Exploring perceptual similarity and its relation to image-based spaces: an effect of familiarity

Rosyl S. Somai, Peter J.B. Hancock

Psychology, Faculty of Natural Sciences, University of Stirling, Stirling, FK9 4LA

Correspondence to Rosyl Somai, rss1@stir.ac.uk

Abstract

One challenge in exploring the internal representation of faces is the lack of controlled stimuli transformations. Researchers are often limited to verbalizable transformations in the creation of a dataset. An alternative approach to verbalization for interpretability is finding image-based measures that allow us to quantify image transformations. In this study, we explore whether PCA could be used to create controlled transformations to a face by testing the effect of these transformations on human perceptual similarity and on computational differences in Gabor, Pixel and DNN spaces. We found that perceptual similarity and the three image-based spaces are linearly related, almost perfectly in the case of the DNN, with a correlation of 0.94. This provides a controlled way to alter the appearance of a face. In experiment 2, the effect of familiarity on the perception of multidimensional transformations was explored. Our findings show that there is a positive relationship between the number of components transformed and both the perceptual similarity and the same three image-based spaces used in experiment 1. Furthermore, we found that familiar faces are rated more similar overall than unfamiliar faces. That is, a change to a familiar face is perceived as making less difference than the exact same change to an unfamiliar face. The ability to quantify, and thus control, these transformations is a powerful tool in exploring the factors that mediate a change in perceived identity.

Keywords

Perceptual space; similarity; image-based spaces; PCA; familiarity

Background

Although there are many debates in face recognition, there tends to be agreement on the existence of some organizational space in which faces are represented, dubbed face space (Valentine, 1991). The “multidimensional space framework” proposed by Valentine (1991) has brought many insights in face recognition, but, as clearly stated by the author, it never included any definitions of the dimensions or even attempted this. As Valentine phrased it, “These aspects have not been specified simply because there is as yet insufficient evidence available to enable informed choices to be made.” Nearly three decades later, we find ourselves in a surprisingly similar situation, with little progression in understanding the important dimensions of face variability.

Most research that explores facial variation, uses the approach of reverse engineering: manipulating the visual input to assess whether this content is used in face recognition. For example, changing facial components such as the nose, mouth and eyes (Abudarham & Yovel, 2016) or displacing them (Goffaux & Rossion, 2007). This approach is useful but needs well understood transformations of the visual input to have a meaningful presentation. For example, a change in eyebrow thickness is easily verbalized and can be varied systematically, but a change based on morphing two faces together is difficult to interpret in terms of dimensional change. Verbalizable transformations are insightful but are only a part of the complete story. The variation between and within faces is what makes up our internal representations of faces (Burton, Kramer, & Ritchie, 2016) and it is clear that this variation goes beyond easy to verbalize features. An alternative approach to verbalization for interpretability is finding image-based measures that allow us to quantify image transformations. Such an approach would need image-based measures that preferably correlate with perceptual space, and well-controlled transformations.

Principal Component Analysis (PCA) was first proposed as a means of face recognition by Kirby and Sirovich (1990), and Turk and Pentland (1991). PCA produces a compact representation of a face but is very sensitive to image properties such as lighting.

Nevertheless, using controlled images, PCA will reveal systematic differences between faces, such as gender (O’Toole et al., 1998). The analysis is greatly improved by first morphing all faces to the same average shape and performing PCA on the shape vectors and the ‘shape free’ faces separately (Craw & Cameron, 1991). This separate analysis yields significant correlations with human perceptions of memorability and distinctiveness of faces (Hancock, Burton, & Bruce, 1996). It can also be used to synthesise novel face images (Hancock, 2000) and by extension, provides a controlled way to produce changes in the shape or surface

appearance of face images. It is this ability to produce controlled transformations that is exploited here.

One aim of this work is to test the utility of PCA-induced changes in face images to probe human face perception. PCA can produce graded changes and it is to be expected that a large PCA change will produce a bigger perceptual difference than a smaller one. However, it would be helpful to have a separate computational measure of the likely perceptual effect of varying different components or combinations of them. If a purely image-based measure (without a biological premise) correlates with perceptual ratings, it can be used in the creation of stimuli to make a prediction of how a human observer will perceive the image. For example, one could predict whether a certain image transformation is strong enough for the altered face and the original face to be perceived as different from one another. Here, we test three different computational measures of image similarity.

The simplest measure of image similarity is pixel-wise Euclidean distance. For this to have any chance of working for face images, they need to be very carefully controlled. A simple misalignment between two copies of the same image will produce a large difference, while changes in lighting will also result in a large difference that may be barely noticed by a human observer. Because of this, the measure is little tested in face perception. However, since our images are generated by a PCA model, they are inherently well-controlled. Meytlis (2011) used carefully controlled images and found an almost perfect correlation between distinctiveness ratings and Pixel Euclidean distance.

A more biologically plausible measure commonly used in computer vision (Lades et al., 1993; Wiskott, Fellous, Krüger, & Von der Malsburg, 1997) and neurocomputational models of face recognition (Dailey & Cottrell, 1999; O'Reilly & Munakata, 2000) is Gabor filtering. Gabor filters are found to be a good representation of how cells in the primary cortex respond to a certain scale and orientation (De Valois & De Valois, 1990; Jones & Palmer, 1987; Ringach, 2002). As Gabor filtering reflects very early processing and not any higher order processing, it would seem unlikely that Gabor distances would relate to perceptual similarity of faces. However, this is exactly what Yue, Biederman, Mangini, Malsburg and Amir (2012) found: Gabor-based image measures strongly correlate with human performance in a discrimination task with complex figures. For this reason, we included Gabor space to quantify image transformations.

