Putative Ratios of Facial Attractiveness in a Deep Neural Network

Song Tong^a, Xuefeng Liang^{b,*}, Takatsune Kumada^{a,*}, Sunao Iwaki^c

^a*IST, Graduate School of Informatics, Kyoto University, Kyoto, Japan*

^b*School of Artificial Intelligence, Xidian University, Xi'an, P. R. China*

^c*Information Technology and Human Factors, AIST, Tsukuba, Japan*

Abstract

Empirical evidence has shown that there is an ideal arrangement of facial features (ideal ratios) that can optimize the attractiveness of a person's face. These putative ratios define facial attractiveness in terms of spatial relations and provide important rules for measuring the attractiveness of a face. In this paper, we show that a deep neural network (DNN) model can learn putative ratios from face images based only on categorical annotation when no annotated facial features for attractiveness are explicitly given. To this end, we conducted three experiments. In Experiment 1, we trained a DNN model to recognize the attractiveness (female/male $\times$ high/low attractiveness) of face in the images using four category-specific neurons (CSNs). In Experiment 2, face-like images were generated by reversing the DNN model (e.g., deconvolution). These images depict the intuitive attributes encoded in CSNs of the four categories of facial attractiveness and reveal certain consistencies with reported evidence on the putative ratios. In Experiment 3, simulated psychophysical experiments on face images with varying putative ratios reveal changes in the activity of the CSNs that are remarkably similar to those of human judgements reported in a previous study. These results show that the trained DNN model can learn putative ratios as key features for the representation of facial attractiveness.

*Corresponding author

Email addresses: `tong.song.7s@kyoto-u.ac.jp` (Song Tong), `xliang@xidian.edu.cn` (Xuefeng Liang), `t.kumada@i.kyoto-u.ac.jp` (Takatsune Kumada), `s.iwaki@aist.go.jp` (Sunao Iwaki)

This finding advances our understanding of facial attractiveness via DNN-based perspective approaches.

Keywords: Facial Attractiveness, Deep Neural Network, Deconvolution Method, Golden Ratio, Psychophysical, Face Perception

1. Introduction

Attractiveness is a facial feature that conveys significant biological advantages in terms of mating success (Pashos & Niemitz, 2003), earning potential (Frieze et al., 1991), and longevity (Henderson & Anglin, 2003) across age groups
5 and cultures (Perrett et al., 1994, 1998; Rubenstein et al., 2002). A large body of studies has suggested that people use similar criteria for determining facial attractiveness, although some individual and cross-cultural differences have been noted. One criterion of facial attractiveness is defined using ideal ratios (Valenzano et al., 2006), such as neoclassical canons (Schmid et al., 2008; Jayaratne
10 et al., 2012), golden proportions (Borissavlievitch & Hautecoer, 1958; Jefferson, 2004), facial thirds (Farkas & Schendel, 1995; Farkas & Kolar, 1987), and new golden ratios (Pallett et al., 2010; Bóo et al., 2013). These putative rules define a physical measure for facial attractiveness using spatial relations between parts of the face. For example, the ancient Greeks believed that the golden ratio
15 represents the essence of beauty (Rhodes, 2006). Leonardo da Vinci’s painting *Mona Lisa* also illustrates the golden ratio of facial beauty (Jefferson, 2004; Atiyeh & Hayek, 2008).

Empirical evidence provides support for the association between putative ratios and the human perception of facial attractiveness. For example, the
20 golden ratio is an irrational value (0.618) and a commonly accepted metric for measuring the harmoniousness of proportions in aesthetics. The golden ratio has also been proposed as a universal standard of facial attractiveness (Jefferson, 2004), especially as a standard in plastic surgery (Holland, 2008), and has been used to compute facial attractiveness (Schmid et al., 2008). In addition, Pallett
25 et al. (2010) has reported that ideal facial ratios (called the new golden ratios)

can optimize the attractiveness of face images. These new golden ratios consist of two ratios: a vertical ratio of the eye–mouth distance to the height of the face that is approximately 0.36, and a horizontal ratio of the distance between the eyes to the width of the face that is approximately 0.46. Furthermore, Shen et al. (2016) found that the facial attractiveness induced by facial proportions evokes a neural response in some regions of the brain in the reward system, such as the orbitofrontal cortex and amygdala. This suggests that the putative ratios contain high-level information about facial attractiveness.

The rules of putative ratios have also been used to construct computer models for the assessment of the attractiveness of face images (Gunes & Piccardi, 2006; Schmid et al., 2008; Chen & Zhang, 2014). Some of these models relied on a set of putative ratios such as the golden proportions and facial thirds as features of the face images and used them to calculate attractiveness scores. For instance, Schmid et al. (2008) used a feature vector comprising 77 putative ratios to build a computer model for facial attractiveness. Results of experiments revealed a significant correlation between the predicted facial attractiveness scores and those rated by human subjects. Other approaches (Kagian et al., 2008; Zhang et al., 2017) used only geometric features based on the automatic detection of facial landmarks and all the ratios among them. For example, Kagian et al. (2008) used 6,972 distance vectors of 84 facial landmarks as features to predict the attractiveness of faces. They reported that these geometric features implicitly capture the basic human psychophysical characteristics for the perception of facial attractiveness, such as averageness (Langlois & Roggman, 1990) and symmetry (Rhodes et al., 1999). However, these ratio-based models require the annotation of the facial landmarks, which incurs high labor costs. These landmarks also lack hierarchical, high-level semantic information (Wang et al., 2014).

To solve the above problems, researchers have applied deep neural networks (DNNs) to compute the attractiveness of face images (LeCun et al., 2015; Wang et al., 2014; Rothe et al., 2016). In contrast to methods using handcrafted features, DNNs learn higher level features from a large number of face images to

yield more precise predictions (Wang et al., 2014; Chen et al., 2016). For example, Rothe et al. (2016) constructed a convolutional neural network (CNN) to learn facial attractiveness from thousands of images by relying on millions
60 of ratings from the Internet. The computational performance of DNNs has led to their use in smartphones, such as for face beautification (Li et al., 2015) and facial makeup recommendation (Nguyen & Liu, 2017). This indicates that, compared with handcrafted features, better features contributing to facial attractiveness can be learned by DNN models.

65 Some striking similarities have recently been observed between DNNs and biological vision systems in terms of their operation (Cadieu et al., 2014; Yamins & DiCarlo, 2016; Cichy et al., 2016a,b; Wang & Cottrell, 2017; Seeliger et al., 2018; O’Toole et al., 2018; Kietzmann et al., 2019). The work in Cichy et al. (2016a) showed that the DNN can capture the stages of human visual processing
70 in both time and space, from the early visual areas to the dorsal and ventral streams. Seeliger et al. (2018) reported that DNNs can capture hierarchical representations of color for object recognition that rivals the performance of primates. For instance, the middle layers of the DNN show a denser sampling of hue that suggests some correlation with hue maps in V2 (Conway, 2003).
75 Thus, some research (Kriegeskorte, 2015; VanRullen, 2017) has proposed using DNNs as a new tool to gain insight into biological visual perception. However, the above studies mainly concentrated on object recognition tasks to explain the association between DNNs and biological visual perception. In addition, a little research has explored the implicit feature representations learned by DNNs that
80 are associated with high-level perception. The studies of McCurrie et al. (2018) and Parde et al. (2019) are examples in which implicit feature representations learned by DNNs are used to better understand subjective social traits of face perception.

This study examines whether the putative ratios can be learned by DNNs
85 based only on categorical annotation, i.e. where no annotated facial features for attractiveness are explicitly given. To investigate this issue, we addressed the following methodological problems: (1) how to interpret the feature representa-

tions of facial attractiveness learned by DNNs, and (2) how to evaluate that implicit feature representations contain the putative ratios of facial attractiveness.

90 However, it should be noted here that we examined feature representations of faces only based on photos. These do not necessarily cover all inherent features of individual faces.

To investigate these questions, we applied a data-driven approach and conducted three experiments. In Experiment 1, we trained a DNN model for facial
95 attractiveness (FADNN) with four category-specific neurons (CSNs) from thousands of face images annotated according to high/low attractiveness and gender (Figure 1a). In Experiment 2, the implicit features of the CSNs were projected to four face-like images by reversing FADNN (i.e. deconvolution) (Yosinski et al., 2015). The face-like image shown in Figure 1b, is an example generated by the
100 neuron representing the category of highly attractive females. The face-like images not only make the implicit features of attractiveness explicit, but also help us to examine whether the putative ratios are encoded in the implicit features. In Experiment 3, psychophysical experiments were conducted on the trained FADNN to examine whether the responses of the CSNs match the results of
105 human judgments of the putative ratios (e.g. Pallett et al. (2010)).

