



OPEN

SUBJECT AREAS:

PSYCHOLOGY

COMPUTATIONAL
NEUROSCIENCE

COGNITIVE NEUROSCIENCE

Own-race and own-species advantages in face perception: a computational view

Christoph D. Dahl³, Chien-Chung Chen³ & Malte J. Rasch^{1,2}Received
30 June 2014Accepted
10 September 2014Published
17 October 2014

Correspondence and requests for materials should be addressed to C.D.D. (christoph.d.dahl@gmail.com) or M.J.R. (malte.rasch@bnu.edu.cn)

¹State Key Laboratory of Cognitive Neuroscience and Learning & IDG/McGovern Institute for Brain Research, Beijing Normal University, China, ²Center for Collaboration and Innovation in Brain and Learning Sciences, Beijing Normal University, China, ³Department of Psychology, National Taiwan University, Roosevelt Road, Taipei, Taiwan (ROC).

The frequency to which an organism is exposed to a particular type of face influences recognition performance. For example, Asians are better in individuating Asian than Caucasian faces, known as the own-race advantage. Similarly, humans in general are better in individuating human than monkey faces, known as the own-species advantage. It is an open question whether the underlying mechanisms causing these effects are similar. We hypothesize that these processes are governed by neural plasticity of the face discrimination system to retain optimal discrimination performance in its environment. Using common face features derived from a set of images from various face classes, we show that maximizing the feature variance between different individuals while ensuring minimal variance within individuals achieved good discrimination performances on own-class faces when selecting a subset of feature dimensions. Further, the selected subset of features does not necessarily lead to an optimal performance on the other class of faces. Thus, the face discrimination system continuously re-optimizes its space constraint face representation to optimize recognition performance on the current distribution of faces in its environment. This model can account for both, the own-race and own-species advantages. We name this approach Space Constraint Optimized Representational Embedding (SCORE).

Humans have more difficulty to identify individuals of a different race than their own¹. For example, to Asians with little exposure to Caucasians all Caucasian faces look alike. A large body of scientific evidence exists describing discrimination advantages for the own face class as opposed to the other face classes, an effect commonly known as the own-race advantage of face processing (also referred to as the other-race effect)^{2–5}, for an overview see^{6,7}. Similarly, visual recognition performance of humans^{8,9}, apes^{10–12} and monkeys^{8,13–17} is biased toward their own as opposed to another species' faces, referred to as the own-species advantage (or the other-species effect)⁶. Inferior recognition performance of other faces arises at a subjective level ('all faces look alike'), and can also be objectively measured in terms of decreased accuracy and increased response latency. A computational explanation for the own-race advantage has been up for debate¹⁸, while the nature of the own-species advantage and its relationship to the own-race advantage have been unaddressed.

According to O'Toole¹⁹, a few basic elements characterize the core concept of a computational model under the assumption that experience fundamentally shapes the representations of faces. In an attempt to build a unified model for both the own-race and the own-species effects, we here take into account the following assumptions (partly cited from¹⁹): (1) Face of different races [and species] comprise different statistical categories of faces. (2) Within a given category of faces, a set of differentially weighted "features" is optimal for encoding faces in a manner that makes faces within the category most discriminable. (3) The size and number of features, which can be simultaneously employed for discriminating faces, is limited by neural resources. (4) The face feature representation slowly learns and optimizes to the combined distribution of exposed faces experienced over a longer time frame. Thus, with exposure to many faces of a given race [or species] and a smaller number of faces of other races [or species], perceptual learning enables observers to make optimal use of the features that are best for processing faces from the category with which they have had the most experience, typically faces of their own race [and species]. Hence, we explore in this study to what extent the difficulties in perceiving other races and species is reflected in the deviation of selecting the optimal features for distinguishing faces of the own race and species from the selection of other face classes, as it has been done previously with faces of different races²⁰. This further sheds light not only on whether the same underlying computational principle can explain both the own-race and the



own-species advantages, but also to what extent a system is capable to classify other-class when tuned to own-class face exemplars.

As proposed in an earlier model¹⁰, we assume that the neural machinery for face discrimination has access to complex non-linear face features, representing the facial features as extracted from high-dimensional face space by means of sensory processing. While in this previous model the facial features are artificially generated normal distributions of feature samples, we here extract facial features using topological methods²¹ and test the model predictions under the scenarios of the own-race and own-species advantages. We here show that under assumption that the visual system is generically plastic and under consideration of the exposure history, the own-race and own-species advantages can be explained with the same underlying computational mechanism. We thus conclude that face perception underlies a life-long learning process that continuously optimizes the recognition system to the changing environment.

Results

We used four sets of face pictures. The first was a collection of 66 chimpanzee photographs from 22 individuals and the second set contained 54 photographs from 18 human (Asian) individuals. For each individual we used three viewpoints (first and second sets: face database 1, see methods). The third and fourth sets consisted of rendered images from 3-D reconstructions of faces²². These face images were of Caucasian and Asian races and contained 60 individuals from three viewpoints (third and fourth sets: face database 2, see methods). In the following we use face database 1 to evaluate a learning process toward another species and face database 2 to evaluate a similar effect toward another race. From all face data sets we extracted face features by implementing the topological methods suggested by²¹. In Figures 1A and H typical face images are shown together with the extracted features. Before feature extraction, faces were vertically aligned. The features were computed on the face image and then averaged arbitrarily along the horizontal dimension to enhance the processing efficacy (Figures 1B, I), however, the exact type of the face features used is not crucial for our method. Note that eye, nose, and mouth regions are captured by characteristic feature shapes.

We first tested whether the distribution of features in both were indeed different between classes of faces (Figures 1D, K), as stated in the assumptions above. A simple (and sufficient) way to show this is by comparing the projections onto the first principle components of the both data samples against each other (Figures 1C, J). It can be seen clearly in the case of human and chimpanzee faces (species comparison) that both data samples spread out toward different directions (Figure 1C), indicating that the feature distributions of chimpanzees and humans are indeed different at the principle component explaining the greatest variance of feature distribution. A closer look at the average profiles reveals that humans and chimpanzees show clear morphological differences in their facial feature characteristics (Figure 1D). On the other hand, the differences across races (Asians and Caucasians) are rather subtle (Figure 1K). However, when using a more elaborated statistical method for testing whether two multi-dimensional distributions are significantly different (Maximum Mean Discrepancy²³), we indeed found that both the features derived from Caucasians versus Asians as well as Asians versus Chimpanzees are significantly different (alpha level .001), in agreement with our initial assumptions.

Further, the neural system has to optimize its representation to achieve best discrimination of the type of faces it was already exposed to ('Optimized Representational Embedding'). Although the exact neural implementations are of course not known, a reasonable way would be to maximize the discrimination power of individuals while taking the limitation of neural resources into account ('Space Constraint'). A simple way to achieve this is by selecting those potential feature dimensions which reduce the feature variance for face

images belonging to the same individuals (within-individual variance) and simultaneously maximizing the variance between individual faces (between-individual variance). This approach is known as the Fisher's Linear Discriminant (LFD) (hence "Fisher-faces") and was shown to be superior to an "Eigenface" method (using the first few principle components for dimension reduction instead)²⁴. We name the general approach proposed here 'Space Constraint Optimized Representational Embedding' (SCORE). In Figures 1E, F, L and M both data sets are projected onto the first two components using LFD. One notes that images of the same faces/identities (same colors) tend to group together, potentially easing the discrimination (in respect to other projections). Computing the distances of face exemplars of the same identities ('within') and of exemplars of various identities ('across'), we found a significantly closer average distance for 'within' than 'across' in the species comparison (chimpanzee versus human faces) (*within-across*: $F(1,279) = 4320, p < .001$; interaction *within-across* x *classes*: $F(1,279) = 259, p < .001$; chimpanzee: *within*: $.09 \pm .02$, *across*: $.19 \pm .01$ (mean \pm std), Figure 1G, left panel; human: *within*: $.05 \pm .02$, *across*: $.21 \pm .01$ (mean \pm std), Figure 1G, right panel), as well as for the race comparison (Asian versus Caucasian faces) (*within-across*: $F(1,719) = 62591, p < .001$; interaction *within-across* x *classes*: $F(1,719) = 48.56, p < .001$; Asian: *within*: $.03 \pm .01$, *across*: $.20 \pm .001$ (mean \pm std), Figure 1N, left panel; Caucasian: *within*: $.03 \pm .01$, *across*: $.19 \pm .001$ (mean \pm std), Figure 1N, right panel).

