

Simulating Pareidolia of Faces for Architectural Image Analysis

Stephan K. Chalup and Kenny Hong

Newcastle Robotics Laboratory
School of Electrical Engineering and Computer Science
The University of Newcastle, Callaghan 2308, Australia
Stephan.Chalup@newcastle.edu.au, Kenny.Hong@uon.edu.au

Michael J. Ostwald

School of Architecture and Built Environment
The University of Newcastle
Callaghan 2308, Australia
Michael.Ostwald@newcastle.edu.au

Abstract

The hypothesis of the present study is that features of abstract face-like patterns can be perceived in the architectural design of selected house façades and trigger emotional responses of observers. In order to simulate this phenomenon, which is a form of pareidolia, a software system for pattern recognition based on statistical learning was applied. One-class classification was used for face detection and an eight-class classifier was employed for facial expression analysis. The system was trained by means of a database consisting of 280 frontal images of human faces that were normalised to the inner eye corners. A separate set of test images contained human facial expressions and selected house façades. The experiments demonstrated how facial expression patterns associated with emotional states such as surprise, fear, happiness, sadness, anger, disgust, contempt or neutrality could be identified in both types of test images, and how the results depended on preprocessing and parameter selection for the classifiers.

1. Introduction

It is commonly known that humans have the ability to 'see faces' in objects or random structures which contain patterns such as two dots and line segments that abstractly resemble the configuration of the eyes, nose, and mouth in a human face. Photographers François and Jean Robert collected a whole book of photographs of objects which seem to display face-like structures [110]. Simple abstract typographical patterns such as emoticons in email messages are not only associated with faces but also with emotion categories. The following emoticons using Western style typography are very common.

:)	happy	:D	laugh
:-)	sad	:?	confused
:O	surprised	:3	love
;-)	wink	:	concerned

Pareidolia is a psychological phenomenon where a vague or diffuse stimulus, for example, a glance at an unstructured background or texture, leads to simultaneous perception of the real and a seemingly unrelated unreal pattern. Examples are faces, animals or body shapes seen in walls, clouds, rock formations, or trees. The term originates from the Greek 'para' (παρά = beside or beyond) and 'eidōlon' (εἶδωλον = form or image) and describes the human visual system's tendency to extract patterns from noise [92]. Pareidolia is a form of apophenia which is the perception of connections and associated meaning of unrelated events [43]. These phenomena were first described within the context of psychosis [21, 42, 49, 124] but are regarded as a tendency common in healthy people [3, 108, 128] and can explain or inspire associated visual effects in arts and graphics [86, 92].

Various aspects of architectural design analysis have contributed to questions such as: How do we perceive aesthetics? What determines whether a streetscape is pleasant to live in? What visual design features influence our well-being when we live in a particular urban neighbourhood? Some studies propose, for example, the involvement of harmonic ratios, others calculate the fractal dimension of façades and skylines to determine their aesthetic value [11, 13, 14, 22, 57, 96, 95, 135]. Faces frequently occur as ornaments or adornments in the history of architecture in different cultures [9].

The present study addresses a hypothesis that is inspired by the phenomenon of pareidolia of faces and by recent results from brain research and cognitive science which show that large areas of the brain are dedicated to face processing, that the perception of facial expressions involves the emotional centres of the brain [31, 141, 142], and that (in contrast to other stimuli [79]) faces can be processed sub-cortically, non-consciously, and independently of visual attention [41, 71]. Recent brain imaging studies using magnetoencephalography suggest that the perception of face-like objects has much in common with the perception of real faces. Both occur (in contrast to non-face objects) as a relatively early processing step of signals in the brain [56]. Our hypothesis is that abstract face expression features that ap-

pear in the architectural design of house façades trigger via a pareidolia effect emotional responses in observers. These may contribute to inducing the percept of aesthetics in the observer. Some pilot results of our study have been presented previously [12]. Related ideas have attracted attention in the area of car design where associations of frontal views of cars with emotions or character are claimed to influence sales volume [145]. A recent study investigated how the human ability to detect faces and to associate them with emotion features transfers to objects such as cars [147].

The topic of face recognition traditionally plays an important role in cognitive science, particularly in research on object perception and affective computing [17, 32, 78, 89, 103, 125, 126, 153]. A widely accepted opinion is that face recognition is a special skill, distinct from general object recognition [91, 78, 105]. It has frequently been reported that psychiatric and neuropathological conditions can have a negative impact on the ability to recognise facial expression of emotion [58, 72, 73, 74, 138]. Farah and colleagues [33, 34, 36] suggested that faces are processed holistically and in specific areas of the human brain, the so-called fusiform face areas [37, 77, 78, 107]. Later studies confirmed that activation in the fusiform gyri plays a central role in the perception of faces [50, 101] and that a number of other specific brain areas also showed higher activation when subjects were confronted with facial expressions than when they were shown images of neutral faces [31]. It was shown that the fusiform face areas maintain their selectivity for faces independently of whether the faces are defined intrinsically or contextually [25]. From recognition experiments using images of faces and houses, Farah [34] concluded that holistic processing is more dominant for faces than for houses.

Prosopagnosia, the inability to recognise familiar faces while general object recognition is intact, is believed by some to be an impairment that exclusively affects a subject's ability to recognise and distinguish familiar faces and may be caused by damage of the fusiform face areas of the brain [55]. In contrast, there is evidence which indicates that it is the expertise and familiarity with individual object categories which is associated with holistic modular processing in the fusiform gyrus and that prosopagnosia not only affects processing of faces but also of complex familiar objects [45, 46, 47, 48]. These contrasting opinions are the subject of on-going discussion [35, 44, 90, 98, 109]. Although the debate about how face processing works is far from over, a developmental perspective suggests that 'the ability to recognize faces is one that is learned' [94, 120]. It is assumed that the 'learning process' has ontogenetic and phylogenetic dimensions and drives the development of a complex neural system dedicated for face processing in the brain [26, 91, 100].

In order to parallel nature's underlying concept of 'learning' and/or 'evolution', the first milestone of the present project was to design a simple face detection and facial expression classification system purely based on pattern recognition by statistical learning [59, 140] and train it on images of faces of human subjects. After optimising the system's learning parameters using a data set of images of human facial expressions, we assumed that the system represented a basic statistical model of how human subjects would detect faces and classify facial expressions. An evaluation of the system when applied to images of selected house façades should allow us to test under which conditions the model can detect facial features and assign façade sections to human facial expression categories.

There is quite a large body of work on computational methods for automatic face detection and facial expression classification. For face detection a variety of different approaches have been successfully applied, for example, correlation template matching [8], eigenfaces [102, 139] and variations thereof [143, 151], various types of neural networks [17, 112, 123], kernel methods [20, 24, 60, 65, 69, 70, 97, 106, 111, 150, 151] and other dimensionality reduction methods [18, 27, 54, 66, 131, 148, 154]. Some of the methods focus specifically on improvements under difficult lighting conditions [23, 113, 133], non-frontal viewing angles [5, 19, 81, 84, 119, 121, 155], or real-time detection [121]. More details can be found in survey papers on various aspects of face detection and face recognition [1, 6, 17, 61, 80, 83, 87, 125, 126, 152, 155, 156, 157]. Other papers specifically highlight affect recognition or facial expression classification [38, 62, 63, 68, 99, 114, 122, 123, 130, 136, 149, 153]. Related technology has been implemented in some digital cameras such as the Sony DSC-W120 with Smile Shutter (TM) technology. This camera can analyse facial features such as lip separation and facial wrinkles in order to release the shutter only if a smiling face is detected [67]. Some recent face detection methods aim at detecting and/or tracking particular individual faces [2, 7, 132] and some of the methods are able to estimate gender [24, 52, 53, 54, 88] or ethnicity [64, 158]. Multimodal approaches [93, 115, 144] and techniques for dynamic face or expression recognition [51, 137] appear to be particularly powerful. Recent interdisciplinary studies demonstrated how a computer can learn to judge the beauty of a face [75] or how to perform facial beautification in simulation [82].

The remainder of the present paper is structured as follows. In Section 2 a description of the system is given, which includes modules for preprocessing, face detection and facial expression classification. The experimental results are presented and discussed in Section 3. Section 4 contains a summarising discussion and conclusion.

