors are rendered from the same dataset, distractors are selected in such a way that distance to the target be above the smallest channel's measured just noticeable difference (JND, as described in Section 5.5) to guarantee discriminability. Similar to Task 1, we generate a total of 90 trials corresponding to 3 repetitions per 2 grid sizes for each channel and glyph.

Glyph. To assess the relative performance of each glyph type, the same set of data is used for each glyph. There are thus 30 different sets of target and grid. Each set is rendered three times with each of the three glyph types used once. The randomized order of the subtasks is controlled so that a glyph rendered in a subtask is not followed by one of the same type. The order of subtasks is arranged so that the same data set is not used in consecutive subtasks.

Grid Size. Within this task we use two grid sizes 5×5 and 6×6, and vary the position of the target within the grid. Chernoff's original hypothesis was that the documented ability of humans to recognize facial expressions would improve the representation of multidimensional data. One possible hypothesis for any improved representation is pre-attentive attraction [7, 23]. A consequence of pre-attentive attraction is that the pop-out effect should not be affected by the number of other objects vying for attention. This pre-attentive aspect of Chernoff faces was investigated by Sivagnanasundaram [39]. We followed Sivagnanasundaram's method of testing for pre-attentive attraction by varying the size of our glyph grid using grid sizes of 5×5 and 6×6.

5.3 Task 3: Pattern Detection

The third of our synoptic tasks is pattern detection. Participants are asked to locate one set of three glyphs embedded within a larger set. Each participant is presented with a 6×6 glyph grid. As aforementioned the 36 glyphs were rendered from an artificially created dataset, data were chosen at random with the condition that glyphs within a grid would be separated by a channel-wise JND. Each participant was presented with a vector of glyphs that occur

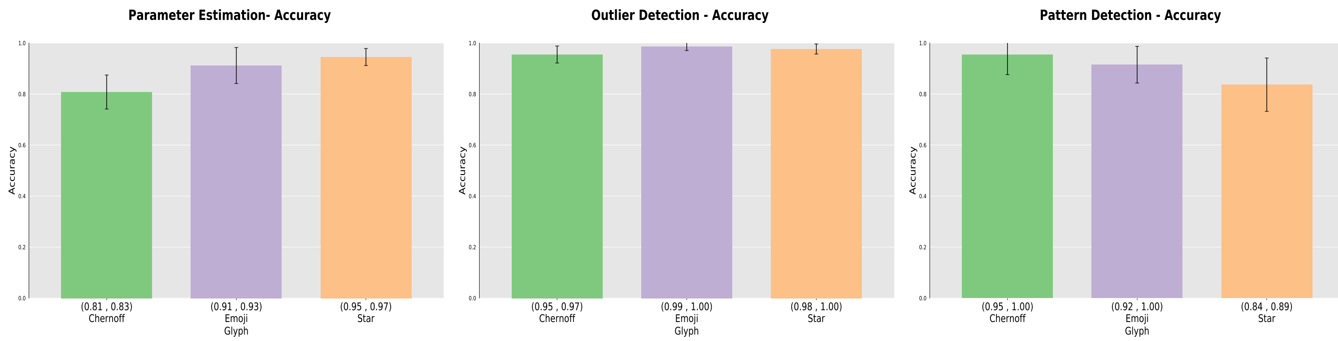

Figure 4: Analysis of accuracy results for Task 1, 2 & 3 respectively. Results shown for each visual encoding, (mean, median) values are indicated below each bar. Error bars show 95% confidence intervals.

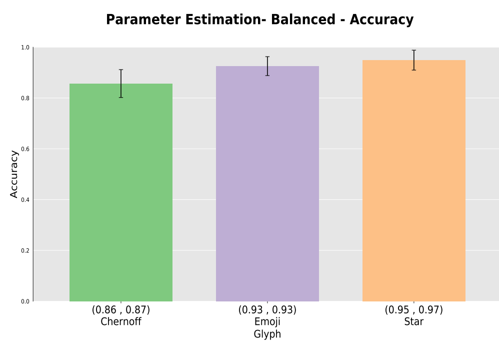

Figure 5: Analysis of accuracy results for Task 4 – Parameters Estimation. Results shown for each visual encoding, (mean, median) values are indicated below each bar. Error bars show 95% confidence intervals.

as a horizontal sequence in the grid. The task is to locate the vector within the grid and select any of the glyphs that match, as shown in Fig. 15. For each set of three glyphs distractors are randomly generated to enforce participants to focus on a triplet as a whole and not on a single glyph within the sequence, e.g. a three glyphs sequence cannot be always identified uniquely by a single glyph within the full grid.

The originally randomized order of the trials is controlled so that a glyph rendered in a subtask is not followed by one of the same type or position. Each participant saw identical trials in the same sequence. For Task 3 we generated 10 possible glyph arrangements for a total of 30 trials (i.e. 10 per glyphs).

5.4 Task 4: Parameters Estimation (Balanced Emoji)

In our fourth task we aim at measuring changes in performance of the Emoji representation after balancing its five channels (see Section 4). To achieve our goal we repeat Task 1 but with the new balanced Emoji in place of the original one. Apart from the use of the newly balanced encoding, the task remains identical to Task 1.

5.5 Just Noticeable Differences

To guarantee discriminability, target and distractors where designed to be above just noticeable difference (JND) for their respective visual encoding. Glyphs are by definition composite images therefore to compute JNDs we chose as threshold value for each glyph the maximum JND value of all the visual channel employed within its design. If we indicate with $glyphA_{ci}$ and $glyphB_{ci}$ the i^{th} channel of $glyphA$ and $glyphB$ respectively then:

$$glyphA_{ci} \neq glyphB_{ci} \forall i \in \{1, \ldots, 5\}$$

$$|(glyphA_{ci} - glyphB_{ci})| \geq max(jnd(glyphA_{ci}), jnd(glyphB_{ci}))$$

In our designs the main visual channels to be considered are size, shape, orientation and curvature. For size, shape and orientation we follow the same approach to compute JND thresholds as used by Chung *et al*. [13]. Size is therefore mapped to circle radius, following results which demonstrate how perception of size (e.g., area) is logarithmic and can be modelled using Weber-Fechners Law [3]. Shape JND is computed in function of the number of spikes of the star-shaped glyph. Shape differences are measured using image moment statistics [27]. Orientation follows the more conservative measurement proposed by Chung *et al*. [13], with changes between consecutive elements of at least 11.3°. We similarly avoid ambiguous orientations. Curvature JND threshold was computed according to work by Foster [19, 20], we relied on Weber fraction for contour-curvature discrimination and favoured the viewpoint invariant features of the curvature. Color is not used as a discriminant between target and distractors, however in the design of Emoji we selected a shade whose hue and value could be easily distinguished from the background as previously suggested [6].