Current artificial face recognition deploys deep neural networks (DNNs) and their face matching ability easily surpasses human performance in many situations. Hancock, Somai and Mileva (2020) compared six state-of-the-art DNNs with human performance on four face

matching tasks designed to be challenging for human participants. The humans averaged 73.3%, while the best DNN achieved an accuracy of 98.6%. Despite the difference in accuracy, there were significant correlations between the six DNN similarity scores for a pair of faces and the human accuracy on that pair. This was especially true for matching pairs, where the two images do show the same person. Humans and the DNNs seem to agree on which faces look similar. However, the correlation for mismatch face pairs was lower, indicating that humans and the DNNs show less agreement on what makes faces look different. Indeed, the DNNs worked in distinctly non-human like ways, often reporting that faces which differed in apparent race or sex were the same identity. Given these discrepancies, it is valuable to include DNN similarity measures in the present analysis of controlled image changes.

Hancock, Bruce and Burton (1998) looked at correlations between human similarity ratings and PCA on existing faces images. Here, we invert that process and ask humans to rate the similarity of face images that have been altered in a controlled way by varying principal components. We compare the human data with similarity scores obtained from Gabor, Pixel and DNN analysis of the same images. In experiment 1, we test the effects of single component transformations of both colour and shape, using unfamiliar faces. In Experiment 2, we test the effect of multiple component transformations on both familiar and unfamiliar faces. This allows us to test the effect of familiarity on perceptual similarity.

Experiment 1 - exploring the effects of single component transformations

Background

One major issue in exploring facial variation is finding an image transformation that reflects a likely dimension in face space which can also be interpreted in some meaningful way. PCA transformations can be difficult to interpret, as they reflect variation in pixel values that commonly occur together. For example, a PCA component that changes eye colour will also change hair colour, because people with darker eyes often have darker hair than people with lighter eye colours. Thus, verbalizing the change of PCA components can be difficult. On few occasions, a component can be classified after visual review, such as a component that varies masculinity. However, for most components the variation is difficult to classify. In this experiment, we will explore the effect of single component transformations on Gabor, Pixel and DNN space, and (human) perceptual space in order to assess their usefulness in quantifying image transformations.

Methods

Participants

We tested 25 participants (11 female, 13 male, 1 other, mean age = 39.7 years, 2 missing values, SD = 16.3 age range = 12 - 74 years. Participants were naive to the purpose of the experiment and took part on a completely voluntarily basis without any monetary compensation. The experiment was approved by the local General University Ethics Panel of the University of Stirling.

Apparatus and stimuli

The experiment was programmed through the online platform Qualtrics and participation was allowed via PC, laptop, phones and tablets. Face image stimuli were generated from a PCA-based face model, written in Matlab (MATLAB, 2020). The PCA space was calculated from 72 eye-aligned images, each an average of 10 images of a famous (mostly Caucasian) female. Use of average images helps to reduce variations due to lighting and viewpoint. Even so, the first few components encode image properties such as lighting and head movements, so we selected components 10 – 16 to alter. Three different types of transformations were used: shape-only, colour-only and both. To explore the effect of the size of the transformation, two transformation sizes were used: 3 and 6 SD. The SD for each component was scaled according to the variance in the (72- eye-aligned) original images. This gave 42 pairs in total: 7 shape, 7 colour and 7 both, two transformation sizes each. For each pair a different face image was generated by setting all the shape and colour components to a random normal value, mean zero, sd 1, though truncated at ± 2 sd to prevent faces looking too strange. An example of the transformation is shown in Figure 1. In addition to the PCA transformations, three single feature transformations derived from Independent Component Analysis (ICA) were added. However, these were purely exploratory and will not be analysed or further discussed. The images produced were all 420 by 595 pixels, size in the experiment was browser and screen dependent.

Figure 1 Example of the transformation used. The top row shows the left end of the scale which is -3 SD from the original and the bottom row is the right end of the scale which is 3 SD in this example. The first column shows shape transformations only, the middle column shows colour transformations only and the last column shows shape and colour transformations.

Procedure

Participants were shown an information sheet and then gave consent online. The face matching task consisted of 45 trials (14 alterations for three conditions and the 3 exploratory ICA transformations) that were presented in a random order. In each trial, two faces were presented along with the question “Could these be two pictures of the same person? 1 – certain no; 2 – think no; 3 – guess no; 4 – guess yes; 5 – think yes; 6 – certain yes”. An example of a trial is shown in **Error! Reference source not found..**

Figure 2 Example of a trial

The three computational spaces

Gabor

We used the Gabor Features in Signal and Image Processing Toolbox (Kämäräinen, Kyrki, & Kälviäinen, 2002b, 2002a), written in Matlab (2019). This requires a 256 square input image, so the original images were cropped to just below the chin, resized and converted to monochrome, Figure 3. We used the default settings of the toolbox, which computes phase

and magnitude for 8 orientations and 5 scales across a 10x10 grid, giving a vector of 8000 values for each image. The difference between two images is given by the Euclidean distance between the two vectors.

