UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The face inversion effect and the anatomical mapping from the visual field to theprimary visual cortex

Authors

Gahl, Martha
Yuan, Meilu
Sugumar, Arun
et al.

Peer reviewed

The face inversion effect and the anatomical mapping from the visual field to the primary visual cortex

Martha Gahl (mgahl@eng.ucsd.edu)
Department of Computer Science and Engineering, University of California San Diego
9500 Gilman Dr 0404, La Jolla, CA 92093 USA

Meilu Yuan (m2yuan@illinois.edu)
Department of Computer Science, University of Illinois at Urbana-Champaign
601 E John St, Champaign, IL 61802 USA

Arun Sugumar (asugumar@eng.ucsd.edu)
Department of Computer Science and Engineering, University of California San Diego
9500 Gilman Dr 0404, La Jolla, CA 92093 USA

Garrison Cottrell (gary@eng.ucsd.edu)
Department of Computer Science and Engineering, University of California San Diego
9500 Gilman Dr 0404, La Jolla, CA 92093 USA

Abstract

The face-inversion effect, or the drastic decrease in accuracy seen when a participant is asked to identify inverted faces when compared to upright faces, is an effect that is not found in object inversion. Here we suggest a new explanation of this effect using computational models to show that the phenomenon can be explained by the anatomical mapping from the visual field to primary visual cortex. We propose that the way inverted faces are mapped onto the cortex is fundamentally different from the way upright faces are mapped. Our work first shows the advantages of this mapping due to its scale and rotation invariance when used as input to a convolutional neural network. We train the network to perform recognition tasks and show it exhibits scale and realistically constrained rotation invariance. We then confirm that the decline in accuracy seen when a participant is asked to identify inverted faces is not seen in the network with inverted object recognition tasks. With the support of these two findings, we test the face-inversion effect on our network and are able to show the unique decline in accuracy, suggesting that the way the visual field is mapped onto the primary visual cortex is a key facet in the manifestation of this effect.

Keywords: Face-inversion effect; object inversion; scale invariance; rotation invariance; log polar transformation.

Introduction

A well known, well defined, but still poorly understood phenomenon in perception is the face-inversion effect. Early experiments showed the difficulty people have both with identifying and remembering upside down faces. The effect is even more pronounced when compared with people's abilities to do these same tasks with objects (Yin, 1969). In subsequent years, many studies have explored the questions that still remain about our cognitive abilities in facial recognition. Some of these studies have focused on creating cognitive hypotheses or new models to understand the processing that occurs in the brain (Rakover, 2013; Schwaninger et al., 2003). Others have focused on the cognitive processes that occur and the way they manifest differently when the stimulus is an inverted face (Schwaninger & Mast, 2005; Rezlescu et al., 2017; Rock, 1988). Still others have explored the time course of the differences when processing right side up faces versus upside down faces during recognition (Taubert et al., 2011; Freire et al., 2000). We propose an additional method for approaching these questions by studying the anatomical basis for the way faces in the visual field are mapped onto primary visual cortex.

Previous work studying the face-inversion effect using an anatomical approach has provided evidence for the link between inverted face processing and different cortical areas. Much of the work focuses on cortical regions such as the fusiform face area (FFA) and the occipital face area (OFA), and how they respond differently to upright and inverted faces (Pitcher et al., 2011; Yovel and Kawisher, 2005; Kanwisher et al., 1998). This occurs late in processing and in specialized brain regions. Our work instead uses an anatomical approach to study the inversion of a face as it is perceived at the level of the primary visual cortex (V1). We are interested in the way faces, both upright and inverted, are mapped onto V1 and how that mapping plays a role in the face-inversion effect.

The mapping of the visual field onto the visual cortex (V1, V2, and V3) can be approximately described by geometric transformations of the visual field (J.R. Polimeni et al., 2006). Polimeni describes a "Wedge-Dipole model" of this mapping, which has two component parts: 1) "a (quasiconformal) wedge mapping with" 2) "a (conformal) dipole mapping." Together these two transformations map

Figure 1. Scale is a horizontal translation.

Figure 2. Rotation is a vertical translation

the visual field onto a 2D representation of cortical space. In this map, a wedge map is a compression of the image in the visual field along the azimuth, and the conformal dipole mapping "is an extension of the standard log-polar or complex logarithm mapping" (J.R. Polimeni et al., 2006).