6 Study Procedure

6.1 Pilot

Three pilot studies were performed to guide the design of the present study. In the first pilot participants were asked to rate the overall level of difficulty of the study as a whole and of each task. They were also asked for any general feedback and to detect any display errors or logic problems. Feedback from the first pilot were used to tune the study procedure in terms of number of trials and length of tasks. In the second pilot study we measured performances when varying grid sizes (for Tasks 2 and 3) and sequence length (for Task 3). We experimented with 3 grid sizes 4×4, 5×5, 6×6, results showed no significant effect for grid sizes below 5. In the case of sequence length we experimented with sequences of 2, 3 and 4, results showed no significant effect for sequence length below 3, while post-hoc comparison of results of sequences of length 3 and 4 showed no significant difference. Moreover during the debriefing interview the majority of participants (78%) reported to be focusing on either the first 3 or last 3 elements in a sequence to locate the correct answer. The findings guided our decision to design sequences of length 3 for Task 2, and restrict grid sizes to 5×5 and 6×6. The third pilot study was used to collect interim results to guide our balancing strategy as described in Section 4.

6.2 Participants

Participants for the first two pilots and the main study were recruited from students and staff at Swansea University. All participants were compensated with £5 vouchers. For the first pilot study we recruited 10 participants, all male below age 25. All pilot participants were Computer Science graduates with experience in data visualization. For the second pilot study we recruited 16 participants, 4 females

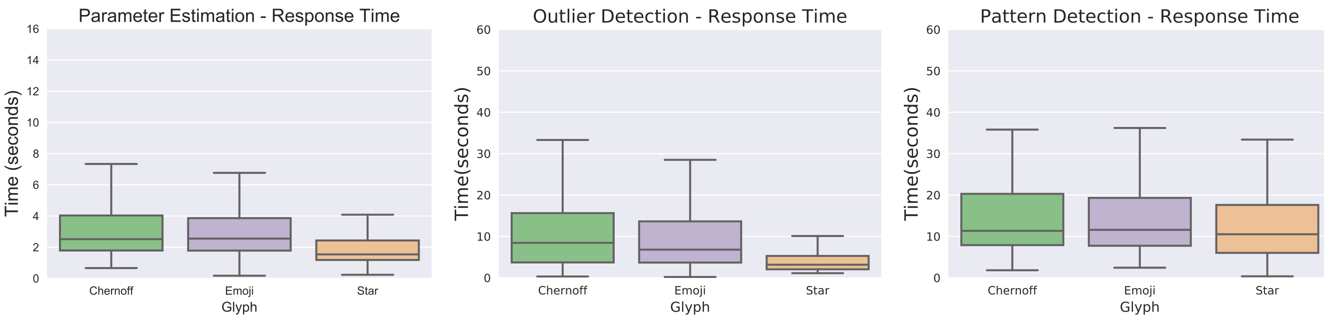

Figure 6: Analysis of response time results for Task 1, 2, & 3 respectively. Results shown for each visual encoding.

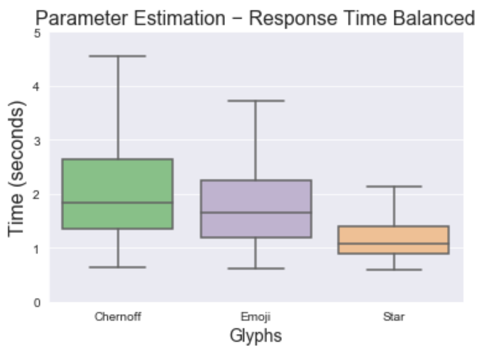

Figure 7: Analysis of response time results for Tasks 4 – Parameters Estimation with balanced Emoji glyph. Results shown for each visual encoding.

and 12 males all below age 25. All pilot participants were studying for a Computer Science degree. 12 described themselves as knowledgeable of data visualization. For the third pilot study we recruited 10 participants, 1 female and 9 males between the age of 23 and 45. All pilot participants were visualization experts at either academic or postgraduate level. Participants of all three pilot studies were automatically excluded from the main study. For the main study we recruited 18 participants – 5 females and 13 male. Ages ranged from 20 to 44 (Mean=23.24, SD=8.27). All participants had education at or above high school. 17 either had or were studying for a Computer Science degree. 1 described himself as knowledgeable of data visualization. All participants had normal or corrected to normal vision and were not informed about the purpose of the study prior to the beginning of the session.

6.3 Experimental Setting

The study took place in a university computer laboratory. The computers were all Viglen Omnino 5 computers running windows 10. The screen size was 24 inches with native resolution of 1,920 x 1,080. The maximum brightness was 300cd/m2. The contrast ratio was 1000:1 and the response time was 5ms. As described in Section 4.5 the tasks were presented as web pages. The browser that was used throughout was Google Chrome version 54.0.

