Recognition of identity and expressions as integrated processes

Kathryn C. O’Neill, Rebecca Saxe, and Stefano Anzellotti*

February 2019

Abstract

Recent neuroimaging evidence challenges the traditional view that face identity and facial expressions are processed by segregated neural pathways, and shows that information about identity and expressions is encoded within the same brain regions. This article proposes that integrated representations of identity and expressions arise naturally within neural networks, and demonstrates that deep networks trained to recognize facial expressions spontaneously develop representations of identity. Furthermore, features distinguishing between identities and features distinguishing between expressions become increasingly orthogonal from layer to layer, indicating that face perception is a process of disentanglement of representational subspaces corresponding to different generative sources.

1 Introduction

Recognizing face identity and facial expressions is a critical component of our ability to navigate the social environment. Identity recognition enables us to acquire knowledge about specific individuals that we can retrieve in future encounters (5, 31). Expression recognition helps us to infer the emotional states of an individual (50, 51, 44) and predict their future actions and reactions.

1.1 The dominant view: separate processing of identity and expressions

The dominant view on the recognition of face identity and facial expressions proposes that facial expressions and face identity are processed by distinct pathways (5, 19). In this view, face identity is recognized by a ventral temporal pathway (19), including the occipital face area (OFA, (14)) and the fusiform face area (FFA, (24)). By contrast, facial expressions are recognized by a lateral pathway (19), including the posterior superior temporal sulcus (pSTS, (20)).

In support of this view, several lines of evidence show that the pSTS is a key region for the recognition of facial expressions. The pSTS responds selectively to faces (1) and shows greater responses to moving faces than static faces (40).

Furthermore, the patterns of activity in this region encode information about the valence of facial expressions (38, 45). Finally, patients with pSTS damage have deficits for facial expression recognition (13), providing causal evidence in support of the involvement of pSTS in facial expression recognition.

Faces showing dynamic expressions do not evoke increased responses in OFA and FFA to the same degree as in pSTS (40). In addition, several studies have shown that information about individual faces can be decoded in FFA (32, 49) and even OFA (3).

Taken together, this evidence supports the involvement of OFA and FFA in the recognition of identity, and the involvement of pSTS in the recognition of expressions. However, the evidence does not rule out a role for OFA and FFA in expression recognition, nor does it rule out a role for pSTS in the recognition of identity. In sum, this evidence is not sufficient to demonstrate a separation between the recognition of identity and expressions.

1.2 Challenges to the dominant view

Previous work has noted the lack of strong evidence in support of the dominant view (7). Prosopagnosics often show some amount of impairment for expression recognition (7). Deficits of identity recognition can spare expression recognition (12, 54), but such double dissociations might occur at later stages of processing and do not exclude substantial integration between expression and identity recognition at earlier stages (7). These observations further highlight that support for the existence of separate processing systems for identity and expressions is not strong, but they do not provide direct evidence against the dominant view.

Recent evidence, however, indicates that recognition of face identity and facial expressions might be more integrated than previously thought. The valence of facial expressions can be decoded from ventral temporal regions including OFA and FFA (45, 28). At the same time, identity information can be decoded from the face-selective pSTS (2, 18, 11). In fact, in one study identity could be decoded with higher accuracy from pSTS than from OFA and FFA ((11), Figure 6), and in two other studies identity could be decoded in pSTS across faces and voices (2, 18). Furthermore, pSTS damage leads to impairments for recognizing faces across different expressions (13), suggesting that pSTS plays a causal role for identity recognition too.

This evidence is part of a broader context of findings suggesting that dynamic information contributes to the recognition of person identity (55). Dynamic information is used to recognize face identity (33), and it is given higher weight when the shape of the face is less reliable (10) and when the face is familiar (6). Person identity can also be recognized from gait presented with point light displays (34).

Dynamic faces also evoke stronger responses in the face-selective pSTS (40). Furthermore, neighboring regions of the pSTS respond to biological motion (39) and rTMS to the pSTS disrupts biological motion recognition (16). However, the view that pSTS is specialized for the processing of dynamic stimuli (4) does not fully account for the empirical data. Identity information can also be decoded

from the face-selective pSTS after presentation of static face images and voices (2, 18). Rather than encoding exclusively dynamic information, pSTS appears to integrate form, motion, and sound (voices).