Pixel

A simple way to compare two images is Euclidean distance in pixel space. Since our images are all aligned by the eyes, pixel distance might plausibly capture the perceptual changes caused by variations purely in the colour space. However, shape transformations will result in different face features being compared, or face pixels in one image compared with background in another. Meytlis (2011) normalized each image to the same average grey level and cropped them to a tight oval. Since we are comparing pairs of images that differ by adding a principal component that is inherently zero mean we did not normalize the colour values, but did crop tightly around the face with an oval 378 x 277 pixels, see Figure 3. The difference between images was computed by concatenating the RGB values and calculating the Euclidean distance between them.

Figure 3 Example of cropped images, differing in shape component 10. Left, for Pixel comparison, right for Gabor.

DNN

The DNN used was an experimental system from the University of Surrey. It produces an output vector of length 512. The difference between two images was given by the cosine of the angle between the two output vectors.

Results

Our human participants were not asked to rate the similarity of a pair of faces as such but to make a decision about identity: do the two images show the same person? On a 6-point scale, 3 was ‘guess different’ and 4 was ‘guess same’. The boundary for same or different identity

was therefore at 3.5. We report the agreement between the average human classification and that offered by our 3 computational spaces. The deep neural network has a specified decision boundary at a similarity score of 0.4. For the two image-based approaches, we computed the minimum distance between all 72 of the images used to generate the PCA space. This provides a conservative estimate of ‘different’: any pair less similar than this threshold should be different people but it is possible that images of different identities might be more similar than this threshold. It should be obvious that a threshold could be set that more accurately describes the data based on the findings. However, this is exactly why we wanted to find a rule of thumb for setting a threshold: creating a suitable stimulus set can be a laborious task that requires repeated piloting and it is specific to the set of faces the transformations were based on, the exact images, the parameters of the transformations, etc. For this reason, we explore the usefulness of a generic threshold, the minimum distance in a linear image-based space between all 72 of the images used to generate the PCA space.

Gabor

We found a significant negative correlation between the human similarity ratings and the Gabor distances ($r = -.87$, $p = <.001$), shown in Figure 4. The smallest distance between different identities in the Gabor space was 185. Using this threshold, shown as a vertical dotted line in Figure 4, results in 8 points that are misclassified, relative to the human decision boundary at 3.5. As expected, the threshold of 185 is somewhat conservative: a lower threshold of around 160 would produce a very similar classification to the human data, differing only in 3 points. Interestingly, 5 out of 8 of the ‘incorrectly classified’ items are colour transformations. Despite the effect of a change in colour components not crossing the Gabor threshold, the effect of the change is apparently strong enough to decrease perceptual similarity to a point of a ‘dissimilar’ classification. This shows that Gabor distance can be used to quantify PCA transformations and can even be used to predict whether or not a participant is likely to consider two images to depict the same person.

Figure 4 Correlation between perceptual similarity and Gabor distance. The 6 SD transformations are indicated by a bold outline.

Pixel

We found a significant negative correlation between the human similarity ratings and the Pixel distances ($r = -.85$, $p = <.001$). To avoid non-face image information inflating pixel values, Meytlis (2011) cropped the images to an oval size. In this experiment, a comparable crop has been applied (described above). Using the minimum Pixel distance between the 72 cropped images that make up our PCA space, the threshold is 7350 Pixel distance. The Pixel threshold performed slightly worse than the Gabor threshold, by only correctly identifying 30 out of 42 points (see Figure 5). Again, the colour transformations are the majority of the misclassified points, but only by 1 more error than the transformation of both (colour and shape transformed). Similar to the Gabor threshold and its conservative prediction of dissimilar, the Pixel threshold would benefit from a lower value. For example, lowering the threshold to 4500 would leave only 4 misclassifications.

Figure 5 Correlation between perceptual similarity and Pixel distance. The 6 SD transformations are indicated by a bold outline.

DNN

We found a significant correlation between the human similarity ratings and the DNN distances ($r = .94$, $p = <.001$). This specific DNN uses a value of 0.4 or higher to classify two faces as similar. This threshold provides a slightly less useful guide for predicting perceptual similarity compared to the Gabor threshold and performs the same as the Pixel threshold, as it only correctly identifies 30 out of 42 data points. However, it is important to note that this threshold was not based on our stimulus set but is the threshold used by the DNN. The value of the smallest distance between identities in our PCA set (of the average images of 72 actresses) was 0.18, which is clearly not a useful guideline, as it would not consider anything a match using that threshold.

Figure 6 Correlation between perceptual similarity and DNN distance. The 6 SD transformations are indicated by a bold outline.

Discussion experiment 1

The linear relationship between the DNN similarity space and human similarity is remarkable. The human response scale is, strictly speaking, ordinal: there is no reason to think that the differences between ‘guess’, ‘think’ and ‘certain’ should be the same number, and the same as the difference between ‘guess same’ and ‘guess different’. The DNN similarity space is given by the cosine of the angle between the feature vectors. Cosine is clearly non-linear, although most of the non-linearity occurs in the region above 0.8 and most of our points are below that, so within the relatively linear region. Yet these two non-linear spaces, averaged over a number of observations, contrive to give an almost perfect straight-line relationship.

