

How Diagnostic are Spatial Frequencies for Fear Recognition?

Martial Mermillod

Psychology and NeuroCognition Laboratory
CNRS UMR 5105
University Pierre Mendès France
(martial.mermillod@upmf-grenoble.fr)

Patrik Vuilleumier

Laboratory of Neurology and Imaging of Cognition
University of Geneva
(patrik.vuilleumier@medecine.unige.ch)

Nathalie Guyader

Psychology and NeuroCognition Laboratory
CNRS UMR 5105
University Pierre Mendès France
(nguyader@yahoo.fr)

David Alleysson

Christian Marendaz

Psychology and NeuroCognition Laboratory
CNRS UMR 5105
University Pierre Mendès France
(David.Alleysson@upmf-grenoble.fr)
(Christian.Marendaz@upmf-grenoble.fr)

Abstract

Vuilleumier, Armony, Driver & Dolan (2003) have shown that amygdala cells to fearful expressions of human faces seem to be more activated by intact or low spatial frequency (LSF) faces than high spatial frequency (HSF) faces. These fMRI results may suggest that LSF components might be processed by a subcortical pathway that is assumed to bypass the striate cortex in order to process LSF components faster than HSF components of visual stimuli. The purpose of the present paper is to test the usefulness of LSF information as compared to HSF information in a visual classification task performed by an artificial neural network and a statistical classifier. Our results show that visual information, conveyed by LSF faces, allows the statistical and connectionist models to better recognize or categorize fearful faces amongst neutral faces than HSF faces. These results suggest that high-speed connections from the magnocellular layers to the amygdala might be a fast and efficient way to perform classification of human faces with respect to their emotional expressions.

Introduction

Neuropsychological results have shown “blindsight” for fearful faces in a hemianopic patient (with unilateral destruction of primary visual cortex) when he was exposed to emotional stimuli in his blind visual hemifield (de Gelder, Vroomen, Pourtois & Weiskrantz, 1999; Rossion, de Gelder, Pourtois, Guérit & Weiskrantz, 2000). This has led to the hypothesis that a neural route, by-passing the striate cortex, might reach the amygdala using a subcortical visual pathway from the lateral geniculate nucleus (LGN) through the pulvinar and superior colliculus.

Enroth-Cugell & Robson (1966) reported the spatiotemporal characteristics of X (responding to high-resolution stimuli) and of Y (responding to low-resolution stimuli) retinal ganglion cells; they showed that, following retinal processing, there is a difference between high and low spatial frequencies. Hubel & Wiesel (1977) reported that this distinction remains for the lateral geniculate nucleus: the

magnocellular layers receiving preferentially projections from Y retinal ganglion cells, whereas X cells project to both parvo and magnocellular layers.

Formally, in the visual thalamus, the magnocellular layer is equivalent to a high-pass filter in the temporal frequency domain and a low-pass filter in the spatial frequency domain. Thus, magnocellular neurons mainly provide rapid but low spatial frequency (LSF) information encoding configural features, as well as brightness and motion of objects; whereas the parvocellular neurons provide slower but high spatial frequency (HSF) information about local shape features, color, and texture.

Testing the role of magnocellular inputs in fearful face recognition, Vuilleumier, Armony, Driver & Dolan (2003) conducted a functional magnetic resonance imaging (fMRI) experiment in which human observers were exposed to different spatial frequency components of faces (i.e. LSF only, HSF only, or the integral broad spatial frequency (BSF) images), with either a fearful or a neutral expression. Results showed that HSF and BSF faces produced more activation of the fusiform cortex than LSF faces, irrespective of expression; this suggests predominant contribution of the parvocellular information to the ventral visual stream for face identification. In contrast, the amygdala and subcortical tecto-pulvinar areas were “blind” to the difference of expressions conveyed by HSF information, but selectively activated by fearful relative to neutral faces seen in LSF or BSF images; this suggests an important role of magnocellular information for the activation of amygdala-related circuits in face emotion recognition.