Procedure. The study began with a brief overview read by the experimenter using a predefined script. Detailed instructions were then given through a self-paced slide presentation and video of the user interface. The presentation included a briefing on how to interpret each visual design. Participants completed a number of training trials for each task. Each training trial included a feedback to the participants regarding the correct answer. None of the stimuli used during training were repeated in the main tasks. All four tasks were completed in sequential order. Maintaining the same section order for each participant meant that each participant experienced similar experimental conditions. This increased robustness of the analysis of the collected data. To compensate for possible learning effects Task 1 through 3 required the user to perform ontologically different actions with different level of complexity, stimuli within each tasks were randomly generated to avoid repetition on stimuli across tasks. Each task featured a different level of complexity in terms of number and size of distractors to compare against and size of target. Randomness was introduced at trial level. Within a task, trials were randomized to avoid learning effects. The study was closely monitored and participants abode to the study requirements. At the end of each task a short multiple choice questionnaire was presented to collect qualitative information from the participants. At the end of the study each participant completed a post-study debriefing interview and questionnaire to collect demographic and further qualitative information. All three visual designs were at all times presented as valid options, especially during post-processing interview, to maintain unbiased judgment and preserve validity of the collected qualitative feedback. Participants were given no time limit to complete each trial but were encouraged to be as accurate and fast as possible.

7 Evaluation

7.1 Summary Statistics

In our analysis we mainly consider the effect of task vs. visual encoding. We perform the analysis of the collected results in two stages. First, we consider the effect of visual encoding overall, as this is our primary research question. To check for normality, we run a Shapiro-Wilk test on each distribution, for normally distributed data we use a repeated measure analysis of variance (ANOVA), for data not always normally distributed we use a non-parametric Friedman's test both with standard statistical level $\alpha = 0.05$ to determine the statistical significance between conditions. Post-hoc analysis is conducted using Wilcoxon paired-samples test for all conditions that pass Friedman's test. In a second stage we then analyze differences in performances of visual channels within visual encoding and grid size, due to the different semantic of each encoding dimensions a channel-wise comparison across visual encoding is not feasible. To compensate for measuring multiple outcome from the same conditions we adjust the p-value, e.g. α value, applying a Bonferroni adjustment to the significance value. We choose a more conservative approach to provide stronger control of the condition-wise error rate. When separating the data by channel we apply a Bonferroni correction, reducing the significance level to $\alpha = 0.01$ obtained by subdividing the alpha value by the number of channels. When separating the data by grid size we apply a Bonferroni correction, reducing the significance level to $\alpha = 0.025$ obtained by subdividing the alpha value by the number of grid sizes. Post-hoc analysis is conducted as above. Unless explicitly reported no trade-off effects (i.e., less time leading to more errors) were detected. Due to space limitation we report mean correctness and response time results for the channel-wise analysis in the supplemental material. Since grid size dealt no significance effect, i.e. no cases where found in which either time or accuracy data produced significant results, we omit those results. Table 1 summarizes the significance values for

Table 1: Task 1, 2 & 3 results – Post-hoc p-values results for parameter estimation, outliers detection and pattern detection overall. Significant differences are highlighted in red.

Task 1: Parameters Estimation – Significance		
Pair-wise Test	Accuracy	Time (sec.)
Emoji vs. Chernoff	0.01	0.248
Emoji vs. Star	0.206	0.002
Chernoff vs. Star	0.003	0.001

Task 2: Outliers Detection – Significance		
Pair-wise Test	Accuracy	Time (sec.)
Emoji vs. Chernoff	0.027	0.084
Emoji vs. Star	0.184	$\ll 0.001$
Chernoff vs. Star	0.085	$\ll 0.001$

Task 3: Pattern Detection – Significance		
Pair-wise Test	Accuracy	Time (sec.)
Emoji vs. Chernoff	0.234	0.463
Emoji vs. Star	0.07	0.093
Chernoff vs. Star	0.008	0.055

Table 2: Task 4 results – Post-hoc p-values results parameter estimation overall with balanced Emoji (Emoji_B). Significant differences are highlighted in red.

Task 4 – Parameters Estimation – Significance [Balanced Emoji]		
Pair-wise Test	Accuracy	Time (sec.)
Emoji_B vs. Chernoff	0.01	0.001
Emoji_B vs. Star	0.13	0.002
Chernoff vs. Star	0.009	0.001

Table 3: Comparison by visual design of results from Task 1 & Task 4 – T1 refers to data from Task 1 and T4 refers to data from Task 4, we indicate with Emoji_B the balanced Emoji. Post-hoc p-values results for parameter estimation, outliers detection and pattern detection overall. Significant differences are highlighted in red.

Task 1 vs Task 4 – Parameters Estimation – Significance		
Pair-wise Test	Accuracy	Time (sec.)
Emoji vs. Emoji_B (T4)	0.7	$\ll 0.001$
Star(T1) vs. Star(T4)	0.83	$\ll 0.001$
Chernoff(T1) vs. Chernoff(T4)	0.11	$\ll 0.001$

post-hoc p-values for the first three tasks and three different visual encodings. Tables 2 summarizes the significance values for post-hoc p-values for Task 4 and the three different visual encodings with a balanced Emoji. Table 3 summarizes the significance values for post-hoc p-values in the pairwise comparison of the three different visual encodings in Task 1 and Task 4.

Tables in supplementary material summarize the significance values for post-hoc p-values with respect to each visual channel.

Task 1 – Parameters Estimation – Non Balanced Glyphs. Performances for Task 1, as a function of visual design, are summarized in Fig. 4 and Fig. 6. We found a significant difference overall of Emoji and Star versus Chernoff faces, with Emoji being significantly more accurate than Chernoff, and Star being both more accurate and faster than Chernoff faces.

Post-hoc analysis on channels revealed an effect of channel in Emoji and Chernoff. In Chernoff faces Eye Width was significantly more accurate than Brow Width and Nose, while Eye Slant significantly faster than Eye Width and Month. In Emoji Smile was significantly faster than Left Brow. Performances for Task 1, as a function of visual design and channel, are summarized in Fig. 9 and Fig. 10.

Task 2 – Outliers Detection. Performances for Task 2, as a function of visual design, are summarized in Fig. 4 and Fig. 6. We found a significant difference overall of Emoji versus Chernoff faces, with Emoji being significantly more accurate than Chernoff faces while Star being significantly faster than both Emoji and Chernoff faces.

Table 4: Feedback received from participants, each cell shows how many participants selected a level of difficulty for a specific task.