2. Experiment 1: Training category-specific neurons (CSNs) for the attractiveness of face images

In Experiment 1, we trained a DNN model for facial attractiveness classification (FADNN). The goal of FADNN is to extract feature representation across
110 categories using its four CSNs. The flowchart of this experiment is shown in Figure 1a. To let FADNN learn features of facial attractiveness, the training dataset must have sufficient representative instances and be diverse, i.e. large amounts of high and low attractive face images that contain a variety of attributes specific to each category. High and low attractive samples ensure DNN model focusing on learning the intended features of facial attractiveness. The diverse data
115 would prevent some shortcut features (Geirhos et al., 2020) from learning for de-

Figure 1: Summary of the methods. (a) Four category-specific neurons (CSNs) were trained using thousands of face images and their attractiveness-related annotations. (b) The feature representations learned by the CSNs were projected to four face-like images by reversing the DNN using deconvolution.

cision. Here, the shortcut features denote the image attributes are correlated to the task while irrelevant to facial attractiveness. However, most open datasets are either too small (Altwaijry & Belongie, 2013; Eisenthal et al., 2006; Kagian et al., 2008) or contain mostly neutrally attractive faces (Rothe et al., 2016; Zhang et al., 2017). To address this issue, we collected the high attractive set from celebrities, while collected the low attractive set from mugshots. This provides sufficient representative instances and a diverse data sources. The dataset is composed of 4,512 face images divided into four categories (female/male $\times$ high/low attractiveness).

However, images collected from two distinct data sources may introduce some confounding factors to the FADNN model. To alleviate this issue, we

take following three steps to mitigate the influence of potentially confounding factors. Firstly, a strict selection process has been carried out to exclude low-quality images, such as cropping faces using a unified identification method (face localization) and three steps of subjective evaluations (see section 2.1.1 for details). Secondly, we applied transfer learning based on the VGG-Face model (Parkhi et al., 2015), which was trained on millions of face images for face identification. VGG-Face model can be used as a feature extractor, and is similar to the traditional scoring keys that accompany psychometric tests. A traditional scoring key involves converting responses to test questions into a set of psychometric scores, such as the Big Five personality scores. VGG-Face translates a face image into a score with 4,096 dimensions. Each dimension might subsume differences in multiple features in the convolution layers. VGG-Face offers two main advantages in the context of this study. 1) Successful facial recognition depends on robust facial features. VGG-Face aims to represent a given face as a feature vector of identity that is invariant under changes in facial expression, background, lighting, head orientation, and such properties of the image as brightness and contrast. 2) Transferring the parameters of a well-trained DNN to a different but related task reduces the risk of overfitting (Wang & Kosinski, 2018). Thirdly, we tested the FADNN on the dataset out of the training data distribution in Experiments 1 and 3 to verify if FADNN learns the intended features contributing to facial attractiveness.

We designed the output layer of FADNN to consist of four neurons corresponding to the four categories (CSNs). Because VGG-Face translates a face image into a feature vector, the training process on the collected dataset involved optimizing the weights of the connections between the feature vector and the four CSNs. Furthermore, the CSNs need to be strongly activated by samples from the corresponding categories. The feature representation of FADNN was compared with two commonly accepted features of facial attractiveness, i.e. geometric (Kagian et al., 2008) and Gabor features (Chen et al., 2010). Geometric features describe the arrangement of facial features and are composed of pairwise distances between facial landmarks. Gabor features represent detailed

information about faces (such as texture).

160 2.1. Methods

2.1.1. Dataset

To collect candidates for the high attractiveness sets, we focused on celebrities and their publicly available images (Zhang et al., 2017). They were collected from two websites, TC Candler ¹ and TB World ². These sites list millions of suggestions in terms of images of attractive people every year. In this study, 165 a list consisting of images of faces of 406 females and 247 males was created according to the rankings of these websites in recent years, and 65,300 photos were crawled via Google Image Search (100 photos/person).

For candidate images for the low attractiveness sets, we used mugshots. The stereotypical notion that "ugly is bad" makes people think the mugshots of criminals are unattractive (Valla et al., 2011). Some empirical evidence indicates that visually unattractive people pursue economic and sexual gains through illegal means, lacking the advantages and options of people in the set of attractive face images (Cavior & Howard, 1973; Agnew, 1984; Cavior et al., 1974); thus, we 175 used a large number of mugshots as the source of unattractive faces. MacLin & MacLin (2004) claimed that the faces of criminals are less attractive than neutral faces on average. We collected 39,764 mugshots from the St. Louis Arrests and Mugshots website ³.

Because the locations of the facial regions varied in the images, they were 180 detected and cropped using the method proposed in (Mathias et al., 2014) and normalized to 224×224 . To ensure image quality, we cleaned up the raw dataset using the following two steps: (1) Two authors (male: aged 25 and 41 years) scrutinized all face images and filtered out erroneously tagged, side-view, and duplicate photos. (2) Three males and three females (average age: 27.3 years, 185 $SD = 2.7$) rated their agreement (agree or disagree) with the attractiveness of

¹<https://www.facebook.com/independentcritics/>

²<https://www.facebook.com/topbeautyworld.d/>

³<http://www.stlmugshots.com/>

the faces in the sets of highly attractive and unattractive face images. Finally, we separated the two sets into four categories: highly attractive females (HAFs), unattractive females (UAFs), highly attractive males (HAMs), and unattractive males (UAMs). For each category, we chose 1,128 face images that had secured
190 the agreement of at least three participants to construct a dataset of facial attractiveness consisting of 4,512 face images.

To ensure that the images are extremely (un)attractive, we conducted a subjective rating of the attractiveness of these data. First, 240 images (60 per category) were randomly selected from the dataset. Each image was displayed
195 on the screen for 1.5 s, and then participants were given 3.5 s to rate it (as “very attractive”, “attractive”, “neutral”, “unattractive”, or “very unattractive”). Fifteen participants were recruited for this experiment (all male, average age: 22.9 years, SD = 2.4). The rating results of 13 participants were valid, and these results are shown in Figure 7. A significant difference (unpaired t-test: $t(118) =$
200 52.002 , $p = 3.387 \times 10^{-83}$) exists between attractive females (Mean $\pm$ SD: 4.073 ± 0.308) and unattractive females (1.576 ± 0.208). Analogously, a significant difference (unpaired t-test: $t(118) = 39.769$, $p = 3.399 \times 10^{-70}$) also exists between attractive males (4.110 ± 0.359) and unattractive males (1.731 ± 0.293). This study was approved by the Ethics Committee in the Unit for Advanced
205 Studies of the Human Mind, Kyoto University. Participants provided their written informed consent to participate in this study.

2.1.2. Training the CSNs

VGG-Face, which was originally designed for facial recognition, includes six convolutional layers and three fully connected layers. Hierarchical computation
210 was used to repeatedly stack the following operations: (a) a set of learned two-dimensional (2D) convolution operators that extract features hierarchically from the input stimuli and the output of previous convolutional layers; (b) spatial maximum pooling applied to a small local region of the feature maps obtained from the convolutional operation, which reduces their dimensionality and gains
215 some degree of translational invariance; and (c) nonlinear activation functions

applied to the pooled responses. The learned low-level features (i.e. edges, corners, and contours) are represented in the initial layers whereas the subsequent layers represent high-level object-related representations (i.e. parts of objects or entire objects) (Zeiler & Fergus, 2014). The fully connected layers (fc6, fc7, and fc8) are typically on top of these stacked convolutional operations and generate more abstract and task-dependent features. The fc6 layer transforms the two-dimensional (2-D) feature map into a 4,096-D feature vector. The fc7 layer is a hidden layer with the same number of dimension, and could be considered as a middle-level feature vector. The fc8 layer has 2,622 neurons, each of which represents an identity. Finally, the output layer is a softmax layer with a cross-entropy objective function that calculates the probability of the target category of the input. In this study, we replaced the fc8 layer of VGG-Face with a layer consisting of four neurons for attractiveness classification, that is, FADNN.