For our model, we have chosen to simplify Polimeni's model map to the visual cortex in two ways. The first is in the visual cortex itself. Instead of mapping to V1, V2, and V3, we are only concerned with primary visual cortex (V1). Our focus is on studying the mapping from the visual field onto the visual cortex as a possible factor in the explanation of the face-inversion effect. Because of this, we do not model the mapping of the visual field onto higher visual areas, nor did we model anatomically where the stimuli are

Figure 3. The map to V1 starts at the vertical meridian and proceeds clockwise from there.

processed later in the temporal lobe. The second simplification we made was with respect to the 100:1 image compression from retina to V1 (100M photoreceptors compared to 1M synapses onto the LGN). We chose to only use log polar transformations on our images as a simplification. We believe this is a valid approximation, and not an oversimplification, because the compression of the visual field is consistent across the visual field. Even if the size of the final mapping differs, the proportions of the visual field mapped into the cortex are maintained for different areas of the visual field.

By using a log polar transformation on the images, we are still able to capture important features of the mapping. Because of the log transformation of the radial axis, the pixels at the center of the image are more highly represented than those in the periphery. This mirrors the greater representation of the fovea on the cortical surface when compared with the periphery. In the primate visual system, this "cortical magnification" is due to the number and arrangement of receptors in the retina: densely packed in the fovea, and dropping off logarithmically in the radial direction. Another simplification is that we did not use the anatomical constraint that only half of the visual field is represented in each hemisphere.

Our work includes two components. We first wanted to test how well a log polar mapping could perform a recognition task with faces at multiple scales. This provides insight into how well the anatomically inspired mapping will work in tandem with the convolutional neural network (CNN) that we use for our recognition tasks. Hence, images are preprocessed by the log polar transformation before being input to the network. We chose scale as the transformation to test first to verify that the combination of the log-polar input with a translation-invariant convnet would be scale invariant, as has been verified independently in an unpublished paper (Remmelzwaal et al., 2019). We tested the network's ability to generalize to new scales through interpolation, or generalizing to a scale not seen in testing that lies between scales seen in testing, and extrapolation, generalizing to scales not seen in testing and that do not lie between scales seen in testing. These experiments allowed us to verify that our mapping allowed

CNNs to become invariant to scaled images, and validated our use of them with rotated images.

The second aspect of the work was to test the network's ability to generalize to rotated images. To parallel our scaled experiments, we tested the ability of the network to generalize to new rotational orientations through interpolation and extrapolation. This verified that the network, with images mapped with log polar transformations, could not only become invariant to scale, but also to rotation. We then tested the network on upside down faces to explore the face-inversion effect with log polar images.

Note that there is a topological difference between scale and rotation. Whereas scale is just a horizontal shift, rotation requires wrap-around due to the fact that 360° is the same as 0°. Indeed, due to the vertical meridian in the retina, the top of the V1 representation corresponds to 270° and the bottom to 90° (see Figure 3). This results in a configuration of the input that has the features in a different relationship to one another (see Figure 2, bottom row). Now the eyes are above the nose and mouth, versus in upright position, the log-polar representation has the eyes below the mouth. We believe that this configural difference is what contributes to the inverted face effect.

Methods

Model

For each of our experiments, we used ResNet-18 to perform the identification tasks. This network has an initial convolutional layer followed by four blocks with four convolutional layers in each, and makes use of global average pooling before the final, fully connected layer (He et al., 2016).

Data

We used images of faces for training and testing. All images of faces were previously collected by members of the lab as part of a new face dataset. The dataset includes 200 different people, each with approximately 200 unique images. These images portray each person in a variety of contexts, with differing backgrounds, lighting conditions, orientations, and facial expressions. The training set including 30,030 total images, or approximately 150 images per person. The test set includes 3,236 total images, or approximately 16 images per person (Figure 4).

We also used a dataset of cars, The Comprehensive Cars (CompCars) Dataset (Yang et al., 2015), for our experiments with object inversion, as the inversion effect for objects is much smaller than it is for faces (Yin, 1969). The dataset includes cars of different makes, models, and years. In addition, the cars in each image are oriented differently. We used 16,200 images in our training set and 1,800 images in our test set. The images were evenly distributed across nine categories. We chose to use the labels of car types, such as "Sedan" or "Pickup", for the categories in our object

inversion task, as these can be thought of as non-trivial basic level categories (Figure 5).