Participant Perceived Difficulty			
	Easy	Medium	Hard
Task 1 – Parameter Estimation	7	11	0
Task 2 – Outliers Detection	6	10	2
Task 3 – Pattern Detection	0	5	13

Table 5: Feedback received from participants, table shows how many selected a visual encoding as good (g), or bad (b), to be used in a specific task. Participants were allowed select multiple glyphs.

Participant Glyph Preference				
	Chernoff g/b	Emoji g/b	Star g/b	N/A
Task 1 – Parameter Estimation	1/6	8/1	7/5	1/0
Task 2 – Outliers Detection	1/6	8/1	8/2	1/1
Task 3 – Pattern Detection	1/4	9/0	4/4	1/1

Post-hoc analysis on channels revealed an effect of channel across all visual encodings. Chernoff faces Brow Width was significantly faster than all other Channels, Eye Slant and Mouth were significantly faster than Nose. Emoji channels featured also significant differences with Right Eye and Smile being faster than Left Eye, Left Brow, Right Brow.

Post-hoc analysis on grid size revealed no significant effect of grid size on performances and reconfirmed Star as significantly faster than Emoji and Chernoff faces across both grid sizes.

Task 3 – Pattern Detection. Performances for Task 3, as a function of visual design, are summarized in Fig. 4 and Fig. 6. We found a significant difference of Chernoff faces which proved to be significantly more accurate than Star.

Post-hoc analysis revealed no effect of visual layout between Emoji and Star. No effect was detected on response time across all three visual layouts.

Task 4 – Parameters Estimation – Balanced Glyphs. Performances for Task 4, as a function of visual design, are summarized in Fig. 5 and Fig. 7. We found a significant difference overall of Emoji and Star versus Chernoff faces, with Star being still both more accurate and faster than Chernoff faces. An increase in Emoji accuracy was measured with Emoji being now both more accurate and faster than Chernoff faces.

Post-hoc analysis on channels revealed an effect of channel on Chernoff faces with Eye Slant significantly faster than Channel Eye Width. Performances for Task 4, as a function of visual design and channel, are summarized in Fig. 9 and Fig. 10.

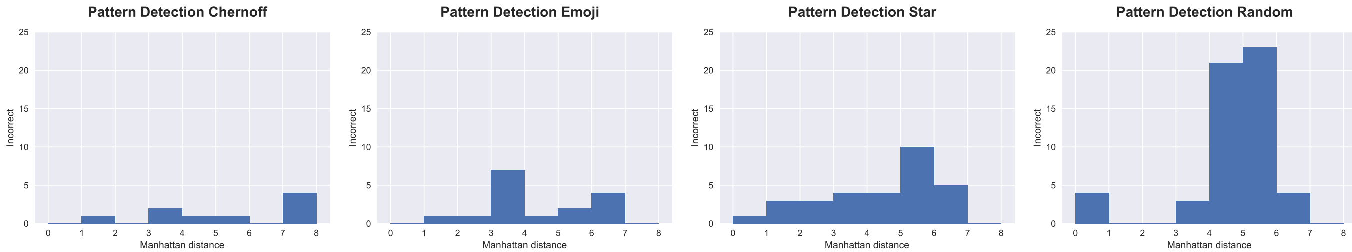

Figure 8: Manhattan distances of target versus answer for Task 3. Data shows the number of incorrect answers and distance from target.

Task 1 and 4 – Parameter Estimation – Balanced vs. Non-balanced Glyphs. Comparison of performances for Tasks 1 and 4, as a function of visual design, are summarized in Fig. 11 and Fig. 12. Comparison of glyphs across Task 1 and 4 revealed a significant decrease in response time across all visual encodings with balanced Emoji featuring the highest decrease. No trade-off effect was detected.

User survey. At the end of each task the participants were asked to rate the difficulty of the task and to choose the most and least suited glyph. Tables 4 and 5 summarize the results. Participants selected Emoji as the preferred Glyph for two of the tasks, and equal preference with Star for the other (Table 5).

7.2 Further Analysis

For Task 3, pattern detection, we performed a topology-based analysis by comparing distances between target and selected answer. In the task each sequence was composed by 3 glyphs of 5 dimension each (i.e. the 5 channels), a sequence can therefore be interpreted as a 15-dimensional vector. When the user selects the correct sequence from within the grid, the vector that represents the selected sequence and the vector that represents the target vector will be the same. However when a user selects an incorrect sequence the two vectors will be different. A quantitative measure of the difference can be obtained by calculating the Manhattan distance between the two vectors, where the Manhattan distance computes the distance between target and distractor and is measured as the sum of the absolute difference between respective channel coordinates. The Manhattan distance was chosen as distance metric as it shows higher reliability for high dimensional data where it is not straightforward to compare dimensions [1].

The smaller the Manhattan distance the more similar the chosen and target vector will be in appearance. The Manhattan distance thus becomes a measure of how good is the guess of the participant.

Results of the pattern detection task showed a total of 55 incorrect attempts over a total of 540. We calculated Manhattan distance for each incorrect answer and broke down the results by glyph type. We also created 55 totally random answers for the same data set and calculated the Manhattan distance between these random answers and the target sequence. Results are shown in Fig. 8, histogram of the Manhattan distance. Results for all the glyphs differ from the random answers featuring lower Manhattan distance values than the random ones. Chernoff glyph has fewest incorrect answers as reflected in the flatness of its histogram. Star glyph has a similar profile to the random histogram while Emoji is more skewed towards the lower Manhattan distance values.

8 Discussion

Collected results both met and confounded our expectations. In our second pilot study Star glyphs outperformed other visual layouts especially with respect to response time, we therefore expected a similar behavior. Nevertheless in both the parameter estimation and outlier detection tasks, Emoji glyphs outperformed Chernoff faces, while performed equally well to Star with respect to accuracy. Emoji glyphs carry a much higher cognitive load than Star and Chernoff faces due to the extra color channel. The wireframe nature of both Star and Chernoff faces reduces the load in terms of information processing. Emoji performances can be explained if we consider the boost provided by familiarity and frequency of use of the encoding. Despite our Emoji design following a minimalistic approach it still abides by the Unicode standard [14].