Similar to the findings of Meytlis (2011) on Euclidean distance, our results show that all three image-based distance values correlate significantly with perceptual similarity. Moreover, we find that a threshold based on the PCA set can be used as a rough prediction of the likelihood that a pair will be considered the same or different identity. Taken together, the strong correlations between the image-based spaces and the informative nature of a threshold based on the stimulus set, provide a controlled way to alter the appearance of a face. Although the results are straightforward, it remains unclear if these linear relationships are also present when multidimensional transformations are used. In experiment 2, multidimensional transformations are explored by using transformations with a varying number of principal components changed.

Experiment 2: the relationship between the number of principal components transformed, perceptual similarity and familiarity

Background

This experiment is similar to experiment 1, but with multidimensional instead of single-dimensional transformations. Specifically, the effect of the number of PCs changed on perceptual similarity is explored, using four different transformations that increase in image-based dissimilarity and their effect in perceptual, Gabor, Pixel and DNN space. Based on the findings of Experiment 1, we hypothesize that the relationship between the number of PCs changed and any of the four image spaces is linear. Since it is well known within face research that people are better at processing familiar than unfamiliar faces (e.g. Johnston & Edmonds, 2009) and representations of familiar faces are more robust (Burton, Jenkins, Hancock, & White, 2005), we included familiarity to explore the effects of familiarity on perceptual similarity. Therefore, transformations in this experiment were performed on famous faces, as opposed to the unfamiliar, randomly generated faces in Experiment 1. A familiarity check was added in the beginning of the experiment.

Although there is an abundance of research on the effect of familiarity, it remains unclear to what underlies this enhanced performance to recognize familiar over unfamiliar faces. A common proposal is that the representation and processing of familiar faces is different than that of unfamiliar faces (Johnston & Edmonds, 2009; Ramon & Gobbini, 2018). Recent work from Ritchie & Burton (2017) shows that variability is important for learning a new face. Additionally, the variability within images of a specific identity appears to be idiosyncratic information that is key for learning new faces (Burton et al., 2016).

That two images of a person look more similar when familiar is almost the definition of familiarity. Someone familiar with the person can identify them both as being the same individual, while someone unfamiliar has to rely on what they can infer from the appearance (Jenkins, White, Van Montfort, & Burton, 2011). Here, we are applying fixed changes to faces, some of which will be familiar to a given participant and others not. It is plausible that in this situation, viewers familiar with a face will be better able to identify a change: e.g. ‘that’s X, but the nose is wrong’. They may be able to use their existing memory of the face to identify the changes.

The results from Beale and Keil (1995) make a more subtle prediction. They used morph sequences between different individuals’ faces to test the effect of familiarity on the decision boundary. They found that familiar face pairs had a sharper decision boundary.

Moving along the morph continuum, the face looked like the first identity until nearly half way, when it rapidly shifted to looking like the other identity. For unfamiliar face pairs, the shift was more linear, with a gradual change from one identity to the other. The prediction from this for our imposed changes in appearance is that, for a familiar face, the changed images should look like the original up to some critical point, where it will rather suddenly no longer look like them. If unfamiliar with the original face, then the changes should produce a more linear change in similarity. Therefore, the prediction is that changed faces will initially look more similar when familiar than unfamiliar but that, if the change becomes large enough, this will reverse, with changed familiar faces looking more different.

In these experiments, we are asking participants to decide about probable identity, rather than about similarity directly. So, while those familiar with the face may be able to say ‘the nose is wrong’, we expect participants to say ‘same identity’ with more confidence if they are familiar with the face. Whether we see the reversal predicted by the results of Beal and Keil will depend on whether the changes are sufficiently large.

Methods

Participants

We tested 61 participants (43 females, 16 males, 1 other, 1 missing value, mean age = 29.9 years, SD = 9 age range = 20 - 68 years. Testing was done through an online platform, Testable. Participants were naive to the purpose of the experiment. All participants signed an informed consent form after a written explanation of the procedure. Participation was on a completely voluntarily basis without any monetary compensation. The experiment was approved by the local General University Ethics Panel of the University of Stirling.

Apparatus and stimuli

The experiment was programmed through the online platform Testable and participation was only allowed via a PC or laptop. A calibration to control stimulus size was performed in form of matching the size of an on-screen bar to a credit card, to ensure proper visualisation throughout the experiment. Stimuli preparation and creation were done using Matlab (2019), using the same PCA space as described in experiment 1. Eight of the 72 average images that went into the PCA were chosen for use in this experiment. The stimuli consisted of these 8 average images and variations that had either 2, 3, 5 or 7 colour components (components 10 and 11; 10 – 12; 10 – 14; 10 – 16 respectively) in the calculated PCA space changed. The size of the change was 2 SD. To avoid any directional effects from the PCA transformation, for

half the trials the negative transformation and the other half the positive transformation was used. As shown in experiment 1, any of the three spaces can be used for insight into the size of the perceptual change. We applied this principle in this experiment to ensure that the effect of the transformation was similar for all identities. Figure 7 shows the Gabor distance between each original image and its four variations. It can be seen that the pattern of differences is very similar for each of the identities, suggesting that we should not expect any marked differences between items in our set. An example of the transformation can be found in Figure 8.

Figure 7 Gabor distance between the original and altered image to check if the transformations had a similar effect in Gabor space.

Figure 8 An example of the transformation: the left image is -2 SD, the middle image is the original and the right image is $+2$ SD.