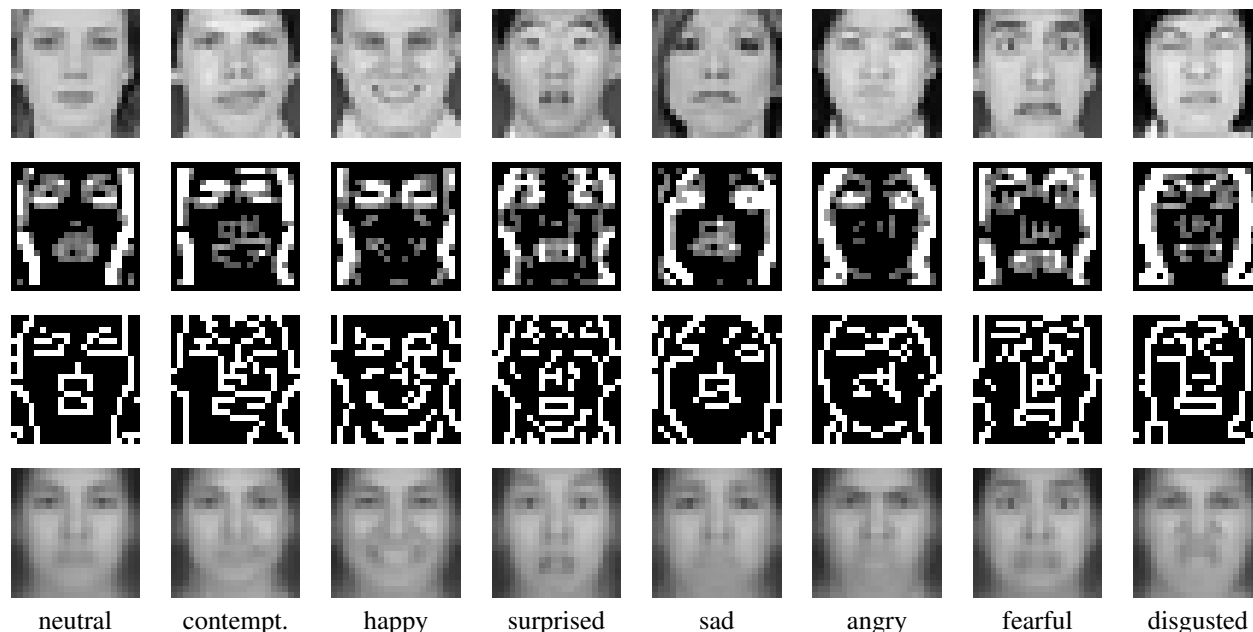


Figure 1. Training data normalised to the inner eye corners in 22×22 pixel resolution; *First row:* Examples of greyscale images, *second row:* Sobel edge images, *third row:* Canny edge images. *Fourth row:* For each expression category the averages of all associated greyscale images in the training set are displayed. The underlying images stem from the JACFEE and JACNeuF image data sets (©Ekman and Matsumoto 1993) [28, 29].

2. System and Method Description

The aim was to design and implement a clearly structured system based on a standard statistical learning method and train it on human face data. In a very abstract way this should simulate how humans learn to recognise faces and assess their expressions. The system should not rely on domain-specific techniques from human face processing, such as eye and lip detection used in some of the current systems for biometric human face recognition.

A significant part of the project addressed data selection and preparation. The final design was a modular system consisting of a preprocessing module followed by two levels of classification for face detection (one-class classifier) and facial expression classification (multi-class classifier).

2.1. Face Database

The set of digital images for training the classifiers for face detection and facial expression classification consisted of 280 images of human faces taken from the research image data sets of Japanese and Caucasian Facial Expressions of Emotion (JACFEE) and Japanese and Caucasian Neutral Faces (JACNeuF) (©Ekman and Matsumoto 1993) [28, 29]. All images in the training set were cropped

and resized to 22×22 pixels so that each showed a full individual frontal face in such a way that the inner eye corners of all faces appeared in exactly the same position. This normalisation step helped to reduce the false positive rate. Profiles and rotated views of faces were not taken into account.

For training the expression classifier, the images were labelled according to the following eight expression classes: neutral, contemptuous, happy, surprised, sad, angry, fearful, and disgusted, as shown by representative sample images in Figure 1. Half of the training data (140 images) showed neutral expressions. The other half of the training set was composed of images of the remaining seven expression classes, each represented by 20 images.

The images of human faces for testing generalisation (shown in Figure 2) were selected from a separate database, the Cohn-Kanade human facial expression database [76]. None of these images was used for training. The test images for house façades in Figures 3 to 6 were sourced from the author's own image database.

2.2 Preprocessing Steps

The preprocessing module converts all images into greyscale. This can be followed by histogram equalisation and/or application of an edge filter.




Figure 2. The trained SVMs for face detection ($\nu = 0.1$) and expression classification ($\nu = 0.1$) were applied to a squared test image which was assembled from four images that were taken from standard database face images [76] (©Jeffrey Cohn). Face detection was based on equalised greyscale (left column), Sobel edge images (middle column), or Canny edge images (right column). The upper left face within each test image was classified as ‘disgusted’ = green, the upper right face was classified as ‘angry’ = red, and the bottom right face was identified as ‘happy’ = white. The bottom left face was not detected in the equalised greyscale test image. Otherwise, the dominant class of the bottom left face was ‘neutral’ = grey. In the case of the Canny edge filtering, additional face boxes were detected with relatively high decision values including two smaller face boxes in which the nose openings were mistakenly recognised as eyes. That is, the system performs as desired, with a tendency to false negatives in the case of equalised greyscale and a tendency to false positives in the case of additional Canny edge filtering.

Histogram equalisation [129] compensates for effects owing to changes in illumination, different camera settings, and different contrast parameters between the different images. In many (but not all) cases, histogram equalisation can have a significant impact on edge detection and system performance.

Equalised or non-equalised greyscale images were either directly used for training and testing or they were converted into edge images with Sobel [127] or Canny [10] edge filters. Examples are shown in Figures 1 and 2. Sobel and Canny edge operators require several parameters to be cho-

sen that can have significant impact on the resulting edge image. We used the ‘Filters’ library v3.1-2007_10 [40]. The selection of the Canny and Sobel filter parameters was based on visual evaluation of ideal facial edges in selected training images. For both filters we used a lower threshold of 85 [0-255] and an upper threshold of 170 [0-255]. Additional parameters for the Sobel filter were blur = 0 [0-50] and gain = 5 [1-10].




Figure 3. Example where one dominant face is detected within a façade but the associated facial expression class depends on the size of the box. The small black boxes represent the category ‘fearful’, blue boxes denote a ‘sad’ expression while the larger violet boxes denote ‘surprise’. This is mostly consistent between the two different aspects of the same house in the left and right images. Only the boxes with the highest decision values are displayed. The bottom row shows the Canny edge images of the above equalised greyscale pictures. Both SVMs, for face detection and expression classification, used $\nu = 0.1$. This is the same parameter setting and preprocessing used for the right top result shown in Figure 2.

2.3 Support Vector Machines for Classification

The present study employed ν -support vector machines (ν -SVMs) with radial basis function (RBF) kernel [116, 118] as implemented in the libsvm library [16]. The ν parameter in ν -SVMs replaces the C parameter of standard SVMs and can be interpreted as an upper bound on the fraction of margin errors and a lower bound on the fraction of support vectors [4]. Margin errors are points that lie on the wrong side of the margin boundary and may be misclassified. Given training samples $x_i \in \mathbf{R}^n$ and class labels $y_i \in \{-1, +1\}$, $i = 1, \dots, k$, SVMs compute a binary de-

cision function. In case of a one-class classifier, it is determined if a particular sample is a member of the class or not. Platt [104] proposed a method to employ the SVM output to approximate the posterior probability $Pr(y = 1|x)$. An improvement of Platt’s method was proposed by Lin et al. [85] and has been implemented in libsvm since version 2.6 [16]. In the experiments of the present study the posterior probabilities output by the SVM on test samples are interpreted as decision values that indicate the ‘goodness’ of a face-like pattern.

In pilot experiments, visual evaluation was used to select suitable values for the parameter ν within the range from

0.001 to 0.5. The width parameter γ of the RBF kernel was left at libsvm's default value of 0.0025. It was found that $\nu = 0.1$ was a suitable value for the SVMs of both, face detection and facial expression classification stages, in greyscale, Sobel, and Canny filtered images that were first equalised. This parameter setting was used to obtain all of the reported results, except the results shown in Figure 5.

2.4 Face Detection

The central component for the face detection module is a one-class support vector machine (SVM) with radial basis function (RBF) kernel [16, 117, 134]. Input to the classifier is an image array of size $22 \times 22 = 484$ where pixel values ranging from zero to 255 were normalised into the interval $[-1, 1]$. Output of the classifier is a decision value which, if positive, indicates that the sample belongs to the learned model class (i.e. it is a face). Support Vector Machines (SVMs) were previously successfully employed for face detection by several authors, for example, [60, 65, 70, 97, 106, 111].