Emoji glyphs also held their ground against Star glyphs in the pattern detection task with respect to both accuracy and response time. In the same task Chernoff faces, still a face-like representation, outperformed Star in accuracy but not response time. Chernoff face's better performance with respect to accuracy might corroborate the hypothesis of an effect of semantic appealing of face like visual layouts in search task. In this respect it is interesting to compare results between Tasks 2 and 3, both visual search tasks at the core, but with different target lengths. Results would indicate that a "group" of faces within a crowd is easier to identify than a single individual. Both tasks required conjunction search (target and distractors differed by more than a single feature), psychology literature provides substantial evidence [16, 35] that conjunction search, in its early stages, favors a bottom-up process when looking for features within a stimulus. In later stages conjunction search moves towards a serial process of the pre-defined stimuli to identify the stimulus that best represents the targets.

Feature integration theory (FIT) [40, 42] suggests that features like luminance, color, orientation, and simple aspects of form are registered early and rapidly coded in parallel across the visual field using pre-attentive processes. FIT also suggests that to integrate two or more visual features belonging to the same object, a later process involving integration of information from different brain areas is needed and is coded serially using focal attention. For Emoji this would entail integration of both color and shape information to locate the target. For Chernoff faces this would entail integration of shape information to locate the target. Eriksen et al. [18] report how visual attention can either be focused on high resolution smaller areas or spread over wider areas but with loss of detail. Treiseman et al. [40] extends the analogy, suggesting that attention can either be narrowed to focus on a single feature, when there is a need to see what other features are present and how they contribute to form an object, or distributed over a whole group of items sharing the same relevant feature. Our results seem to suggest that, when dealing with a group of items, not only pre-attentiveness of features in face-like representations but also semantic meaning/empathic feedback and familiarity with the visual encoding support the cognitive process involved in the visual search.

Analysis of Manhattan distances in Task 3 revealed also how, despite the complexity of the task, response did not exhibit a random pattern and, when presented with face like representations, participants selected responses much closer to the correct answers even when incorrect responses were selected.

Effects of Glyph Balancing. The symmetric nature of Star glyphs makes it a balanced visual representation by design, this is supported by our results which do not show any effect of channels. More complex glyph representations however might introduce channel bias due to different dominance of features. Bias can in-

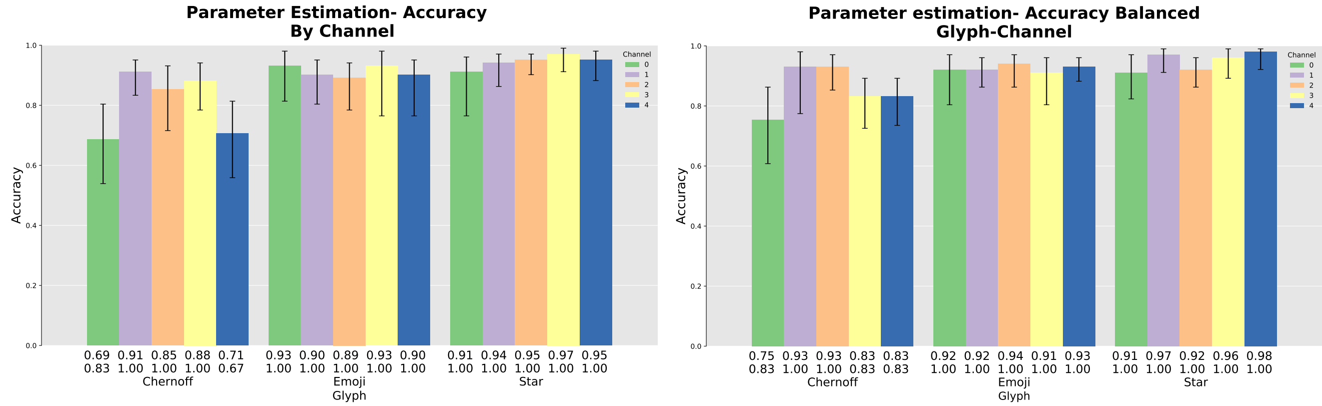

Figure 9: Analysis of accuracy results for Task 1 & 4 – Parameters Estimation (unbalanced vs. balanced). Results are broken down by channel for each visual encoding, (mean, median) values are indicated below each bar. Error bars show 95% confidence intervals.

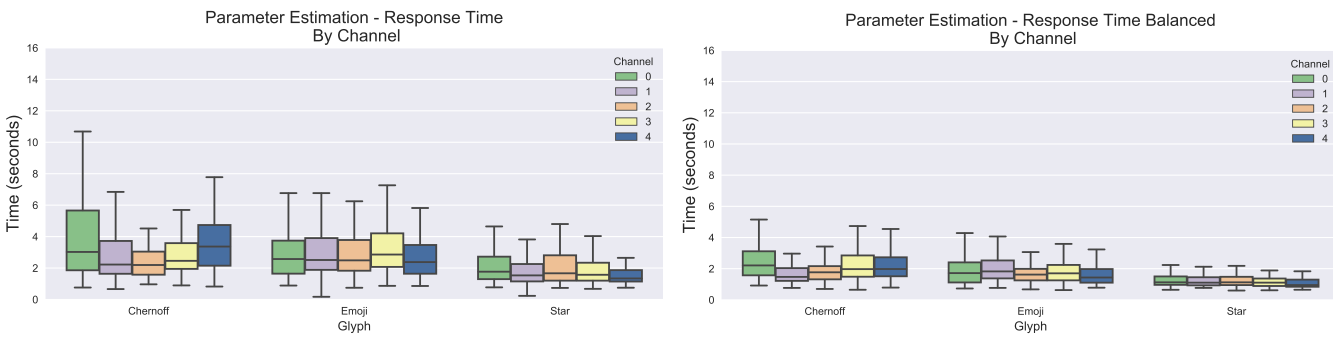

Figure 10: Analysis of response time results for Task 1 & 4 – Parameters Estimation (unbalanced vs. balanced). Results are broken down by channel for each visual encoding.

