

ZN-Face: A system for access control using automated face recognition

Wolfgang Konen and Ekkehard Schulze-Krüger
Zentrum für Neuroinformatik GmbH, Bochum
Universitätsstr. 160
D-44801 Bochum
Germany
E-mail: wolfgang@zn.ruhr-uni-bochum.de

Abstract

We present a biometric access control device which is based on the identification of human faces. The system combines a console for semi-automated image acquisition with the necessary algorithms for face recognition. Facial features are stored in a relatively compact data format (1.6 kB). The system uses an extension of the Elastic Graph Matching Algorithm of [1] which has been ported here to a standard PC (Pentium 90 without any special accelerator hardware) where it performs image acquisition, face localization and identification in 3.5 seconds. The system not only allows robust identification of stored persons (despite changes in facial expression or size), but also reliable rejection of unknown persons. With an acceptance criterion which safely rejects all unknown persons we achieve an identification rate above 96% (FRR=4%) on a database of 87 persons.

1 Introduction

Face recognition is a remarkable example for the ability of humans to perform with great reliability a complex visual recognition task. We are able to recognize thousands of faces which we have learned during our lifetime. Our visual performance is very robust against a variety of factors like changes in facial expression, head posture or size, illumination, background, facial aging or partial occlusion of a face. It is one of the challenges of artificial information processing to achieve at least partially a similar performance on systems for automated visual recognition. The ease with which the problem is solved by the human brain, i. e. a biological neural system, may serve as a guideline in the sense that neural information processing paradigms can be important also for the construction of automated face recognition systems.

In this contribution we will report on the development of the face recognition system ZN-Face which offers a turn-key solution for an access control system. ZN-Face is based on the neural algorithms of [1] which are briefly reviewed in Section 2 (Elastic Graph Matching). In Section 3 we describe the extensions necessary to migrate the algorithm into an automated access control system.

Numerous approaches for solving the face recognition problem can be found in the literature from which we can mention here only a few. A good and more comprehensive survey of the state of the art can be found in [2, 3]. Early algorithms [4] use feature-based techniques (e. g. features like localization or thickness of eyebrows) or template matching for recognition [5]. Recently, a systematic comparison of feature-based vs. template-based algorithms

has been undertaken by Brunelli and Poggio [6]. Gilbert and Yang [7] presented a real-time face recognition system using custom VLSI hardware. The hardware allows fast template correlation and the system is able to perform an identification from a 34-person database in 2 to 3 seconds. Another approach uses the decomposition of facial images into an 'eigenface' expansion (similar to Karhune-Loeve expansion) [8, 9]. These expansions achieve a very compact representation of faces, they are, however, rather sensitive against variations in imaging conditions.

Another group of algorithms (including this work) uses neural processing techniques for face recognition [10, 11] or gender recognition [12]. One of the few systems which does not only present a face recognition algorithm but also a complete access control booth has been proposed by Bichsel and Seitz [13]. This direction is also the aim of the current work, where a complete hardware-setup has been developed as a biometric access control device for real-world security applications.

2 Face recognition by graph matching

The basic recognition algorithm of ZN-Face is an extension of the Elastic Graph Matching Algorithm [1]. In contrast to most other face recognition algorithms which require distinct processing steps like localization, separation, standardization and finally recognition of faces, the current algorithm is more coherent in the sense that *one* basic principle is used to perform the above steps simultaneously. Here, faces are stored as flexible graphs or grids (see Fig. 1) with characteristic visual features attached to the nodes of the graph (*labeled graphs*). Features are obtained by a convolution with Gabor wavelets

$$G_{\vec{k},\sigma}(\vec{x}) = C_{\vec{k},\sigma} \exp\left(\frac{\vec{k}^2 x^2}{2\sigma^2}\right) \exp(i\vec{k}\vec{x}), \quad (1)$$

where the length of \vec{k} characterizes the size of the wavelet and the direction of \vec{k} its orientation sensitivity. The convolution is computed at the location of the graph node, thus a feature contains information about the image from the graph node and its vicinity. Such an information distribution over units with localized and partially overlapping sensitive fields is typical for neural systems.¹

¹In biological systems, the sensitive region of a neuronal unit is known as *receptive field*. The use of Gabor wavelets is further motivated by results from neurobiology [14] which indicate that the visual cortex of mammals contains neurons with receptive fields very similar in shape to Gabor wavelets.