Our basic face detection module performs a pixel-by-pixel scan of the image to select boxes and then tests if they contain a face. The procedure can be described as follows:

Step 1: Given a test image, select a centre point (x, y) for a box within the image. Start at the top left corner of the image at pixel $(x, y) = (11, 11)$ (i.e. distance to the boundary is half of the diameter of the intended 22×22 box). In later iterations scan the image deterministically column by column and row by row.

Step 2: For each centre point select a box size starting with 22×22 . In each later iteration increase the box size by one pixel as long as it fits into the image.

Step 3: Crop the image to extract the interior of the box generated around centre point (x, y) and rescale the interior of the box to a 22×22 pixel resolution.

Step 4: At this step histogram equalisation and/or Canny or Sobel edge filters can be applied to the interior of the box. Note that an alternative approach with possibly different results would be to apply the filters first to the whole image and then extract and classify the candidate face boxes.

Step 5: Feed the resulting 22×22 array into the trained one-class SVM classifier to decide if the box contains a face. If the box contains a face store the decision value and colour the centre pixel yellow.

Step 6: Continue the loop started in Step 2 and increase box size until the box does not fit into the image area. Then continue the outer loop that was started in

Step 1 by progressing to the next pixel to be evaluated as centre point of a potential face box.

At the completion of the scan, each of the evaluated box centre points can have assigned several positive decision values for differently-sized associated face boxes. If a pixel was assigned several values, only the box with the highest decision value for that pixel was kept.

The procedure up to this point generated a cloud of candidate solutions (shown as yellow clouds in Figures 2 to 7) consisting of centre points of boxes with the highest positive decision values output by the one-class SVM. Note that every pixel within a 'face cloud' had a positive decision value (if the value was negative it meant that the pixel was not associated with a face box).

Within the yellow face clouds local peaks of decision values can be identified and highlighted by means of the following filter procedure:

1. Randomly select a pixel with positive decision value and examine a 3×3 area around it.
2. If the centre pixel has the highest decision value, flag it as a local peak. Otherwise, move to the pixel within the group of nine which has the highest decision value and evaluate the new group.
3. Repeat until all pixels with positive decision values (i.e. those in the yellow clouds) have been examined.

The resulting coloured pixels displayed within the yellow face clouds indicate faces associated with local peaks of high decision values. The colours indicate the associated facial expression classes as explained further below.

2.5 Facial Expression Classification

Affect recognition has become a wide field [153]. Good results can be obtained through multi-modal approaches; Wang and Guan [144] combined audio and visual recognition in a system capable of recognising six emotional states in human subjects with different language backgrounds with a success rate of 82%. The purpose of the present study was to evaluate architectural image data using a clearly structured statistical learning system. Therefore, a purely vision based approach had to be adopted and good classification accuracy was not the highest priority.

As the facial expression classifier, an eight-class ν -SVM [116] with radial basis function (RBF) kernel was trained on the labelled data set of 280 images (from Section 2.1). Eight classes corresponding to the facial expression classification system's (FACS) eight emotional states were distinguished [28]. Face expressions were colour coded via the frames of the boxes which were determined to contain a face by the face detection module in the first stage of the



Figure 4. Example where the system detects several dominant ‘faces’ within the same façade and there is some consistency in detection and classification (all SVMs used $\nu = 0.1$) between the different aspects of the same house in the left and right images. Only boxes with the highest decision values are displayed. The bottom row shows the associated Sobel edge images.

system. The following list describes which colours were assigned to which facial expressions of emotion:

sad	=	blue
angry	=	red
surprised	=	violet
fearful	=	black
disgusted	=	green
contemptuous	=	orange
happy	=	white/yellow
neutral	=	grey

Figure 2 shows how the system was applied to example test images each of which contains four human faces.

Classification accuracies for facial expression classification in the training set were determined by ten-fold cross-

validation. In order to determine which preprocessing steps deliver the best classification accuracies we compared the results obtained for greyscale, Sobel, and Canny filtered images, each of them with and without equalisation. The best correct classification accuracy was about 65% and was achieved when non-equalised greyscale images were used for training a ν -SVM with $\nu = 0.1$. This result was an improvement of about a 10% over our pilot tests with the same dataset before its images were normalised to the inner eye corners. Note that the class averages of the greyscale training images (as shown in the bottom row of Figure 1) show clearly recognisable differences. Some of the differences are expressed by the direction and shape of the eyebrows which are, quite recognisable owing to the inner eye corner normalisation [30].




Figure 5. Consistently in all four shown results a central violet (‘surprised’) facebox was detected. The four images used different filter and parameter settings as follows; *Top*: Canny filter on equalised grayscale, *bottom*: Canny filter on non-equalised grayscale, *left*: $\nu = 0.4$, *right*: $\nu = 0.05$. The non-equalised results show more local peaks, and the smaller the ν , the more local peaks are detected. Equalisation has more impact than changing ν .

The test image shown in Figure 2 was composed of four face images from the Cohn-Kanade data set [76]. All four faces were detected as the dominant face pattern by the face detection module, except the bottom left face in the case of equalised greyscale filtering. For the equalised greyscale, Sobel, and Canny filtered versions, the facial expression classification module consistently assigned the same sensible emotion classes to all detected faces. The top left face was classified as ‘disgusted’ (green), the top right face as ‘angry’ (red), and the bottom right face was classified as ‘happy’ (white). Outcomes of processing of the bottom left face showed some instability between the different filtering options. The face was detected and was classified

as ‘neutral’ with the Sobel and Canny filters. It was not detected with equalised greyscale as input. In the case of Canny filtering, additional smaller faces were detected for the emotion categories ‘sad’ (blue), ‘disgusted’ (green), and ‘surprised’ (violet). The boxes associated with the latter two categories were so small that they only contained the mouth and the bottom part of the nose. This indicates that the classifier interpreted the nose openings as small ‘eyes’ above the mouth. The yellow face clouds in Figure 2 also show that the desired face pattern was exactly detected as expected in the case of equalised greyscale and Sobel filtering (with the exception of the bottom left image in the case of equalised greyscale).

For the examples in the left and middle columns of Figure 2, the face boxes for all local peaks (coloured dots within the yellow clouds) are displayed. For the result with Canny filtering in the right column, the yellow face cloud was larger and several local peaks were detected. Many of these can be regarded as false positives but often there is some room for interpretation about what exact emotion is expressed by a face. The final selection of the displayed boxes was made by the experimenter using an interactive viewer. The interactive viewer is part of the software system we have developed and allows the display of coloured face boxes and associated decision values by mouse click on the associated local peak (coloured pixel) at the centre of the box. All displayed boxes are associated with the highest decision values found by the SVM for face detection in stage 1. The experimenter had to decide how many boxes should be displayed if several local peaks were detected. A full automation of this last step of the procedure is still a work in progress. We found that in most cases the decision values are a very good indicator for selecting sensible face boxes. We also observed, however, that the decision values depend heavily on preprocessing and parameter selection, and for results with many local peaks the decision value should not be taken as the only and absolute measure. The interactive viewer became even more useful when the system was tested on images of selected house façades.

3. Experimental Results with Architectural Image Data

After the system was tuned and trained on face detection and facial expression classification using only human face data following the above described approach, it was applied to selected images of house façades. Figures 3 to 6 show characteristic results of these experiments.

The images in Figure 3 show the façade of a house on Glebe Road in Newcastle. The face detection system indicated that the house façade contains a dominant pattern that can be classified as a face. Two images of the same house, taken at different distances and at slightly different angles, were compared (left and right images in Figure 3). The facial expression classifier consistently delivered high decision values for ‘surprised’ (violet box) if the box contained the full garage door and ‘fearful’ (black box) or ‘sad’ (blue box) if the box only contained a section of the upper part of the garage door. The yellow face cloud contains several other local peaks of lower decision values. These are typical of our approach using Canny filtering which was applied to equalised face boxes in this example. Preprocessing and SVM parameter settings for this example were exactly the same as used for the top right image in Figure 2, which for human test images had a tendency to show false positives.

In Figure 4 the Sobel filter was applied without prior

equalisation of the greyscale image. Within the façade the black bottom right face box was classified as ‘fearful’ and was consistently detected in a straight frontal view and a slight side view of the same building. Similarly, several of the indicated violet (‘surprised’) face boxes were detected in both views. The example in Figure 4 also shows that the yellow face cloud can have several components. The violet (‘surprised’) face patterns could be detected at several structurally similar parts of the house façade.