The collected dataset (4,512 face images) was split into a training set and a testing set, where the former consisted of 4,000 images, (1,000 per category), and the latter contained 512 images (128 per category). FADNN was fine-tuned on the last fully connected layer while other parameters were fixed. The model was trained by minibatch gradient descent with a batch size of 20 (owing to memory-related constraints) for 30,000 iterations to ensure convergence. The learning rate of the last layer was set to 0.001, the weight decay was set to 5×10^{-4} , and the momentum was 0.9. The process was conducted using the Caffe deep learning framework (Jia et al., 2014) on an NVIDIA Tesla K40 GPU 12 GB.

2.2. Results and discussion

To evaluate the learned representation of FADNN, it was compared with traditional feature-based models for facial attractiveness, including those based on geometric and appearance-related features (Tong et al., 2017). Details of these models are provided in Appendix 7.3. We then used principal component analysis (PCA) to visualize the high-dimensional features extracted by the models for an intuitive analysis of the separability of the 512 test images (Barshan

et al., 2011). These test images were not presented to FADNN in the training phase. Figure 2 shows the distributions of the extracted features in the 2-D space of PCA. In the model that used geometric features, the four categories were mixed with one another, and thus have low separability. The model using
250 Gabor features has better separability. In contrast, the VGG-Face feature (fc6) exhibits the best separability, which indicates it has the ability of classifying different attractiveness categories.

In addition, the within-class and between-class distances (WD and BD) were used to quantitatively evaluate the separability of these features. WD and BD
255 denote the distances between features within a category and between categories, respectively. Thus, smaller the values of WD, larger the values of BD, more separable the features are. Detailed definitions of WD and BD are provided in Appendix 7.2. Table 1 shows the quantitative results of the models. The VGG-Face feature model obtains the smallest WD and largest BD, which suggests that
260 it is more discriminating than the other feature models and more representative of the attractiveness categories.

Figure 2: Comparison of the separability of models using three kinds of features through PCA-based dimensional reduction. The visualization of the extracted features illustrates the distribution of the facial features in the four categories (marked by different colors). Blue dots represent highly attractive female faces; green dots represent unattractive female faces; yellow dots represent highly attractive male faces; and red dots represent unattractive male faces.

We compared the models in terms of the predictive accuracy of the 512 test images as shown in Table 2. Both the geometric feature-based and Gabor feature-based models achieved accuracies over 90%, but FADNN outperformed
265 them with an accuracy of 98.1%. These results indicate that FADNN can learn

Table 1: Comparison of separability of the three feature models using WD and BD,

Methods	Within-class distance (WD)	Between-class distance (BD)
Geometric feature	1.660	0.633
Gabor feature	1.284	1.527
VGG-Face feature	0.923	1.910

more discriminative information from the categories than the traditional feature-based models.

Table 2: Comparison of the classification accuracy of the models.

Methods	Accuracy
Geometric features model	90.6%
Gabor features model	92.8%
Proposed model (FADNN)	98.1%

In addition, to further investigate whether FADNN learned robust features contributing to attractiveness, we tested it on the SCUT-FBP5500 dataset (Liang et al., 2018) with attractiveness scores, which has been used for emerging studies of facial attractiveness analysis in both computer vision (Xu et al., 2019; Shi et al., 2019) and the psychological community (Zhao et al., 2019). The dataset contains 5,500 face images. Each image was rated by 60 volunteers, and the average score was used as the ground-truth to remove personal preference bias. The dataset was divided into four subsets with diverse races and both genders, and includes 2,000 Asian females, 2,000 Asian males, 750 Caucasian females, and 750 Caucasian males. Because the dataset for FADNN training has few Asian persons, only 750 Caucasian females and 750 Caucasian males were used for the test. Note that [the SCUT-FBP5500 dataset is out of the distribution of the training data](#).

To evaluate the robustness of our model, we used PCA to visualize the high-dimensional features extracted by the model on the data in the SCUT-FBP5500 dataset. Because the SCUT-FBP5500 lacks categorical labels, which is different

from our constructed dataset, we chose the 10% highest/lowest scoring faces for
 285 testing, i.e., 75 images for each category. Figure 8 shows the distribution of the
 VGG-Face features of these 300 images, which is very similar to the distribution
 shown in Figure 2.

In addition, we compared the correlation between the activation of the CSNs
 and the attractiveness scores. However, we found the distribution of the attrac-
 290 tiveness scores of the SCUT-FBP5500 dataset is unbalanced (see Figure 9),
 which may affect the correlation analysis. Thus, we re-balanced the dataset as
 follows. First, the female and male facial sets were each divided into 20 subsets
 according to the attractiveness scores (i.e., from lowest score to highest score
 in 20 steps). Second, 10 face images were randomly selected from every sub-
 295 set. All images would be selected if the subset had no more than 10 faces. We
 performed 10 repeated sampling phases. In each sampling phase, 173 female
 face images and 180 male images were selected. Figure 9 (right) shows the
 distribution of the attractiveness scores of a sampling phase (353 face images).
 Finally, the face stimuli were fed to FADNN, which outputted the scores of
 300 the four CSNs in the last fully connected layer (fc8): the highly attractive fe-
 male neuron (Neuron-HAF), unattractive female neuron (Neuron-UAF), highly
 attractive male neuron (Neuron-HAM), and unattractive male neuron (Neuron-
 UAM). For each sampling phase, the scores of Neuron-HAF and Neuron-UAF
 were recorded for the 173 Caucasian females. Analogously, the scores of Neuron-
 305 HAM and Neuron-UAM were recorded for the 180 Caucasian males.

Pearson correlation was used as the metric to evaluate the relationship be-
 tween the CSN scores and human-rated attractiveness scores. For the female
 faces, the average correlation coefficient between the scores of Neuron-HAF and
 the attractiveness scores was $r = 0.7583 \pm 0.0198$, all $p < 0.00001$ (see Figure 3a),
 310 whereas the average correlation coefficient between the scores of Neuron-UAF
 and the attractiveness scores was $r = -0.5655 \pm 0.0217$, all $p < 0.00001$. For
 the male faces, the average correlation coefficient between the scores of Neuron-
 HAM and the attractiveness scores was $r = 0.6653 \pm 0.0161$, all $p < 0.00001$
 (see Figure 3b), whereas the average correlation coefficient between the scores

315 of Neuron-UAM and the attractiveness scores was $r = -0.6221 \pm 0.0202$, all
 $p < 0.00001$. Xu et al. (2019) found that the handcrafted feature-based models
 can predict attractiveness scores with Pearson correlations of $[0.6, 0.7]$ whereas
 the DNN-based models can predict attractiveness scores with Pearson corre-
 lations of $[0.8, 0.9]$. Unlike these models, even though that our FADNN was
 320 not trained on the SCUT-FBP5500 dataset, Neuron-HAF and Neuron-HAM
 achieve performances that are comparable with those of handcrafted feature-
 based models. This result demonstrates that the four trained CSNs subsume
 the main features associated with facial categories and can be used as metrics
 to determine the membership of an image in a category.

Figure 3: (a) Correlation between the Neuron-HAF scores and attractiveness scores of 173
 female images. (b) Correlation between the Neuron-HAM scores and attractiveness scores of
 180 male images.

325 **3. Experiment 2: Rules of putative ratios in neuron interpretation**

Experiment 1 shows that the FADNN learned features contribute to the
 facial attractiveness classification. However, these feature representations are
 not intuitively interpretable. To address this issue, we used the deconvolution

method (Zeiler & Fergus, 2014; Mahendran & Vedaldi, 2015) to visualize the
 330 learned representations. A flowchart of the visualization is shown in Figure 1b.

The deconvolution method was proposed to interpret the intrinsic functions
 of DNNs. This method can generate an image that maximizes the activation of a
 category-specific neuron by inputting a content-free (noise) image (Liang et al.,
 2019). Because the CSNs are located in the fully connected layer, the position
 335 information of the preferred patterns is lost (Horikawa & Kamitani, 2017). In-
 stead of a content-free image, we thus use the mean of all faces as input. Thus,
 the features learned by a CSN can be projected back into a facial space (Grant
 et al., 2016). Using this method, because the generated images contain essen-
 tial features (such as the mouth, eyes, and chin), we can assess the information
 340 using FADNN on the basis of physical measures of facial attractiveness such as
 the golden ratio.