In addition, a data representation with labeled graphs has computational advantages:

robustness: Gabor features are invariant against intensity changes and contrast changes in the image. Furthermore, Gabor features are less affected by changes in head posture, size and facial expression than raw grey level features.

data compression: compared to the raw image which has in our case a size of 128×128 pixels, the size of the graph (about 1.6 kB) is smaller by a factor of 10.

scaling: a sparse graph can be readily adjusted to changes in the geometry (size, perspective, see Fig. 1B,C). Change in size in a pixel-based data representation, for example, would require a more complicated transformation.

distribution: Compared to a high-level feature description (e.g. 'eye', 'nose') the graph contains sufficient information distribution on simple, but numerous features. Even if the information at a single graph node is missing due to occlusion, recognition is usually still possible due to the information at the remaining nodes.

When comparing a stored graph of a person with an actual image, the graph is shifted and elastically deformed on the image with the aim to optimize the similarity between graph node features and corresponding features in the actual image. This optimization process follows the elastic graph matching scheme of [1]. As a result of the match, a similarity value $\mathcal{M} \in [0, 1]$ is obtained which characterizes the goodness of the match (1 = perfect match). In extension to [1] we found that it is also possible to optimize the stored graph to a different size in the actual image (Fig. 1B) where the size differences may be as large as ± 20 -30%. This is remarkable for two reasons:

- Optimizing the graph simultaneously with respect to two spatial (shift) directions and one scale space direction increases the search space by one dimension, thus increasing the possibility of a mismatch (i. e. wrong local minimum of the similarity value). However, in experiments with facial images the optimization algorithm has proven to be very robust: In almost all cases the procedure locates the face with the correct size and this behaviour is stable within a wide range of parameters.
- Changing the size of the graph would in principle require also a transformation of the Gabor features, which would increase the computational cost of the optimization process. However, it was found that for size changes up to $\pm 30\%$ the untransformed Gabor features can be used instead as a good approximation.

The latter can be understood from the following argument: If the image contains a step edge $\theta(\vec{x})$, the Gabor wavelets are normalized in such a way that the convolution coefficients are the same for all scales. Although faces are of course not build from step edges, dominant features are at least similar to a step edge so that the wavelet coefficients change only weakly as the scale changes.

In summary, the face recognition algorithm consists of three basic steps:

- (1) Convolution of the image with each Gabor wavelet used in the labeled graph representation.
- (2) A stored graph is matched to the image by an optimization procedure: Position, size and inner structure of the graph are varied in order to maximize the similarity between graph node features and corresponding image features. Penalty terms inhibit too large deformations of the graph. As a result, a new graph with the actual image features as node labels can be extracted.
- (3) In order to decide, whether stored graph and image show the face of the same person, the extracted graph is compared also to a number of reference graphs \mathcal{G}_{ref} (similarity values \mathcal{M}_{ref}). Only if certain significance conditions

$$c_i(\mathcal{M}, \mathcal{M}_{ref}) > t_i \quad (2)$$

are fulfilled (for example, the stored graph's \mathcal{M} has to be significantly better than the best of the other \mathcal{M}_{ref} 's), the match is accepted.

3 Access Control System ZN-Face

ZN-Face offers an automated access control system which performs a biometric identification of persons from their facial images. To achieve this, several extensions are necessary to transform the Elastic Graph Matching algorithm of Sec. 2 into a system which is easy to use in real-world security applications:

- *Fast computation on standard hardware:* In [1] the graph matching was performed on a transputer-based platform and required about 12 seconds for the identification. Now, the algorithm has been optimized and ported to a standard PC (Pentium PC/90 Mhz *without* special accelerator hardware). Computation time for a full identification is 3.5 seconds.
- *Semi-automated image acquisition:* No operator is available to adjust the camera (persons may differ in their size and their distance to the camera) and to trigger the image acquisition. An appropriate semi-automated procedure is described below.
- *Easy-to-use system administration:* A Windows-based² GUI has been developed which allows simple control by an authorized system administrator (e.g. adding or deleting access rights, monitoring the access protocol).
- *Possibility to handle large databases (1000 persons and more).*

In order to achieve the last point without unacceptable increase in computation time, we perform with ZN-Face *verification* instead of recognition. Verification has the additional advantage of higher security because in addition to a standard personal identification via secret number (PIN) or card, the face is used as independent control of the person's identity.