Transformations

We preprocessed our images to create an anatomically-realistic mapping of the visual field onto V1. The visual field is mapped into polar coordinates in V1. In addition, there is a log transformation that accounts for more receptors and greater representation in cortical regions for central vision as compared to peripheral vision (Polimeni, 2006). We used a log polar transformation on our images as a 2D approximation of this mapping onto the cortex.

We conducted experiments studying scale and rotation in facial recognition to verify the effects discussed above. For scale, we used unscaled images (224 x 224), images that were scaled to 80% of the original image size (180 x 180), and images that were scaled to 60% of the original image size (134 x 134). For each of the two different scaled conditions, we padded the images back to the original image size of 224 x 224. For each scale of image, we first found a bounding box around the face in the image and then took the log polar transformation from the center of the box. This allowed us to take the log polar transformation from a point that represented the fixation point on the face. After the images have undergone a log polar transformation, changes in scale appear as translation changes in shifts left and right (Figure 1).

In the rotation experiment, we rotated the training images to commonly observed head orientations: 20°, 10°, 0°, -10°, and -20°. To study the face inversion effect, we also rotated testing images to 180° for testing (not included in training). For each case of rotation, we first found a bounding box around the face and then took the log polar transformation from the center of the face, similarly to the scaled images. After the images have undergone a log polar transformation,


Figure 4. Images from face dataset

Figure 5. Images from The Comprehensive Cars Dataset

the changes in degrees of rotation appear as changes in translation in shifts up and down. However, in contrast to the shifts that occur when scaled images undergo a log polar transformation, when an image shifts up or down the pixels that "fall off" the edge of the image wrap around to the opposite side of the image. Instead of appearing as a simple translation, changes in rotation in images that have undergone a log polar transformation result in a rearranging of features, particularly the eyes and the mouth.

Experiments and Results

To test the validity of using log polar images as an approximation for the mapping from the visual field to primary visual cortex, we tested the network's scale invariance on log polar images. Convolutional neural networks are invariant to shift, and when an image has undergone a log polar transformation, changes in scale appear as changes in translation. Therefore, the network should be invariant to scale. We tested scale invariance under two different conditions: interpolation and extrapolation.

We trained with our preprocessed images for 90 epochs. We then tested with two different conditions. Both test conditions used images that depicted the same 200 people seen in training, but were all novel images. The testing images had different backgrounds, lighting, and poses than the training images. The first test condition used the novel images with the same preprocessing and transformations used in training. This is the accuracy for familiar orientations, which we call the "training condition." The second test condition used the same novel images as were used in the first test condition, but with novel transformations performed. This we call the "testing condition." If these two accuracies are similar, then invariance to the transformation performed in the second test condition is shown. For all experiments, the difference between the training condition and the testing condition is the primary result being reported.

Scaling

We first trained a network on only one scale and then tried to generalize to different scales. These experiments resulted in large differences between the training and testing conditions; accuracy was much lower for the test condition. We then tried training using two scales, and tested using extrapolation to larger and smaller scales. This treatment resulted in similar accuracies between train and test conditions. The results are shown in Table 1.

Table 1. Scale: training condition/testing condition

	Trained with one scale	Trained with two scales
Generalizing to larger images	81.3% / 68.6%	89.1% / 88.8%
Generalizing to smaller images	80.8% / 64.2%	89.5% / 87.2%

Interpolation Interpolation in scale invariance is the ability of the network to generalize to unseen scales when it has been trained on a combination of scales that are larger and scales that are smaller than the testing scale. We trained a network using unscaled images and images that were scaled to 60% of the original image size. We then tested the network on the facial recognition task with images that were scaled to 80% of the original size. Table 2 shows the results of this experiment. By the end of training, the training condition accuracy reached 87.8% and the testing condition accuracy reached 87.5%. This is only a 0.3% difference in training condition accuracy and testing condition accuracy.