Figure 5 shows results where the non-equalised images generated larger yellow clouds than the equalised version. A decrease of ν for the one-class SVM for face detection could also lead to more boxes being detected. The results in Figure 5 used $\nu = 0.4$ on Canny filtered non-equalised greyscale images (left column) and $\nu = 0.05$ on Canny filtered equalised greyscale images (right column). It appears that preprocessing has greater impact than the selection of ν . In all of the four shown results a central violet (‘surprised’) facebox, which is the largest box in the top row, was consistently detected with a high decision value. In the bottom two examples, additional red (‘angry’), blue (‘sad’), and green (‘disgusted’) face boxes could be detected with relatively high decision values. Several different emotions could be detected within the same house façade and the type of filtering had substantial impact on the outcome.

The results so far show that for detecting face-like patterns in façades the combination of greyscale equalisation and Canny filtering (Figure 3) performs similarly well as if a Sobel filter is applied to a non-equalised greyscale image (Figure 4). The Canny filtered images tend to have larger face clouds than the Sobel filtered images but greyscale equalisation seems to compensate and shrink the clouds.

The example in Figure 6 shows a house façade which allows the detection of several face-like patterns. In contrast to Figure 3 it is not clear which should be declared the most dominant pattern. Results based on our standard 22×22 resolution face boxes in the first row are compared with results that used a 44×44 resolution shown in the second row. The underlying image for all results was a non-equalised greyscale image and all SVMs used $\nu = 0.1$. The left column shows the results for greyscale, the middle column for Sobel filtering and the right column for Canny filtering. The different sizes of the yellow face clouds are typical of the different filter settings. The presented results with the 44×44 resolution have smaller face clouds than the corresponding results with 22×22 resolution. The examples show that a change of resolution can lead to a different outcome but not necessarily to an ‘improvement’ of the pareidolia effect. The highest decision values were obtained for the ‘angry’ (red) and ‘surprised’ (violet) boxes. Other faces, some of them with similarly high decision values, could be detected, but in different parts of the image. Sometimes additional faces were detected in clouds in the sky.

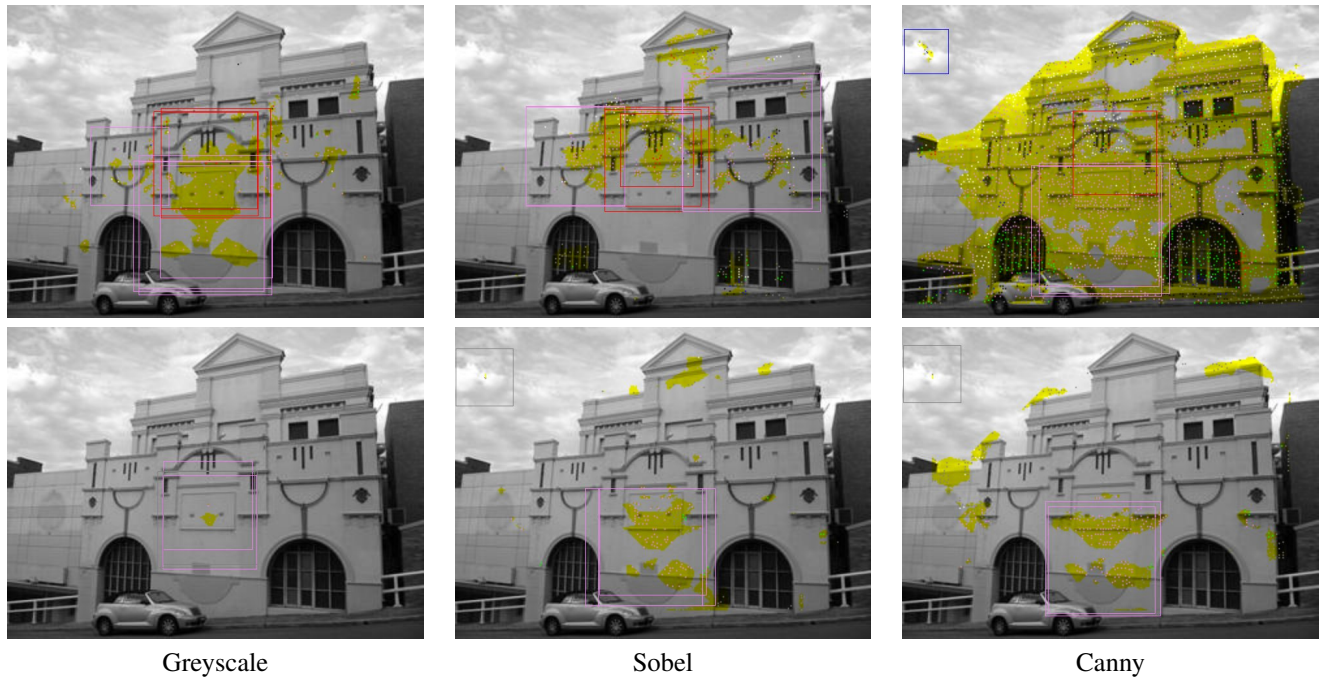


Figure 6. Results in the first row used our standard 22×22 resolution for the face boxes while the results shown in the second row used a 44×44 resolution. The underlying image for all results was a non-equalised greyscale image and all SVMs used $\nu = 0.1$. For the results in the middle and right columns additional Sobel or Canny filtering was applied, respectively.

4. Discussion and Conclusion

A combined face detection and emotion classification system based on support vector classification was implemented and tested. The system was trained on 280 images of human faces that were normalised to the inner eye corners. This allowed for a statistical model that emphasised details around the eye region [30]. The system detected sensible face-like patterns in test images containing human faces and assigned them to the appropriate categories of facial expressions of emotion. The results were mostly stable if filter types were changed moderately, avoiding extreme settings.

Using preprocessing and parameter settings that had a slight tendency to generate false positives in face detection on human test images (e.g. right column in Figure 2) we demonstrated that the system was also able to detect face-expression patterns within images of selected house façades. Most ‘faces’ detected in houses were very abstract or incomplete and often allowed the assignment of several different emotion categories depending on the choice of the centre point and the box size. Slight changes in viewing angle seemed not to have much impact on the outcome.

Sometimes face-like patterns, some of which had simi-

larly high decision values, could be detected in other parts of the image. Alternative face structures could originate from the texture of other façade structures but could also be caused by artefacts of the procedure, which includes box cropping, resizing, antialiasing, histogram equalisation, and edge detection. If the order of the individual processing steps is changed, this can also have an impact on the outcome of the procedure.

Overall, the experiments of the present study indicate that for selected houses a face pattern associated with a dominant emotion category is identifiable if appropriate filter and parameter settings are applied.

A limitation of the current system is that its statistical model learned geometric features of the human face data. That includes, for example, height–width proportions inherent in the training data shown in Figure 1. Consequently the system had difficulties in assigning sensible emotion categories to face-like patterns that do not have the same geometrical properties as the learned data but still have the topological properties required to be identified as face patterns by humans. For example, if the system is tested on images of ‘smileys’, as in Figure 7, the result is not always as expected. Inclusion of ‘smileys’ in the training dataset is one possibility to address this issue. This could, however,




Figure 7. Test image assembled of four ‘smileys’; The system detected all four face-like patterns but did not always assign the expected emotion categories. Left top: ‘neutral’ (grey), right top: ‘happy’ (white), left bottom: ‘disgusted’ (green), right bottom: ‘surprised’ (violet). These tests used a Canny filter on a non-equalised greyscale image and $\nu = 0.1$ for the SVM face detector.

lead to lower accuracy of the human test data, as previously observed in [12].

The present study demonstrated that a simple statistical learning approach using a small dataset of cropped and normalised face images can to some degree simulate the phenomenon of pareidolia. The human visual system, however, is much more sophisticated and consists of a large number of processing modules that interact in a complex manner [15, 39, 126]. Humans are able to process rotated and distorted face-like patterns and to recognise emotions utilising subtle features and micro-expressions. The scope and resolution of the human visual system are far beyond the simulation which was employed in the present study. Future research may investigate other compositions and normalisations of the training set and extensions of the software system which allow, for example, combinations of holistic approaches with component-based approaches for face detection and expression classification.