1.3 A new hypothesis

On the basis of this evidence, we propose that recognition of facial expressions and identities are implemented by integrated mechanisms, and not by separate neural pathways. Unlike classical models, which suggest that information relevant to identity recognition should be shed as emotion representations develop, we propose that representations of identity contribute to expression recognition, and vice versa. In our account, identity and expressions are represented within the same brain regions for computational efficiency. Furthermore, we hypothesize a mechanism by which representations of identity and of expressions contribute respectively to the recognition of expressions and identity. The appearance of a face image is the result of both the identity and the expression. Recognizing what aspects of an image are due to face identity could prevent errors of expression recognition that would result from misattributing those aspects of the image to the expression. Likewise, recognizing what aspects of an image are due to the expression can prevent errors of identity recognition that would result from misattributing those aspects to the identity. This hypothesis leads to two non-trivial predictions. First, if recognition of facial expressions and face identity are mutually beneficial, training an algorithm to recognize facial expressions might lead to the spontaneous formation of representations that encode face identity information. Second, this phenomenon would not occur because representations of identity and expression rely on common features. On the contrary, features important for the recognition of expressions and features important for the recognition of identity should become increasingly disentangled and orthogonal along the processing stream.

Deep neural networks have been used to achieve remarkable accuracy on image recognition tasks (29, 36), and features extracted from deep network layers have been successful at predicting responses to visual stimuli in temporal cortex in humans (26) and monkeys (52) (see (53, 27) for reviews). To test the hypothesis that recognition of facial expressions and identity are mutually beneficial, we asked whether features from hidden layers of deep networks trained to recognize facial expressions could be used successfully to recognize face identity. As an even stronger test of the hypothesis we investigated whether, as accuracy for expression recognition increased from layer to layer, accuracy for identity recognition increased correspondingly. This prediction is in stark contrast with classical theories of abstraction in visual processing suggesting that information about task-orthogonal information is progressively discarded (42, 47, 25).

To test this prediction, we trained a densely-connected deep convolutional neural network (DenseNet, (22)) to recognize facial expressions using a database of over 28000 facial expression images ('fer2013'). The DenseNet architecture was chosen because of its excellent performance on a variety of visual recognition tasks (22), and because, like the human visual system, it features connections

between non-adjacent stages of the processing hierarchy (48). After training, the DenseNet achieved human level performance at facial expression labeling on the fer2013 database. We then tested whether features from the DenseNet trained to recognize facial expression could be used to recognize face identity. If identity information were discarded, we would expect to see decreasing performance on identity classification across layers. If, as we predict, identity representations grow alongside emotion representations, we would expect to see increasing performance on identity classification across each layer of the net.

2 Methods

2.1 Stimuli

The DenseNet was trained to label emotions using the face images in the facial expression recognition 2013 (fer2013) database (15). The database contains 28,709 images for training and an additional 3,589 images labeled as ‘public test’ (these latter images were used to test the performance of the network after training and to compare it to human performance). All images are sized 48×48 pixels and grayscale. The Karolinska Directed Emotional Faces (KDEF) database consists of 4900 images depicting 70 individuals showing 7 different expressions from 5 different angles, each photographed twice (30). Images are sized 562 (width) $\times$ 762 (height) and in color (RGB). In order to match the format of the fer2013 images, all KDEF images were converted to grayscale, cropped to squares, and downsampled to 48×48 pixels. The images were converted to grayscale by averaging pixel values of the red, green, and blue channels. As the positioning of the face within the image was consistent across KDEF images, the rectangular images were all cropped to the same 388×388 pixel region around the face. Example face images from each dataset can be seen in Figure 1. Visual inspection confirmed that the entire face was visible in each image after cropping. For both sets of analyses, we used the frontal views and views rotated by 45 degrees on both sides (left and right).

2.2 Network architecture

Using Pytorch (37), we implemented a deep DenseNet consisting of 1 convolutional layer, 3 dense blocks, and 1 fully connected linear layer. The convolutional layer consisted of 64 channels of 2D convolutions using a 3×3 kernel and padding=1. Each dense block consisted of 3 densely connected convolutional layers with kernel size=3, stride=1, and padding=1. Each layer in the dense block produced 32 channels of output. Therefore, the number of input channels for the first layer in a dense block was equal to the number of output channels of the previous layer outside the dense block (i.e., for the first layer of the first dense block it was equal to 64: the number of output channels of the first convolutional layer). The number of input channels for each subsequent layer in each dense block increased by 32. This choice is widely used and featured on pub-

Figure 1: Face image stimuli. Top: selected images from fer2013 database. Bottom: selected images from KDEF dataset (after gray-scaling and cropping; before spatial downsampling).

licly available DenseNet implementations (i.e. <https://github.com/pytorch/vision/blob/master/torchvision/models/densenet.py>).