Extrapolation We then explored extrapolation in scale invariance. In order for a network to show scale invariance in extrapolation it would have to be able to generalize to testing scales that were either larger than all scales seen in training or smaller than all scales seen in training. We first trained a network on larger images to test the ability to generalize to smaller images. The training and testing conditions were then flipped to see if a network trained on smaller images was able to generalize to larger images.

To test the ability to generalize to smaller images, we trained a network on unscaled images and images that had been scaled to 80% of the original image size. We then tested the network on images that were scaled to 60% of the original image size. The results of this experiment are found in Table 2. We found the training condition accuracy with these conditions to be 89.5% and the testing condition accuracy to be 87.2%. This gives us a difference of 2.3% between the two accuracies.

For the experiment to generalize to larger images, we trained a network on images that were scaled to 80% of the original image size and images that were scaled to 60% of the original image size. We then tested on unscaled images. From Table 2, the training condition accuracy at the end of training was 89.1% and the testing condition accuracy was 88.8%, which gave a final difference of 0.3%.

Rotation

We then explored how effective using log polar transformations on images to map the visual field to V1 is when the transformation being performed is rotation instead of scaling. Once images have undergone a log polar transformation, changes in rotation appear as a translation of the image up or down. However, when the log polar image shifts, the pixels wrap around to the other side of the image. We wanted to determine the effect of this wrapping first with more realistic changes in head orientation, and then with inverted faces. Because of our results from scaling experiments in which multiple examples of orientations in training proved more robust to transformation, we performed all rotation experiments using multiple orientations in training. Similarly to the scaling experiments, we looked at both interpolation and extrapolation conditions when testing the realistic changes in head orientation.

Interpolation Interpolation in rotation experiments is the ability of the network to generalize to an unseen head orientation when it has been trained on head orientations to the right and to the left of the testing orientation. To test interpolation in rotation, we trained a network on faces rotated 10° and faces rotated -10°. We then tested the network on faces that were not rotated. The results are in Table 3. The network achieved a training condition accuracy of 87.2% and a testing condition accuracy of 84.9%, which gives a difference in accuracy of 2.3%.

Extrapolation Extrapolation in our rotation experiment is how well the network can generalize to unseen head orientations that do not lie between head orientations seen in training. Because we are using multiple training orientations and rotation is cyclic, any testing orientation can be considered to be between the training head orientations. For the purpose of our experiments, extrapolation means using a testing head orientation that is directly adjacent to all training head orientations, either clockwise or counterclockwise.

To test rotation invariance with extrapolation we tested both clockwise and counterclockwise rotation. For clockwise extrapolation, we trained the network on faces rotated 10°, faces rotated 0°, and faces rotated -10°. We then tested the network on faces rotated -20°. The training condition accuracy at the end of training was 89.5% and the testing condition accuracy was 86.9%. The difference in accuracies for clockwise rotation was 2.6%. For counterclockwise extrapolation we used the same training examples as were used in clockwise extrapolation. The only difference was that we tested on faces that were rotated 20°. For the counterclockwise experiment we achieved a training condition accuracy of 89.9% and a testing condition accuracy of 87.9%. The total difference in accuracy was 2.0%. The results of these two experiments can be seen in Table 3.

Rotation of Objects As a means of validating that inversion affects the perception of faces in a different way than it affects the perception of objects, we tested our network on inverted cars. We trained for 50 epochs on images of cars in nine categories. The cars were rotated to five orientations for training: 20°, 10°, 0°, -10°, and -20°. We then tested with two different testing conditions. The first testing condition, or the training condition, tested novel images of the cars at orientations seen in training. The second testing condition, or the testing condition, tested novel images of cars that were rotated 180°. This allowed us to demonstrate the effect of object inversion on the network when used in combination with log polar images. These results are shown in Table 4. For cars, the training condition accuracy at the end of training was 50.7% and the testing condition accuracy at the end of training was 19.4% for a total difference in accuracies of 31.3%. This demonstrates that rotated images have a significant effect on the network's performance.