It may be argued that detecting and classifying face-like structures in house façades is an exotic way of design evaluation. However, as mentioned in the introduction, recent results in psychology found that the perception of faces is qualitatively different from the perception of other patterns. Faces, in contrast to non-faces, can be perceived non-consciously and without attention [41, 56]. These findings support our hypothesis that the perception of faces or face-like patterns [71, 146] may be more critical than previously thought for how humans perceive the aesthetics of the environment and the architecture of house façades of the buildings they are surrounded by in their day-to-day lives.

Acknowledgements

This project was supported by ARC discovery grant DP0770106 “Shaping social and cultural spaces: the application of computer visualisation and machine learning techniques to the design of architectural and urban spaces”.

References

- [1] Andrea F. Abate, Michele Nappi, Daniel Riccio, and Gabriele Sabatino. 2d and 3d face recognition: A survey. *Pattern Recognition Letters*, 28(14):1885–1906, 2007.
- [2] F. Al-Osaimi, M. Bennamoun, and A. Mian. An expression deformation approach to non-rigid 3d face recognition. *International Journal of Computer Vision*, 81(3):302–316, 2009.
- [3] Michael Bach and Charlotte M. Poloschek. Optical illusions. *Advances in Clinical Neuroscience and Rehabilitation*, 6(2):20–21, May/June 2006.
- [4] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, New York, 2006.
- [5] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3d morphable model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(9):1063–1074, 2003.
- [6] K. W. Bowyer, K. I. Chang, and P. J. Flynn. A survey of approaches and challenges in 3d and multi-modal 3d-2d face recognition. *Computer Vision and Image Understanding*, 101(1):1–15, January 2005.
- [7] Alexander M. Bronstein, Michael M. Bronstein, and Ron Kimmel. Three-dimensional face recognition. *International Journal of Computer Vision*, 64(1):5–30, 2005.
- [8] R. Brunelli and T. Poggio. Face recognition: Features versus templates. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 15(10):1042–1052, 1993.
- [9] Ernest Burden. *Building Facades: Faces, Figures, and Ornamental Details*. McGraw-Hill Professional, 2nd edition, 1996.
- [10] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:679–714, 2001.
- [11] Stephan K. Chalup, Naomi Henderson, Michael J. Ostwald, and Lukasz Wiklendt. A computational approach to fractal analysis of a cityscape’s skyline. *Architectural Science Review*, 52(2):126–134, 2009.
- [12] Stephan K. Chalup, Kenny Hong, and Michael J. Ostwald. A face-house paradigm for architectural scene analysis. In Richard Chbeir, Youakim Badr, Ajith Abraham, Dominique Laurent, and Fernando Ferri, editors, *CSTST 2008: Proceedings of The Fifth International Conference on Soft Computing As Transdisciplinary Science and Technology*, pages 397–403. ACM, 2008.

- [13] Stephan K. Chalup and Michael J. Ostwald. Anthropocentric biocybernetic computing for analysing the architectural design of house façades and cityscapes. *Design Principles and Practices: An International Journal*, 3(5):65–80, 2009.
- [14] Stephan K. Chalup and Michael J. Ostwald. Anthropocentric biocybernetic approaches to architectural analysis: New methods for investigating the built environment. In Paul S. Geller, editor, *Built Environment: Design Management and Applications*. Nova Science Publishers, 2010.
- [15] Leo M. Chalupa and John S. Werner, editors. *The Visual Neurosciences*. MIT Press, Cambridge, MA, 2004.
- [16] Chih-Chung Chang and Chih-Jen Lin. *LIBSVM: a library for support vector machines*, 2001. Software available at www.csie.ntu.edu.tw/~cjlin/libsvm.
- [17] R. Chellappa, C. L. Wilson, and S. Sirohey. Human and machine recognition of faces: a survey. *Proceedings of the IEEE*, 83(5):705–741, May 1995.
- [18] Jie Chen, Ruiping Wang, Shengye Yan, Shiguang Shan, Xilin Chen, and Wen Gao. Enhancing human face detection by resampling examples through manifolds. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 37(6):1017–1028, 2007.
- [19] Shaokang Chen, C. Sanderson, S. Sun, and B. C. Lovell. Representative feature chain for single gallery image face recognition. In *19th International Conference on Pattern Recognition (ICPR) 2008, Tampa, Florida, USA*, pages 1–4. IEEE, 2009.
- [20] Tat-Jun Chin, Konrad Schindler, and David Suter. Incremental kernel svd for face recognition with image sets. In *FGR '06: Proceedings of the 7th International Conference on Automatic Face and Gesture Recognition*, pages 461–466, Washington, DC, USA, 2006. IEEE Computer Society.
- [21] Klaus Conrad. *Die beginnende Schizophrenie. Versuch einer Gestaltanalyse des Wahns*. Thieme, Stuttgart, 1958.
- [22] J. C. Cooper. *The potential of chaos and fractal analysis in urban design*. Joint Centre for Urban Design, Oxford Brookes University, Oxford, 2000.
- [23] Mauricio Correa, Javier Ruiz del Solar, and Fernando Bernuy. Face recognition for human-robot interaction applications: A comparative study. In *RoboCup 2008: Robot Soccer World Cup XII*, volume 5399 of *Lecture Notes in Computer Science*, pages 473–484, Berlin, 2009. Springer.
- [24] N. P. Costen, M. Brown, and S. Akamatsu. Sparse models for gender classification. *Proceedings of the Sixth IEEE International Conference on Automatic Face and Gesture Recognition (FGR04)*, pages 201–206, May 2004.
- [25] David Cox, Ethan Meyers, and Pawan Sinha. Contextually evoked object-specific responses in human visual cortex. *Science*, 304(5667):115–117, April 2004.
- [26] Christoph D. Dahl, Christian Wallraven, Heinrich H. Bülthoff, and Nikos K. Logothetis. Humans and macaques employ similar face-processing strategies. *Current Biology*, 19(6):509–513, March 2009.
- [27] Chunhua Du, Jie Yang, Qiang Wu, Tianhao Zhang, and Shengyang Yu. Locality preserving projections plus affinity propagation: a fast method for face recognition. *Optical Engineering*, 47(4):040501–1–3, 2008.
- [28] Paul Ekman, Wallace V. Friesen, and Joseph C. Hager. *Facial Action Coding System, The Manual*. A Human Face, Salt Lake City UT, 2002.
- [29] Paul Ekman and David Matsumoto. Combined Japanese and Caucasian expressions of emotion (JACFEE) and Japanese and Caucasian neutral faces (JACNeuF) datasets. www.mettonline.com/products.aspx?categoryid=3, 1993.
- [30] N. J. Emery. The eyes have it: the neuroethology, function and evolution of social gaze. *Neuroscience & Biobehavioral Reviews*, 24(6):581–604, 2000.
- [31] Andrew D. Engell and James V. Haxby. Facial expression and gaze-direction in human superior temporal sulcus. *Neuropsychologia*, 45(14):323–341, 2007.
- [32] Michael W. Eysenck and Mark T. Keane. *Cognitive Psychology: A Student's Handbook*. Taylor & Francis, 2005.
- [33] M. J. Farah. *Visual Agnosia: Disorders of Object Recognition and What They Tell Us About Normal Vision*. MIT Press, Cambridge, MA, 1990.
- [34] M. J. Farah. Neuropsychological inference with an interactive brain: A critique of the ‘locality assumption’. *Behavioral and Brain Sciences*, 17:43–61, 1994.
- [35] M. J. Farah, C. Rabinowitz, G. E. Quinn, and G. T. Liu. Early commitment of neural substrates for face recognition. *Cognitive Neuropsychology*, 17:117–124, 2000.
- [36] M. J. Farah, K. D. Wilson, M. Drain, and J. N. Tanaka. What is “special” about face perception? *Psychological Review*, 105:482–498, 1998.
- [37] Martha J. Farah and Geoffrey Karl Aguirre. Imaging visual recognition: PET and fMRI studies of the functional anatomy of human visual recognition. *Trends in Cognitive Sciences*, 3(5):179–186, May 1999.
- [38] B. Fasel and J. Luetttin. Automatic facial expression analysis: a survey. *Pattern Recognition*, 36(1):259–275, 2003.
- [39] D. J. Felleman and D. C. Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1:1–47, 1991.
- [40] Filters. Filters library for computer vision and image processing. <http://filters.sourceforge.net/>, V3.1-2007.10.
- [41] M. Finkbeiner and R. Palermo. The role of spatial attention in nonconscious processing: A comparison of face and non-face stimuli. *Psychological Science*, 20:42–51, 2009.
- [42] Leonardo F. Fontenelle. Pareidolias in obsessive-compulsive disorder: Neglected symptoms that may respond to serotonin reuptake inhibitors. *Neurocase: The Neural Basis of Cognition*, 14(5):414–418, October 2008.
- [43] Sophie Fyfe, Claire Williams, Oliver J. Mason, and Graham J. Pickup. Apophenia, theory of mind and schizotypy: Perceiving meaning and intentionality in randomness. *Cortex*, 44(10):1316–1325, 2008.