Each dense block except the last was followed by a transition layer receiving as input the outputs from all layers of the dense block plus the layer preceding the dense block, and producing an output with half the number of channels and with a max pooling using a 2x2 kernel. The last dense block was followed by an average pooling with an 8x8 kernel and then by a fully connected linear layer. In sum, the number of input and output channels for the 13 layers of the network were: (1,64); (64,32); (96,32); (128,32); (160,80); (80,32); (112,32); (144,32); (176,88); (88,32); (120,32); (152,32); (32,1).

All layers used rectified linear units (ReLU) as nonlinearity (except the fully connected layer which was linear), and all layers in dense blocks and all transition layers used 2D dropout with a dropout probability $p = 0.1$ (8). All convolutional layers were followed by batch normalization (23).

2.3 Training and validation

Given 48x48 grayscale images in the fer2013 database, the DenseNet produced as output 7 values, one for each emotion label in the dataset. The network was trained minimizing the cross-entropy loss between the output and the true labels using stochastic gradient descent, with a learning rate starting at 0.1 and halving every 30 epochs. The training was run for 200 epochs, and images were presented to the network in batches of 16. After training, the accuracy of the

Figure 2: DenseNet architecture. Top: The DenseNet consists of a convolutional layer, three dense layers, and a linear classifier. Bottom: Single dense block; red lines represent connections that would exist in a typical convolutional neural net, purple lines represent connections that are unique to the densely-connected net.

network was validated using an independent subset of the fer2013 that was not used for training (the images marked as ‘PublicTest’). The network achieved an accuracy of 63.5%, matching closely reported human accuracy on the fer2013 stimuli (65%, (15)).

2.4 Labeling expression across identity and identity across expression

The weights of the network were then fixed to prevent further training. Henceforth we will refer to the network with the weights fixed after training with fer2013 as the ‘pretrained network’. To test identity labeling, we used a new database of images: the Karolinska Directed Emotional Faces database (KDEF, (30)) that has identity labels in addition to emotion labels. First, we tested whether the network could accurately label the emotions in the new dataset as a sanity check, then we tested identity labeling.

The output that the network needs to produce for identity labeling and for emotion labeling is different. For instance, the number of identity labels is different from the number of emotion labels (70 v 7). To accommodate this, we attached a fully connected linear layer after the end of the last dense block

of the pretrained network that was re-trained to produce the desired output (emotion or identity) using a subset of the KDEF database. Since we wanted to probe the features learned by the pretrained network in different layers, we attached fully connected layers to different layers of the pretrained network, producing as output the emotion or identity labels. Critically, these added fully connected layers were linear, and achieved very different performance depending on the layer of the network that they were attached to (that is, depending on the nonlinear features that they received as an input). Accuracy at emotion and identity labeling was tested using the left-out portion of the KDEF database.

Partitioning of the KDEF database was performed as follows. In a first analysis, we used frontal images as well as images with the head rotated 45 degrees to the left and right. For emotion labeling, 8 identities were left out. For identity labeling, one of the emotions was left out, and the other 6 were used for training. Accuracy was averaged across all cross-validation folds. In a second analysis, the linear layer was trained using all but one of the viewpoints (left, straight, and forward), and accuracy was tested using the left-out viewpoint (as in (3)). Accuracy values for both emotion and identity were then averaged across averaged across the three conditions. This choice was made to provide a more stringent test of identity recognition: facial expressions might leave some parts of the face unchanged, but rotation in depth alters all parts.