The Face-Inversion Effect

Our main focus in studying rotation under these conditions was to test the face-inversion effect on our network knowing that taking a log polar mapping of our images provided an approximation to the mapping from the visual field to V1 that allowed our network to be invariant to scale. To test this effect, we trained on normally oriented faces and tested on upside down faces. We used our conclusion from the scaling experiments that networks are more robust to transformations in recognition tasks when they are trained on more examples of transformations to inform our face-inversion effect experiment. Because of this, we chose to train our network on five degrees of rotation: 20°, 10°, 0°, -10°, and -20°. While this provides a training range of only 40°, the degrees of rotation are frequently seen orientations of the head and face. This also allowed for direct comparison with the object inversion experiment. We then tested using the same two testing conditions used in the experiment with rotation of objects. The first testing condition was the training condition, in which we tested the network on novel faces at the same orientations used in training. The second testing condition was the testing condition, in which we tested on only the one orientation, faces rotated 180°. The results of this experiment are shown in Table 4. Our training condition accuracy after training for 50 epochs was 90.6%. Our testing condition accuracy was 20.4%. The difference in our training condition accuracy and testing condition accuracy when testing the face-inversion effect was 70.2%. This difference is significantly higher than any other difference seen in our experiments, including the experiment on the inversion of objects. While together the experiments showed that inversion is a much more difficult task for the network, they also show that the network behaves differently when the stimuli are faces.

Table 2. Scaling: training condition/testing condition

Interpolation	87.8% / 87.5%
Extrapolation to smaller scales	89.5% / 87.2%
Extrapolation to larger scales	89.1% / 88.8%

Table 3. Rotation: training condition/testing condition

Interpolation	87.2% / 84.9%
Extrapolation to clockwise rotation	89.5% / 86.9%
Extrapolation to counterclockwise rotation	89.9% / 87.9%

Table 4. Face-inversion effect and object inversion experiments: training condition/testing condition

Face-inversion effect	90.6% / 20.4%
Object inversion	50.7% / 19.4%

Discussion

We used an anatomically-inspired log polar transformation to preprocess the input images to our network as a parallel to the mapping that occurs from the visual field onto the primary visual cortex. This transformation accounts for a number of aspects of the visual system. The log transformation gives greater representation in the final input images to the pixels in the region at the center of the transformation. This is analogous to the center of fixation on an image. This greater representation of central pixels mirrors the cortical magnification of central vision as compared to peripheral vision, which occurs because of the increased number of receptors for central vision. This transformation also allows the network to be invariant to the scale of faces, which is an important component in face and object recognition. After a log polar transform, changes in scale appear as changes in horizontal translation. CNNs are invariant to translation, which allows them to be invariant to scaling performed on these preprocessed images. The log polar transformation is unique in being able to facilitate the replication of aspects of biological vision.

Through our experiments with scaled faces, we were able to show that, with log polar images, scale invariance in a CNN can be achieved. This can be shown for interpolation and extrapolation, provided the network is trained on multiple scales. We believe this validates the use of the log polar transformation as a 2D approximation of the mapping of the visual field onto V1.

The success of introducing scale invariance to the network led us to perform experiments with rotated images. Because of the difference in behavior of log polar images that are rotated as compared to scaled, we wanted to again verify the validity of using log polar transformations as 2D approximations of the mapping of the visual field onto V1. We used multiple head orientations in training to increase the robustness of the network to unseen orientations. In both interpolation and extrapolation of head orientations, the network was able to generalize to unseen head orientations. This supports the conclusions from the scaled experiments that log polar transformations provide a valid approximation for the mapping of the visual field onto V1, because the network was also able to demonstrate rotation invariance to commonly observed head orientations in addition to scale invariance.

With these results, we proceeded to test the face-inversion effect in our network. Even when training on multiple head orientations meant to increase robustness to unseen orientations, the difference between the training condition accuracy and the testing condition accuracy was significantly larger than any other difference seen in all experiments. The network was not invariant to a rotation of 180°. Instead of a simple transformation, a rotation of 180° appears as a fundamental rearrangement of the features as they are mapped onto the cortex. It is this rearrangement that we believe plays a part in the face-inversion effect.

As a final validation that the network was demonstrating this behavior uniquely for faces, we trained the network on multiple orientations of cars and then tested it on inverted cars. The network was able to generalize to inverted objects better than it was able to generalize to inverted faces, strengthening the argument that the effect of the log polar mapping onto V1 causing a decline in recognition accuracy is unique to faces.