To estimate the association between the generated images and human per-
 ception of facial attractiveness, the putative ratios of facial attractiveness, in-
 cluding the golden ratio and new golden ratios, were employed to quantify the
 345 generated images. The rules of the golden ratio can be used measure ratios
 between facial elements to obtain the local putative ratios of facial attractive-
 ness (Jefferson, 2004; Holland, 2008), whereas the new golden ratios give global
 features of these putative ratios (Pallett et al., 2010; Bóo et al., 2013).

3.1. Methods

350 3.1.1. Generating mean face images

The mean face images were automatically generated by three sequential pro-
 cesses: face landmark detection, coordinate transforms, and averaging. First,
 the facial landmarks of all faces were extracted by the active shape model as
 shown in Figure 10. Second, landmarks on the coordinate system of the input
 face were transferred to the average face coordinate system by the following
 similarity transform matrix:

$$S = \begin{bmatrix} s_x \cos(\theta) & \sin(\theta) & t_x \\ -\sin(\theta) & s_y \cos(\theta) & t_y \end{bmatrix}, \quad (1)$$

where the first two columns of this matrix encode rotation (θ) and scale (s_x, s_y), respectively, and the last column encodes translation (t_x, t_y) in the x and y directions. Third, the mean face image was generated by taking the mean of the pixel intensities of all warped face images.

355 3.1.2. Generating face-like images for CSNs

The deconvolution method (Mahendran & Vedaldi, 2016; Nguyen et al., 2016) was employed to visualize the learned representation for the CSNs of FADNN. The technique is used to project the learned representation of a particular neuron onto an input image by maximizing the activation of this neuron
 360 (Erhan et al., 2009) (as an optimization problem). A mean face image was first propagated through FADNN, and the derivative of a target CSN activation was then calculated with respect to each pixel of the input image. This describes the way to change the pixel values of the input image to increase the activation of the CSN. In addition, total variation and jitter (described in Mahendran &
 365 Vedaldi (2016)) were used in the optimization process to improve the quality of the generated images. The optimizer was implemented using the source code provided by Nguyen et al. (2016). Finally, four face-like images corresponding to the four CSNs (categories) were generated by the deconvolution method.

3.2. Results and discussion

370 The four generated face-like images are shown in Figure 4 that mainly present facial features. In addition, these images illustrate different visual features of the four categories. For instance, the generated image for Neuron-HAF has a pointed chin and standard arched eyebrows. On the contrary, the generated image for Neuron-UAF has a short chin and short eyebrows. These features are
 375 consistent with the shape analysis in terms of the attractiveness of female faces (Valenzano et al., 2006; Feser et al., 2007). In addition, the image generated for Neuron-HAM has a square chin, which correlates with the analysis of attractive male faces (Rhodes, 2006).

Figure 4: Generated face-like images illustrating the discriminating features learned from raw face data. The images were generated using methods described in 3.1.2 using different CSNs: (a) highly attractive female neuron (Neuron-HAF), (b) unattractive female neuron (Neuron-UAF), (c) highly attractive male neuron (Neuron-HAM), and (d) unattractive male neuron (Neuron-UAM). The boxes and lines represent the rules of the golden ratio and new golden ratios, respectively. A detailed description of the measurement of the rules of the golden ratio is provided in Table 3.

The putative ratios—the golden ratio and new golden ratios—were calculated for the generated images as to provide quantitative measures of facial attractiveness. To measure the rules of the golden ratio for facial attractiveness, the six ratios (r_m) described in Table 3 were measured, as shown by the boxes in Figures 4a and 4d. In addition, the new golden ratios of facial attractiveness included a horizontal and a vertical ratio, mathematically defined as $R_h = \frac{0.5(H_1+H_2)}{H_3}$ and $R_v = \frac{V_1}{V_2}$, respectively. The measurements of H_1 , H_2 , and H_3 , and V_1 and V_2 are shown in Figures 4b and 4c. Tables 4 and 5 list the

results of the measurements according to the rules of the golden and new golden ratios, respectively, for the generated images.

Table 3: Rules of the golden ratio for attractive faces (Schmid et al., 2008). The colored names in brackets represent the colors of the rectangles in Figures 4d and 4a. Here, r_m represents the m -th ratio for an attractive face with the golden ratio.

No.	Description of Ratio
r_1	Distance between the top and bottom of eyebrows to that between eyebrows and hairline (White)
r_2	Distance between the mouth and bottom of nose to that between bottom of nose and pupil of eye (Blue)
r_3	Distance between the pupil of eye and bottom of nose to that between bottom of nose and chin (Blue)
r_4	Distance between the chin and mouth to that between mouth and pupil of eye (Blue)
r_5	Distance between the side of face side and inside of eye to that between inside of eye and side of face (Golden)
r_6	Distance between the center of nose and inside of eye to that between inside of and outside of eye (Black)

As shown in Table 4, the numbers in the brackets describe the absolute error
 390 between the golden ratio and the measured values. These five ratios (except for
 r_5) for the generated image for Neuron-HAF are closer to the golden ratio than
 those of the image generated for Neuron-UAF. An analogous difference was
 found between the generated images for Neuron-HAM and Neuron-UAM. The
 mean absolute error (MAE), which reflects the difference between the measured
 395 ratios and the golden ratio, was used to describe the rules of the golden ratio
 for attractive faces in the generated images. Table 4 shows that the generated
 image for Neuron-HAF has a much lower MAE (0.066) than that generated for
 Neuron-UAF (0.183). Thus, the image for Neuron-HAF is closer to the putative
 ratios of facial attractiveness according to the rules of the golden ratio. This
 400 tendency also holds for the image of Neuron-HAM (0.175), which has ratios
 closer to the golden ratio than does Neuron-UAM (0.258).

Table 5 shows that the values of R_v of Neuron-HAF and Neuron-HAM are
 close to the standard value (0.36) of the new golden ratio. Significantly, both
 R_h and R_h for the image for Neuron-HAM (0.470, 0.354) are close to the new
 405 golden ratios (0.46, 0.36). An analogous consistency was observed in the image
 of Neuron-HAM (0.494, 0.357). However, neither of these consistencies appears
 in the images for Neuron-UAF and Neuron-UAM. For instance, the image for
 Neuron-UAF has ratios of (0.339, 0.393), which are far from the new golden
 ratios.

Table 4: Measured ratios for generated images for different CSNs. The values in the brackets indicate the absolute error between the measured values and the golden ratio. MAE represents the mean absolute error.

Golden Ratio	r_1	r_2	r_3	r_4	r_5	r_6	MAE
Neuron-HAF	0.681 (0.063)	0.536 (0.082)	0.761 (0.143)	0.619 (0.001)	0.585 (0.033)	0.694 (0.076)	0.066
Neuron-UAF	0.755 (0.137)	0.416 (0.202)	1.032 (0.415)	0.551 (0.067)	0.639 (0.021)	0.360 (0.258)	0.183
Neuron-HAM	0.346 (0.272)	0.381 (0.237)	0.660 (0.042)	0.426 (0.192)	0.668 (0.050)	0.363 (0.255)	0.175
Neuron-UAM	1.020 (0.402)	0.907 (0.289)	0.379 (0.239)	1.094 (0.476)	0.530 (0.088)	0.669 (0.051)	0.258

Table 5: Results of the new golden ratios for generated images for the CSNs. Here, R_h and R_v respectively represent the horizontal and vertical ratios defined by the new golden ratios for attractive faces. The values in the brackets indicate the absolute error between the measured values and the new golden ratios.

New Golden Ratio	$R_h(0.46)$	$R_v(0.36)$
Neuron-HAF	0.494 (0.034)	0.357 (0.003)
Neuron-UAF	0.339(0.121)	0.393(0.033)
Neuron-HAM	0.470(0.010)	0.354(0.006)
Neuron-UAM	0.415(0.045)	0.367(0.007)

410 The generated images intuitively show the learned feature representation of
the CSNs for different categories of attractiveness. Some shape features consis-
tent with the findings of psychophysical research (Valenzano et al., 2006; Feser
et al., 2007; Rhodes, 2006) are observed in the generated images. Furthermore,
the quantitative results show that the images for Neuron-HAF and Neuron-
415 HAM contain arrangements of facial parts close to the golden ratio and new
golden ratios. However, we did not explicitly include these shape features and
rules of the putative ratios in the learning process. [Therefore, this result sug-
gests that the rules of the putative ratios were learned by FADNN without any](#)

cues.