²Windows is a registered trademark of Microsoft

The basic verification procedure can be described as follows: The hardware setup of the system consists of a standard PC with framegrabber and a ZN-Face console (a prototype is shown in Fig. 2) containing camera, ID acquisition device (card reader or PIN pad) and LC-display. The camera is positioned behind a semi-permeable mirror which is tilted about $30-50^\circ$ against the vertical line. Camera and mirror are tilted in order to allow users of different height to view and position themselves centrally in front of the camera. (at the expense of some size variation in the acquired facial images). The user triggers the image acquisition and identifies himself (PIN pad or card). Now ZN-Face starts the verification whether the graph stored under the given PIN in the database fits to the acquired image. The final decision 'yes' or 'no' is based on the significance conditions mentioned above under step (3). This does not require to search the whole database, but only the stored graphs to the given PIN and a fixed number of reference graphs have to be searched.

It has to be mentioned that even with the user's cooperation during image acquisition, there is considerable variation in the images with respect to head posture, size and position. This is due to the user-dependence of the automated image acquisition. Most of the variation can be handled successfully by the robustness of the elastic graph matching. Also, variations like different facial expressions or wearing glasses are handled very well by the algorithm. To cope with the remaining cases, the system offers the possibility to store more than one image for a given person ("teach-in", see below).

Adding new persons to the database is made very simple for an authorized system administrator by using the graphical user interface (GUI) under Windows. An image of the new user is acquired in the same way as described above, and the system automatically locates the face and extracts the new graph (Fig. 3). In addition, the GUI offers a facility to browse through the verification protocol (Fig. 4). This control panel also offers the possibility of "teach-in": If an image acquired during verification is rejected although it shows the correct person (perhaps because stored and acquired image differ too strongly in head posture, thus leading to a low significance), the authorized system administrator has the option to add the picture into the person's record in the database. The system can handle an arbitrary number of images in a person's record, although normally 1-3 images are sufficient for reliable verification.

4 Results

Nearly all papers on face recognition algorithms – with the exception of [1] – usually report only on the recognition rate achieved by the algorithm, i. e. how successfully a person is accepted whose image is in the database. But this is only one side of the coin: Often even more important for biometric devices is the question, how successfully persons are rejected whose image is *not* in the database. Biometric research names errors of the former kind as *False Rejection Rate (FRR)* while errors of the latter kind contribute to the *False Acceptance Rate (FAR)*. Of course, there is a trade-off between both types of

errors and consequently it is much more difficult to build a system with a high recognition rate *and* a low FAR than a system with a high recognition rate alone.

The system ZN-Face offers the possibility to balance optimally the trade-off between FAR and FRR. With a simple neural network (perceptron) the significance conditions of Eq. (2) are combined into one scalar acceptance value in such a way that an optimal discrimination between acceptance and rejection cases is achieved. 'Optimal' is defined here as the simultaneous minimization of FAR and FRR. In the verification, the actual acceptance value is compared against an acceptance threshold. If needed, the user has the possibility to shift the acceptance threshold, e. g. in the direction of lower FAR, usually at the expense of a somewhat larger FRR.

We conducted two experiments with a database of 87 persons: In the first one (Fig. 5A), the database contained one image per person. For each person, the verification was tested with a different image \mathcal{I} of the same person. Unsuccessful verifications (acceptance value too low) contribute to the FRR in Fig. 5A. Then, the person was removed from the database and it was tested, whether the image \mathcal{I} was now rejected, as it should. False acceptances contribute to the FAR. As can be seen from Fig. 5A, the combined minimum of FAR and FRR is achieved at a level of about 3.5%. In Fig. 5B the same experiment is repeated, but now with 2 images per person in the database. In this case, comparing with a person's third image \mathcal{I} usually yields higher acceptance values. Here, it is possible to shift the acceptance threshold in such a way, that FAR and FRR simultaneously drop to zero. (This is of course only true for the specific database, more comprehensive tests with larger data material have to be undertaken to estimate with higher accuracy, how low FAR and FRR actually are). Setting the threshold to the 'safe' value 0.5 still produces a recognition rate of 96% (FRR=4%).