Two key observations need to be made about the methods. First, the network trained with the fer2013 database is nonlinear, and its weights are fixed after training (they are never modified using the KDEF images). By contrast, the fully connected layers attached to the different layers of the pretrained network – which are trained with a subset of the KDEF images for the purpose of producing output in the correct format – are linear. Since the added fully connected layers cannot implement nonlinear transformations, their performance is heavily dependent on the nonlinear features they receive as inputs. Second, if the added linear layers trained with a subset of the KDEF images could achieve high accuracy without needing the features from the pretrained network, this should be evident also when they are attached early layers of the pretrained network. Accuracy improvements when using features from later layers as compared to features from early layers of the pretrained network cannot be due to the attached linear layer trained with a subset of KDEF images, because the attached linear layer was used in both cases in the same way.

The pretrained net, training code, datasets, and code used to crop and label images can be found at <https://osf.io/p8jm2/>.

3 Results

The DenseNet trained with the fer2013 database labeled emotions with an accuracy of 63.5% on the left-out ‘public test’ images: close to the human performance of 65%. Importantly, the accuracy at emotion labeling on the frontal and 45 degree angle view images of the KDEF database was high (85.4%), showing that the DenseNet generalized successfully to the new database. Labeling accu-

racy was higher for the KDEF database than for the fer2013. This is likely due to the fact that images in the KDEF database are centered and posed, while images in the fer2013 exhibit a much larger degree of natural variation in illumination and viewpoint. In addition images in the fer2013 database often show only part of the face, and in some cases the faces are covered by occluders (i.e. the hands).

As expected, labeling accuracy increased from layer to layer of the DenseNet. A linear layer applied directly to the KDEF images obtained an accuracy of 66.4% for emotion classification, but subsequent layers were necessary to reach the final accuracy of 85.4%. A large increase in accuracy was observed in the second dense block, yielding an accuracy of 83%.

Identity classification from the frontal and 45 degree angle view images of the KDEF database was also very accurate (98.8%). As detailed in the Methods section, in order to generate the 70 identities as output (instead of the 7 emotion labels), a linear layer was attached to the second-last layer of the DenseNet pretrained with the fer2013 database, and trained with images from all but one emotion label. Critically, the DenseNet weights were fixed at this stage: only the weights of the newly attached linear layer could change.

Accuracy was consistently high (above 98%) regardless of the emotion that was left out. Importantly, the high accuracy was not due to the newly attached linear layer: a linear layer alone achieved an accuracy of 68.7%, far below the 98.8% achieved by a linear layer attached to the features from the DenseNet. Interestingly, a large increase in accuracy was observed in the second dense block (leading to an accuracy of 98.5%), parallel to the large increase in accuracy observed for emotion labeling at the same processing stage.

Having established that the features learned by the DenseNet using an emotion labeling task can be used to successfully label face identity, we aimed to investigate the invariance of the face representations across image transformations. Recognition of face identity across changes in viewpoint is notoriously difficult (41, 3). The DenseNet was not trained to recognize identity across different viewpoints, but it was trained to label emotions across viewpoints. Could the features it developed for emotion labeling be used to support view-invariant identity recognition?

To address this question, we used images from the KDEF database showing a frontal view as well as 45 degree views (left and right) of the faces. We tested whether the DenseNet could label emotions and identities when the newly attached linear layer was trained with two of the three views, and tested with the left-out view.

Emotion labeling across viewpoint in the KDEF database was more challenging than across identity, leading to an accuracy of 63.7%. This value is close to the accuracy of 63.5% obtained with the ‘public test’ subset of the fer2013 database. Interestingly, emotion labeling was more accurate when the two 45 degree views were used for training and the front view was left-out for testing (88%), than when the left view was left-out (44.7%) or when the right view was left-out (58.4%). In this case, accuracy also increased from layer to layer in the network, and a linear layer alone achieved low accuracy (14.3%).

Figure 3: Classification accuracy for identity (top) and emotion (bottom) for a linear layer attached to successive sections of the pre-trained DenseNet. Dotted line represents performance at chance. Leftmost bar represents performance of unattached linear classifier.

Identity labeling was also more challenging across viewpoint, with an average accuracy of 31.4%. However, the DenseNet features were critical to achieve this level of accuracy: a linear layer trained directly on the images labeled identity with an accuracy of only 2.3%. As in the cases that did not require generalization across viewpoints, features from the second dense block allowed for large increases in accuracy for both emotion labeling (from 16.4% to 41%) and for identity labeling (from 7.2% to 23.7%).