The purpose of the present paper is to examine the usefulness of LSF cues in fearful face recognition by comparing the performance of a distributed neuronal and a statistical models of visual processing exposed to different spatial frequency information. We tested how facial information provided by LSF and HSF images influenced two different computational models for an emotional classification task of face images.

Neuro-computational models

Our simulations were based on two computational models.

Computational model of vision

Several recent advances in computer vision for the categorization of facial emotions have been made during the last decade. Some models have used a feature-based approach (Brunelli & Poggio, 1993), or a more holistic approach based on principal component analysis (Turk & Pentland, 1991; Abdi, Valentin, Edelman & O'Toole, 1995; Cottrell, Branson & Calder, 2002), or non-linear neural-network (Cottrell, 1990). These different techniques, promising at a computational level, do not explore the role of spatial frequency (SF) information. However, some connectionist simulations of visual processes have permitted successful categorization and recognition tasks using Gabor wavelet coding of visual inputs (Cottrell, Branson & Calder, 2002). Dailey & Cottrell (1999) used this technique to differentiate faces from objects. Moreover, Dailey, Cottrell, Padgett & Ralph (2002) have shown by means of Gabor wavelet filtering the possibility to provide good classification performance on database of facial expressions.

Gabor functions provide an efficient way to describe the content of the frequency domain while losing the minimum of information in the spatial domain (Gabor, 1946). Therefore, it was shown that visual information is reliably compressed by Gabor wavelet decomposition. For example, for face recognition, Wiskott (1997); Wiskott, Fellous, Krüger & Von der Malsburg (1999) proposed applying several jets of Gabor wavelets to extract different orientation and spatial frequency information at specific locations. Moreover, at both the computational and behavioral levels, it has been shown that accurate categorization can be achieved using the energy spectrum of natural images (Ginsburg, 1986; Guyader, Chauvin, Peyrin, Hérault & Marendaz, 2004; Hughes, Nozawa & Kitterle, 1996; Hérault, Oliva & Guerin-Dugué, 1997; Mermillod, Guyader & Chauvin, 2004; Torralba & Oliva, 2003).

Our model describes images by sampling their energy spectrum. It is divided into the following steps. First, an Hanning window is applied to avoid an over-representation of vertical and horizontal orientations (due to image edges) in the Fourier domain. After this pre-processing, images were transferred into the Fourier domain using a two-dimensional Fast Fourier Transform algorithm and, then, filtered by a set of Gabor filters. Filter sizes were normalized with respect to a $1/f$ decreasing of the amplitude spectrum for natural images (Field & Brady, 1997). We applied a bank of fifty-six Gabor filters corresponding to seven different spatial frequency bands (one octave per spatial frequency channel) and eight different orientations (each 22.5 deg of visual angle). Then the mean energy at each filter output is measured. An image is then described by 56 different values that correspond to the image energy in different orientation and frequency bands.

Statistical and connectionist models of categorization

We tested two different models in categorization tasks.

The connectionist network involves a distributed model of categorization based on a 3-layer back-propagation neural network. We used the standard hetero-association training algorithm, whose function is to associate each of the different category exemplars with a specific output vector coding for them. This training algorithm is completely supervised because each category is associated with a unique label coding for it. Previous simulations have shown that the combination of these two artificial models allows reliable categorization capacities with respect to empirical data (French, Mermillod, Quinn, Chauvin & Mareschal, 2002; Mermillod, Guyader & Chauvin, 2004).

The statistical model is based on supervised classifier. Using a Principal Component Analysis we reduce the dimension of our data and describe each category by its mean vector and its eigenvectors. Then, test data are projected into the "training" eigenspace where a Mahalanobis distance is applied in order to classify the data. The combination of PCA and Mahalanobis distance is often used for classification purposes. This was also used in Face recognition (Sirovich & Kirby, 1987). The difference here is that, following Dailey & Cottrell (1999), we applied PCA not to the face images but to the Gabor responses to the face images.