We next evaluated the discrimination performances the SCORE system would achieve using a discrimination task. We divided each data set into train and test sets, computed the LFD projection only on the train set and selected the optimal features. The ratio of faces stemming from the two classes of faces was constantly varied in the train set to simulate various levels of experience with one and the other class of faces. This process simulates long-term exposure and adaption to a particular predominant face class in a given system (Asians, Caucasians, humans, chimpanzees), where the features space is optimized to the discrimination of behaviorally relevant faces. To evaluate discrimination performance on the test set we draw – analogous to a real experiment^{10,25} – two faces representing two individuals, A and B, together with a third face image depicting one of the first two individuals. The hypothetical subject then had to report whether A or B was depicted on the third image. This judgment was done on the basis of the similarity in the face feature representation in the hypothetical neural system. In particular, all three pictures were projected into the feature space (using the learned weights from above) and if the distance in the feature space between the first and third picture was nearer than between the second and the third pictures, individual A was reported or otherwise B. In this manner the expected performance could be computed on the whole test set. This analysis revealed that indeed testing and training on two different face classes reduced the system performances as opposed to testing and training on the same classes (Figures 2A, D), reproducing both the own-race and own-species effect, respectively.

We next quantified the size of the effect when using SCORE. We define two factors: (1) the *testing class*, which is either identical with the training class or different, and analogously (2) the *training class*. We found own-race advantages for both Asian and Caucasian faces above the other face class (Caucasian faces for a system tuned to Asian faces and vice-versa), reflected in a significant effect of the factor *testing class* $F(1,399) = 234, p < .001$ and a non-significant effect in the factor *training class*: $F(1,399) = 1.07, p = .30$ (Figure 2D). Further we show that the performance clearly depends on the exposure history, reflected in an increasing discrimination rate with increasing own-face training ratio in both Asian (Figure 2E) and Caucasian faces (Figure 2F). Similarly, we found an own-species advantage for both human and chimpanzee faces above the other face classes (chimpanzee faces for a system tuned