Procedure

The experiment consisted out of a familiarity check and a face matching task. In the first 8 trials, participants were asked “Would you recognise this person on the street? yes(y)/ no(n)” for each identity (presented in a random order). Afterwards, participants were presented with the face matching task instructions and had 2 initial practice trial with an identity not included in the actual experiment. This was followed by 32 trials in a random order. In each trial, the original was shown, followed by a mask, the altered image and another mask. Finally, the participants were presented with the question “Could these be two pictures of the same person? 1 – certain no; 2 – think no; 3 – guess no; 4 – guess yes; 5 – think yes; 6 – certain yes”. An overview of a trial is shown in Figure 9.

Figure 9 Overview of a trial.

Data analysis

To explore the effects of familiarity, we constructed a linear mixed effects model with the similarity rating as dependent variable. The model included Familiarity and Transformation level as categorical factor (Familiarity: familiar and unfamiliar rescaled to -0.5 and 0.5 respectively to allow for the analysis of the interaction; Transformation Level: 2, 3, 5, or 7), and the intercept, slope and interaction of familiarity and Transformation level in the performance per participant as a random effect. Due to convergence issues, the main effects were excluded from the random effects. We report the parameter estimate (b), standard error (SE), t value and p value. The linear mixed effect analysis was performed in R (R Core Team, 2019) with *lmerTest* (Kuznetsova, Brockhoff, & Christensen, 2017). The threshold for

significance was set at $\alpha = 0.05$. Normality of the residuals on the model were inspected visually and showed no violation.

Results and discussion

In the current study, we explored the effect of multidimensional transformations on the three image-based spaces and perceptual similarity. Similar to experiment 1, we explore how informative the use of a threshold based on the stimulus set is to predict whether or not the two images are perceived as the same identity. Additionally, we investigated the effect of familiarity on similarity ratings.

Gabor

We found a significant negative correlation between the human similarity ratings and the Gabor distances ($r = -0.77$, $p < 0.001$).

The same threshold as Experiment 1 was used, which was a Gabor value of 185.33 out of 42 points were correctly classified (see Figure 10).

Especially when five components were changed, the estimation was incorrect: 5 out of the 9 incorrect estimations

using the threshold were made in this condition. However, in this experiment only colour components were changed, which were the least accurately predicted in Experiment 1 and this could also explain the misclassifications. In other words, it could be that it is not a certain size of change that is poorly predicted but it might be a poor fit for predicting colour transformations. Similar to Experiment 1, lowering the threshold (to about 167) would increase predictive value.

Figure 10 Correlation between perceptual similarity and Gabor distance. The trendline is fitted to the data of all change size groups.

Pixel

We found a significant negative correlation between the human similarity ratings and the Pixel distances ($r = -0.76$, $p = <.001$). The same threshold as in Experiment 1 was used, which was 7350 Pixel distance. Similar to the Gabor threshold, this provides a decent prediction whether or not a face pair will be perceived as similar or not (misclassified points: 33 out of 42; see Figure 11). This accuracy

is similar to that in Experiment 1 (30 out of 42). Again, the estimation using the threshold was poorest in the condition where 5 principal components were changed, whereas the other conditions are quite well estimated. In contrast to the Gabor space, it does not appear that one Pixel threshold can be set to increase accuracy of prediction for the different types of transformations (single- and multidimensional).

Figure 11 Correlation between perceptual similarity and Pixel distance. The trendline is fitted to the data of all change size groups.

DNN

We found a significant correlation between the human similarity ratings and DNN distances ($r = .79$, $p = <.001$). As mentioned before, this specific DNN uses a value of 0.4 or higher to classify two faces as similar and a threshold based on our image set was uninformative. Contrary to the first experiment in which the DNN threshold was less useful compared to the other two spaces, this threshold provides an equally useful guide for predicting perceptual similarity by correctly identifying 33 out of 42 data points (see Figure 12).

Figure 12 Correlation between DNN distance and perceptual similarity. The trendline describes all datapoints.

The effect of size of change and familiarity on perceptual similarity

The descriptive statistics of experiment 2 are shown in **Error! Reference source not found.** Table 1 in the supplementary materials shows the distribution of familiarity of the 8 identities. The results show that increasing the number of components transformed decreased the similarity rating for all levels compared to the base change (2 components changed, all p-values are $p < 0.001$, see Table 2 - Table 4 in supplementary materials). Interestingly, the similarity rating between 5 and 7

Figure 13 Average similarity rating per level of change grouped by familiarity.

components did not significantly differ (estimate = -0.14, SE = 0.09, $t = -1.62$, $p = 0.106$). Secondly, we found that participants rated familiar faces as more similar overall (estimate = 0.34, SE = 0.13, $t = 2.59$, $p = 0.010$; see Figure 13). Lastly, there was no interaction found between familiarity and level of change (see Table 2 - Table 4 in supplementary materials).

The lack of a significant difference in similarity rating between 5 and 7 components (see figure 13) is surprising because the difference in Gabor space between 5 and 7 components changed is the biggest in the whole stimuli set (see Figure 10). Therefore, one might expect the difference in perceptual distance would be similarly large, but this was not found. As one could see, the relationship is linear until 5 components. This could indicate that Gabor space and perceptual space do not relate linearly for bigger transformations. Thus, the lack of a significant difference could be an indication that more multidimensional transformations are perceived differently. Alternatively, it could be simply an effect of encoding rate and limited presentation duration: the presentation of all faces was 1500 ms so participants might not have had enough time to perceive all the changes.