- [44] I. Gauthier and C. Bukach. Should we reject the expertise hypothesis? *Cognition*, 103(2):322–330, 2007.
- [45] I. Gauthier, T. Curran, K. M. Curby, and D. Collins. Perceptual interference supports a non-modular account of face processing. *Nature Neuroscience*, (6):428–432, 2003.
- [46] I. Gauthier, P. Skudlarski, J. C. Gore, and A. W. Anderson. Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, 3:191–197, 2000.
- [47] I. Gauthier and M. J. Tarr. Becoming a “greeble” expert: Exploring face recognition mechanisms. *Vision Research*, 37:1673–1682, 1997.
- [48] I. Gauthier, M. J. Tarr, A. W. Anderson, P. Skudlarski, and J. C. Gore. Activation of the middle fusiform face area increases with expertise in recognizing novel objects. *Nature Neuroscience*, 2:568–580, 1999.
- [49] M. Gelder, D. Gath, and R. Mayou. Signs and symptoms of mental disorder. In *Oxford textbook of psychiatry*, pages 1–36. Oxford University Press, Oxford, 3rd edition, 1989.
- [50] Nathalie George, Jon Driver, and Raymond J. Dolan. Seen gaze-direction modulates fusiform activity and its coupling with other brain areas during face processing. *NeuroImage*, 13(6):1102–1112, 2001.
- [51] Shaogang Gong, Stephen J. McKenna, and Alexandra Psarrou. *Dynamic Vision: From Images to Face Recognition*. World Scientific Publishing Company, 2000.
- [52] Arnulf B. A. Graf, Felix A. Wichmann, Heinrich H. Bülthoff, and Bernhard H. Schölkopf. Classification of faces in man and machine. *Neural Computation*, 18(1):143–165, 2006.
- [53] Abdenour Hadid and Matti Pietikäinen. Combining appearance and motion for face and gender recognition from videos. *Pattern Recognition*, 42(11):2818–2827, 2009.
- [54] Abdenour Hadid and Matti Pietikäinen. Manifold learning for gender classification from face sequences. In Massimo Tistarelli and Mark S. Nixon, editors, *Advances in Biometrics, Third International Conference, ICB 2009, Alghero, Italy, June 2-5, 2009. Proceedings*, volume 5558 of *Lecture Notes in Computer Science (LNCS)*, pages 82–91, Berlin / Heidelberg, 2009. Springer-Verlag.
- [55] Nouchine Hadjikhani and Beatrice de Gelder. Neural basis of prosopagnosia: An fMRI study. *Human Brain Mapping*, 16:176–182, 2002.
- [56] Nouchine Hadjikhani, Kestutis Kveraga, Paulami Naik, and Seppo P. Ahlfors. Early (N170) activation of face-specific cortex by face-like objects. *Neuroreport*, 20(4):403–407, 2009.
- [57] C. M. Hagerhall, T. Purcell, and R. P. Taylor. Fractal dimension of landscape silhouette as a predictor of landscape preference. *The Journal of Environmental Psychology*, 24:247–255, 2004.
- [58] Jeremy Hall, Heather C. Whalley, James W. McKirdy, Liana Romaniuk, David McGonigle, Andrew M. McIntosh, Ben J. Baig, Viktoria-Eleni Gountouna, Dominic E. Job, David I. Donaldson, Reiner Sprengelmeyer, Andrew W. Young, Eve C. Johnstone, and Stephen M. Lawrie. Over-activation of fear systems to neutral faces in schizophrenia. *Biological Psychiatry*, 64(1):70–73, July 2008.
- [59] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*. Springer, New York, 2nd edition, 2009.
- [60] Bernd Heisele, Purdy Ho, and Tomaso Poggio. Face recognition with support vector machines: Global versus component-based approach. In *Proceedings of the Eighth IEEE International Conference on Computer Vision, 2001. ICCV 2001*, volume 2, pages 688–694, 2001.
- [61] E. Hjelmås and B. K. Low. Face detection: A survey. *Computer Vision and Image Understanding*, 83:236–274, 2001.
- [62] H. Hong, H. Neven, and C. Von der Malsburg. Online facial expression recognition based on personalized galleries. In *Proceedings of the Third International Conference on Face and Gesture Recognition, (FG98), IEEE, Nara, Japan*, pages 354–359, Washington, DC, USA, 1998. IEEE Computer Society.
- [63] Kenny Hong, Stephan Chalup and Robert King. A component based approach improves classification of discrete facial expressions over a holistic approach. *Proceedings of the International Joint Conference on Neural Networks (IJCNN 2010)*, IEEE 2010 (in press).
- [64] Satoshi Hosoi, Erina Takikawa, and Masato Kawade. Ethnicity estimation with facial images. *Proceedings of the Sixth IEEE International Conference on Automatic Face and Gesture Recognition (FGR04)*, pages 195–200, May 2004.
- [65] Kazuhiro Hotta. Robust face recognition under partial occlusion based on support vector machine with local gaussian summation kernel. *Image and Vision Computing*, 26(11):1490–1498, 2008.
- [66] Haifeng Hu. ICA-based neighborhood preserving analysis for face recognition. *Computer Vision and Image Understanding*, 112(3):286–295, 2008.
- [67] Sony Electronics Inc. Sony Cyber-shot(TM) digital still camera DSC-W120 specification sheet. www.sony.com, 2008. 16530 Via Esprillo, San Diego, CA 92127, 1.800.222.7669.
- [68] Spiros Ioannou, George Caridakis, Kostas Karpouzis, and Stefanos Kollias. Robust feature detection for facial expression recognition. *EURASIP Journal on Image and Video Processing*, 2007(2):1–22, August 2007.
- [69] Xudong Jiang, Bappaditya Mandal, and Alex Kot. Complete discriminant evaluation and feature extraction in kernel space for face recognition. *Machine Vision and Applications*, 20(1):35–46, January 2009.
- [70] Hongliang Jin, Qingshan Liu, and Hanqing Lu. Face detection using one-class-based support vectors. *Proceedings of the Sixth IEEE International Conference on Automatic Face and Gesture Recognition (FGR04)*, pages 457–462, May 2004.