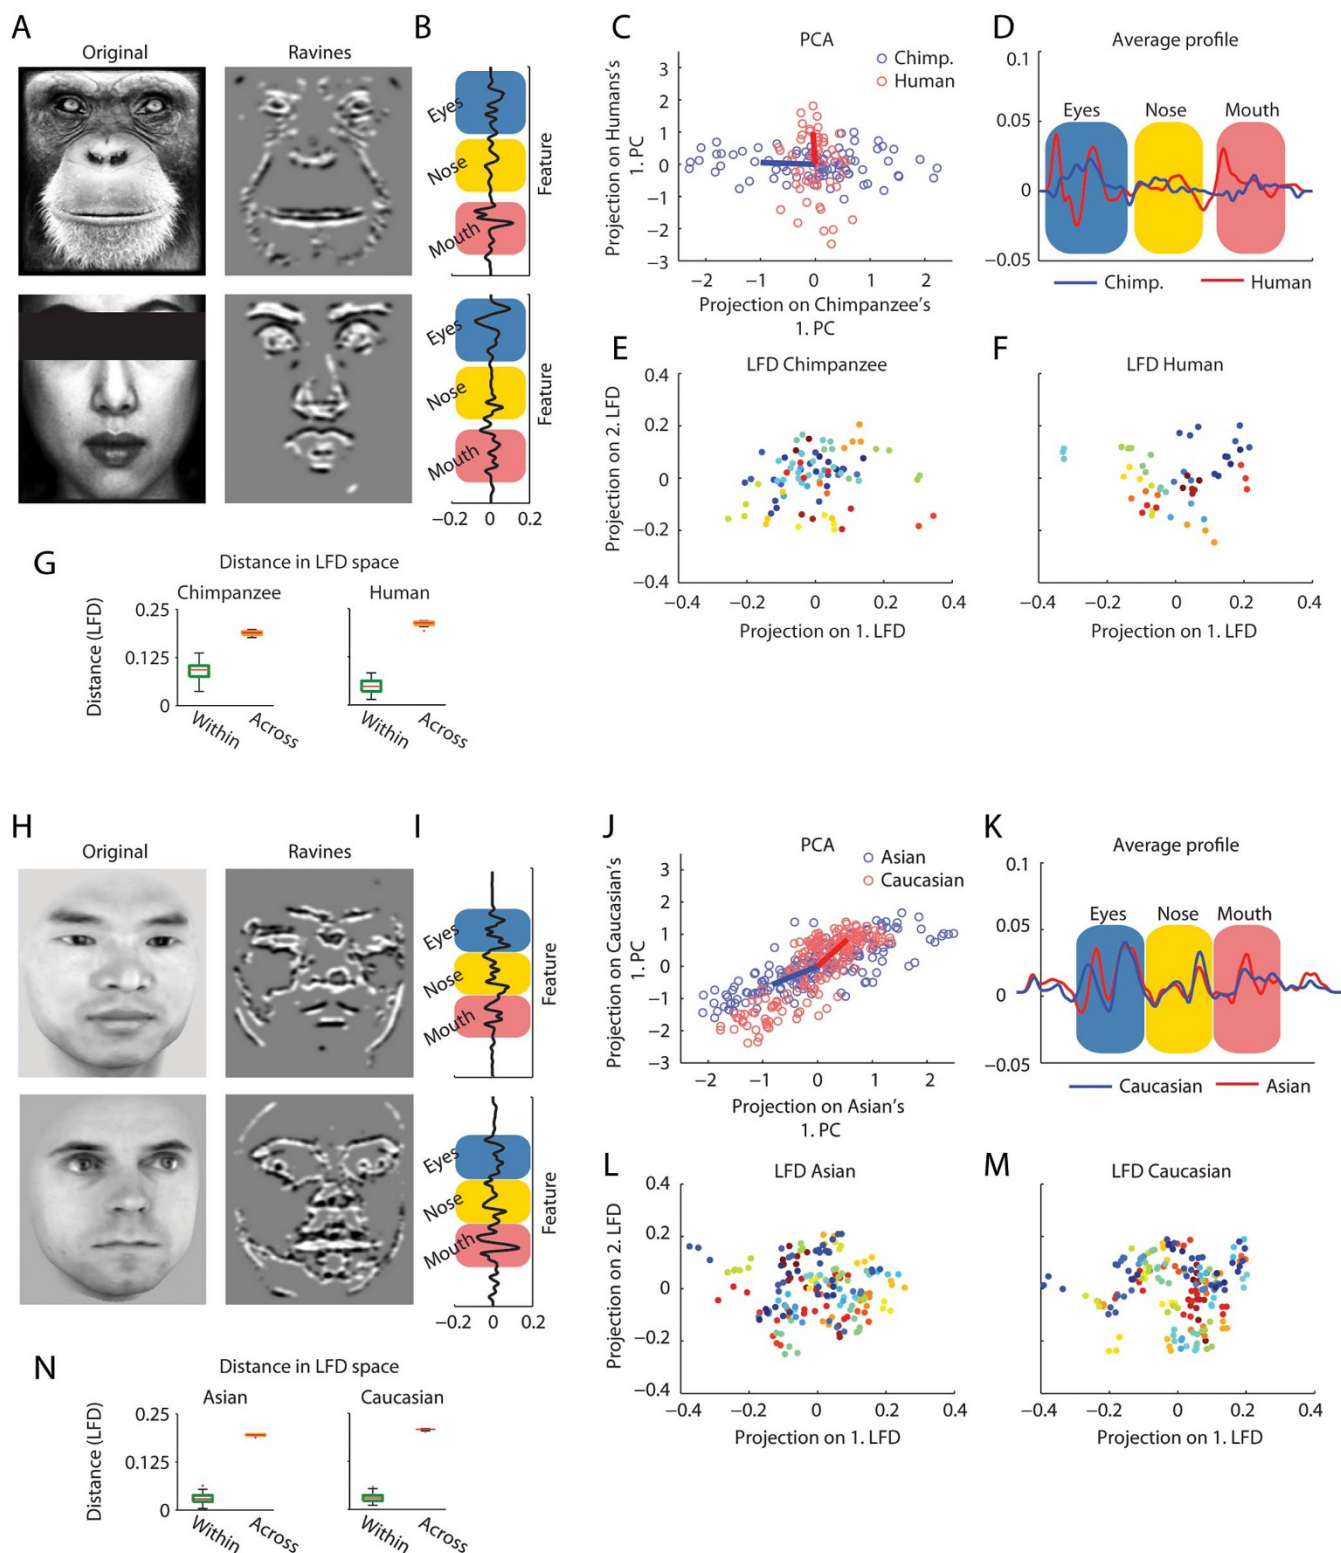


Figure 1 | Species and race comparisons. (A,H). Exemplar faces. The chimpanzee face was contributed by the Great Ape Research Institute (GARI), the human faces were contributed by Heinrich Bulthoff and the ScanLab at the Max Planck Institute for Biological Cybernetics, Tübingen, Germany (see Methods and Acknowledgements for details). (A,B,H,I). Feature extraction using topological features. (A,H). The original faces (left column) and the extracted features (right column). (B,I). The horizontally average profile. (C–G, J–N.) Projections in features space. (C,J). Principle components of human and chimpanzee (C) and Asian and Caucasian faces (J). Distributions are different, thus a system has to adapt if either of the classes are new to the system. (D,K.) Means (extracted in PCA) are also different. (E, F, L, M.) Chimpanzee and human (E, F) and Asian and Caucasian faces (L, M) projected with LFD. Note that individuals are grouped together (same color). (G, N.) Distance in LFD space for chimpanzee and human (G) and Asian and Caucasian faces (N).