Next, we tested whether features that show large variation across different identities and features that show large variation across different emotions are overlapping or orthogonal. If features that are most useful for classifying

Figure 4: Classification accuracy for identity (top) and emotion (bottom) for a linear layer attached to successive sections of the pretrained DenseNet when one of three face directions is held out. Dotted line represents performance at chance. Leftmost bar represents performance of unattached linear classifier.

emotion and identity are shared in each layer of the net, we would expect for these features to overlap. If, on the other hand, features that define emotion and identity are disentangled by the net, we would expect to see them become more orthogonal. We differentiated between these two hypotheses by applying a principal component analysis to activations for each emotion and identity and finding the angles between emotion principal components and identity principal components. First, we ran each face image from the the KDER dataset through our DenseNet, extracted activations from the convolutional layer and each of the three dense layers, averaged the activations for each layer over individual emotions and identities, then found the top 5 principal components (PCs) for

emotion and identity in each of the four layer’s activation spaces. We then found the angle between each emotion PC and identity PC per layer, and averaged them to get a mean angle. In the first iteration, we only used the top two PCs, then three, increasing the number of PCs until we reached five. The emotion and identity PCs started with low angles in the convolutional layer, then grew more orthogonal across the net (see Figure 5A). We also projected each face activation onto the top two identity and emotion PCs for each layer (see Figure 5B,C). In each case, the relevant aspect (emotion or identity) visibly clusters in deeper layers of the net, while the other aspect does not.

4 Discussion

We have argued that the existing empirical evidence cannot be reconciled by the view that recognition of facial expressions and identity are implemented by separate neural mechanisms. Furthermore, we have proposed that the mechanisms for the recognition of expressions can contribute to the recognition of identity, and vice versa. Based on this hypothesis, we predicted that features generated for the purpose of recognizing facial expressions would also support accurate recognition of identity. We trained a dense deep convolutional neural network (DenseNet) to label facial expressions, and found that as the labeling of expressions improved from layer to layer, the labeling of identity also correspondingly improved – despite the fact that the features of the DenseNet were never explicitly trained for identity recognition.

While our findings do not guarantee that all neural network architectures will show the same effect, they demonstrate that identity representations arise naturally in a simple, standard architecture trained to achieve near-human accuracy at emotion labeling. This result lends support to the view that the neural mechanisms for identity and expression recognition in the human brain might not be integrated by accident, but to improve computational efficiency.

Why do features relevant to classification of facial expressions contribute to recognition of identity? One possibility is that recognition of both identity and expressions are characterized by image patches with similar statistical properties. In such case, the same ‘vocabulary’ of features would contribute to both expression and identity recognition because it captures statistical properties of image patches (i.e. spatial frequency, pixel patterns) that are important for both identity and expressions (i.e. a feature detecting an low-curvature u shaped pattern could be used both to represent a person’s jaw line or a smile). However, our principal component analysis suggests that this is not the case. In early layers of the net, we captured low level features relevant to both identity and emotion which overlapped in activation space. However, the late layers of the net seem to represent emotion and identity in increasingly distinct features. This suggests that, while features relevant to emotion can allow for identity classification, the features most relevant to emotion and identity are not the same.

A second possibility is that successful labeling of expressions relies on avoid-

Figure 5: Emotion and identity principal components. A) Mean angle between principal components derived from activations averaged over either emotion or identity. B) Face activations labeled by emotion projected into emotion and identity principal component spaces for each layer. C) Face activations labeled by identity (only 7 of 70 identities are displayed for clarity) projected into emotion and identity principal component spaces for each layer.

ing errors due to considering expression-irrelevant aspects of an image as expression-relevant. From layer to layer, the network might develop increasingly accurate representations of identity that can be used by downstream layers to isolate what aspects of the image are due to the facial expression.

A third possibility is that features relevant to classifying expressions might aid in the detection of identity-specific patterns of expression-dependent changes due to individual variation in facial musculature, that could be used to aid recognition of person identity. This suggestion is in line with a growing literature on the contributions of facial motion to the recognition of person identity (33, 9), and with a broader range of research examining the contribution of motion cues to identity recognition, including information about body movements and gait (46, 34, 43, 55, 17).

This third possibility does not seem sufficient to explain the full pattern of empirical results. If pSTS only extracts identity information that is specific to particular facial movements, it should be difficult to classify identity across different expressions using pSTS representations. By contrast, neuropsychological evidence indicates that damage to pSTS leads to deficits for recognition of identity generalizing across expressions (13). However, extraction of identity-specific variation in expressions can be one of several factors accounting for the integration between the recognition of identity and expressions.