420 4. Experiment 3: Rules of putative ratios for neuron activation

The images generated by FADNN show that the rules of putative ratios may be encoded in feature representations of four CSNs. To further explore the association between the putative ratios and the CSNs, we conducted a psychophysical experiment using FADNN. In a psychophysical experiment with
425 human observers, Pallett et al. (2010) distorted the same face using different horizontal and vertical ratios, and found the optimal ratios (i.e. new golden ratios) for facial attractiveness. Similarly, we distorted the horizontal and vertical ratios of the faces while keeping other factors constant, and used these faces as stimuli for FADNN. Figure 5 shows examples of facial stimuli with differ-
430 ent degrees of expansion. The test stimuli were distorted automatically from a new face dataset. When the modified face images were fed into FADNN, if it had learned the golden ratio as feature for attractiveness, its CSNs would be expected to generate lower attractiveness scores for distorted face images.

Figure 5: Examples of face stimuli. Image of the face of a male (upper) and female. Original face image and four degrees of expanded face images are shown.

4.1. Methods

4.1.1. Method of distorting face shapes

To obtain the stimuli of the distorted faces, we used a computer vision method to automatically manipulate the horizontal and vertical proportions. First, facial landmarks were relocated using the active shape model. Second, new landmarks were computed by simultaneously increasing the eye–mouth distance and the distance between the eyes to a certain extent. Third, Delaunay triangulation (Hong et al., 2015) was employed to distort the faces in the photographs according to the new landmarks.

For each original photograph, the eye–mouth distance and the distance between the eyes were simultaneously expanded by 4%, 8%, 12%, and 16% compared with those in the original image. The samples of the expanded faces are shown in Figure 5. In addition, we contracted the distance between the eyes by 16%, as shown in Figure 11. We selected photographs from two public face image datasets: colorFERET⁴ and GTAV (Tarrés, 2012). After removing low-quality images, 887 images (527 males and 360 females) were selected. In addition, we measured the R_h and R_v (described in Experiment 2) for every original image. The mean values and standard deviations of the R_h and R_v of the female faces are 0.465 ± 0.024 and 0.358 ± 0.020 , while those of the R_h and R_v of the male faces are 0.465 ± 0.026 and 0.365 ± 0.020 , respectively. Both were approximated as 0.46 and 0.36, respectively, which is analogously consistent with the new golden ratios. Each photograph was cropped and resized to 224×224 pixels.

Although the ColorFERET dataset consists of 994 face images, some were images of the same person. We chose one image per person to obtain 504 images of males and 347 images of females. The GTAV face dataset comprises 44 face images. Again, by discarding multiple images of the same person, 23 images of males and 13 of females were obtained. In total, 887 face images were used in this experiment.

⁴<http://face.nist.gov/colorferet/request.html>

4.1.2. Experimental procedure

The face stimuli were input into FADNN. For each photograph, the model
465 assigned scores (activations) for each of the four CSNs. The scores were used
as metrics to evaluate the extent to which an input image belongs to one of
the experimental categories. The expanded stimuli consist of five degrees of
distortion: the original photograph and four levels of expansion (887 original
images: 527 males and 360 females). The contracted stimuli include two degrees,
470 the original photograph and a level of contraction. The dependent variables
were scores in the CSNs. The results were statistically analyzed using SPSS 25
statistical software.

4.2. Results and discussion

Figure 6 shows the scores (mean values and standard deviations (SDs)) of
475 the four CSNs for different degrees of expansion. As the faces were expanded,
both the scores of Neuron-HAF and Neuron-HAM decreased whereas those of
Neuron-UAF and Neuron-UAM increased.

For the original faces, the paired t-test was used to compare the scores
of Neuron-HAF with those of Neuron-UAF, as well as those of Neuron-HAM
480 with those of Neuron-UAM. For male faces, there was no significant difference
between either pair (Neuron-HAF vs. Neuron-UAF: $t(526) = 1.385$, $p=0.167$;
Neuron-HAM vs. Neuron-UAM: $t(526) = 1.522$, $p=0.129$). This indicates that
the original images of the faces of the males are close to neutral faces, and belong
to neither HAM nor UAM. For female faces, there were significant differences
485 in both pairs (Neuron-HAF vs. Neuron-UAF: $t(359) = -7.587$, $p < 10^{-12}$;
Neuron-HAM vs. Neuron-UAM: $t(359) = 13.483$, $p < 10^{-33}$). The scores for
Neuron-UAF are 10.44 higher than those for Neuron-HAF on average, which
indicates that the original images for of females were mostly predicted to be
unattractive.

490 To examine the effects of the degree of expansion, an analysis of variance
(ANOVA) was used independently for each CSN. The main effects of the de-
gree of expansion were significant in all neurons ($F(4, 4425) = 126.815$, 102.354 ,

Figure 6: Scores (mean values and standard deviations) of the four CSNs for different degrees of expansion. The top row illustrates the neuron scores (shown by four dependent sub-figures) of 2,635 male facial stimuli. The bottom row illustrates the neuron scores of 1,800 female facial stimuli.

132.332, and 116.690, respectively; $p < 0.00001$). Multiple comparisons (Bonferroni) of the degree of expansion were significantly different in all neurons
495 for the original photographs, and those expanded by 4%, 8%, and 12% (all $p < 0.00001$). In particular, in a comparison of images expanded by 12% and 16%, the activations of Neuron-HAM and Neuron-UAM changed significantly ($p = 0.011$ and $p = 0.029$, respectively), whereas the changes in Neuron-HAF and Neuron-UAF were more moderate ($p = 0.624$, and $p = 0.720$, respectively). The
500 scores of Neuron-HAF and Neuron-HAM decreased when the faces expanded, whereas those of Neuron-UAF and Neuron-UAM increased, indicating that FADNN is more likely to recognize a face as unattractive when its measures are more than the putative ratios of facial attractiveness.

In addition, the interaction (degree of expansion $\times$ gender) is significant in
505 Neuron-HAF, Neuron-HAM, and Neuron-UAM ($F(4, 4425) = 11.405$, $p < 10^{-8}$, $F(4, 4425) = 3.935$, $p = 0.003$, and $F(4, 4425) = 17.331$, $p < 10^{-13}$, respectively).

However, in Neuron-UAF, the interaction (degree of expansion $\times$ gender) is only marginally significant ($F(4, 4425)=2.113, p=0.077$).

Figure 12 compares the scores of the four CSNs for contracted stimuli with those of the original photographs. For male stimuli, the scores of Neuron-HAM, Neuron-HAF, and Neuron-UAF change significantly ($t(526) = 3.322, 13.231$ and -11.698 , respectively; all $p<0.00001$). Specifically, the contracted stimuli obtain lower Neuron-HAM scores than do the original photographs, which means the confidence that the image was of an attractive male decreased as distance between the eyes contracted. For the female stimuli, the scores of Neuron-HAF, Neuron-UAF, Neuron-HAM, and Neuron-UAM change significantly ($t(359) = 15.780, -6.953, -3.662$, and -4.892 respectively; all $p<0.00001$). Analogously, the scores of Neuron-HAF decrease and the scores of Neuron-UAF increase when the distances between eyes are contracted. These results demonstrate that FADNN recognizes a face as unattractive when its measures are less than the putative ratios of facial attractiveness.

Pallett et al. (2010) reported that human judgement of attractiveness decreases when the facial arrangement strays from the new golden ratios. The above experiments demonstrate a similar trend in the attractiveness determination of FADNN, although it was not explicitly trained using putative ratios. This result confirms that the putative ratios were encoded into FADNN. Moreover, the new golden ratios had a significantly stronger effect for female faces than male faces, which suggests that rules of the putative ratios impact judgement of faces of females more than those of males in FADNN. This result might have been obtained because of the dataset used in this study. Although some psychophysical studies of putative ratios with human observers have shown that the new golden ratios can be employed to judge both female and male faces (Bóo et al., 2013; Yoo et al., 2013), to the best of our knowledge, the difference in size effect between female and male face images has not been studied.