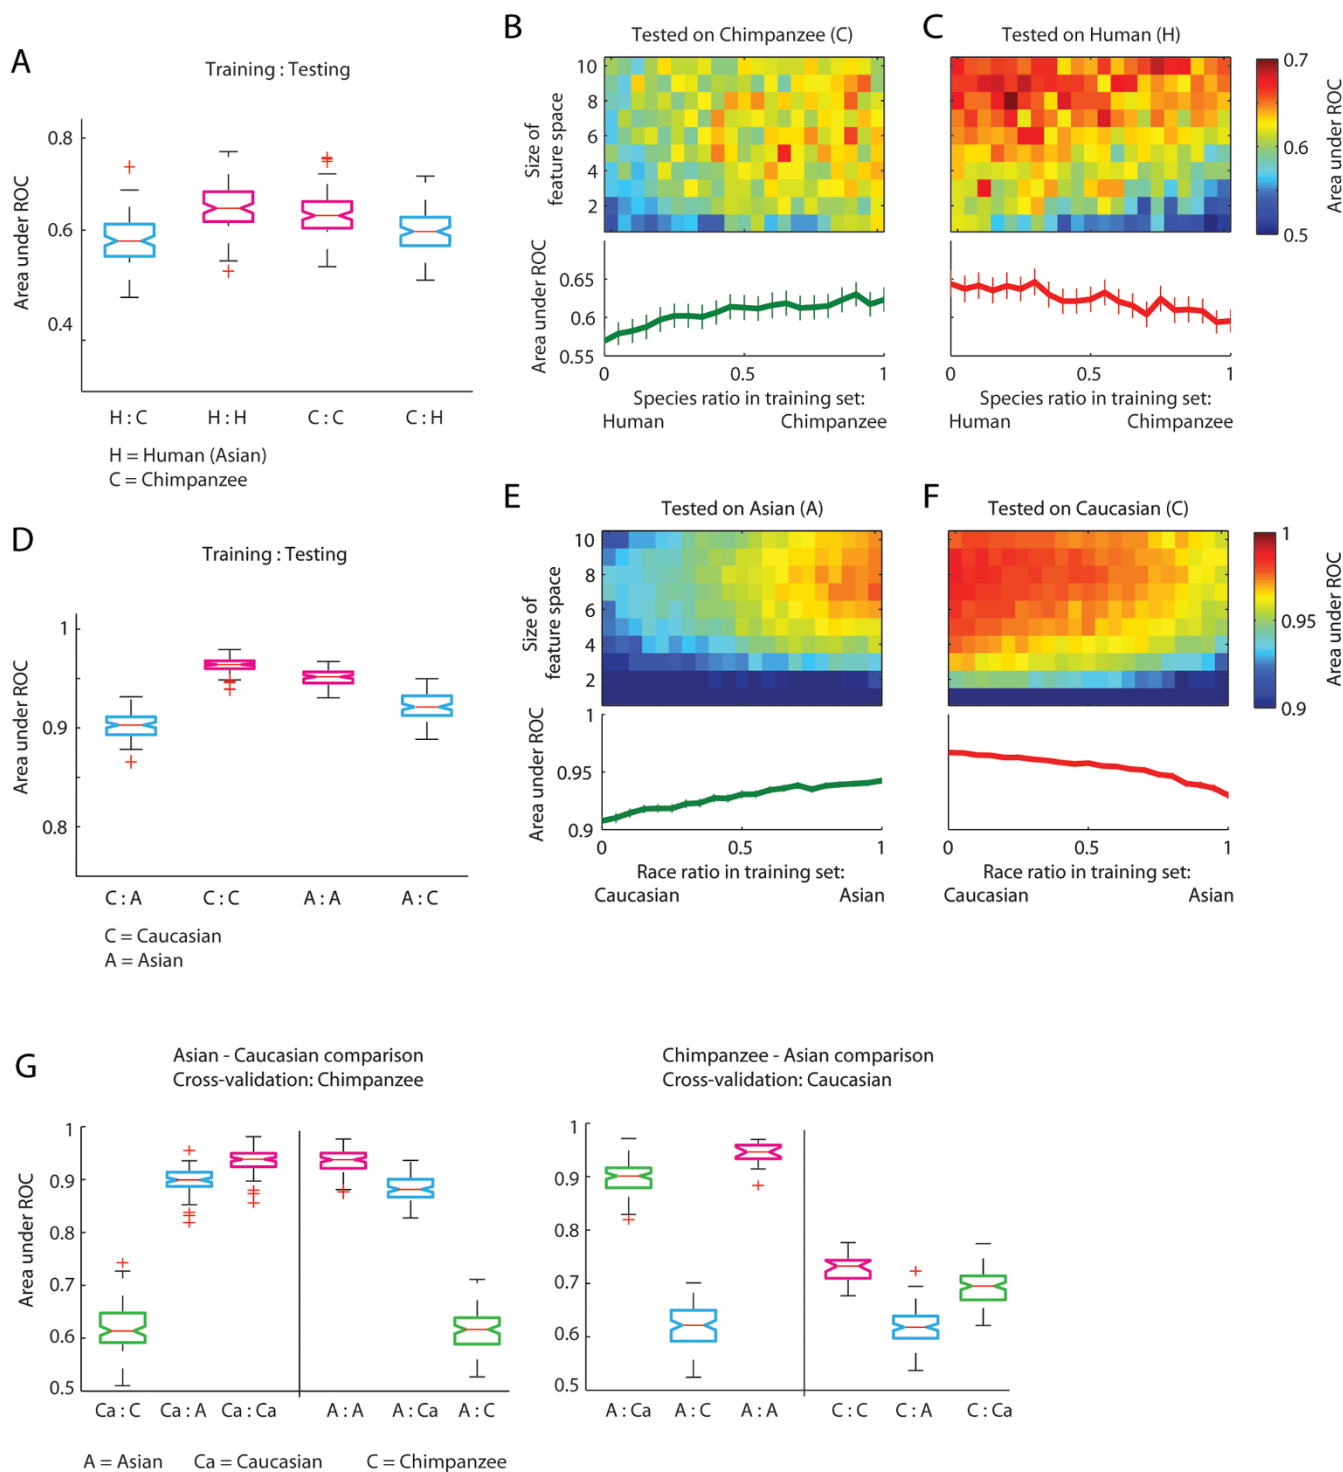


Figure 2 | Discrimination performances. (A,D.) The performances on the human and the chimpanzee (A) as well as the Asian and Caucasian face data sets (D) are plotted. Notably, as expected, if the training set contained examples from only one species (A), the performance tested on the other species degraded. A similar effect was observed for races (D). (B,C.) Change of the species ratio in the train set (from 1 (only chimpanzees) to 0 (only humans)) versus number of used dimensions in the feature space. Performances in color-code. (B.) Tested on chimpanzee faces, (C.) Tested on human faces. Note that performance changes with the species ratio in the train set signifying the adaption of the recognition system to a change in the input distribution. “Old” faces get partly forgotten by this process (i.e. trained on humans, tested on chimpanzees is worse than trained on chimpanzees and tested on chimpanzees). Accordingly, figures (E,F) show the change of race ratio in the train set versus number of used dimensions in the feature space. (G.) Cross-validation: The performances on the race comparison (left panel) as well as the species comparison (right panel) are plotted. In addition, performances on second other-class faces are shown in green boxplots. Notably, face samples come from both face databases (see methods) and y-axis label of (A, D and G) do not correspond.

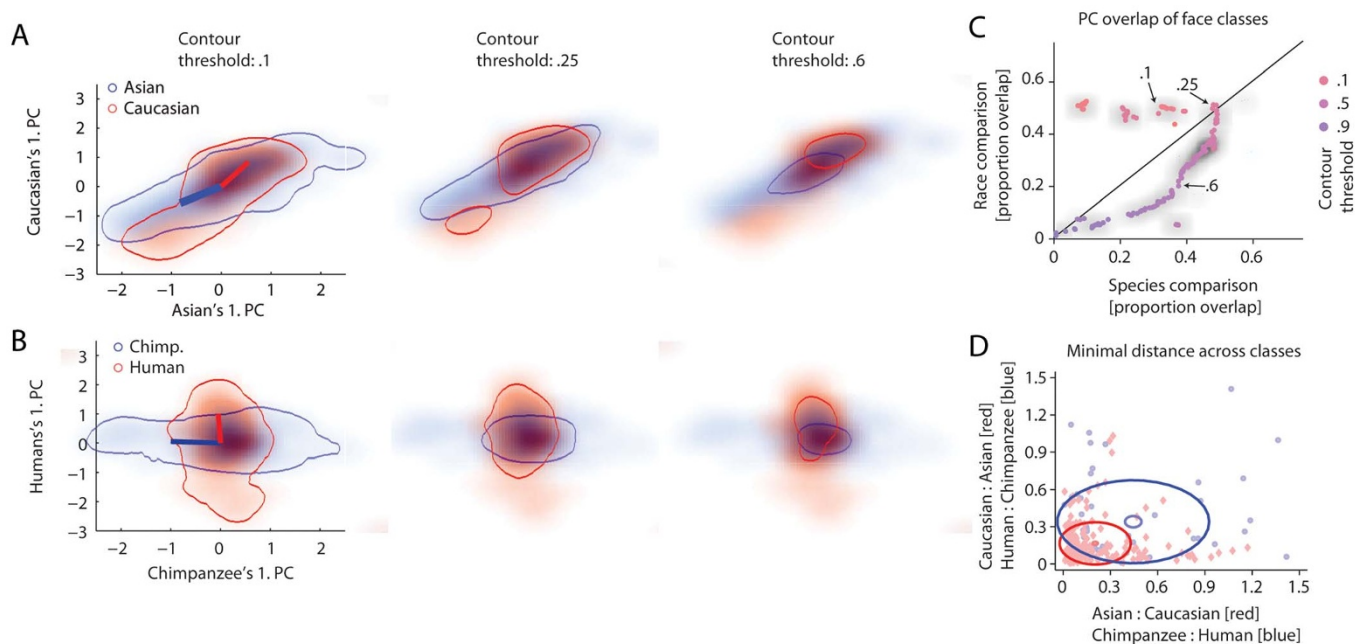


Figure 3 | (A–D.) Feature transfer effect. (A.) Asian and Caucasian face feature distributions (1st PCs) are compared by calculating the proportional overlap while varying the contour threshold of the distributions. (B.) Same as A with chimpanzee and human (Asian) faces. (C.) Direct comparison of PC overlap of face classes between races and species. The proportional overlap of both the species and race comparisons are plotted against each other. In a red-blue color graduation the contour threshold is indicated, ranging from .05 to .95. (D.) Minimal distance across classes. The minimal distance was calculated across classes for both the species and the race comparison in both directions (from class 1 to 2 and from class 2 to 1, see x and y axes). Notably, the minimal distance for the species comparison (blue) was larger than for the race comparison (red).

to human faces and vice-versa) reflected in a significant effect of the factor *testing class* $F(1,399) = 9.7, p < .01$ and a non-significant effect in the factor *training class*: $F(1,399) = 1.96, p = .16$ (Figure 2A). Similarly to the race comparison, the simulated recognition performance of faces from different species is critically shaped by the exposure history: We found an increase in performance scores with increasing ratio of own-species exposure (Figures 2B, C). When relaxing the space constraint of the neural system and allowing for more feature dimensions to be present simultaneously, we see performance increases with increasing size of the face feature space for both, the own-race and the own-species advantages. Performance saturates around eight feature dimensions, however, the exact number necessary to reach the saturation level will likely depend on the details of the implementation, e.g. the chosen face feature extraction method and the number of represented faces.

Further, we cross-validated our model to demonstrate its efficacy: The idea is to show that when the system is tuned to one particular class (e.g. Asian faces) it not only exhibits a discrimination advantage of own-race above other-race faces (e.g. Caucasian faces), but also above other-species (chimpanzee) faces (factor *testing class*). Moreover, we expect that with changing training ratio, e.g. a shift from a predominant exposure to Asian faces to a predominant exposure to Caucasian faces, discrimination of the other-species face class (chimpanzee faces) remains unaffected. To test this, we compared the factor *testing class*, here reflected as own-race (1), other-race (2) and other-species (3) (corresponding to the color code in Figure 2G left panel) and analogously the factor *training class* (right versus left boxplots in Figure 2G, left panel). We found significant effects for *testing class* ($F(2,599) = 6568, p < .001$) and the factor *training class* ($F(1,599) = 6.73, p < .01$; Figure 2G, left panel). Analogously, when the system is tuned to own-species (chimpanzee) faces, the discrimination performances of human face classes should deteriorate (right versus left boxplots in Figure 2G, right panel). We found an own-species advantage above both other-species face classes. In addition, with changing training ratio from a predominant exposure to chim-

panzee to Asian faces we found increased discrimination performance for Asian faces and a decreased performance for chimpanzee faces, analogous to the above-mentioned case. Importantly, the second other-species face class (Caucasian faces) remained unaffected by the introduced change in exposure (chimpanzee to Asian). Statistically, this is reflected in a significant effect of the factor *testing class* ($F(2,599) = 584, p < .001$) and a significant effect in the factor *training class* ($F(1,599) = 480, p < .001$; Figure 2G, left panel).