5 Conclusion and outlook

Elastic Graph Matching has proven to be a powerful algorithm for the recognition of human faces, especially because it achieves at the same time a reliable rejection of unknown faces. This makes it an appropriate tool for a biometric access control device which has been presented in this work. In addition to the results presented in the preceding section, the system has also been tested recently 'live' on several exhibitions (VISION'94 and SECURITY'94) where it operated successfully under real-world conditions with 'untrained' visitors. The fast identification of faces on standard hardware opens new application possibilities on the biometry market.

There do remain, of course, further development areas where human face recognition is still superior to automated systems: The present algorithm – as most others algorithms, too – is still sensitive to large changes in lumination (e. g. high contrast shadows). Additionally, further developments will be made in the direction of fully automated acquisition of facial images and/or face localization in more complex scenes.

References

- [1] M. Lades, J. Vorbrüggen, J. Buhmann, J. Lange, C.v.d. Malsburg, R.P. Würtz, and W. Konen. Distortion invariant object recognition in the dynamic link architecture. *IEEE Transaction on Computers*, 42:300–311, 1993.
- [2] V. Bruce and M. Burton. *Processing Images of Faces*. ABLEX Publishing-Corporation, Norwood, NJ, 1992.
- [3] S. Toelg and T. Poggio. Towards an example-based image compression architecture for video-conferencing. Technical report, Center for Automation Research, Univ. of Maryland, CAR-TR-723, 1994.
- [4] T. Kanade. *Picture Processing System by Computer Complex and Recognition of Human Faces*. Unpublished Ph.D. thesis, Dept. of Information Science, Kyoto Univ., 1973.
- [5] R. J. Baron. Mechanisms of human facial recognition. *International Journal of Man Machine Studies*, 15:137–178, 1981.
- [6] R. Brunelli and Tomaso Poggio. Face recognition: Features versus templates. Technical Report TR 9110-04, Istituto per la Ricerca Scientifica e Tecnologica, October 1992.
- [7] J. M. Gilbert and Woodward Yang. A real-time face recognition system using custom VLSI hardware. In *Proc. of Computer Architectures for Machine Perception Workshop*, December 1993.
- [8] M. Kirby and L. Sirovich. Application of the Karhunen-Loeve procedure for characterization of human faces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(1):103–108, January 1990.
- [9] M. Turk and A. Pentland. Face recognition using eigenfaces. In *IEEE Proc. of CVPR*, pages 586–591, Maui, Hawaii, June 1991.
- [10] T. Kohonen, E. Oja, A. Kortekangas, and K. Mäkisara. In *Proc. Intl. Conf. on Cybernetics and Society*, Washington D.C., 1977.
- [11] G. Cottrell, P. Munro, and D. Zipser. Learning internal representations of grey scale images: An example of extensional programming. In *Proc. Ninth Annual Cognitive Science Society Conference*, Seattle, WA, 1987.
- [12] B. A. Golomb, D. T. Lawrence, and T. J. Sejnowski. SEXNET: A neural network identifies sex from human faces. In D. S. Touretzky and R. Lippman, editors, *Advances in Neural Information Processing Systems 3*. Morgan Kaufmann, San Mateo, 1991.
- [13] M. Bichsel and P. Seitz. Der elektronische Pförtner: Automatisches Erkennen und Identifizieren von menschlichen Gesichtern. In R.E. Grosskopf, editor, *Mustererkennung 1990, 12. DAGM-Symposium*. Springer, 1990.
- [14] R. and K. de Valois. *Spatial Vision*. Erlbaum, NY, 1990.