535 5. General discussion

Our study offers a demonstration that features learned by a DNN for facial attractiveness include the putative ratios [while the attractiveness category is given as a supervisory signal](#). In Experiment 1, we trained FADNN only using gender and attractiveness (2×2) categories, and it accurately classified im-
540 ages into these categories of facial attractiveness. FADNN used four CSNs as features of facial attractiveness and gender in each category. In Experiment 2, the CSNs projected the facially attractive features of different categories onto four face-like images. Through quantitatively analyses, we found that the generated images representing high attractiveness contain the putative ratios whereas
545 images representing low attractiveness do not. To validate this finding, we designed a psychophysical experiment for FADNN in Experiment 3. face images were distorted away from the new golden ratios and used as input to FADNN. The scores given by the four CSNs illustrate tendencies that are similar to those observed in human judgements of [facial attractiveness in Pallett et al. \(2010\)](#).
550 [These results of Experiments 2 and 3 suggest that the rules of putative ratios are contained in the latent distinctive features that represent facial attractive categories.](#)

FADNN learned the putative ratios similar to the human perception of attractiveness for two reasons. First, attractiveness was annotated by human
555 participants in the training images. The rules of the putative ratios are shared criteria for the perception of facial attractiveness across humans of different ages. They are a set of distinguishable features across categories of attractiveness. Second, the DNN has the capacity to learn abstract features from raw inputs (Nguyen et al., 2016). For example, Zeiler & Fergus (2014) found that,
560 in object recognition, the initial layers of the DNN capture low-level features (i.e. color, edges, and corners) whereas subsequent layers capture high-level object-related representations. In particular, the last layer captures the representation of the entire object. The rules of the putative ratios represent a high-level representation of facial attractiveness (Shen et al., 2016). Thus, it

565 is natural to argue that FADNN can learn high-level representations similar to these rules in the human perception of attractiveness.

Many factors that can be perceived in human faces, such as first impressions (Vernon et al., 2014), health (Russell et al., 2016), and sexual orientation (Wang & Kosinski, 2018), are otherwise difficult to measure. The neuronal interpretation (especially the generation of face-like images) and selections with
570 known links to such factors and psychological traits can provide an avenue for generating hypotheses that can be verified in experimental studies. For example, in Experiment 3, we found that the new golden ratios impact female face in images more than those of males in FADNN. However, there is no empirical evidence for this in human behavior, even with the new golden ratios employed for
575 both female and male face images (Bóo et al., 2013; Yoo et al., 2013). Further research is required to clarify this issue.

There is a limitation in the dataset. It is composed of non-standardized images and the (un)attractive categories are collected from two sources (celebrity
580 and mugshots) even though we excluded low-quality images before the training phase. These may introduce confounding factors to FADNN, because recent studies (Colón et al., 2020; Hill et al., 2019) showed that Deep Neural Networks retain certain features irrelevant to the tasks (annotations). Thus, the FADNN may learn some shortcut features, which could influence the results of
585 this study. For example, the quality of makeup in the face images of attractive females might have been different from those of unattractive females. To alleviate learning shortcut features from each source, we applied transfer learning on VGG-face model, specifically, fine tune FADNN based on the high-level representation of face identity. Then, we tested FADNN on the face images out
590 of the distribution of training data in Experiments 1 and 3, and importantly, a highly controlled dataset with varied putative ratios was used in Experiment 3. Under such complicated conditions, the results still suggest that FADNN has learned putative ratios (in Experiment 3) from the dataset. However, it doesn't mean that the attractive scores are solely determined by the putative
595 ratios. The other features that have been learned by FADNN require further

investigation.

It has been suggested that some machine learning models can implicitly capture basic human psychophysical characteristics, such as averageness, symmetry (Kagian et al., 2008), and sexual dimorphism (Said & Todorov, 2011). However, they are based on either handcrafted features or artificial stimuli. For example, the input facial stimuli in Said & Todorov (2011) were generated by specific features (50 dimensions for shape and 50 dimensions for reflectance), which were well controlled. In this paper, we showed that FADNN can learn the putative ratios without knowing any prior features annotated in the face images. This opens up opportunities for elucidating the origins of facial attractiveness, from basic perceptual mechanisms to high-level human perception, via a DNN-based perspective approach.

6. Acknowledgment

The authors thank Dr. Yuen Peng Loh (Multimedia University) and Dr. Yuyu Zhang (Kyoto University) for their great comments on the earlier version of this manuscript. We would also like to thank NVIDIA Corporation for providing the Tesla K40 for this research. Finally, we would like to thank Dr. Anh Nguyen (Auburn University) for kindly letting us use their software.

This work was supported in part by the Fundamental Research Funds of the Central Universities of China (No. JC1904), and in part by the China Scholarship Council under Grant No. 201608050017.

References

- Agnew, R. (1984). Appearance and delinquency. *Criminology*, 22, 421–440.
- Altwaijry, H., & Belongie, S. (2013). Relative ranking of facial attractiveness. In *2013 IEEE Workshop on Applications of Computer Vision* (pp. 117–124). IEEE.

- Atiyeh, B., & Hayek, S. (2008). Numeric expression of aesthetics and beauty. *Aesthetic Plastic Surgery*, *32*, 209–216.
- Barshan, E., Ghodsi, A., Azimifar, Z., & Jahromi, M. Z. (2011). Supervised
625 principal component analysis: Visualization, classification and regression on
subspaces and submanifolds. *Pattern Recognition*, *44*, 1357–1371.
- Bóo, F. L., Rossi, M. A., & Urzúa, S. S. (2013). The labor market return to an
attractive face: Evidence from a field experiment. *Economics Letters*, *118*,
170–172.
- 630 Borissavlievitch, M., & Hautecoer, L. (1958). The golden number and the sci-
entific aesthetics of architecture.
- Cadiou, C. F., Hong, H., Yamins, D. L., Pinto, N., Ardila, D., Solomon, E. A.,
Majaj, N. J., & DiCarlo, J. J. (2014). Deep neural networks rival the repre-
sentation of primate it cortex for core visual object recognition. *PLoS Comput*
635 *Biol*, *10*, e1003963.
- Cavior, H. E., Hayes, S. C., & Cavior, N. (1974). Physical attractiveness of
female offenders: Effects on institutional performance. *Correctional Psychol-*
ogist, *1*, 321–331.
- Cavior, N., & Howard, L. R. (1973). Facial attractiveness and juvenile delin-
640 quency among black and white offenders. *Journal of Abnormal Child Psy-*
chology, *1*, 202–213.
- Chang, C.-C., & Lin, C.-J. (2011). Libsvm: A library for support vector ma-
chines. *ACM Transactions on Intelligent Systems and Technology*, *2*, 27.
- Chen, F., Xiao, X., & Zhang, D. (2016). Data-driven facial beauty analy-
645 sis: prediction, retrieval and manipulation. *IEEE Transactions on Affective*
Computing, *9*, 205–216.
- Chen, F., & Zhang, D. (2014). Evaluation of the putative ratio rules for facial
beauty indexing. In *2014 International Conference on Medical Biometrics*
(pp. 181–188). IEEE.

- 650 Chen, Y., Mao, H., & Jin, L. (2010). A novel method for evaluating facial attractiveness. In *2010 International Conference on Audio, Language and Image Processing* (pp. 1382–1386). IEEE.
- Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A., & Oliva, A. (2016a). Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific Reports*, *6*, 27755.
- 655 Cichy, R. M., Pantazis, D., & Oliva, A. (2016b). Similarity-based fusion of meg and fmri reveals spatio-temporal dynamics in human cortex during visual object recognition. *Cerebral Cortex*, *26*, 3563–3579.
- 660 Colón, Y. I., Castillo, C. D., & O’Toole, A. J. (2020). Facial expression is retained in deep networks trained for face identification, .
- Conway, B. R. (2003). Colour vision: a clue to hue in v2. *Current Biology*, *13*, R308–R310.
- Cootes, T. F., Taylor, C. J., Cooper, D. H., & Graham, J. (1995). Active shape models-their training and application. *Computer Vision and Image Understanding*, *61*, 38–59.
- 665 Eishental, Y., Dror, G., & Ruppel, E. (2006). Facial attractiveness: Beauty and the machine. *Neural Computation*, *18*, 119–142.
- Erhan, D., Bengio, Y., Courville, A., & Vincent, P. (2009). Visualizing higher-layer features of a deep network. *University of Montreal*, *1341*, 1.
- 670 Farkas, L. G., & Kolar, J. C. (1987). Anthropometrics and art in the aesthetics of women’s faces. *Clinics in Plastic Surgery*, *14*, 599–616.
- Farkas, L. G., & Schendel, S. A. (1995). Anthropometry of the head and face. *American Journal of Orthodontics and Dentofacial Orthopedics*, *107*, 112–112.
- 675