The aim of these simulations was to test the role of low spatial frequency content in faces on the expression recognition performance of a distributed classifier network. In the case of a failure of the neural network to categorize emotions based on LSF images only, the hypothesis of an important functional role of coarse (subcortical) magnocellular inputs to the amygdala would have to be seriously questioned.

Simulation 1: Testing spectral information in a connectionist network

Network

We used a standard 24-6-2 feedforward backpropagation hetero-associator (learning rate: 0.1, momentum: 0.9).

Stimuli

For all simulations, stimuli were the original stimuli used in the neuro-imaging study by Vuilleumier et al. (2003). These included 160 human faces from two categories (80 neutral face exemplars and 80 fearful face exemplars). Each of 80 different individuals appeared with the two emotional expressions (fearful vs. neutral), always in a frontal viewpoint. Face images were grey level photographs with an average stimulus luminance, on a 256 gray-level scale, of 112, 118, and 115 for BSF, HSF, and LSF stimuli, respectively, and of 117 and 114 for the neutral and fearful face categories, respectively. These average luminance values did not significantly differ across the different stimulus conditions (Vuilleumier et al., 2003). The size of all images was squared to the same frame for computational reasons, by applying an area of 198×198 pixels on the centre of each face, in

such a way to retain a similar amount of information for each stimulus and to keep all internal facial details from the original images (from the base of the chin to the top of the forehead). And as we described in the first part, each image is then described by 56 different values that correspond to the image energy in 8 different orientation and 6 different frequency bands.

In their fMRI study, Vuilleumier et al. (2003) used a high-pass cut-off >24 cycles per image for HSF faces and a low-pass cut-off <6 cycles per image for LSF faces. In order to reproduce this cut-off in the connectionist simulations, we removed the 4 lowest spatial frequency channels (or Gabor filters), coding for HSF faces, and the 4 highest spatial frequency channels (Gabor filters), coding for LSF faces. Thus, we kept the three highest SF bands for the HSF face inputs, and the three lowest SF bands for the LSF face inputs. This method allowed us to remove spectral information that was not relevant for one or the other simulation, while keeping the same vector-size for both simulations. Consequently, the input vector size was 24 units (3 spatial frequency bands by 8 orientations).

Procedure

The procedure included two phases: a training phase with a subset of fearful and neutral faces, and a testing phase in which the neural network was tested on its ability to categorize new facial expressions. The procedure described below concerns LSF faces. Exactly the same procedure as described below was then applied on HSF faces.

Training phase. Twenty LSF fearful and twenty corresponding LSF neutral faces were randomly extracted from the categories of emotional expression. Then, the 24-dimensional energy vector was associated with the corresponding output vector by the 3-layer back-propagation network. Then, a new image from the training set was coded and associated by the neural network in an iterative process. Each run began with a random selection of 40 training exemplars (20 exemplars per category). Then the training consisted of associating each of the 20 exemplars with the suitable output (0 1 coding for “fearful” face, 1 0 coding for “neutral” face) for a fix number of 500 epochs.

Test phase. The neural network was trained on the two expression categories and then tested on the 60 remaining exemplars from the trained category versus the corresponding 60 exemplars from the other category. Results were averaged over 50 runs of the above training-test procedure. After applying a winner-take-all on the output nodes, the dependent measure was the correctness of the outputs produced by the tested vector.

Results

After training on the low spatial frequency of natural images, the neural network produced an average of 94.3 % of correct response when tested on new fearful faces. When tested on new neutral faces, the network produced 94.4 % of correct responses. The difference between the two test conditions is not significant.

After training on the high spatial frequency of natural images, the neural network produced an average of 87.2 % of correct response when tested on new fearful faces. When tested on new neutral faces, the network produced 92.1 % of correct responses. More importantly, the difference between generalization performance produced by LSF fearful faces compared to HSF fearful faces was significant ($\chi^2(1)=90.36$, $p<.001$). Similarly, the difference of performance between LSF neutral faces compared to HSF neutral faces was significant ($\chi^2(1)=12.6$, $p<.001$).