As described above, the work of Beale and Keil (1995) predicts that changed faces will initially look more similar when familiar than unfamiliar but that, if the change becomes large enough, this will reverse, with changed familiar faces looking more different. If this prediction became true, there would be a main effect of familiarity and an interaction effect with transformation level. As expected, we observed greater perceptual similarity for familiar over unfamiliar faces. However, the expected interaction with the transformation

level was not found. Based on the current experiment alone, it is impossible to say if the prediction were wrong because the change could have simply been not great enough to change the identity, i.e., cross the border of a category. The current paradigm is an important first step to gaining more insight is how familiarity influences perceptual similarity and it offers a clear-cut way of comparing other transformations that will be used in future research.

General discussion

In this paper, we used transformations based on PCA to assess the relationship between perceptual similarity and Gabor, Pixel and DNN distance values. Experiment 1 explored the effects of shape and colour changes on the size of the perceptual change and the three image-based spaces. The results showed that the PCA transformations used result in a linear relationship between the three spaces and human similarity ratings. In addition to assessing the relationship between perceptual similarity and Gabor, Pixel and DNN distance values, Experiment 2 explored the effect of altering different numbers of principal components on the size of the perceptual change, and the effect of familiarity. Similar to experiment 1, the relationship between each of the spaces and perceptual similarity was linear. Additionally, our findings show that there is a significant relationship between the number of components changed and both the perceptual similarity and the three image-based spaces and that participants rated familiar faces as more similar overall. These findings combined indicate that there is a clear relationship between perceptual and our image-based spaces.

Our findings supports those of Yue et al. (2012) that Gabor-based image measures correlate highly with human performance in a discrimination task with complex figures and Meytlis (2011) that there is a relationship between distinctiveness ratings and pixel-wise Euclidean distance. Both of these studies and our current studies all had different tasks. In Yue et al. (2012), participants performed a match-to-sample task and in Meytlis (2011) participants were asked to rate single faces on distinctiveness and face pairs on perceptual distance (rating how different they looked). In both experiments the stimuli were simultaneously on-screen. Since it is known that availability of stimuli reduces the use of visual working memory and promotes the use of eye movements to gather visual input when needed (Somai, Schut, & Van der Stigchel, 2019), one could argue the observed relationship between perceptual and image-based spaces is due to the ongoing presence of the stimuli that serves as an external representation. Since the external representation is an image and the spaces are image-based, this could mean the findings do not reflect face space but simply the

use of visual input. However, in this paper, Experiment 2 includes a memory component because there is a delay between the two faces rated and the findings are similar. This indicates that the internal representation of faces does indeed relate to image-based spaces.

In the current study, we have proposed the use of a threshold based on the face set used to create the PCA space as an aid in stimulus creation. Overall, the three thresholds all gave a decent prediction whether or not a face pair will be perceived as similar or not for each experiment. However, the predictive value of the thresholds varied between the two experiments. The DNN and Pixel threshold made more correct classifications when multi-dimensional rather than single-dimensional transformations were used (single-dimensional: 30/42, multi-dimensional: 33/42), while the Gabor threshold was almost equally informative in both experiments (single-dimensional: 34/42, multi-dimensional: 33/42). The equal performance of the DNN and Pixel threshold might look surprising but is most likely to be consequent of how well-controlled the PCA transformations used are. It is noteworthy that the Gabor threshold was consistent in its mistakes over different types of transformations and could be set lower to increase predictive value, where the Pixel and DNN threshold could not be changed to describe all transformations better. Some might argue that setting a threshold in such a fashion is arbitrary because the threshold would then be different for each stimulus set. However, the goal was never to find the best ‘one-size-fits-all’ threshold, it is merely a suggestion to assist in stimulus creation. Nevertheless, it will be interesting to see how big the variation in the predictive value of the thresholds is between different stimulus sets.

It must be noted that we do not want to imply that the 3 chosen image-based spaces are superior to any others. This is simply unknown at the time of this paper, and these spaces were chosen based on personal experience with them. Another argument for using these spaces over others is the existing literature on the measures, although the literature is very limited. Additionally, the Gabor and Pixel distance values are easily obtained and understood. The DNN used in this particular instance is not openly available, and other systems with different architecture and training might produce different results (although the 6 DNNs tested in Hancock, Somai and Mileva (2020) showed similar, significant correlations). In the context of creating a human-like neural network, it will be interesting to compare the differences in DNN spaces and how they relate to perceptual measures.

Apart from exploring the use of image-based threshold to predict perceptual similarity, another aim of this paper was to test the utility of PCA-induced changes in face images to probe human face perception. Our findings show that the graded changes produced by PCA translate to a graded change in perceptual space. In other words, large PCA change indeed

produced a bigger perceptual difference than a smaller one. Additionally, the bigger the number of principal components that were altered, the bigger the perceptual change (up until 5 components). Another argument for the use of PCA-induced changes is the potential insights that can be gained from understanding why it describes perceptual similarity quite accurately in some cases and not in others. For example, in experiment 2, 5 out of 8 faces with 5 principal components changed were rated as dissimilar and 3 similar by participants. One could wonder why this variation exists: what makes some faces appear different with 5 principal components changed and how are other faces still perceived the same? Did the specific components changed have a different effect on the different faces? For example, if someone has especially distinctive eyes, then a change to the eye region may have more perceptual effect than on a face with more typical eyes. In the current study, the faces had all the same components changed and no specific selection was made in identities. The use of thresholds could give more insight in facial distinctiveness.