troduce distortion in the data representation and user performance may be influenced depending on the channel that was used for depiction [15, 31, 32]. Emoji and Chernoff faces both report effect of channels in accordance with Henderson *et al*. [25] findings with channel 4 (i.e. Smile) being significantly faster. Henderson study also highlights how focal attention varies when looking at faces, given the importance of focal attention in visual search tasks this can introduce a bias when mapping values to facial features. Results from Task 4 appear to support our initial hypothesis. Emoji is the only glyph undergoing balancing, and results show a significant increase in accuracy and decrease in response time, while effects of channels disappear. It is fair to report that an improvement in performances with respect to Task 1 was reported across all three visual encoding, as shown in Fig. 5 and Fig. 7. This could be due to Task 4 being a repetition of Task 1, participants becoming more familiar with the interface and therefore learning effects might start to surface. It is important to highlight that balanced Emoji stimuli were seen for the first time in Task 4, therefore learning effects, if present, should have had a more visible positive impact on Chernoff faces and Star, encodings unchanged through all four tasks. Balanced Emoji however feature the highest improvement overall, as shown in Table 2. Performance results also showed more uniform responses across channels for Emoji (as shown in Fig. 9 and Fig. 10), a result which we do not detect in Chernoff faces, where effect of channel is still significant (significance values reported in supplemental material). The controlled design of our study would imply that any learning effect would impact equally all visual encodings and, as such, should generate equal improvements, again no uniform improvement across all encodings can be detected from the results. In designing our stimuli we also ensured an equal distribution of face expressions to avoid any confounding effect introduced by people favoring "happy" versus "sad" faces as reported in literature [5, 26].

Aesthetic Influence. Tables 4 and 5 summarize the feedback received from participants at the end of each task. Task 3 was ranked as the most difficult followed by Task 2. Participants expressed clear preference for Emoji overall but selected also Star as the optimal visual encoding for Task 2. The latter partially reflects the study results which sees Star outperforming all visual encodings within Task 2 in response time. Preference for Emoji can be explained by both familiarity and aesthetic appeal. In post-study debriefing participants reported Emoji as their favorite encoding and how both color, *"finally color instead of boring black and white"*, and design, *"looking at happy faces"* made performing a task more enjoyable. It is worth noting that not all Emoji stimuli were "smiling" but they still appeared more friendly than their Chernoff counterparts, this shows how familiarity with the visual encoding has built an empathic link between participants and visualization.

9 Lessons Learned

The primary objective of our study has been to investigate the use of emoji to encode multidimensional data, as previously attempted by Chernoff [11]. A by-product of our investigation has been the need to control for effects of the different visual channels, more prominent when dealing with antropomorphic encodings. We have called the formalization of such attempt as Balancing. We have chosen to be conservative in our approach to maintain control over as many parameters as it was feasible, however dealing with human factors has highlighted interesting aspects of balancing as a problem with no absolute solution, but rather converging to what can be defined as local minima determined by the surrounding space shape. Balancing is inherently an optimization problem with a set

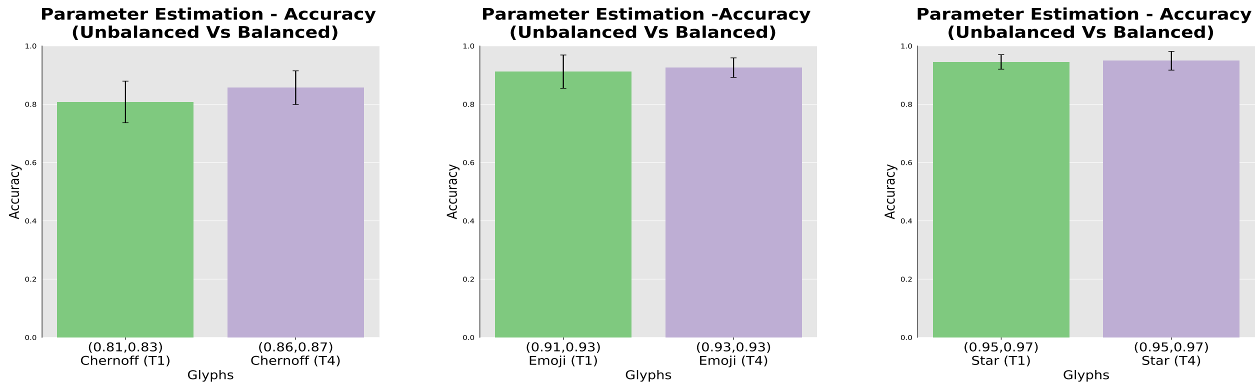

Figure 11: Pairwise analysis of performance results of Chernoff, Emoji and Star glyphs for Task 1 & 4.

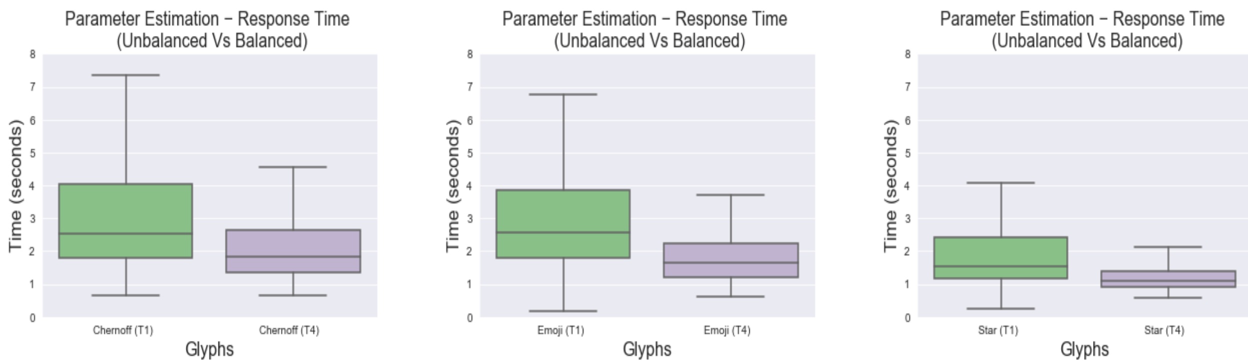

Figure 12: Pairwise Analysis of response time results of Chernoff, Emoji and Star glyphs for Task 1 & 4.