Discussion

This first simulation has important implications for the neurobiological and cognitive underpinnings of emotional face recognition, particularly fearful expressions. The lower performance produced by the network after training on HSF information suggests a problem for a distributed classifier in this task.

Therefore, we suggest that the statistical distribution of LSF and HSF faces in terms of their spectral energy vector may provide a clear explanation for human imaging and connectionist results. The orientation and spatial frequency decomposition occurring in the human visual system is able to provide a pattern of responses that clearly distinguish between fearful and neutral expressions at the level of LSF information, whereas HSF information is worse for this particular classification task. Therefore, the LSF information provided by the magnocellular layers may be capable of providing the necessary information to do the classification of fearful expressions.

Simulation 2: Testing spectral information in a statistical model

Statistical model

We used, for this second simulation, a classical Principal Component Analysis (PCA) in order to reduce the dimensionality of our data. Then, to classify images we used the Mahalanobis distance.

Stimuli

The stimuli were exactly the same as the one used in the simulation 1. (cf. Simulation 1). We had two sets of data: one set for LSF and another one for HSF. Each set corresponded to a 160×24 matrix (160 different faces, 80 neutral and 80 fearful, each described with 24 values).

Procedure

The procedure also included two phases: a training phase with a subset of fearful and neutral faces, and a testing phase in which the neural network was tested on its ability to categorize new facial expressions. The procedure described below concerns LSF faces. Exactly the same procedure as described below was then applied on HSF faces.

Training phase. The 2 principal eigenvectors of the 30 LSF data were computed using a Principal Component Analysis

(PCA). These 2 principal vectors preserved around 90% of the total variance. Then each face category: neutral and fearful is described by its gravity center in this reduced space.

Test phase. The remaining data, the test ones, were projected on the 2- eigenvector space computed during the training phase. Then, in order to attribute a class to each test data we computed the Mahalanobis distance between each test data and the learnt gravity center of each class. The Mahalanobis distance is proved to be a good distance for classification purpose (Yambor, W., Draper, B. & Beveridge, 2000).

Results

This classification reaches a percentage of correct classification of approximately 89% for LSF data, and only 58% for HSF data.

Discussion

Globally, these results imply that the whole spatial frequency spectrum available from the retinal image may not be entirely needed to perform a visual categorization of human faces in terms of basic emotional expressions, such as fear vs. neutral. Therefore, it might indeed be useful for a distributed system (i.e. visual and emotional pathways in the brain) to exploit the most rapid neuronal pathways conveying LSF information (i.e., the magnocellular channel), in order to achieve sufficiently reliable but also fast categorization of fearful stimuli.

Conclusions

The purpose of that paper was to explore the computational basis in support of the hypothesis that LSF pathways within the visual system may be preferentially responsible for carrying visual information to the amygdala about the emotional (fearful) expression of faces. However, our network model did not make any definite assumptions about the anatomical neuronal stream potentially involved in this visual processing pathway (i.e., the geniculostriate cortical stream vs. the tectopulvinar subcortical stream). Taking our different simulations together, the main results suggest that an artificial model of categorization can perform a more reliable categorization of faces in terms of emotional expression based on their LSF content rather than their HSF content. This provides indirect support for a computational advantage of extracting LSF cues from faces, as previously hypothesized for the amygdala on the basis of brain imaging results showing greater activation to LSF than HSF fearful faces (Vuilleumier et al., 2003).

In the human visual system, LSF inputs from magnocellular visual neurons project to a wide range of different areas including subcortical tectopulvinar regions (Schiller et al. 1979; Orban, 1984) and frontoparietal cortical areas (Bullier, 2001; Bar, 2004), but also to ventral temporal cortex (Livingstone & Hubel, 1988; Merigan & Maunsell, 1993). By contrast, HSF information from parvocellular neurons predominantly projects to the ventral temporal cortex. The

amygdala may therefore receive LSF inputs from either subcortical or cortical pathways, although preserved activation by fearful faces during masked presentations or in blind patients (Morris et al., 2001; de Gelder et al., 1999; Pegna et al., 2004) may suggest an important role of the subcortical tectopulvinar pathways known to carry LSF inputs (Vuilleumier et al., 2003).