To our knowledge, very little research has been done on the role of familiarity in the relationship between (systematic) changes to a face and perceptual similarity. The paradigm proposed in this paper was used to explore this and successfully found a main effect of familiarity and more surprisingly, no interaction effect. Although this work is just the first step in exploring the factors that mediate a change in perceived identity, future work can be easily related to these findings because of the relationship between the image-based spaces. Regardless of the transformations, datasets and images used, it can be quantified and compared.

In conclusion, our study shows the predictive nature of Gabor, Pixel and DNN distances in the estimation of perceptual changes in faces. This relationship can be used in the exploration of the dimensionality of face space and facilitate in stimulus creation. For example, in research related to variability and learning, the experimenter is often limited in creating a high and low variability set based on observation. Any of the three image spaces can be used to quantify variability and study the effects on a more gradual scale. However, each space has its own limitations and advantages, so the best suitable space to base a threshold on will be dependent on the face image set and pre-processing tools one has available. Lastly, Experiment 2 presents a framework on how to use controlled multidimensional face transformations to explore differences in the perception of familiar and unfamiliar faces. The ability to quantify, and thus control, these transformations is a powerful tool in exploring the factors that mediate a change in perceived identity.

References

- Abudarham, N., & Yovel, G. (2016). Reverse engineering the face space: Discovering the critical features for face identification. *Journal of Vision*, 16(3), 40.
<https://doi.org/10.1167/16.3.40>
- Beale, J. M., & Keil, F. C. (1995). Categorical effects in the perception of faces. *Cognition*, 57(3), 217–239. [https://doi.org/10.1016/0010-0277\(95\)00669-X](https://doi.org/10.1016/0010-0277(95)00669-X)
- Burton, A. M., Jenkins, R., Hancock, P. J. B., & White, D. (2005). Robust representations for face recognition: The power of averages. *Cognitive Psychology*, 51(3), 256–284.
<https://doi.org/10.1016/j.cogpsych.2005.06.003>
- Burton, A. M., Kramer, R. S. S., & Ritchie, K. L. (2016). Identity From Variation : Representations of Faces Derived From Multiple Instances. *Cognitive Science*, 40, 202–223. <https://doi.org/10.1111/cogs.12231>
- Craw, I., & Cameron, P. (1991). Parameterising images for recognition and reconstruction. *BMVC91*, 367–379. <https://doi.org/10.5244/c.5.52>
- Dailey, M. N., & Cottrell, G. W. (1999). Organization of face and object recognition in modular neural network models. *Neural Networks*, 12, 1053–1073.
- De Valois, R. L., & De Valois, K. K. (1990). *Spatial vision*.
- Goffaux, V., & Rossion, B. (2007). Face Inversion Disproportionately Impairs the Perception of Vertical but not Horizontal Relations Between Features. *Journal of Experimental Psychology: Human Perception and Performance*, 33(4), 995–1002.
<https://doi.org/10.1037/0096-1523.33.4.995>
- Hancock, P. J. B. (2000). Evolving faces from principal components. *Behavior Research Methods, Instruments, and Computers*, 32(2), 327–333.
<https://doi.org/10.3758/BF03207802>
- Hancock, P. J. B., Bruce, V., & Burton, A. M. (1998). Testing Principal Component Representations for Faces. *4th Neural Computation and Psychology Workshop, London, 9–11 April 1997*, 84–97. https://doi.org/10.1007/978-1-4471-1546-5_8
- Hancock, P. J. B., Burton, A. M., & Bruce, V. (1996). Face processing: Human perception and principal components analysis. *Memory & Cognition*, 24(I), 26–40.
<https://doi.org/10.3758/BF03197270>
- Hancock, P. J. B., Somai, R. S., & Mileva, V. R. (2020). Convolutional neural net face recognition works in non-human-like ways. *Royal Society Open Science*, 7, 1–7.
- Jenkins, R., White, D., Van Montfort, X., & Burton, A. M. (2011). Variability in photos of the

- same face. *Cognition*, 121(3), 313–323. <https://doi.org/10.1016/j.cognition.2011.08.001>
- Johnston, R. J., & Edmonds, A. J. (2009). Familiar and unfamiliar face recognition: A review. *Memory*, 17(5), 577–596. <https://doi.org/10.1080/09658210902976969>
- Jones, J. P., & Palmer, L. A. (1987). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58(6), 1187–1211. <https://doi.org/10.1152/jn.1987.58.6.1187>
- Kämäräinen, J., Kyrki, V., & Kälviäinen, H. (2002a). Fundamental Frequency Gabor Filters. *Proceedings of ICPR2002 16th International Conference on Pattern Recognition*, 1, 628–631.
- Kämäräinen, J., Kyrki, V., & Kälviäinen, H. (2002b). Gabor features for invariant object recognition. Lappeenranta University of Technology.
- Kirby, M., & Sirovich, L. (1990). Application of the Karhunen-Loève procedure for the characterization of human faces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(1), 103–108.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, 82(13), 1–26. <https://doi.org/10.18637/jss.v082.i13>
- Lades, M., Vorbrüggen, J. C., Buhmann, J., Lange, J., Malsburg, C. von der, Würtz, R. P., & Konen, W. (1993). Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions on Computers*, 42(3), 300–311.
- MATLAB. (2019). MATLAB r2019a. Natick, Massachusetts: The Mathworks Inc.
- MATLAB. (2020). MATLAB r2020a. Natick, Massachusetts: The Mathworks Inc.
- Meytlis, M. (2011). A model of face space. *Visual Cognition*, 19(1), 13–26. <https://doi.org/10.1080/13506285.2010.512507>
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational Explorations in Cognitive Neuroscience : Understanding the Mind by Simulating the Brain*.
- O'Toole, A. J., Deffenbacher, K. A., Valentin, D., McKee, K., Huff, D., & Abdi, H. (1998). The perception of face gender: The role of stimulus structure in recognition and classification. *Memory and Cognition*, 26(1), 146–160. <https://doi.org/10.3758/BF03211378>
- R Core Team. (2019). R: A language and environment for statistical computing. *R Foundation for Statistical Computing, Vienna, Austria*. Retrieved from <https://www.r-project.org/>
- Ramon, M., & Gobbini, M. I. (2018). Familiarity matters: A review on prioritized processing