Describing what we are observing in more formal terms can be a starting point to better understand this phenomenon. During training, the weights of the DenseNet are modified to minimize a loss function L , which captures the performance of the network on the training data. The network is a composition of a nonlinear component g trained only with the training subset of the fer2013 database, with a linear fully connected layer attached at the end. We have three versions of the fully connected layer, the f that is trained with the training subset of the fer2013 database, an f_E trained for expression labeling with the training subset of the KDEF database, and an f_I trained for identity with the training subset of the KDEF database.

One way to think about the present findings is that g trained to minimize the loss $L(f \circ g)$ on the training subset of fer2013 also leads to low losses $L(f_E \circ g)$ and $L(f_I \circ g)$ on the testing subset of KDEF. If we plot the losses $L(f \circ g)$, $L(f_E \circ g)$ and $L(f_I \circ g)$ as a function of the weights W_g that characterize g , we have that in regions of parameter space where $L(f \circ g)$ is low, also $L(f_E \circ g)$ and $L(f_I \circ g)$ are low with high probability. If we consider the conditional probability distribution $p(L(f_I \circ g) = y | L(f \circ g) = x)$, this amounts to saying that given low values of x , the distribution is concentrated on low values of y .

In this article, we studied the spontaneous emergence of representations of face identity when deep neural networks are trained to label facial expressions. However, it is unlikely that this type of phenomenon would occur only in this particular case. Recent studies have reported related evidence for the emergence of representations of viewpoint and position in the visual field for deep networks trained to label objects (21). This is also the case for faces: late layers of deep convolutional networks trained to recognize identity encode information about yaw and pitch (35). More broadly, integrated implementation of complementary

computations might be a large-scale principle of organization of human cortex, determining by virtue of computational efficiency what sets of cognitive processes are represented within the same neural systems, and could apply to cases as diverse as the recognition of person identity and word in speech processing, and the inference of mental states and traits in social cognition.

Future work can test whether face representations generated for labeling expressions also support identity recognition with invariance across other kinds of transformations, such as translation, scaling, and the more challenging case of occlusion. In addition, facial expressions are dynamic: extending the present results to neural networks processing dynamic stimuli will be an important step forward to better understand the relationship between features built for expression recognition and features built for identity recognition given moving faces. Finally, while we have discussed one patient study showing impaired recognition of identity across expressions (13), additional causal evidence of the interaction between impairments for the recognition of facial expressions and face identity is needed.

References

- [1] Timothy J Andrews and Michael P Ewbank. Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. *Neuroimage*, 23(3):905–913, 2004.
- [2] Stefano Anzellotti and Alfonso Caramazza. Multimodal representations of person identity individuated with fmri. *Cortex*, 89:85–97, 2017.
- [3] Stefano Anzellotti, Scott L Fairhall, and Alfonso Caramazza. Decoding representations of face identity that are tolerant to rotation. *Cerebral Cortex*, 24(8):1988–1995, 2013.
- [4] Michal Bernstein and Galit Yovel. Two neural pathways of face processing: a critical evaluation of current models. *Neuroscience & Biobehavioral Reviews*, 55:536–546, 2015.
- [5] Vicki Bruce and Andy Young. Understanding face recognition. *British journal of psychology*, 77(3):305–327, 1986.
- [6] Natalie Butcher and Karen Lander. Exploring the motion advantage: Evaluating the contribution of familiarity and differences in facial motion. *The Quarterly Journal of Experimental Psychology*, 70(5):919–929, 2017.
- [7] Andrew J Calder and Andrew W Young. Understanding the recognition of facial identity and facial expression. *Nature Reviews Neuroscience*, 6(8):641, 2005.
- [8] George E Dahl, Tara N Sainath, and Geoffrey E Hinton. Improving deep neural networks for lvsr using rectified linear units and dropout. In *2013*