On the other hand, it has been shown that HSF information can also play an important role in face processing, particularly for the accurate identification of specific exemplars (Morrison & Schyns, 2001). Therefore, depending on the situational constraints (i.e. identify a target or categorize a stimulus in terms of danger), one or the other spatial frequency channel might be preferentially used by the cognitive system to deal efficiently with its visual environment.

Furthermore, our new results do not only support the observations previously made at a neurophysiological level concerning the possible substrates for fast, non-conscious process of fearful faces (de Gelder et al., 1999; Vuilleumier et al., 2003), but also more generally provide additional evidence for the hypothesis of coarse-to-fine processing in visual recognition. The coarse-to-fine hypothesis suggests an advantage of LSF information for the initial categorization of visual objects or scenes (Ginsburg, 1986; Parker, Lishman, & Hughes, 1992, 1997; Parker & Costen, 1999; Schyns & Oliva, 1994), prior to finer visual analysis based on HSF information. These psychological data are supported by anatomical evidence showing faster LSF integration at the level of the magnocellular layers in the lateral geniculate nucleus of the thalamus (Hubel & Wiesel, 1977). In other words, a fast propagation of LSF information within the perceptual system might constitute an efficient mechanism for the fast categorization of visual stimuli into most salient or relevant entities (see also Bullier 2001, Bar, 2004). Thus, rapid connections from the magnocellular visual neurons in early thalamic and other subcortical relays to the amygdala might be in general agreement with the computational demands of a distributed cognitive system. Such a functional architecture would be highly consistent with results from the present simulations showing more efficient visual classification of facial expressions based on their LSF content (rather than HSF alone), and with previous neuro-imaging results showing more robust amygdala activation to fearful faces seen from LSF images (Vuilleumier et al., 2003).

These empirical results reported here provide a first attempt to understand the complex relationships unifying basic visual perceptual processes with higher cognitive and emotional recognition systems. A next step will be to use such computational modeling to simulate and to predict further empirical results. Based on their elementary physical properties, it is possible to generate stimuli for which magnocellular pathways would be completely blind, and then, test the response of recognition systems for different emotional categories and different categorization processes. Schyns & Oliva (1999) have reported psychological evidence showing that different regions of the SF spectrum are used depending on the task: LSF is preferentially used to describe facial emotions in explicit terms

of happy, angry or neutral whereas HSF seems to be used to determine if a face is expressive or not. Future simulation models should therefore also investigate whether training on the same set of faces may lead to specialized processing streams (i.e., at the level of hidden-layer or in different subparts of the network) extracting distinct LSF or HSF components for different task purposes (e.g. emotion recognition based on LSF in some neurons, identity or age recognition based on HSF in other neurons).

Acknowledgments

This work was supported by a post-doctoral grant from the Fyssen Foundation to MM, the French CNRS to DA and CM and a grant from the Swiss National Science Foundation to PV.