- Feser, D. K., Gründl, M., Eisenmann-Klein, M., & Prantl, L. (2007). Attractiveness of eyebrow position and shape in females depends on the age of the beholder. *Aesthetic Plastic Surgery*, *31*, 154–160.
- Frieze, I. H., Olson, J. E., & Russell, J. (1991). Attractiveness and income for
680 men and women in management 1. *Journal of Applied Social Psychology*, *21*, 1039–1057.
- Geirhos, R., Jacobsen, J.-H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., & Wichmann, F. A. (2020). Shortcut learning in deep neural networks. *arXiv preprint arXiv:2004.07780*, .
- 685 Grant, E., Sahm, S., Zabihi, M., & van Gerven, M. (2016). Predicting and visualizing psychological attributions with a deep neural network. In *2016 23rd International Conference on Pattern Recognition (ICPR)* (pp. 1–6). IEEE.
- Gunes, H., & Piccardi, M. (2006). Assessing facial beauty through proportion analysis by image processing and supervised learning. *International Journal*
690 *of Human-computer Studies*, *64*, 1184–1199.
- Henderson, J. J., & Anglin, J. M. (2003). Facial attractiveness predicts longevity. *Evolution and Human Behavior*, *24*, 351–356.
- Hill, M. Q., Parde, C. J., Castillo, C. D., Colon, Y. I., Ranjan, R., Chen, J.-C., Blanz, V., & OâToole, A. J. (2019). Deep convolutional neural networks in
695 the face of caricature. *Nature Machine Intelligence*, *1*, 522–529.
- Holland, E. (2008). Marquardt’s phi mask: pitfalls of relying on fashion models and the golden ratio to describe a beautiful face. *Aesthetic Plastic Surgery*, *32*, 200–208.
- Hong, W., Chen, T.-S., & Chen, J. (2015). Reversible data hiding using delaunay
700 triangulation and selective embedment. *Information Sciences*, *308*, 140–154.
- Horikawa, T., & Kamitani, Y. (2017). Generic decoding of seen and imagined objects using hierarchical visual features. *Nature Communications*, *8*, 15037.

- Jayarathne, Y. S., Deutsch, C. K., McGrath, C. P., & Zwaenlen, R. A. (2012). Are neoclassical canons valid for southern chinese faces? *PloS One*, 7, e52593.
- 705 Jefferson, Y. (2004). Facial beauty-establishing a universal standard. *International Journal of Orthodontics*, 15, 9–26.
- Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., & Darrell, T. (2014). Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM International Conference on Multimedia* (pp. 675–678). ACM.
- 710 Jones, B. C., DeBruine, L. M., & Little, A. C. (2007). The role of symmetry in attraction to average faces. *Perception & Psychophysics*, 69, 1273–1277.
- Kagian, A., Dror, G., Leyvand, T., Meilijson, I., Cohen-Or, D., & Ruppim, E. (2008). A machine learning predictor of facial attractiveness revealing human-like psychophysical biases. *Vision Research*, 48, 235–243.
- 715 Kietzmann, T. C., Spoerer, C. J., Sörensen, L. K., Cichy, R. M., Hauk, O., & Kriegeskorte, N. (2019). Recurrence is required to capture the representational dynamics of the human visual system. *Proceedings of the National Academy of Sciences*, (p. 201905544).
- 720 Kriegeskorte, N. (2015). Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science*, 1, 417–446.
- Langlois, J. H., & Roggman, L. A. (1990). Attractive faces are only average. *Psychological Science*, 1, 115–121.
- 725 LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436.
- Li, J., Xiong, C., Liu, L., Shu, X., & Yan, S. (2015). Deep face beautification. In *Proceedings of the 23rd ACM International Conference on Multimedia* (pp. 793–794). ACM.

- 730 Liang, L., Lin, L., Jin, L., Xie, D., & Li, M. (2018). Scut-fbp5500: A diverse benchmark dataset for multi-paradigm facial beauty prediction. In *2018 24th International Conference on Pattern Recognition (ICPR)* (pp. 1598–1603). IEEE.
- 735 Liang, X., Tong, S., Kumada, T., & Iwaki, S. (2019). Golden ratio: The attributes of facial attractiveness learned by cnn. In *2019 International Conference on Image Processing*. IEEE.
- MacLin, O. H., & MacLin, M. K. (2004). The effect of criminality on face attractiveness, typicality, memorability and recognition. *North American Journal of Psychology*, 6, 145–154.
- 740 Mahendran, A., & Vedaldi, A. (2015). Understanding deep image representations by inverting them. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 5188–5196).
- Mahendran, A., & Vedaldi, A. (2016). Visualizing deep convolutional neural networks using natural pre-images. *International Journal of Computer Vision*, 120, 233–255.
- 745 Mao, H., Jin, L., & Du, M. (2009). Automatic classification of chinese female facial beauty using support vector machine. In *2009 IEEE International Conference on Systems, Man and Cybernetics* (pp. 4842–4846). IEEE.
- Mathias, M., Benenson, R., Pedersoli, M., & Van Gool, L. (2014). Face detection without bells and whistles. In *European Conference on Computer Vision* (pp. 720–735). Springer.
- 750 McCurrie, M., Beletti, F., Parzianello, L., Westendorp, A., Anthony, S., & Scheirer, W. J. (2018). Convolutional neural networks for subjective face attributes. *Image and Vision Computing*, 78, 14–25.
- Nguyen, A., Yosinski, J., & Clune, J. (2016). Multifaceted feature visualization: Uncovering the different types of features learned by each neuron in deep neural networks. *arXiv preprint arXiv:1602.03616*, .

- Nguyen, T. V., & Liu, L. (2017). Smart mirror: Intelligent makeup recommendation and synthesis. In *Proceedings of the 25th ACM International Conference on Multimedia* (pp. 1253–1254). ACM.
- 760 O’Toole, A. J., Castillo, C. D., Parde, C. J., Hill, M. Q., & Chellappa, R. (2018). Face space representations in deep convolutional neural networks. *Trends in cognitive sciences*, 22, 794–809.
- Pallett, P. M., Link, S., & Lee, K. (2010). New “golden” ratios for facial beauty. *Vision Research*, 50, 149–154.
- 765 Parde, C. J., Hu, Y., Castillo, C., Sankaranarayanan, S., & O’Toole, A. J. (2019). Social trait information in deep convolutional neural networks trained for face identification. *Cognitive science*, 43, e12729.
- Parkhi, O. M., Vedaldi, A., Zisserman, A. et al. (2015). Deep face recognition. In *Proceedings of the British Machine Vision Conference*.
- 770 Pashos, A., & Niemitz, C. (2003). Results of an explorative empirical study on human mating in germany: Handsome men, not high-status men, succeed in courtship. *Anthropologischer Anzeiger*, (pp. 331–341).
- Perrett, D. I., Lee, K. J., Penton-Voak, I., Rowland, D., Yoshikawa, S., Burt, D. M., Henzi, S., Castles, D. L., & Akamatsu, S. (1998). Effects of sexual
775 dimorphism on facial attractiveness. *Nature*, 394, 884.
- Perrett, D. I., May, K. A., & Yoshikawa, S. (1994). Facial shape and judgements of female attractiveness. *Nature*, 368, 239.
- Rhodes, G. (2006). The evolutionary psychology of facial beauty. *Annual Review of Psychology*, 57, 199–226.
- 780 Rhodes, G., Sumich, A., & Byatt, G. (1999). Are average facial configurations attractive only because of their symmetry? *Psychological Science*, 10, 52–58.