- IEEE international conference on acoustics, speech and signal processing*, pages 8609–8613. IEEE, 2013.
- [9] Katharina Dobs, Isabelle Bühlhoff, and Johannes Schultz. Identity information content depends on the type of facial movement. *Scientific reports*, 6:34301, 2016.
 - [10] Katharina Dobs, Wei Ji Ma, and Leila Reddy. Near-optimal integration of facial form and motion. *Scientific reports*, 7(1):11002, 2017.
 - [11] Katharina Dobs, Johannes Schultz, Isabelle Bühlhoff, and Justin L Gardner. Task-dependent enhancement of facial expression and identity representations in human cortex. *NeuroImage*, 172:689–702, 2018.
 - [12] Nancy L Etcoff. Selective attention to facial identity and facial emotion. *Neuropsychologia*, 22(3):281–295, 1984.
 - [13] Christopher J Fox, Hashim M Hanif, Giuseppe Iaria, Bradley C Duchaine, and Jason JS Barton. Perceptual and anatomic patterns of selective deficits in facial identity and expression processing. *Neuropsychologia*, 49(12):3188–3200, 2011.
 - [14] Isabel Gauthier, Michael J Tarr, Jill Moylan, Pawel Skudlarski, John C Gore, and Adam W Anderson. The fusiform “face area” is part of a network that processes faces at the individual level. *Journal of cognitive neuroscience*, 12(3):495–504, 2000.
 - [15] Ian J Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, et al. Challenges in representation learning: A report on three machine learning contests. In *International Conference on Neural Information Processing*, pages 117–124. Springer, 2013.
 - [16] Emily D Grossman, Lorella Battelli, and Alvaro Pascual-Leone. Repetitive tms over posterior sts disrupts perception of biological motion. *Vision research*, 45(22):2847–2853, 2005.
 - [17] Carina A Hahn and Alice J O’toole. Recognizing approaching walkers: Neural decoding of person familiarity in cortical areas responsive to faces, bodies, and biological motion. *NeuroImage*, 146:859–868, 2017.
 - [18] Bashar Awwad Shiekh Hasan, Mitchell Valdes-Sosa, Joachim Gross, and Pascal Belin. “hearing faces and seeing voices”: Amodal coding of person identity in the human brain. *Scientific reports*, 6:37494, 2016.
 - [19] James V Haxby, Elizabeth A Hoffman, and M Ida Gobbini. The distributed human neural system for face perception. *Trends in cognitive sciences*, 4(6):223–233, 2000.

- [20] Elizabeth A Hoffman and James V Haxby. Distinct representations of eye gaze and identity in the distributed human neural system for face perception. *Nature neuroscience*, 3(1):80, 2000.
- [21] Ha Hong, Daniel LK Yamins, Najib J Maja, and James J DiCarlo. Explicit information for category-orthogonal object properties increases along the ventral stream. *Nature neuroscience*, 19(4):613, 2016.
- [22] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [23] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [24] Nancy Kanwisher, Josh McDermott, and Marvin M Chun. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of neuroscience*, 17(11):4302–4311, 1997.
- [25] Nancy Kanwisher, Carol Yin, and Ewa Wojciulik. Repetition blindness for pictures: Evidence for the rapid computation of abstract visual. *Fleeting memories: Cognition of brief visual stimuli*, page 119, 1999.
- [26] Seyed-Mahdi Khaligh-Razavi and Nikolaus Kriegeskorte. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS computational biology*, 10(11):e1003915, 2014.
- [27] Tim Christian Kietzmann, Patrick McClure, and Nikolaus Kriegeskorte. Deep neural networks in computational neuroscience. *bioRxiv*, page 133504, 2018.
- [28] Dorit Kliemann, Hilary Richardson, Stefano Anzellotti, Dima Ayyash, Amanda J Haskins, John DE Gabrieli, and Rebecca R Saxe. Cortical responses to dynamic emotional facial expressions generalize across stimuli, and are sensitive to task-relevance, in adults with and without autism. *Cortex*, 103:24–43, 2018.
- [29] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [30] Daniel Lundqvist, Anders Flykt, and Arne Öhman. The karolinska directed emotional faces (kdef). *CD ROM from Department of Clinical Neuroscience, Psychology section, Karolinska Institutet*, 91:630, 1998.
- [31] Peter Mende-Siedlecki, Yang Cai, and Alexander Todorov. The neural dynamics of updating person impressions. *Social cognitive and affective neuroscience*, 8(6):623–631, 2012.