Bibliography

- Abdi, H., Valentin, D., Edelman, B.E., O'Toole, A.J. (1995). More about the difference between men and women: Evidence from linear neural networks and the principal component approach. *Perception*, 24, 539-562.
- Bar, M. (2004). Visual objects in context. *Nature Reviews: Neuroscience*, 5, 619-629.
- Breiter H.C., Etcoff N.L., Whalen P.J., Kennedy W.A., Rauch S.L., Buckner R.L., Strauss M.M., Hyman S.E. & Rosen B.R. (1996). Response and habituation of the human amygdala during visual processing of facial expression. *Neuron* 17(5), 875-887.
- Brunelli, R., & Poggio T. (1993). Face Recognition: Features versus Templates. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 15(10), 1042-1052.
- Cottrell, G.W. (1990). Extracting features from faces using compression networks: Face, identity, emotion, and gender recognition using holons. In *Proceedings of the 1990 Connectionist Models Summer School*, (eds. D. Touretsky, J. Elman, T. Sejnowski & G. Hinton) Kaufman, 328-337.
- Cottrell, G. W., Branson, K. & Calder A. J. (2002). Do expression and identity need separate representations? In *Proceedings of the 24th Annual Cognitive Science Conference*, Fairfax, Virginia. Mahwah: LEA.
- Dailey, M. N. & Cottrell, G. W. (1999). Organization of Face and Object Recognition in Modular Neural Networks. *Neural Networks*, 12(7-8), 1053-1074.
- Dailey, M. N., Cottrell, G. W., Padgett, C., & Ralph A. (2002). EMPATH: A neural network that categorizes facial expressions. *Journal of Cognitive Neuroscience* 14(8), 1158-1173.
- Enroth-Cugell, C., & Robson, J. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *Journal of Physiology*, 187, 517-552.
- Esteves, F., & Ohman, A. (1993). Masking the face: Recognition of emotional facial expressions as a function of the parameters of backward masking. *Scandinavian Journal of Psychology*, 34, 1-18.
- Field, D.J. & Brady N. (1997). Visual sensitivity, blur and the sources of variability in the amplitude spectra of natural scenes. *Vision Research*, 37, 3367-3383.
- French, R. M., Mermillod M., Quinn P. C., Chauvin A. & Mareschal D. (2002). The Importance of Starting Blurry: Simulating Improved Basic-Level Category Learning in Infants Due to Weak Visual Acuity. *Proc. of the 24th Annual Cog. Sci. Society Conference*. NJ:LEA. 322-327.
- Gabor, D. (1946). Theory of Communication. The Journal of the Institution of Electrical Engineers. London: Unwin Brothers. 93(3): 429-457.
- Ginsburg, A. P. (1986). Spatial filtering and visual form perception. In K. Boff, L. Kaufman, & J., Thomas (Eds.), *Handbook of perception and human performance. Volume 2: Cognitive processes and performance* (pp. 34-1 to 34-41). New York: Wiley.
- Guyader, N., Chauvin, A., Peyrin, C., Héroult, J., & Marendaz, C. (2004). Image phase or amplitude? Rapid scene categorization is an amplitude based process. *C. R. Biologies* 327, 313-318.
- de Gelder B., Vroomen, J., Pourtois G. & Weiskrantz, L. (1999). Non-conscious recognition of affect in the absence of striate cortex. *NeuroReport*, 10(18), 3759-3763.
- Héroult, J., Oliva, A., & Guerin-Dugué, A. (1997). Scene Categorisation by Curvilinear Component Analysis of Low Frequency Spectra. 5th European Symposium on Artificial Neural Network., Bruges, Belgium. pp. 91-96.
- Hubel, H. D. & Wiesel, T. N. (1977). Ferrier lecture: Functional architecture of macaque monkey visual cortex. *Proc. Roy. Soc. Lond. [Biol.]*, 98, 1-59.
- Hughes, H.C., Nozawa, G. & Kitterle, F. (1996). Global precedence, spatial frequency channels, and the statistics of natural images. *Journal of Cognitive Neuroscience*, 8, 197-230.
- Jones, J.P. & Palmer L.A. (1987). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58, 1187-1211.
- Jones, J.P., Stepnoski A. & Palmer L.A. (1987). The two-dimensional spectral structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58(6), 1212-1232.
- Livingstone, M., & Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science*, 240,(4853), 740-749.
- Lundqvist, D. & Litton, J.E (1998). The Karolinska Directed Faces (Karolinska Institute).
- Merigan, W. H., & Maunsell, J. H. (1993). How parallel are the primate visual pathways? *Annual Review of Neuroscience*, 16, 369-402.
- Mermillod M., Guyader N. & Chauvin A. (2004). Does the energy spectrum from Gabor wavelet filtering represent sufficient information for neural network recognition and classification tasks? In H. Bowman, C. Labiouse (Eds.) *Connectionist Models of Cognition, Perception and Emotion II. Progress in Neural Processing* (vol. 15). World Scientific,