of feasible solutions contained within a convex polytope. The convex polytope is computed through the intersection of finitely many half spaces, each of which defined by a logarithmic inequality determined by an individual perceptual abilities. Our approach is an attempt at defining an objective function as an iterative algorithm searching for a point in the polytope where the function achieves the smaller value [28]. Due to the influence of elements such as familiarity and/or previous knowledge such function is reliant on the space definition and an optimal unique solution may not exist, but rather a set of user dependent local minima. Nevertheless, results from our study have shown that when such minima is found, significantly positive effects of balancing can be detected. We want to acknowledge that the experimental nature of our work, being a first of its kind for objective and context, and the control imposed on the experimental design, provide a starting point worth pushing forward for validation in more complex settings such as using real datasets ideally at different levels of complexity.

A further point of reflection relates to the antropomorphic nature of face like glyphs, which makes balancing not categorizable as a linear optimization problem. In face like representations channels interact with each other influencing perception in a non-linear way [8]. In our study we have restricted our analysis to 5 channels as in the original Chernoff face implementation [11], and also being one of the most common standard emoji representation. We have favoured 5 channels also as 5 items being the lower limit on our capacity to process information [36]. Based on the results it would be interesting to investigate effects of balancing for higher number of channels within and above the traditional 7 ± 2 suggested boundary [36]. Psychophysical research is a valuable asset to help determine perceptual magnitude of within channels influence, this would also allow to constrain the parameter space of the Balancing optimization problem increasing the velocity of reaching a feasible solution for the cost function.

Finally emoji are becoming a visual language phenomena, based on user post-study feedbacks aesthetic appeal and familiarity emerge as factors worth further investigation.

10 Conclusion

As initially suggested literature has demonstrated that Emoji are being appropriated for tasks other than creating emotional context or intonation in messages [30], we have taken these results a step further by utilizing Emoji for data visualization. Results from our study have shown how Emoji, as a visual encoding, can be as effective for visualising multiple attributes as other more common, and more minimalist, visual designs which in literature have instead outperformed Chernoff faces.

By purposely choosing a neutral dataset we have been able to reduce any confounding effect that might emerge from semantic links between data and its representations. However by doing so we have also restricted the full potential of the expressive power of glyph representation. Participants feedback have highlighted a strong empathic link with Emoji representations, likely due to their common usage. This adds a further level of implicit cognitive processing which could potentially be leveraged to support analytical tasks and complex data interpretation. A further level of analysis would therefore be to exploit semantic link between data and glyph visualization when such a link exists.

We have investigated, and highlighted, possible biases introduced when mapping data to features which might be inherently dominant. We have made a step towards the quantification of such bias and proposed a strategy to compensate for its effect. As future work it would be interesting to pursue this analysis further to inspect any variation with respect to task and difference in compensation strategies which could take the form of either reducing the size of the channel in its visual depiction (the choice we make here), or to transform the underlying data in a preparation stage before being mapped onto a glyph.

References

[1] C. C. Aggarwal, A. Hinneburg, and D. A. Keim. On the surprising behavior of distance metrics in high dimensional spaces. In *Proceedings of the 8th International Conference on Database Theory*, ICDT '01, pages 420–434. Springer-Verlag, 2001.

[2] A. Antonov. *Facing your data with Chernoff faces*, 2016. Available at `http://community.wolfram.com/groups/-/m/t/866748`, accessed on Jan. 10, 2019.

[3] T. Augustin and T. Roscher. Empirical evaluation of the near-miss-to-weber's law: a visual discrimination experiment. *Psychology Science Quarterly*, 50(4):469–488, Dec. 2008.

[4] F. Barbieri, F. Ronzano, and H. Saggion. What does this emoji mean? A vector space skip-gram model for twitter emojis. In *Language Resources and Evaluation conference*, Portoroz, Slovenia, 2016.

[5] C. Bartneck and J. Reichenbach. Subtle emotional expressions of synthetic characters. *International Journal of Human-Computer Studies*, 62(2):179–192, Feb. 2005.

[6] L. D. Bergman, B. E. Rogowitz, and L. A. Treinish. A rule-based tool for assisting colormap selection. In *Proceedings of the 6th Conference on Visualization '95*, VIS '95, pages 118–125, Washington, DC, USA, 1995. IEEE Computer Society.

[7] R. Borgo, J. Kehrer, D. H. S. Chung, E. Maguire, R. S. Laramee, H. Hauser, M. Ward, and M. Chen. Glyph-based Visualization: Foundations, Design Guidelines, Techniques and Applications. In M. Sbert and L. Szirmay-Kalos, editors, *Eurographics 2013 - State of the Art Reports*. The Eurographics Association, 2013.

[8] M. G. Calvo, A. Fernández-Martiín, and L. Nummenmaa. A smile biases the recognition of eye expressions: Configural projection from a salient mouth. *The Quarterly Journal of Experimental Psychology*, 66(6):1159–1181, Jun. 2013.

[9] S. Cappallo, T. Mensink, and C. G. Snoek. Image2emoji: Zero-shot emoji prediction for visual media. In *Proceedings of the 23rd ACM International Conference on Multimedia*, MM '15, pages 1311–1314, New York, NY, USA, 2015. ACM.

[10] J. M. Chambers, W. S. Cleveland, B. Kleiner, and P. A. Tukey. *Graphical Methods for Data Analysis*. Wadsworth, 1983.

[11] H. Chernoff. The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association*, 68(342):361–368, Jun. 1973.

[12] D. H. Chung, P. A. Legg, M. L. Parry, R. Bown, I. W. Griffiths, R. S. Laramee, and M. Chen. Glyph sorting: Interactive visualization for multi-dimensional data. *Information Visualization*, 14(1):76–90, Nov. 2013.