We believe that these results are explained by the log polar transformation that was performed on the images. When using log polar transformations, changes in rotation cause the rearrangement of features. The more severe the rotation, the bigger the impact this on a network's ability to recognize the inverted image. In essence, the inversion of images causes a fundamental rearrangement of features when in conjunction with a log polar transform. When features are rearranged, the configural information about the relationship of features of the object or face in the image is lost. It is the loss of this configural information that causes both cars and faces to have larger drops in accuracy between the training condition and the testing condition. However, we believe this accuracy drop is more significant in facial recognition because of the importance of configural information in recognizing faces. The network is showing that it is relying more heavily on configural information when it is performing facial recognition as opposed to object recognition. The effect of the loss of this information is reflected in the magnitude of the accuracy drop between training and testing conditions.

In addition, we do not believe that these conclusions are in conflict with popular hypotheses. Some studies have suggested that a holistic approach to face perception, or the emphasis of the face as a whole as opposed to the face being viewed as a combination of features, lies at the center of the face-inversion effect because of the way that rotation affects the configuration of features and the mechanisms the brain employs to "undo" this rotation (Shwaninger, 2003; Schwaninger, 2005; Farah et al., 1995). Our work similarly suggests that there is a fundamental rearrangement of features, and that upright faces are not mapped on to the cortex in the same configuration as inverted faces. We believe that these experiments suggest a compelling anatomical basis for the face-inversion effect.

References

Esteves, C., Allen-Blanchette, C., Zhou, X., & Daniilidis, K. (2018). Polar Transformer Networks. *2018 International Conference on Learning Representations (ICLR).*

Farah, M. J., Tanaka, J. W., & Drain, H. M. (1995). What causes the face inversion effect? *Journal of Experimental Psychology: Human Perception and Performance.*

Freire, A., Lee, K., & Symons, L. A. (2000). The Face-Inversion Effect as a Deficit in the Encoding of Configural Information: Direct Evidence. *Perception*.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. *Cognition*.

Krause, J., Stark, M., Deng, J., & Fei-Fei, L. (2013). 3D Object Representations for Fine-Grained Categorization. *4th IEEE Workshop on 3D Representation and Recognition, at the 2013 International Conference on Computer Vision (ICCV)*.

Pitcher, D., Duchaine, B., Walsh, V., Yovel, G., & Kanwisher, N. (2011). The role of lateral occipital face and object areas in the face inversion effect. *Neuropsychologia*.

Polimeni, J., Balasubramanian, M., & Schwartz, E. (2006). Multi-area visuotopic map complexes in macaque striate and extra-striate cortex. *Vision Research*.

Rakover, S. S. (2013). Explaining the face-inversion effect: the face–scheme incompatibility (FSI) model. *Psychonomic Bulletin & Review*.

Remmelzwaal, L., Mishra, A., & Ellis, G. (2019). Human eye inspired log-polar pre-processing for neural networks. *URL https://arxiv.org/abs/1911.01141*.

Rezlescu, C., Susilo, T., Wilmer, J. B., & Caramazza, A. (2017). The inversion, part-whole, and composite effects reflect distinct perceptual mechanisms with varied relationships to face recognition. *Journal of Experimental Psychology: Human Perception and Performance*.

Rock, I. (1974). The Perception of Disoriented Figures. *Scientific American*.

Rock, I. (1988). On Thompsons Inverted-Face Phenomenon (Research Note). *Perception*.

Schwaninger, A., Carbon, C.C., & Leder, H. (2003). Expert face processing: Specialization and constraints. In G. Schwarzer & H. Leder, *Development of face processing*.

Schwaninger, A., & Mast, F. W. (2005). The face-inversion effect can be explained by the capacity limitations of an orientation normalization mechanism1. *Japanese Psychological Research*.

Taubert, J., Apthorp, D., Aagten-Murphy, D., & Alais, D. (2011). The role of holistic processing in face perception: Evidence from the face inversion effect. *Vision Research*.

Wang, P., & Cottrell, G. W. (2017). Central and peripheral vision for scene recognition: A neurocomputational modeling exploration. *Journal of Vision*.

Yang, L., Luo, P., Loy, C. C., & Tang, X. (2015). A large-scale car dataset for fine-grained categorization and verification. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*.

Yovel, G., & Kanwisher, N. (2005). The Neural Basis of the Behavioral Face-Inversion Effect. *Current Biology*.