- pp 148-156.
- Morris, J.S., Frith, C.D., Perrett, D.I., Rowland, D., Young, A.W., Calder, A.J. & Dolan, R.J. (1996). A differential neural response in the human amygdala to fearful and happy facial expressions. *Nature*, 383(6603), 812-5.
- Morris, J.S., Ohman, A. & Dolan, R.J. (1998). Conscious and unconscious emotional learning in the human amygdala [see comments]. *Nature*, 393(6684), 467-70.
- Morrison, D.J. & Schyns, P.G. (2001). Usage of spatial scales for the categorization of faces, objects and scenes. *Psychonomic Bulletin & Review*, 8, 454-469.
- Orban, G.A. (1984). Neuronal operations in the visual cortex. *Studies in brain function*, Vol. II. Berlin: Springer-Verlag.
- Parker, D. M. & Costen, N. P. (1999). One extreme or the other or perhaps the golden mean? Issues of spatial resolution in face processing. *Current Psychology*, 18, 118-127.
- Parker, D. M., Lishman, J. R. & Hughes, J. (1992). Temporal integration of spatially filtered visual images. *Perception*, 21, 147-160.
- Parker, D. M., Lishman, J. R., & Hughes, J. (1997). Evidence for the view that temporospatial integration in vision is temporally anisotropic. *Perception*, 26, 1169-1180.
- Pegna, A.J., Khateb, A., Lazeyras, F. & Seghier, M.L. (2004). Discriminating emotional faces without primary visual cortices involves the right amygdala. *Nature Neuroscience*, 8(1) 24-25.
- Rossion, B., de Gelder B., Pourtois G., Guérit J.M. & Weiskrantz, L. (2000). Early extrastriate activity without primary visual cortex in humans. *Neuroscience Letters*, 279(1), 25-28.
- Schiller, P. H., Malpeli, J.G. & Schein, S. J. (1979). Composition of geniculostriate input to the superior colliculus of the rhesus monkey. *Journal of Neurophysiology*, (42), 1124-1133.
- Schyns, P. G., & Oliva, A. (1994). From blobs to boundary edges: Evidence for time and spatial-scale-dependent scene recognition. *Psychological Science*, 5, 195-200.
- Schyns P.G. & Oliva A. (1999). Dr. Angry and Mr. Smile: when categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition*, 69(3), 243-265.
- Sirovich, L. & Kirby, M. (1987). A low-dimensional procedure for the characterization of human faces. *The Journal of the Optical Society of America*, 4:519 – 524.
- Torrallba, A. & Oliva, A. (2003). Statistics of natural image categories. *Network: Comput. Neural Syst.*, 14, 391–412
- Turk, M. & Pentland, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1), 71-86.
- Vuilleumier, P., Armony, J. L., Driver, J., & Dolan, R. J. (2003). Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature Neuroscience*, 6(6), 624-631.
- Wiskott L. (1997). Phantom Faces for Face Analysis. *Pattern Recognition* 30(6), 837-846.
- Wiskott, L., Fellous, J.M., Krüger, N. & Von der Malsburg C. (1999). Face Recognition by Elastic Bunch Graph Matching. *In Intelligent Biometric Techniques in Fingerprint and Face Recognition*, eds. L.C. Jain et al., CRC Press, 11, 355-396.
- Yambor, W., Draper, B. & Beveridge, R. (2002). Analyzing PCA-based Face Recognition Algorithms: Eigenvector Selection and Distance Measures, in *Empirical Evaluation Methods in Computer Vision*, H. Christensen and J. Phillips (eds.), World Scientific Press, Singapore, 2002