To further quantify the computational origins of the observed own-race and own-species advantages, we evaluated the overlap of the feature distributions for the first few principle components. If the overlap is small, distributions are very different and we expect large own-race and own-species advantages because learned features do not generalize to the other class. We quantified the surface shared between the two distributions: While the overlap of chimpanzee and human facial features remains relatively small (32.5%) at a contour threshold of .1 (due to the perpendicular arrangement of component axes) (Figure 3B, left panel), the overlap of Asian and Caucasian facial features is larger (49.7%) (Figure 3A, left panel). Interestingly, when increasing the contour threshold, and hence carving out the core of the feature distributions, the quantity of overlap equals out between Chimpanzee – Human and Asian – Caucasian comparisons (contour threshold .25, species comparison: 47.3%; race comparison: 50%, Figure 3A,B, middle panels). To estimate the amount to which a system could use features of one face class to classify examples of the other face class, we estimated the minimal distance of each sample to each other sample of the other class (Figure 3D). We found that chimpanzee-human and human-chimpanzee distances are larger than Asian-Caucasian and Caucasian-Asian distances ($t(498) = 7.74, p < .001$; race distances = $.18 \pm .20$, species distances = $.40 \pm .43$ (mean \pm std)), reflecting a higher degree of possible feature transfer between Asian and Caucasian than chimpanzee and human (Asian) faces. In other words, the own-species advantage goes along with a large deviation of the first principal components (angle 84.5 degree) (Figure 3B, left



panel), leading to less feature transfer between chimpanzee and human faces (small overlap at large contour threshold; Figure 3B, left panel, Figure 3C) and hence to a large own-species advantage (Figure 2A–C). In contrast, the own-race advantage is accompanied by a small deviation of the first principal components (angle 22.6 degree) (Figure 3A, left panel), leading to more feature transfer between Asian and Caucasian faces (large overlap at large contour threshold; Figure 3A, left panel) and hence to a small own-race advantage (Figure 2D–F). In addition, we estimated the relative effect sizes of the race and species comparisons by the following formula: $(\text{scores}(\text{training class 1, testing class 1}) - \text{scores}(\text{training class 1, testing class 2})) / (\text{scores}(\text{training class 1, testing class 1}) + \text{scores}(\text{training class 1, testing class 2}))$. It turns out that the species comparison yielded a greater relative effect size ($H:C = 5.74$, $C:H = 4.48$) than the race comparisons ($C:A = 3.43$, $A:C = .79$).

We further quantified to what extent the SCORE approach explained the so-called Mirror Effect²⁶ (Figure 4A,B), i.e. quicker categorical decisions for other-race/species faces as opposed to own-race/species faces. We hypothesize that the categorical decision is made based on the comparison of representational examples in the face space with the probe face to classify. It is possible that the neural system needs more effort to categorize if representational examples are likely to be distributed widely in the face space (and thus are perceived to differ considerably) as opposed to when representational examples of a category are well localized in the face space (and are thus perceived very similar). We thus assume that the mean distance between test samples and all previously trained faces in the face space is proportional to the response time in a categorization task. We quantified the mean distance of a fixed number of test samples with all trained samples for both categories (Asian versus Caucasian) as a function of the changing training-ratio (e.g. 0% Asian, 100% Caucasian to 100% Asian, 0% Caucasian), and normalized the calculated mean distances to the distance for the own-class effect (100% training on the same class as the tested classes). As shown in Figure 4A, relative distance of the test samples with the trained samples decreases for decreasing relative exposure to the tested class during training. In other words, the less the neural system was exposed to a particular face class, the less representational space was allotted to this class. In turn, faces of that category appear more similar (near in face space) and behavioral responses are presumably quicker. We found this to be the case for both Caucasian and Asian faces when trained on a mixture of these two classes (Figure 4A, left panels). One-sample t-tests, comparing each position in the training

ratio (x-axis) and the feature space (y-axis) with the 0% training ratio at given size of feature space (baseline), showed significant differences in mean distances, and thus showing the mirror effect (Figure 4B, upper and lower left panel). However, when tested analogously on the two categories chimpanzee and human faces, ambiguous tendencies occurred: When tested on chimpanzee faces, the mean distances showed an inverted u-shape function along the variation in the training ratio (x-axis) and the size of feature space (y-axis) (Figure 4A, upper right panel). Similarly ambiguous, the mean distances were relatively large, when tested on human faces, irrespective of the training ratio and the size of feature space (Figure 4A, lower right panel). Thus, the species comparison did not reveal a systematic pattern of significance in accordance with the Mirror Effect (Figure 4B, right panels).

Discussion

In the current study, we created a model under the hypothesis that neural plasticity of the face discrimination system retains optimal discrimination performance in its environment, i.e. the face discrimination continuously updates its representation to all the faces exposed to. If neural resources were limited, the face system could be less selective in the choice of represented feature dimensions to allow for the coding of a high variability of faces that then could be discriminate in a similar fashion.

We used common face features derived from face images of various classes and show that a face feature representation in the LFD space achieved solid discrimination performances on own-class faces when selecting a subset of feature dimensions. The selected subset of features, however, does not necessarily lead to an optimal performance on the other classes of faces; effects known as the own-race and own-species advantages. We followed a major trend of computational models on face perception representing faces in a face space framework²⁷. In the literature, a global analysis of faces was implemented with principal components analysis (PCA) directly applied on face images^{28–30} and a simplified model of an own-race advantage was described as an optimal feature selection model³¹. We here first extract facial features via topological methods, then apply PCA on the extracted face profile vectors and optimize the spatial separation of face identities in the representational space with LFD analysis. Conceptually, we followed the idea of Furl²⁰ and O'Toole¹⁹, using both a generic contact hypothesis as well as a developmental contact hypothesis algorithm to simulate the own-race advantage. The authors showed that these experience-based algorithms, in contrast

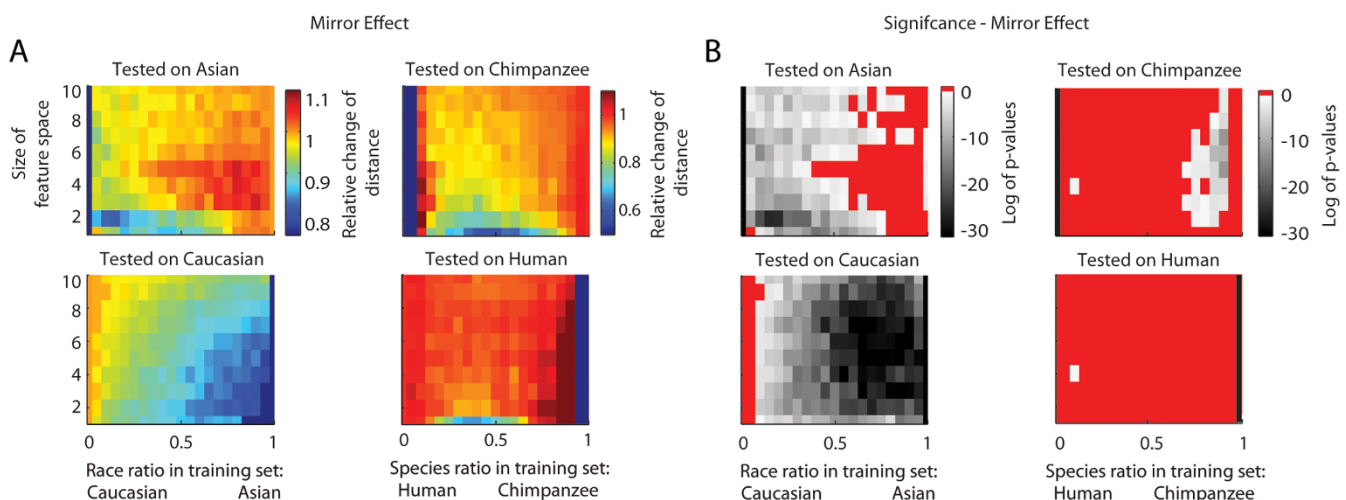


Figure 4 | Mirror Effect. (A.) The relative change of distances between training and testing faces at given training ratio (x-axis) and size of feature space (y-axis). (B.) Axes as in A. Data reflects the significant positions of relative changes as opposed to the baseline (0% training position at given size of feature space, indicated by black vertical lines).



to non-contact hypothesis algorithms, result in an own-race advantage. Following the idea of the contact hypothesis, a recent study³² compares a variety of trainable and non-trainable algorithms across a variety of face classes, age classes and gender. The conclusion of this study is that training the face recognition system on dataset well distributed across all races is critical to reduce deficits for specific demographic cohorts and that face discrimination performance on races improves when trained exclusively on the same race (and age) class. In another recent study, the mechanism of the underlying algorithm effect was investigated, pointing out the need to understand how the ethnic composition (ratio of faces belonging to different race classes) of a training set impacts the performance of the algorithm.