- Rothe, R., Timofte, R., & Van Gool, L. (2016). Some like it hot-visual guidance for preference prediction. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition* (pp. 5553–5561).
- 785 Rubenstein, A. J., Langlois, J. H., & Roggman, L. A. (2002). What makes a face attractive and why: The role of averageness in defining facial beauty. In G. Rhodes, & L. A. Zebrowitz (Eds.), *Advances in visual cognition, Vol. 1. Facial attractiveness: Evolutionary, cognitive, and social perspectives*. Westport, CT, US: Ablex Publishing.
- 790 Russell, R., Porcheron, A., Sweda, J. R., Jones, A. L., Mauger, E., & Morizot, F. (2016). Facial contrast is a cue for perceiving health from the face. *Journal of Experimental Psychology: Human Perception and Performance*, 42, 1354.
- Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., & Pantic, M. (2013). 300 faces in-the-wild challenge: The first facial landmark localization challenge. In
795 *Proceedings of the IEEE International Conference on Computer Vision Workshops* (pp. 397–403).
- Said, C. P., & Todorov, A. (2011). A statistical model of facial attractiveness. *Psychological Science*, 22, 1183–1190.
- Schmid, K., Marx, D., & Samal, A. (2008). Computation of a face attractiveness
800 index based on neoclassical canons, symmetry, and golden ratios. *Pattern Recognition*, 41, 2710–2717.
- Seeliger, K., Fritsche, M., Güçlü, U., Schoenmakers, S., Schoffelen, J.-M., Bosch, S., & van Gerven, M. (2018). Convolutional neural network-based encoding and decoding of visual object recognition in space and time. *NeuroImage*,
805 180, 253–266.
- Shen, H., Chau, D. K., Su, J., Zeng, L.-L., Jiang, W., He, J., Fan, J., & Hu, D. (2016). Brain responses to facial attractiveness induced by facial proportions: evidence from an fmri study. *Scientific Reports*, 6, 35905.

- Shen, L., & Bai, L. (2006). A review on gabor wavelets for face recognition.
810 *Pattern Analysis and Applications*, 9, 273–292.
- Shi, S., Gao, F., Meng, X., Xu, X., & Zhu, J. (2019). Improving facial attractiveness prediction via co-attention learning. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 4045–4049). IEEE.
- 815 Tarrés, F. (2012). Gtav face database. <http://gps-tsc.upc.es/GTAV/ResearchAreas/UPCFaceDatabase/GTAVFaceDatabase.html>,
.
- Tong, S., Liang, X., Kumada, T., Iwaki, S., & Tosa, N. (2017). Learning the cultural consistent facial aesthetics by convolutional neural network. In *2017*
820 *International Conference on Culture and Computing (Culture and Computing)* (pp. 97–103). IEEE.
- Valenzano, D. R., Mennucci, A., Tartarelli, G., & Cellerino, A. (2006). Shape analysis of female facial attractiveness. *Vision Research*, 46, 1282–1291.
- Valla, J. M., Ceci, S. J., & Williams, W. M. (2011). The accuracy of inferences
825 about criminality based on facial appearance. *Journal of Social, Evolutionary, and Cultural Psychology*, 5, 66.
- VanRullen, R. (2017). Perception science in the age of deep neural networks. *Frontiers in Psychology*, 8, 142.
- Vernon, R. J., Sutherland, C. A., Young, A. W., & Hartley, T. (2014). Modeling first impressions from highly variable facial images. *Proceedings of the*
830 *National Academy of Sciences*, 111, E3353–E3361.
- Wang, P., & Cottrell, G. W. (2017). Central and peripheral vision for scene recognition: a neurocomputational modeling exploration. *Journal of Vision*, 17, 9–9.

- 835 Wang, S., Shao, M., & Fu, Y. (2014). Attractive or not?: Beauty prediction with attractiveness-aware encoders and robust late fusion. In *Proceedings of the 22nd ACM international conference on Multimedia* (pp. 805–808). ACM.
- Wang, Y., & Kosinski, M. (2018). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of*
840 *Personality and Social Psychology*, 114, 246.
- Xu, L., Fan, H., & Xiang, J. (2019). Hierarchical multi-task network for race, gender and facial attractiveness recognition. In *2019 IEEE International Conference on Image Processing (ICIP)* (pp. 3861–3865). IEEE.
- Yamins, D. L., & DiCarlo, J. J. (2016). Using goal-driven deep learning models
845 to understand sensory cortex. *Nature neuroscience*, 19, 356.
- Yang, M., Zhang, L., Shiu, S. C., & Zhang, D. (2013). Gabor feature based robust representation and classification for face recognition with gabor occlusion dictionary. *Pattern Recognition*, 46, 1865–1878.
- Yoo, J.-Y., Kim, J.-N., Shin, K.-J., Kim, S.-H., Choi, H.-G., Jeon, H.-S., Koh,
850 K.-S., & Song, W.-C. (2013). Centralization or decentralization of facial structures in korean young adults. *Journal of Craniofacial Surgery*, 24, 1007–1010.
- Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., & Lipson, H. (2015). Understanding neural networks through deep visualization. *arXiv preprint*
855 *arXiv:1506.06579*, .
- Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. In *European Conference on Computer Vision* (pp. 818–833). Springer.
- Zhang, L., Zhang, D., Sun, M.-M., & Chen, F.-M. (2017). Facial beauty analysis
860 based on geometric feature: Toward attractiveness assessment application. *Expert Systems with Applications*, 82, 252–265.

- Zhao, J., Zhang, M., He, C., & Zuo, K. (2019). Data-driven research on the matching degree of eyes, eyebrows and face shapes. *Frontiers in Psychology*, 10.
- 865 Zou, J., Ji, Q., & Nagy, G. (2007). A comparative study of local matching approach for face recognition. *IEEE Transactions on Image Processing*, 16, 2617–2628.

7. Appendices

7.1. Attractiveness Rating

Figure 7: Attractiveness ratings of images in the constructed dataset. NS: ****: $p < 0.00001$.

The DW and DB are defined as follows:

$$\begin{aligned} DW &= \frac{\sum_{j=1}^J \sum_{k=1}^K (\bar{d}(x_{j,k}, c))}{J \times K}, \\ DB &= \frac{\sum_{j=1}^J d(\bar{x}_j, c)}{J}, \end{aligned} \tag{2}$$

where $x_{j,k}$ represents the k -th instance in the j -th class, $\bar{d}(x_{j,k}, c)$ represents the average Euclidean distance between instance $x_{j,k}$ and all other $k - 1$ instances in the j -th class, and $d(\bar{x}_j, c)$ represents the mean Euclidean distance between the center of the j -th class and the center of the other $j - 1$ classes.

875 *7.3. Traditional feature-based models for facial attractiveness*

Geometric features are defined by a set of landmarks on the face (Zhang et al., 2017) that include shape information. The face landmarks were automatically extracted using the active shape model (Cootes et al., 1995; Jones et al., 2007; Sagonas et al., 2013). Figure 10 shows the output of the active shape
880 model in a graphical format. Here, the geometrical features are generated from the extracted landmarks, and are composed of 2,346 pairwise distances. A support vector machine (SVM) is then used as classifier to distinguish images using these extracted features (Mao et al., 2009). The SVM is a robust classifier that constructs an n-dimensional hyperplane to optimally separate the data into dif-
885 ferent categories. The radial-basis function (RBF) kernel (Chang & Lin, 2011) was used as well, $K(x, y) = \exp(-\gamma \|x - y\|)$, $\gamma > 0$, where γ is a kernel parameter. Gabor features are appearance features that encode facial shape and texture information over a range of spatial scales (Yang et al., 2013) and are formed by convolution of the image with a family of Gabor kernels. These ker-
890 nels have five scales and eight orientations with a size of 39×39 pixels (Shen & Bai, 2006; Zou et al., 2007). The results of convolution are vectorized as a 2,560-D vector (Yang et al., 2013) and classified by the SVM.

7.4. Testing FADNN on SCUT-FBP5500

Figure 8: Distribution of the facial features of the SCUT-FBP5500 dataset. Blue dots represent the 75 female images with the highest attractiveness scores (top 10%); green dots represent the 75 female images with the lowest attractiveness scores (bottom 10%); yellow dots represent the 75 male images with the highest attractiveness scores (top 10%); and red dots represent the 75 male images with the lowest attractiveness scores (bottom 10%).

Figure 9: Distribution of the attractiveness scores of the SCUT-FBP5500 dataset. The left figure illustrates the distribution of the attractiveness scores (3,000 face images). The right figure illustrates the distribution of the attractiveness scores after re-balancing (353 face images).

7.5. Active Shape Model

Figure 10: Graphical illustration of the facial landmarks (green dots, $n=68$) and their indices produced by the active shape model. (The sample image, used for illustration only, is of T.S. and is presented with his full consent.)

Figure 11: Examples of original photographs and contracted stimuli.

Figure 12: Scores of the four CSNs obtained by testing the contracted stimuli compared with the original photographs. The left figure illustrates the neuron scores of 527 male photographs and 527 contracted male stimuli. The right figure illustrates the neuron scores of 360 female photographs and 360 contracted female stimuli. NS: $p > 0.05$; ****: $p < 0.00001$.