[13] D. H. S. Chung, D. Archambault, R. Borgo, D. J. Edwards, R. S. Laramee, and M. Chen. How ordered is it? on the perceptual orderability of visual channels. *Computer Graphics Forum*, 35(3):131–140, Jul. 2016.

[14] M. Davis. *Unicode®Emoji*. Available at `http://unicode.org/emoji/`, accessed on Jan. 10, 2019.

[15] B. Duffy, J. Thiyagalingam, S. Walton, D. J. Smith, A. Trefethen, J. C. Kirkman-Brown, E. A. Gaffney, and M. Chen. Glyph-based video visualization for semen analysis. *IEEE Transactions on Visualization and Computer Graphics*, 21(8):980–993, Aug. 2015.

[16] M. Eimer and A. Grubert. The gradual emergence of spatially selective target processing in visual search : From feature-specific to object-based attentional control. *Journal of Experimental Psychology : Human Perception and Performance.*, 40(5):1819–1831, Oct. 2014.

[17] Emojipedia. *Google®Emoji*, 2017 (accessed March, 2017).

[18] C. W. Eriksen and J. E. Hoffman. Temporal and spatial characteristics of selective encoding from visual displays. *Perception & Psychophysics*, 12(2):201–204, Mar. 1972.

[19] D. H. Foster. The cue for contour-curvature discrimination. *Vision research*, 33(3):329–341, Feb. 1993.

[20] D. H. Foster and J. Wagemans. Viewpoint invariant weber fractions and standard contour-curvature discrimination. *Biological Cybernetics*, 70(1):29–36, Nov. 1993.

[21] J. Fuchs, P. Isenberg, A. Bezerianos, F. Fischer, and E. Bertini. The Influence of Contour on Similarity Perception of Star Glyphs. *IEEE Transactions on Visualization and Computer Graphics*, 20:2251 – 2260, Dec. 2014.

[22] J. Fuchs, P. Isenberg, A. Bezerianos, and D. Keim. A systematic review of experimental studies on data glyphs. *IEEE Transactions on Visualization and Computer Graphics*, 23(7):1863 – 1879, Jul. 2017.

[23] C. G. Healey, K. S. Booth, and J. T. Enns. High-speed visual estimation using preattentive processing. *ACM Transanction on Computer-Human Interaction*, 3(2):107–135, Jun. 1996.

[24] J. M. Henderson, R. Falk, S. Minut, F. C. Dyer, and S. Mahadevan. Gaze control for face learning and recognition in humans and machines. In *From fragments to objects: Segmentation processes in vision*, pages 463–481. Elsevier, 2001.

[25] J. M. Henderson, C. C. Williams, and R. J. Falk. Eye movements are functional during face learning. *Memory & Cognition*, 33(1):98–106, Jan. 2005.

[26] U. Hess, S. Blairy, and R. E. Kleck. The intensity of emotional facial expressions and decoding accuracy. *Journal of Nonverbal Behavior*, 21(4):241–257, Dec. 1997.

[27] M.-K. Hu. Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory*, 8(2):179–187, Feb. 1962.

[28] N. I. M. Gould and S. Leyffer. *An Introduction to Algorithms for Nonlinear Optimization*, pages 109–197. Springer, Berlin, Heidelberg, 2003.

[29] E. Kandogan. Star coordinates: A multi-dimensional visualization technique with uniform treatment of dimensions. In *In Proceedings of the IEEE Information Visualization Symposium, Late Breaking Hot Topics*, pages 9–12, 2000.

[30] R. Kelly and L. Watts. Characterising the inventive appropriation of emoji as relationally meaningful in mediated close personal relationships. In *Experiences of Technology Appropriation: Unanticipated Users, Usage, Circumstances, and Design Workshop*, Sept. 2015.

[31] P. A. Legg, D. H. S. Chung, M. L. Parry, M. W. Jones, R. Long, I. W. Griffiths, and M. Chen. Matchpad: Interactive glyph-based visualization for real-time sports performance analysis. *Computer Graphics Forum*, 31(3pt4):1255–1264, 2012.

[32] P. A. Legg, E. Maguire, S. Walton, and M. Chen. Glyph visualization: A fail-safe design scheme based on quasi-hamming distances. *IEEE Computer Graphics and Applications*, 37(2):31–41, Mar. 2017.

[33] A. E. Lie, J. Kehrer, and H. Hauser. Critical design and realization aspects of glyph-based 3d data visualization. In *Proceedings of the Spring Conference on Computer Graphics (SCCG 2009)*, pages 27–34, 2009.

[34] E. Maguire, P. Rocca-Serra, S. A. Sansone, J. Davies, and M. Chen. Taxonomy-based glyph design – with a case study on visualizing workflows of biological experiments. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2603–2612, Dec. 2012.

[35] B. McElree and M. Carrasco. The temporal dynamics of visual search: Evidence for parallel processing in feature and conjunction searches. *Journal of Experimental Psychology: Human Perception and Performance*, 25(6):1517–1537, Dec. 1999.

[36] G. A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2):81–97, Apr. 1956.

[37] I. Moschini. The face with tears of joy emoji. a socio-semiotic and multimodal insight into a japan-america mash-up. *HERMES - Journal of Language and Communication in Business*, 55:11–25, Mar. 2016.

[38] T. Ropinski and B. Preim. Taxonomy and usage guidelines for glyph-based medical visualization. In *Proceedings of the 19th Conference on Simulation and Visualization (SimVis08)*, pages 121–138, 2008.

[39] N. Sivagnanasundaram, A. Chaparro, and E. Palmer. Evaluation of the presence of a face search advantage in chernoff faces. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 57, pages 1610–1614. SAGE Publications, 2013.

[40] A. M. Treisman and G. Gelade. A feature-integration theory of attention. *Cognitive Psychology*, 12(1):97 – 136, Jan. 1980.

[41] M. O. Ward. A taxonomy of glyph placement strategies for multidimensional data visualization. *Information Visualization*, 1(3/4):194–210, Dec 2002.

[42] J. M. Wolfe. What can 1 million trials tell us about visual search? *Psychological Science*, 9(1):33–39, Jan. 1998.