In our study, when quantifying the relative deterioration of other-class as opposed to own-class faces, the own-species advantage is of greater effect size than the own-race advantage. We followed the question whether the own-race and own-species advantages can be explained by the same underlying mechanism. We show that the relative deterioration of performance scores is present in both, the race as well as the species comparisons. The relative impairment, indeed, reflects what has been recently reported in the literature²⁵. Hence, it is plausible to assume that the own-race and own-species advantages are caused by the same fundamental mechanism.

We found absolute performance differences when comparing race classes as opposed to species classes (Figures 2A–G). This can be explained by the general quality of the stimuli: while the race faces were based on 3D-models and rendered under controlled and standardized conditions (Figure 2D–G), the species faces were photographs that include a wider degree of variation of factors (Figure 2A–C,G), such as the exact perspective, lighting condition, alignment of the face inside the image canvas, and others. What looks like a limitation in the methodological procedure at first glance, turns out to be an interesting way of validating the model under more realistic, less controlled conditions, like the species comparison. Further, given certain restrictions, like the fact that there are no 3D models of chimpanzee faces available, there is simply no way to account for improvements in that aspect.

It can be argued if there is some sort of feature transfer between two classes of faces. We accounted for this question by calculating relative distances between individual faces of one class and all other faces of the other class and by estimating the overlap of feature distribution for two classes, showing that the distributions between race classes are generally closer than those of species classes, which seems plausible, given the morphological (and evolutionary) greater distance of human and chimpanzee than Asian and Caucasian faces. Given these facts, we might assume that there is indeed some sort of feature transfer from Asian to Caucasian faces and vice-versa, while it might be a rather challenging task for the face processing system to use human-face tuning to account for chimpanzee faces and vice-versa. In the literature, there are some interesting cases to illustrate exactly this assumption: Captive chimpanzees, a rare case of individuals with an immense exposure to non-conspecific faces over decades of years, but only little exposure to conspecific faces, were tested on discriminating chimpanzee and human faces¹⁰: Young chimpanzees (around 10 years of age) showed a clear advantage for chimpanzee faces above human faces; however, the advantage turned into a disadvantage with increasing exposure to human faces and limited exposure to chimpanzee faces: older chimpanzees (around 30 years of age) showed an advantage for human above chimpanzee faces. In other words the sensitivity toward one class of faces in early life adapted toward another class of faces more strongly exposed to over decades. These same chimpanzees showed a more pronounced face inversion effect for the category of expertise¹² and a more pronounced left-chimeric face bias, reflecting more right hemispheric processing, for the category of expertise¹¹. Hence, tuning a system to two morphologically strongly different face classes seems to be rather chal-

lenging and seems not to occur in biological systems. Indeed, according to our SCORE approach, neural resources are limited and thus features have to be updated. This implicates that the representation of previously learned faces might change as well when the current exposure changes dramatically. In other words, discrimination performance of the own-face class might decrease after an almost exclusive long-term exposure to other-species faces. This view is supported by the findings in the captive chimpanzees^{10–12}.

On the other hand, face feature transfer seems likely across human race classes, as illustrated by³³: Caucasian participants showed a larger whole-face advantage (in comparison to using individual facial parts) for own-race as opposed to other-race faces, while Asian participants, living in a society predominantly populated by Caucasians, showed an equal whole-face advantage for both types of faces³³. Accordingly, training with own-race faces improves processing of other-race faces in patients with developmental prosopagnosia³⁴ as it does in healthy participants^{35,36}. Hence, with increasing amount of experience with the other-race class, the own-race advantage decreases^{37,38}, however, does not flip over into the other-race advantage, possibly pointing to an additional strong priming component in early years of life.

Our findings for the own-race advantage are generally in accordance with the perceptual expertise hypothesis: The more exposure with faces of different races, the more familiar the system becomes with different races and the more the own-race advantage diminishes^{39–43}. However, since the own-race and own-species advantages seem to be based on a morphological basis, as previously shown for races⁴⁴, and morphological species differences are generally larger, it is questionable whether experimentally the own-species advantage diminishes with increasing experience. According to¹⁰, the face perception system competes for one of the two face classes. Importantly, in this latter case, the exposure to the own-face class was extremely limited, potentially supporting a complete cross-over effect from the conspecific to the heterospecific face class.

The own-race advantage (other-race effect) has been well-established in the literature and appears very consistent. Aside from this effect, there is a somehow paradoxical finding that, when categorizing faces by the race, other-race faces show an advantage^{45,46}. This effect has been demonstrated mainly in measuring response latencies. It has to be noted that in order to discriminate faces of a specific race or species, the system relies on identity-specific information; in other words, the system solves the task at the subordinate-level of categorization. To tell whether a face belongs to one race/species or to another, requires race/species-specific information, and hence, operates at the basic-level of categorization. To address the above-described mirror effect, the model needs to be interpreted from the basic-level point of view: We define a categorical decision as the product of a comparison with representational examples of that face class in the Fisher-space. If we change the ratio of own- and other-race/species faces in favor of other-race/species faces (i.e. more other-race/species faces), then the more subtle differences in the other-race/species faces will be coded for and thus the representational space enlarges. The increased dissimilarity between representational examples of the category likely results in an increase of difficulty to make a categorical decision due to, for example, an increase of the search area, while in naïve system (tuned to one face class only, with small sample size of other-race/species faces) the small search area with only a few other-race/species faces is more distinctive. Hence, with only a few representations of other-race/species faces, the variance among these samples is rather small, while with increasing numbers the variance increases. We here showed that, indeed, the distances between representational examples increasingly changes as a function of relative number of other-race/species faces represented in the Fisher-space. Indirectly, the distances reflect the processing speed for a categorical decision: the larger the representational space of a face class, the longer it takes to determine that a face belongs to that class.



This assumption is in accordance with face space models by Valentine⁴⁷.

Together, feature-based approaches, using PCA-features^{19,20} and LFD-representations, as proposed here, are used to explain the own-race and own-species advantages by supervised classification methods. In Wallis⁴⁸, the question about biological relevance has been raised of implementing a layer of supervised training in the neural network model that forces the system to focus on one face class⁴⁹. At the other end, an entirely unsupervised, self-organizing system has been proposed, producing an effect similar to own-race advantage and, interestingly, showed an interplay of the own-race advantage and the degree of holistic processing⁴⁸. Unfortunately, this model uses abstract features, rather than features extracted from facial models or pictures. Here, we chose to extract facial features using topological methods and represented them in a LFD space, conceptually describing the outcome of a self-organized process as mentioned above. We acknowledge that the practical implementation of feature extraction is somewhat arbitrary, especially the generation of profile vectors, and is not without loss of information. We implemented such a processing step to enhance the processing efficacy. However, the main results should be robust to the changes in the face features extraction method, as long as the faces features are rich enough to discriminate faces of all classes. Conceptually, with the features extraction step, the model therefore gains biological relevance, as (1) training is equivalent to the exposure history and not forcefully biased toward one face class, and (2) data samples are actual facial features extracted from a variety of face classes. Hence, we position the current work between models with explicit supervised training and fully self-organized procedures.

In sum, with simple model assumptions we were able to simulate both the own-race¹⁹ and the own-species advantages. We found that experience (training) with one or the other class is the crucial factor; hence the learning history builds the face feature space required for successful discrimination of faces most exposed to.