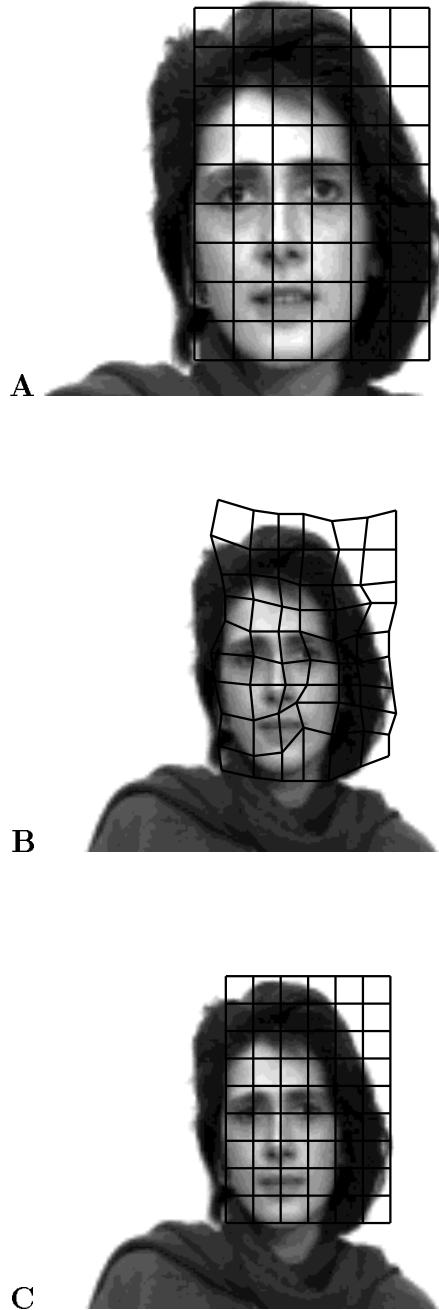


Figure 1: Data representation in *ZN-Face*: Faces are stored as elastic grids or graphs with (here) 7×10 nodes (**A,B**). Such graphs can be shifted or scaled (**C**) efficiently in the image domain.



Figure 2: The access control system *ZN-Face* in a prototype realization, consisting of camera (behind semi-permeable mirror), LC-display and PIN pad.

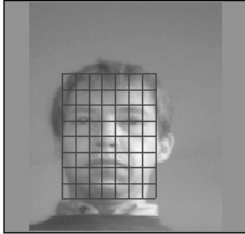
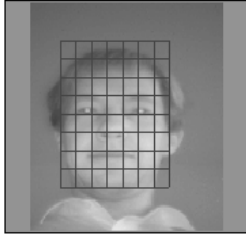
ZNFACE - Protokollinformation	
An Verifikationsstation aufgenommenes Bild	Gespeichertes Bild zur eingegebenen PIN
10012000	konen003
	
An Verifikationsstation eingegebene PIN	Name zur eingegebenen PIN
0012	Konen, Wolfgang
gefundene PIN	Datum / Uhrzeit
0012	17.11.94 15:04
Ergebnis der Verifikation	Station
OK (400)	Haupteingang (2)
OK	Übernahme in Galerie
Bild	
Bild + Graph	

Figure 4: Protocol of the verification: On the left hand side the actually acquired image is shown, on the right hand side the best matching image from the database (to the given PIN).


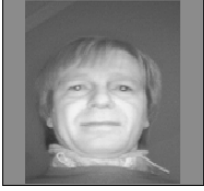
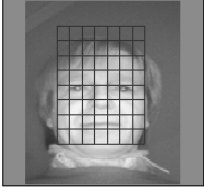


ZNFACE - Datenaufnahme	
 Datenbasis Personen-Aufnahme	
Name <input type="text" value="Mustermann"/>	
Vorname <input type="text" value="Hanspeter"/>	
PIN <input type="text" value="0006"/>	
ID <input type="text" value="123"/>	
GhZ <input type="text" value="0006"/>	
Bilddatei <input type="text" value="muste000"/>	
geladene Galerie <input type="text" value="rexroth"/>	
OK	Abbrechen

Figure 3: GUI for adding new persons to the database. The system automatically locates and extracts the new graph as shown in the lower image.



A



B

Figure 5: False rejection rate (FRR) and false acceptance rate (FAR) for recognition from a database of 87 persons as a function of the adjustable acceptance threshold. The threshold value 0 corresponds to the threshold learned by the neural net as optimal discrimination. The database contains one image per person in (A) and two images per person in (B).