- [71] M. H. Johnson. Subcortical face processing. *Nature Reviews Neuroscience*, 6(10):766–774, 2005.
- [72] Patrick J. Johnston, Kathryn McCabe, and Ulrich Schall. Differential susceptibility to performance degradation across categories of facial emotion—a model confirmation. *Biological Psychology*, 63(1):45–58, April 2003.
- [73] Patrick J. Johnston, Wendy Stojanov, Holly Devir, and Ulrich Schall. Functional MRI of facial emotion recognition deficits in schizophrenia and their electrophysiological correlates. *European Journal of Neuroscience*, 22(5):1221–1232, 2005.
- [74] Nicole Joshua and Susan Rossell. Configural face processing in schizophrenia. *Schizophrenia Research*, 112(1–3):99–103, July 2009.
- [75] Amit Kagian. A machine learning predictor of facial attractiveness revealing human-like psychophysical biases. *Vision Research*, 48:235–243, 2008.
- [76] T. Kanade, J. F. Cohn, and Y. Tian. Comprehensive database for facial expression analysis. In *Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition (FG'00), Grenoble, France*, pages 46–53, 2000.
- [77] N. Kanwisher, J. McDermott, and M. M. Chun. The fusiform face area: A module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience*, 17(11):4302–4311, 1997.
- [78] Nancy Kanwisher and Galit Yovel. The fusiform face area: a cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 361(1476):2109–2128, 2006.
- [79] Markus Kiefer. Top-down modulation of unconscious ‘automatic’ processes: A gating framework. *Advances in Cognitive Psychology*, 3(1–2):289–306, 2007.
- [80] A. Z. Kouzani. Locating human faces within images. *Computer Vision and Image Understanding*, 91(3):247–279, September 2003.
- [81] Ping-Han Lee, Yun-Wen Wang, Jison Hsu, Ming-Hsuan Yang, and Yi-Ping Hung. Robust facial feature extraction using embedded hidden markov model for face recognition under large pose variation. In *Proceedings of the IAPR Conference on Machine Vision Applications (IAPR MVA 2007), May 16–18, 2007, Tokyo, Japan*, pages 392–395, 2007.
- [82] Tommer Leyvand, Daniel Cohen-Or, Gideon Dror, and Dani Lischinski. Data-driven enhancement of facial attractiveness. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2008)*, 27(3), August 2008.
- [83] Stan Z. Li and Anil K. Jain, editors. *Handbook of Face Recognition*. Springer, New York, 2005.
- [84] Yongmin Li, Shaogang Gong, and Heather Liddell. Support vector regression and classification based multi-view face detection and recognition. In *FG '00: Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition 2000*, page 300, Washington, DC, USA, 2000. IEEE Computer Society.
- [85] H.-T. Lin, C.-J. Lin, and R. C. Weng. A note on Platt’s probabilistic outputs for support vector machines. Technical report, Department of Computer Science and Information Engineering, National Taiwan University, Taipei 106, Taiwan, 2003.
- [86] Jeremy Long and David Mould. Dendritic stylization. *The Visual Computer*, 25(3):241–253, March 2009.
- [87] B. C. Lovell and S. Chen. Robust face recognition for data mining. In John Wang, editor, *Encyclopedia of Data Warehousing and Mining*, volume II, pages 965–972. Idea Group Reference, Hershey, PA, 2008.
- [88] Erno Mäkinen and Roope Raisamo. An experimental comparison of gender classification methods. *Pattern Recognition Letters*, 29(10):1544–1556, 2008.
- [89] Gordon McIntyre and Roland Göcke. Towards Affective Sensing. In *Proceedings of the 12th International Conference on Human-Computer Interaction HCI2007*, volume 3 of *Lecture Notes in Computer Science LNCS 4552*, pages 411–420, Beijing, China, 2007. Springer.
- [90] E. McKone and R. A. Robbins. The evidence rejects the expertise hypothesis: Reply to Gauthier & Bukach. *Cognition*, 103(2):331–336, 2007.
- [91] Elinor McKone, Kate Crookes, and Nancy Kanwisher. The cognitive and neural development of face recognition in humans. In Michael S. Gazzaniga, editor, *The Cognitive Neurosciences, 4th Edition*. MIT Press, October 2009.
- [92] David Melcher and Francesca Bacci. The visual system as a constraint on the survival and success of specific artworks. *Spatial Vision*, 21(3–5):347–362, 2008.
- [93] Ajmal Mian, Mohammed Bennamoun, and Robyn Owens. An efficient multimodal 2d-3d hybrid approach to automatic face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11):1927–1943, 2007.
- [94] C. A. Nelson. The development and neural bases of face recognition. *Infant and Child Development*, 10:3–18, 2001.
- [95] M. J. Ostwald, J. Vaughan, and C. Tucker. Characteristic visual complexity: Fractal dimensions in the architecture of Frank Lloyd Wright and Le Corbusier. In K. Williams, editor, *Nexus: Architecture and Mathematics*, pages 217–232. Turin: K. W. Books and Birkhäuser, 2008.
- [96] Michael J. Ostwald, Josephine Vaughan, and Stephan Chalup. A computational analysis of fractal dimensions in the architecture of Eileen Gray. In *ACADIA 2008, Silicon + Skin: Biological Processes and Computation, October 16–19, 2008*, 2008.
- [97] E. Osuna, R. Freund, and F. Girosi. Training support vector machines: An application to face detection. In *Proc. Computer Vision and Pattern Recognition 97*, pages 130–136, 1997.
- [98] T. J. Palmeri and I. Gauthier. Visual object understanding. *Nature Reviews Neuroscience*, 5:291–303, 2004.
- [99] Maja Pantic and Leon J. M. Rothkrantz. Automatic analysis of facial expressions: The state of the art. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(12):1424–1445, 2000.

- [100] Olivier Pascalis and David J. Kelly. The origins of face processing in humans: Phylogeny and ontogeny. *Perspectives on Psychological Science*, 4(2):200–209, 2009.
- [101] K. A. Pelphrey, J. D. Singerman, T. Allison, and G. McCarthy. Brain activation evoked by perception of gaze shifts: the influence of context. *Neuropsychologia*, 41(2):156–170, 2003.
- [102] A. Pentland, B. Moghaddam, and Starner. View-based and modular eigenspaces for face recognition. In *Proceedings IEEE Conference Computer Vision and Pattern Recognition*, pages 84–91, 1994.
- [103] R. W. Picard. *Affective Computing*. MIT Press, Cambridge, MA, 1997.
- [104] J. C. Platt. Probabilistic outputs for support vector machines and comparison to regularized likelihood methods. In A. J. Smola, Peter Bartlett, B. Schölkopf, and Dale Schuurmans, editors, *Advances in Large Margin Classifiers*, Cambridge, MA, 2000. MIT Press.
- [105] Melissa Prince and Andrew Heathcote. State-trace analysis of the face inversion effect. In Niels Taatgen and Hedderik van Rijn, editors, *Proceedings of the Thirty-First Annual Conference of the Cognitive Science Society (CogSci 2009)*. Cognitive Science Society, Inc., 2009.
- [106] Matthias Rätzsch, Sami Romdhani, and Thomas Vetter. Efficient face detection by a cascaded support vector machine using haar-like features. In *Pattern Recognition*, volume 3175 of *Lecture Notes in Computer Science (LNCS)*, pages 62–70. Springer-Verlag, Berlin / Heidelberg, 2004.
- [107] Gillian Rhodes, Graham Byatt, Patricia T. Michie, and Aina Puce. Is the fusiform face area specialized for faces, individuation, or expert individuation? *Journal of Cognitive Neuroscience*, 16(2):189–203, March 2004.
- [108] Alexander Riegler. Superstition in the machine. In Martin V. Butz, Olivier Sigaud, Giovanni Pezzulo, and Gianluca Baldassarre, editors, *Anticipatory Behavior in Adaptive Learning Systems (ABiALS 2006)*. From Brains to Individual and Social Behavior, volume 4520 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 57–72, Berlin/Heidelberg, 2007. Springer.
- [109] R. A. Robbins and E. McKone. No face-like processing for objects-of-expertise in three behavioural tasks. *Cognition*, 103(1):34–79, 2007.
- [110] François Robert and Jean Robert. *Faces*. Chronicle Books, San Francisco, 2000.
- [111] S. Romdhani, P. Torr, B. Schölkopf, and A. Blake. Efficient face detection by a cascaded support-vector machine expansion. *Royal Society of London Proceedings Series A*, 460(2501):3283–3297, November 2004.
- [112] Henry A. Rowley, Shumeet Baluja, and Takeo Kanade. Neural network-based face detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(1):23–38, 1998.
- [113] Javier Ruiz-del Solar and Julio Quinteros. Illumination compensation and normalization in eigenspace-based face recognition: A comparative study of different pre-processing approaches. *Pattern Recognition Letters*, 29(14):1966–1979, 2008.
- [114] Ashok Samal and Prasana A. Iyengar. Automatic recognition and analysis of human faces and facial expressions: A survey. *Pattern Recognition*, 25(1):65–77, 1992.
- [115] Conrad Sanderson. *Biometric Person Recognition: Face, Speech and Fusion*. VDM-Verlag, Germany, 2008.
- [116] B. Schölkopf, A. J. Smola, R. C. Williamson, and P. L. Bartlett. New support vector algorithms. *Neural Computation*, 12(5):1207–1245, 2000.
- [117] Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alex J. Smola, and Robert C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13:1443–1471, 2001.
- [118] Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, 2002.
- [119] Adrian Schwaninger, Sandra Schumacher, Heinrich Bülthoff, and Christian Wallraven. Using 3d computer graphics for perception: the role of local and global information in face processing. In *APGV '07: Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization*, pages 19–26, New York, NY, USA, 2007. ACM.
- [120] G. Schwarzer, N. Zauner, and B. Jovanovic. Evidence of a shift from featural to configural face processing in infancy. *Developmental Science*, 10(4):452–463, 2007.
- [121] Ting Shan, Abbas Bigdeli, Brian C. Lovell, and Shaokang Chen. Robust face recognition technique for a real-time embedded face recognition system. In B. Verma and M. Blumenstein, editors, *Pattern Recognition Technologies and Applications: Recent Advances*, pages 188–211. Information Science Reference, Hersey, PA, 2008.
- [122] Frank Y. Shih, Chao-Fa Chuang, and Patrick S. P. Wang. Performance comparison of facial expression recognition in jaffe database. *International Journal of Pattern Recognition and Artificial Intelligence*, 22(3):445–459, May 2008.
- [123] Chathura R. De Silva, Surendra Ranganath, and Liyanage C. De Silva. Cloud basis function neural network: A modified rbf network architecture for holistic facial expression recognition. *Pattern Recognition*, 41(4):1241–1253, 2008.
- [124] A. Sims. *Symptoms in the mind: An introduction to descriptive psychopathology*. Saunders, London, 3rd edition, 2002.
- [125] P. Sinha, B. Balas, Y. Ostrovsky, and R. Russell. Face recognition by humans: Nineteen results all computer vision researchers should know about. *Proceedings of the IEEE*, 94(11):1948–1962, November 2006.
- [126] Pawan Sinha. Recognizing complex patterns. *Nature Neuroscience*, 5:1093–1097, 2002.