- [32] Adrian Nestor, David C Plaut, and Marlene Behrmann. Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *Proceedings of the National Academy of Sciences*, 108(24):9998–10003, 2011.
- [33] Alice J O’Toole, Dana A Roark, and Hervé Abdi. Recognizing moving faces: A psychological and neural synthesis. *Trends in cognitive sciences*, 6(6):261–266, 2002.
- [34] Alice J O’Toole, P Jonathon Phillips, Samuel Weimer, Dana A Roark, Julianne Ayyad, Robert Barwick, and Joseph Dunlop. Recognizing people from dynamic and static faces and bodies: Dissecting identity with a fusion approach. *Vision research*, 51(1):74–83, 2011.
- [35] Connor J Parde, Carlos Castillo, Matthew Q Hill, Y Ivette Colon, Swami Sankaranarayanan, Jun-Cheng Chen, and Alice J O’Toole. Deep convolutional neural network features and the original image. *arXiv preprint arXiv:1611.01751*, 2016.
- [36] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, et al. Deep face recognition. In *bmvc*, volume 1, page 6, 2015.
- [37] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [38] Marius V Peelen, Anthony P Atkinson, and Patrik Vuilleumier. Supramodal representations of perceived emotions in the human brain. *Journal of Neuroscience*, 30(30):10127–10134, 2010.
- [39] Marius V Peelen, Alison J Wiggett, and Paul E Downing. Patterns of fmri activity dissociate overlapping functional brain areas that respond to biological motion. *Neuron*, 49(6):815–822, 2006.
- [40] David Pitcher, Daniel D Dilks, Rebecca R Saxe, Christina Triantafyllou, and Nancy Kanwisher. Differential selectivity for dynamic versus static information in face-selective cortical regions. *Neuroimage*, 56(4):2356–2363, 2011.
- [41] Tomaso Poggio and Shimon Edelman. A network that learns to recognize three-dimensional objects. *Nature*, 343(6255):263, 1990.
- [42] Michael I Posner. Abstraction and the process of recognition. In *Psychology of learning and motivation*, volume 3, pages 43–100. Elsevier, 1970.
- [43] Allyson Rice, P Jonathon Phillips, Vaidehi Natu, Xiaobo An, and Alice J O’Toole. Unaware person recognition from the body when face identification fails. *Psychological Science*, 24(11):2235–2243, 2013.

- [44] Rebecca Saxe and Sean Dae Houlihan. Formalizing emotion concepts within a bayesian model of theory of mind. *Current opinion in Psychology*, 17:15–21, 2017.
- [45] Amy E Skerry and Rebecca Saxe. A common neural code for perceived and inferred emotion. *Journal of Neuroscience*, 34(48):15997–16008, 2014.
- [46] Sarah V Stevenage, Mark S Nixon, and Kate Vince. Visual analysis of gait as a cue to identity. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, 13(6):513–526, 1999.
- [47] Chris Thornton. Re-presenting representation. *Forms of representation: An interdisciplinary theme for cognitive science*, pages 152–62, 1996.
- [48] David C Van Essen, Charles H Anderson, and Daniel J Felleman. Information processing in the primate visual system: an integrated systems perspective. *Science*, 255(5043):419–423, 1992.
- [49] Sara C Verosky, Alexander Todorov, and Nicholas B Turk-Browne. Representations of individuals in ventral temporal cortex defined by faces and biographies. *Neuropsychologia*, 51(11):2100–2108, 2013.
- [50] HL Wagner, CJ MacDonald, and AS Manstead. Communication of individual emotions by spontaneous facial expressions. *Journal of Personality and Social Psychology*, 50(4):737, 1986.
- [51] Yang Wu and Laura E Schulz. Inferring beliefs and desires from emotional reactions to anticipated and observed events. *Child development*, 89(2):649–662, 2018.
- [52] Daniel L Yamins, Ha Hong, Charles Cadieu, and James J DiCarlo. Hierarchical modular optimization of convolutional networks achieves representations similar to macaque it and human ventral stream. In *Advances in neural information processing systems*, pages 3093–3101, 2013.
- [53] Daniel LK Yamins and James J DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience*, 19(3):356, 2016.
- [54] Andrew W Young, Freda Newcombe, Edward HF de Haan, Marian Small, and Dennis C Hay. Face perception after brain injury: Selective impairments affecting identity and expression. *Brain*, 116(4):941–959, 1993.
- [55] Galit Yovel and Alice J O’Toole. Recognizing people in motion. *Trends in Cognitive Sciences*, 20(5):383–395, 2016.