Methods

We used 66 photographs of 22 chimpanzees and 54 photographs of 18 humans. We used three viewpoints (approximately 0, 10 and 20 degrees) (face database 1). In addition, we rendered images from 3-D reconstructions of faces^{22,50}, consisting of faces of 60 Caucasians and 60 Asians from three viewpoints each (−10, 10 and 20 degrees). These face images were taken from the face database of the Max Planck Institute for Biological Cybernetics (face database 2).

Facial feature extraction. We implemented facial feature extraction according to²¹, making use of the topological features of faces to localize eyes, nose and mouth. With the image taken as surface and values on the z-axis as image intensity, eyes, nose and mouth are singularities in the image, forming valleys and peaks on the luminance surface. The topological features for each face image were computed after luminance histogram equalization across images and averaged horizontally to reduce the size of the feature vector. Horizontally averaging was done to reduce the size of the feature vector and was possible because faces in our database were of the same size and aligned to the image frame. However, the exact method for reducing dimensionality of the features should not affect our main results.

Training procedure. In general, we used 85% randomly selected face images for training and the rest for testing. After feature extraction, each image was represented by its feature vector. As described in the main text, we varied the amount of images from different species/races in the train set to model different exposure ratios. The “training” here corresponded to calculating the Linear Fisher Discriminant [LFD] directions, using the method described by²⁴. In brief, the LFD calculates the directions in the feature space, which minimize the variance (in feature space) between faces of the same individual and simultaneously maximize the variance between faces of different individuals. In this sense, if projecting the feature vectors on the principle LFD direction, faces from the same individual group as close together as linearly possible, while other individuals’ faces are as far apart as linearly possible. Thus, faces from different individuals tend to be well separated. Finding the correct LFD projection corresponds to a simple implementation of the face processing system that optimizes its features to the exposure (experience) to best recognize a known individual. To reduce the dimensions of the face features (in relation to the limited amount of samples), we pre-processed the face features further with a principle component analysis and keeping the first 20 components (as suggested in²⁴). After calculating the first few LFD dimensions (we vary the number of selected LFD components in the figures), we projected the test data on these LFD components

(solely derived from the train set) and subsequently access the discrimination performance of identifying faces from the same individual in a forced choice task.

Testing procedure. We evaluated the discrimination performance in a matching-to-sample fashion, taking one face stimulus (A1) from the test set and comparing the distances in the space of the LFD projections (see above) to another individual’s face (B1) and to another face exemplar belonging to the same individual (A2). The picture having the shorter distance was deemed as the same individual. This choice was either correct or wrong, and the performance of the percentage of correct detections could be evaluated. The test set was held separate and not used for training. The training and testing procedure was repeated 100 times and average performances are reported. We varied the face group ratio in the train set from 0 to 100% as well as the number of LFD components being taken into account for discrimination. The accuracy of the system depends on how well the test separates the two groups (same versus other). Accuracy was measured as the area under the ROC (Receiver operating characteristic) curve, with 1 being a perfect test and .5 being a worthless test.

Testing the Mirror Effect. The mean Euclidean distances among samples of training and testing faces were quantified in the Fisher-space for each point in the training ratio and for each size of feature space. The data was normalized to the overall variance of the classes. We report the relative changes of distances among the samples. We repeated this procedure 100 times for randomized train and test samples. With one-sample t-tests, significance was tested for each position on the training ratio (x-axis) and the size of feature space (y-axis) by comparing the 100 data samples to the baseline (0% training ratio at given size of feature space).

General statistical testing. Dependent variables are the location of samples in the PCA space, the Linear Fisher Discriminant (LFD) space and the area under the ROC curve. We used two-way analysis of variance (ANOVA), one- and two-sampled t-tests with unpaired variance estimate. To determine the similarity of the distribution resulting from the 1st PCs, we calculated the kernel mean discrepancy using bootstrapping (as described in²³), comparing whether two samples can be considered originating from the same underlying high-dimensional probability distribution or not. Note that this test does not make any assumptions on the shape of the distribution.

Additional note. We used face database 1 for the species comparison (Figure 2A–C) and face database 2 for the race comparison (Figure 2D–F). We used the chimpanzee faces of face database 1 and the Asian and Caucasian faces of face database 2 for the cross-validation (Figure 2G).

1. Feingold, C. A. The influence of environment on identification of persons and things. *Journal of Criminal Law and Police Science* 5, 39–51 (1912).
2. Michel, C., Corneille, O. & Rossion, B. Race categorization modulates holistic face encoding. *Cogn Sci* 31, 911–924, doi:10.1080/03640210701530805 (2007).
3. Tanaka, J. W. & Pierce, L. J. The neural plasticity of other-race face recognition. *Cogn Affect Behav Ne* 9, 122–131, doi:Doi 10.3758/Cabn.9.1.122 (2009).
4. Walker, P. M. & Tanaka, J. W. An encoding advantage for own-race versus other-race faces. *Perception* 32, 1117–1125, doi:Doi 10.1068/P5098 (2003).
5. Bothwell, R. K., Brigham, J. C. & Malpass, R. S. Cross-Racial Identification. *Pers Soc Psychol B* 15, 19–25, doi:Doi 10.1177/0146167289151002 (1989).
6. Scott, L. S. & Fava, E. The own-species face bias: A review of developmental and comparative data. *Vis Cog* 1–28, doi:10.1080/13506285.2013.821431 (2013).
7. Anzures, G., Quinn, P. C., Pascalis, O., Slater, A. M. & Lee, K. Development of own-race biases. *Vis Cog* 21, 1165–1182, doi:Doi 10.1080/13506285.2013.821428 (2013).
8. Pascalis, O. & Bachevalier, J. Face recognition in primates: a cross-species study. *Behavioural processes* 43, 87–96 (1998).
9. Pascalis, O., de Haan, M. & Nelson, C. A. Is face processing species-specific during the first year of life? *Science* 296, 1321–1323, doi:10.1126/science.1070223 (2002).
10. Dahl, C. D., Rasch, M. J., Tomonaga, M. & Adachi, I. Developmental processes in face perception. *Sci. Rep.* 3 (2013).
11. Dahl, C. D., Rasch, M. J., Tomonaga, M. & Adachi, I. Laterality Effect for Faces in Chimpanzees (Pan troglodytes). *J Neurosci* 33, 13344–13349, doi:10.1523/JNEUROSCI.0590-13.2013 (2013).
12. Dahl, C. D., Rasch, M. J., Tomonaga, M. & Adachi, I. The face inversion effect in non-human primates revisited - an investigation in chimpanzees (Pan troglodytes). *Sci. Rep.* 3, 2504, doi:10.1038/srep02504 (2013).
13. Dufour, V., Pascalis, O. & Petit, O. It takes one to know one: Species-specific face recognition in humans and monkeys. *J Cogn Neurosci* 69–69 (2002).
14. Dahl, C. D., Logothetis, N. K., Bulthoff, H. H. & Wallraven, C. The Thatcher illusion in humans and monkeys. *Proc Biol Sci* 277, 2973–2981, doi:10.1098/rspb.2010.0438 (2010).
15. Dahl, C. D., Logothetis, N. K., Bulthoff, H. H. & Wallraven, C. Second-Order Relational Manipulations Affect Both Humans and Monkeys. *PLoS one* 6 (2011).
16. Dahl, C. D., Logothetis, N. K. & Hoffman, K. L. Individuation and holistic processing of faces in rhesus monkeys. *Proc Biol Sci* 274, 2069–2076 (2007).
17. Dahl, C. D., Wallraven, C., Bulthoff, H. H. & Logothetis, N. K. Humans and macaques employ similar face-processing strategies. *Curr Biol* 19, 509–513, doi:10.1016/j.cub.2009.01.061 (2009).