- [127] I. E. Sobel. *Camera Models and Machine Perception*. PhD dissertation, Stanford University, Palo Alto, CA, 1970.
- [128] Christopher Summerfield, Tobias Egner, Jennifer Mangels, and Joy Hirsch. Mistaking a house for a face: Neural correlates of misperception in healthy humans. *Cerebral Cortex*, 16(4):500–508, April 2006.
- [129] Kah-Kay Sung. *Learning and Example Selection for Object and Pattern Detection*. PhD dissertation, Artificial Intelligence Laboratory and Centre for Biological and Computational Learning, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 1996.
- [130] M. Suwa, N. Sugie, and K. Fujimora. A preliminary note on pattern recognition of human emotional expression. In *International Joint Conference on Pattern Recognition*, pages 408–410, 1978.
- [131] Ameet Talwalkar, Sanjiv Kumar, and Henry A. Rowley. Large-scale manifold learning. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2008), 24-26 June 2008, Anchorage, Alaska, USA*, pages 1–8. IEEE Computer Society, 2008.
- [132] Xiaoyang Tan, Songcan Chen, Zhi-Hua Zhou, and Fuyan Zhang. Face recognition from a single image per person: A survey. *Pattern Recognition*, 39(9):1725–1745, 2006.
- [133] Li Tao, Ming-Jung Seow, and Vijayan K. Asari. Nonlinear image enhancement to improve face detection in complex lighting environment. *International Journal of Computational Intelligence Research*, 2(4):327–336, 2006.
- [134] D. M. J. Tax and R. P. W. Duin. Support vector domain description. *Pattern Recognition Letters*, 20(11–13):1191–1199, December 1999.
- [135] R. P. Taylor. Reduction of physiological stress using fractal art and architecture. *Leonardo*, 39(3):25–251, 2006.
- [136] Ying-Li Tian, Takeo Kanade, and Jeffrey F. Cohn. Facial expression analysis. In Stan Z. Li and Anil K. Jain, editors, *Handbook of Face Recognition*, chapter 11, pages 247–275. Springer, New York, 2005.
- [137] Massimo Tistarelli, Manuele Bicego, and Enrico Grosso. Dynamic face recognition: From human to machine vision. *Image and Vision Computing*, 27(3):222–232, February 2009.
- [138] Bruce I. Turetsky, Christian G. Kohler, Tim Indersmitten, Mahendra T. Bhati, Dorothy Charbonnier, and Ruben C. Gur. Facial emotion recognition in schizophrenia: When and why does it go awry? *Schizophrenia Research*, 94(1–3):253–263, August 2007.
- [139] Matthew Turk and Alex Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.
- [140] Vladimir Vapnik. *Estimation of dependencies based on empirical data. Reprint of 1982 edition*. Springer, New York, 2nd edition, 2006.
- [141] Patrik Vuilleumier. Neural representation of faces in human visual cortex: the roles of attention, emotion, and viewpoint. In N. Osaka, I. Rentschler, and I. Biederman, editors, *Object recognition, attention, and action*, pages 109–128. Springer, Tokyo Berlin Heidelberg New York, 2007.
- [142] Patrik Vuilleumier, Jorge L. Armony, John Driver, and Raymond J. Dolan. Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature Neuroscience*, 6(6):624–631, June 2003.
- [143] Huiyuan Wang, Xiaojuan Wu, and Qing Li. Eigenblock approach for face recognition. *International Journal of Computational Intelligence Research*, 3(1):72–77, 2007.
- [144] Yongjin Wang and Ling Guan. Recognizing human emotional state from audiovisual signals. *IEEE Transactions on Multimedia*, 10(4):659–668, June 2008.
- [145] Jonathan Welsh. Why cars got angry. Seeing demonic grins, glaring eyes? Auto makers add edge to car ‘faces’; Say goodbye to the wide-eyed neon. *The Wall Street Journal*, March 10 2006.
- [146] P. J. Whalen, J. Kagan, R. G. Cook, F. C. Davis, H. Kim, and S. Polis. Human amygdala responsivity to masked fearful eye whites. *Science*, 306(5704):2061, 2004.
- [147] S. Windhager, D. E. Slice, K. Schaefer, E. Oberzaucher, T. Thorstensen, and K. Grammer. Face to face: The perception of automotive designs. *Human Nature*, 19(4):331–346, 2008.
- [148] T.-F. Wu, C.-J. Lin, and R. C. Weng. Probability estimates for multi-class classification by pairwise coupling. *Journal of Machine Learning Research*, 5:9751005, 2004.
- [149] Xudong Xie and Kin-Man Lam. Facial expression recognition based on shape and texture. *Pattern Recognition*, 42(5):1003–1011, 2009.
- [150] Ming-Hsuan Yang. Face recognition using kernel methods. In T. Diederich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems (NIPS)*, volume 14, 2002.
- [151] Ming-Hsuan Yang. Kernel eigenfaces vs. kernel fisherfaces: Face recognition using kernel methods. In *5th IEEE International Conference on Automatic Face and Gesture Recognition (FGR 2002), 20-21 May 2002, Washington, D.C., USA*, pages 215–220. IEEE Computer Society, 2002.
- [152] Ming-Hsuan Yang, David J. Kriegman, and Narendra Ahuja. Detecting faces in images: A survey. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 24(1):34–58, 2002.
- [153] Zhihong Zeng, Maja Pantic, Glenn I. Roisman, and Thomas S. Huang. A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(1):39–58, 2008.
- [154] Tianhao Zhang, Jie Yang, Deli Zhao, and Xinliang Ge. Linear local tangent space alignment and application to face recognition. *Neurocomputing*, 70(7–9):1547–1553, 2007.
- [155] Xiaozheng Zhang and Yongsheng Gao. Face recognition across pose: A review. *Pattern Recognition*, 42(11):2876–2896, 2009.

- [156] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld. Face recognition: A literature survey. *ACM Computing Surveys*, 35(4):399–458, 2003.
- [157] W. Y. Zhao and R. Chellappa, editors. *Face Processing: Advanced Modeling and Methods*. Elsevier, 2005.
- [158] Cheng Zhong, Zhenan Sun, and Tieniu Tan. Fuzzy 3d face ethnicity categorization. In Massimo Tistarelli and Mark S. Nixon, editors, *Advances in Biometrics, Third International Conference, ICB 2009, Alghero, Italy, June 2-5, 2009. Proceedings*, pages 386–393, Berlin / Heidelberg, 2009. Springer-Verlag.

Author Biographies

Stephan K. Chalup is the Director of the Newcastle Robotics Laboratory and Associate Professor in Computer Science and Software Engineering at the University of Newcastle, Australia. He has a PhD in Computer Science / Machine Learning from Queensland University of Technology in Australia and a Diploma in Mathematics with Biology from the University of Heidelberg in Germany. In his research he investigates machine learning and its applications in areas such as autonomous agents, human-computer interaction, image processing, intelligent system design, and language processing.

Kenny Hong is a PhD Student in Computer Science and a member of the Newcastle Robotics Laboratory in the School of Electrical Engineering and Computer Science at the University of Newcastle in Australia. He is an assistant lecturer in computer graphics and a research assistant in the Interdisciplinary Machine Learning Research Group (IMLRG). His research area is face recognition and facial expression analysis where he investigates the performance of statistical learning and related techniques.

Michael J. Ostwald is Professor and Dean of Architecture at the University of Newcastle, Australia. He is a Visiting Fellow at SIAL and a Professorial Research Fellow at Victoria University Wellington. He has a PhD in architectural philosophy and a higher doctorate (DSc) in the mathematics of design. He is co-editor of the journal *Architectural Design Research* and on the editorial boards of *Architectural Theory Review* and the *Nexus Network Journal*. His recent books include *The Architecture of the New Baroque* (2006), *Homo Faber: Modelling Design* (2007) and *Residue: Architecture as a Condition of Loss* (2007).