- of personally familiar faces. *Visual Cognition*, 26(3), 179–195.
<https://doi.org/10.1080/13506285.2017.1405134>
- Ringach, D. L. (2002). Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of Neurophysiology*, 88(1), 455–463.
<https://doi.org/10.1152/jn.2002.88.1.455>
- Ritchie, K. L., & Burton, A. M. (2017). Learning faces from variability. *Quarterly Journal of Experimental Psychology*, 70(5), 897–905.
<https://doi.org/10.1080/17470218.2015.1136656>
- Somai, R. S., Schut, M. J., & Van der Stigchel, S. (2019). Evidence for the world as an external memory : A trade-off between internal and external visual memory storage. *CORTEX*, 122, 108–114. <https://doi.org/10.1016/j.cortex.2018.12.017>
- Turk, M., & Pentland, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1), 71–86.
- Valentine, T. (1991). A Unified Account of the Effects of Distinctiveness, Inversion, and Race in Face Recognition. *The Quarterly Journal of Experimental Psychology Section A*, 43(2), 161–204. <https://doi.org/10.1080/14640749108400966>
- Wiskott, L., Fellous, J. M., Krüger, N., & Von der Malsburg, C. (1997). Face recognition by elastic bunch graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7), 775–779. https://doi.org/10.1007/3-540-63460-6_150
- Yue, X., Biederman, I., Mangini, M. C., Malsburg, C. von der, & Amir, O. (2012). Predicting the psychophysical similarity of faces and non-face complex shapes by image-based measures. *Vision Research*, 55, 41–46. <https://doi.org/10.1016/j.visres.2011.12.012>

Supplementary materials

Table 1 *Distribution of familiarity.*

ID number	Percentage familiar	Percentage unfamiliar
1	67.2	32.8
2	21.3	78.7
3	62.3	37.7
4	68.9	31.1
5	16.4	83.6
6	31.1	68.9
7	73	39.3
8	72.1	27.9

Table 2 *Fixed effect of the linear mixed effects model for familiarity and transformation level with transformation level 2 as reference for Experiment 2.*

Predictors	Estimates	SE	df	t	p
(intercept)	5.11	.10	110.89	49.46	<.001***
Transformation level 3	-1.05	.09	1584.37	-11.83	<.001***
Transformation level 5	-1.86	.09	1665.97	-20.60	<.001***
Transformation level 7	-2.00	.09	1527.82	-22.48	<.001***
Familiarity	.33	.13	630.66	2.59	.010**
Transformation level 3 x Familiarity	.17	.19	200.70	.92	.359
Transformation level 5 x Familiarity	.10	.18	287.58	.56	.574
Transformation level 7 x Familiarity	.04	.19	167.31	.20	.843

Note * $p < .05$; ** $p < .01$; *** $p < .001$

Table 3 *Fixed effect of the linear mixed effects model for familiarity and transformation level with transformation level 3 as reference for Experiment 2.*

Predictors	Estimates	SE	df	t	p
(intercept)	4.06	.10	116.12	38.83	<.001***
Transformation level 2	1.06	.09	1584.37	11.83	<.001***
Transformation level 5	-.81	.09	1794.75	-9.11	<.001***
Transformation level 7	-.95	.09	1833.78	-10.74	<.001***
Familiarity	.51	.15	79.49	3.35	.001**
Transformation level 3 x Familiarity	-.17	.19	200.70	-.92	.359
Transformation level 5 x Familiarity	-.07	.18	947.76	-.39	.695
Transformation level 7 x Familiarity	-.14	.18	1798.11	-.77	.444

Note * $p < .05$; ** $p < .01$; *** $p < .001$

Table 4 Fixed effect of the linear mixed effects model for familiarity and transformation level with transformation level 5 as reference for Experiment 2.

Predictors	Estimates	SE	df	t	p
(intercept)	3.25	.10	112.56	31.26	<.001***
Transformation level 2	1.86	.09	1665.97	20.96	<.001***
Transformation level 3	.81	.09	1794.75	9.11	<.001***
Transformation level 7	-.14	.09	1813.79	-1.61	.11
Familiarity	.44	.14	114.138	3.06	.002**
Transformation level 2 x Familiarity	-.10	.18	287.58	-.56	.575
Transformation level 3 x Familiarity	.07	.18	947.76	.39	.695
Transformation level 7 x Familiarity	-.07	.18	1146.43	-.37	.712

Note: * $p < .05$; ** $p < .01$; *** $p < .001$