18. O'Toole, A. J. & Natu, V. Computational perspectives on the other-race effect. *Vis Cog* **21**, 1121–1137, doi:Doi 10.1080/13506285.2013.803505 (2013).
19. O'Toole, A. J. Simulating the 'other-race effect' as a problem in perceptual learning. *Connection Science* **3**, 163 (1991).
20. Furl, N., Phillips, P. J. & O'Toole, A. J. Face recognition algorithms and the other-race effect: computational mechanisms for a developmental contact hypothesis. *Cognitive Sci* **26**, 797–815 (2002).
21. Gunduz, A. & Krim, H. Facial feature extraction using topological methods. *2003 International Conference on Image Processing, Vol 1, Proceedings* 673–676 (2003).
22. Blanz, V. & Vetter, T. Face recognition based on fitting a 3D morphable model. *Ieee T Pattern Anal* **25**, 1063–1074, doi:Doi 10.1109/TPAMI.2003.1227983 (2003).
23. Gretton, A., Borgwardt, K. M., Rasch, M. J., Scholkopf, B. & Smola, A. A Kernel Two-Sample Test. *J Mach Learn Res* **13**, 723–773 (2012).
24. Belhumeur, P. N., Hespánha, J. P. & Kriegman, D. J. Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection. *Ieee T Pattern Anal* **19**, 711–720, doi:Doi 10.1109/34.598228 (1997).
25. Dahl, C. D., Rasch, M. J. & Chen, C. C. The other-race and other-species effects in face perception – a subordinate-level analysis. *Frontiers in Psychology* **5**:1068, doi: 10.3389/fpsyg.2014.01068 (2014).
26. Ge, L. Z. *et al.* Two faces of the other-race effect: Recognition and categorisation of Caucasian and Chinese faces. *Perception* **38**, 1199–1210, doi:Doi 10.1068/P6136 (2009).
27. Valentine, T. A Unified Account of the Effects of Distinctiveness, Inversion, and Race in Face Recognition. *Q J Exp Psychol-A* **43**, 161–204 (1991).
28. O'Toole, A. J., Abdi, H., Deffenbacher, K. A. & Valentin, D. Low-Dimensional Representation of Faces in Higher Dimensions of the Face Space. *J Opt Soc Am A* **10**, 405–411, doi:Doi 10.1364/Josaa.10.000405 (1993).
29. Hancock, P. J. B., Burton, A. M. & Bruce, V. Face processing: Human perception and principal components analysis. *Memory & Cognition* **24**, 26–40, doi:Doi 10.3758/Bf03197270 (1996).
30. Turk, M. & Pentland, A. Eigenfaces for recognition. *J Cogn Neurosci* **3**, 71–86, doi:10.1162/jocn.1991.3.1.71 (1991).
31. Haque, A. & Cottrell, G. W. in *In Proceedings of the 27th Annual Cognitive Science Conference* (Mahwah: Lawrence Erlbaum).
32. Klare, B. F., Burge, M. J., Klontz, J. C., Bruegge, R. W. V. & Jain, A. K. Face Recognition Performance: Role of Demographic Information. *Ieee T Inf Foren Sec* **7**, 1789–1801, doi:Doi 10.1109/Tifs.2012.2214212 (2012).
33. Tanaka, J. W., Kiefer, M. & Bukach, C. M. A holistic account of the own-race effect in face recognition: evidence from a cross-cultural study. *Cognition* **93**, B1–B9, doi:Doi 10.1016/j.cognition.2003.09.011 (2004).
34. DeGutis, J., DeNicola, C., Zink, T., McGlinchey, R. & Milberg, W. Training with own-race faces can improve processing of other-race faces: Evidence from developmental prosopagnosia. *Neuropsychologia* **49**, 2505–2513, doi:DOI 10.1016/j.neuropsychologia.2011.04.031 (2011).
35. Elliott, E. S., Wills, E. J. & Goldstein, A. G. Effects of Discrimination Training on Recognition of White and Oriental Faces. *B Psychonomic Soc* **2**, 71–73 (1973).
36. Chance, J., Goldstein, A. G. & McBride, L. Differential Experience and Recognition Memory for Faces. *J Soc Psychol* **97**, 243–253 (1975).
37. Brigham, J. C., Maass, A., Snyder, L. D. & Spaulding, K. Accuracy of Eyewitness Identifications in a Field Setting. *J Pers Soc Psychol* **42**, 673–681, doi:Doi 10.1037//0022-3514.42.4.673 (1982).
38. Chiroro, P. & Valentine, T. An Investigation of the Contact Hypothesis of the Own-Race Bias in Face Recognition. *Q J Exp Psychol-A* **48**, 879–894 (1995).
39. Sporer, S. L. Recognizing faces of other ethnic groups - An integration of theories. *Psychol Public Pol L* **7**, 36–97, doi:Doi 10.1038//1076-8971.7.1.36 (2001).
40. Sporer, S. L. The cross-race effect - Beyond recognition of faces in the laboratory. *Psychol Public Pol L* **7**, 170–200, doi:Doi 10.1037/1076-8971.7.1.170 (2001).
41. Young, S. G., Hugenberg, K., Bernstein, M. J. & Sacco, D. F. Perception and Motivation in Face Recognition: A Critical Review of Theories of the Cross-Race Effect. *Personality and Social Psychology Review* **16**, 116–142, doi:Doi 10.1177/1088868311418987 (2012).
42. Feinman, S. & Entwisle, D. R. Childrens Ability to Recognize Other Childrens Faces. *Child Dev* **47**, 506–510, doi:Doi 10.2307/1128809 (1976).
43. Cross, J. F., Cross, J. & Daly, J. Sex, Race, Age, and Beauty as Factors in Recognition of Faces. *Perception & psychophysics* **10**, 393–&, doi:Doi 10.3758/Bf03210319 (1971).
44. Salah, A. A., Alyuz, N. & Akarun, L. Registration of three-dimensional face scans with average face models. *J Electron Imaging* **17**, doi:Artn 011006Doi 10.1117/1.2896291 (2008).
45. Valentine, T. & Endo, M. Towards an Exemplar Model of Face Processing - the Effects of Race and Distinctiveness. *Q J Exp Psychol-A* **44**, 671–703 (1992).
46. Caldara, R., Rossion, B., Bovet, P. & Hauert, C. A. Event-related potentials and time course of the "other-race" face classification advantage. *Neuroreport* **15**, 905–910 (2004).
47. Valentine, T. Face-space models of face recognition. *J Math Psychol* **42**, 495–495 (1998).
48. Wallis, G. Toward a unified model of face and object recognition in the human visual system. *Frontiers in psychology* **4**, 497, doi:10.3389/fpsyg.2013.00497 (2013).
49. Lewis, M. B. & Johnston, R. A. A unified account of the effects of caricaturing faces. *Vis Cog* **6**, 1–42 (1999).
50. Troje, N. F. & Bulthoff, H. H. Face recognition under varying poses: The role of texture and shape. *Vision Res* **36**, 1761–1771, doi:Doi 10.1016/0042-6989(95)00230-8 (1996).

Acknowledgments

This research was financially supported by the NFSC grants 31371109 and 2014CB846101 (M.J.R.), and the NSC grant 102-2811-H-002-010 (C.C.C.). We thank Heinrich Bulthoff and the ScanLab of the Max Planck Institute for Biological Cybernetics in Tubingen, Germany, for contributing the Asian and Caucasian face pictures and the Great Ape Research Institute (GARI) and the Primate Research Institute of Kyoto University, Japan, for the chimpanzee face photographs.

Author contributions

C.D.D. and M.J.R. conceived, programmed and ran the simulation and analyzed the data; C.D.D., M.J.R. and C.C.C. contributed reagents/materials/analysis tools; C.D.D. and M.J.R. wrote the paper. C.D.D., M.J.R. and C.C.C. discussed the results and commented on the manuscript.

Additional information

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Dahl, C.D., Chen, C.-C. & Rasch, M.J. Own-race and own-species advantages in face perception: a computational view. *Sci. Rep.* **4**, 6654; DOI:10.1038/srep06654 (2014).



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder in order to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-sa/4.0/>