

Efficient Head Pose Estimation with Gabor Wavelet Networks

Volker Krüger, Sven Bruns and Gerald Sommer
Computer Science Institute, Christian-Albrechts University Kiel
Preußstr. 1-9, 24105 Kiel, Germany
Tel: ++49-431-560496, FAX: ++49-431-560481
email: vok@ks.informatik.uni-kiel.de

Abstract

In this article we want to introduce first the Gabor wavelet network as a model based approach for an effective and efficient object representation. The Gabor wavelet network has several advantages such as invariance to some degree with respect to translation, rotation and dilation. Furthermore, the use of Gabor filters ensured that geometrical and textural object features are encoded. The feasibility of the Gabor filters as a model for local object features ensures a considerable data reduction while at the same time allowing *any* desired precision of the object representation ranging from a sparse to a photo-realistic representation. In the second part of the paper we will present an approach for the estimation of a head pose that is based on the Gabor wavelet networks.

1 Introduction

Recently, model-based approaches for the recognition and the interpretation of images of variable objects, like the bunch graph approach, PCA, eigenfaces and active appearance models, have received considerable interest [24; 14; 4; 6]. These approaches achieve good results because solutions are constrained to be valid instances of a model. In these approaches, the term “model-based” is understood in the sense that a set of training objects is given in the form of gray value pixel images while the model “learns” the variances of the gray values (PCA, eigenfaces) or, respectively, the Gabor filter responses (bunch graph). With this, model knowledge is given by the variances of pixel gray values, which means that the actual knowledge representation is given on a pixel basis, that is independent from the objects themselves.

In this work we want to introduce a novel approach for object representation that is based on Gabor Wavelet Networks. Gabor Wavelet Networks (GWN) are combining the advantages of RBF networks with the advantages of Gabor wavelets: GWNs represent an object as a linear combination of Gabor wavelets where the parameters of each of the Gabor functions (such as orientation and position and scale) are optimized to reflect the particular local image structure. Gabor wavelet networks have several advantages:

1. By their very nature, Gabor wavelet networks are invariant to some degree to affine deformations and homogeneous illumination changes,

2. Gabor filters are good feature detectors [13] and the optimized parameters of each of the Gabor wavelets are directly related to the underlying image structure,
3. the weights of each of the Gabor wavelet are directly related to their filter responses and with that they are also directly related to the underlying local image structure,
4. the precision of the representation can be varied to *any* desired degree ranging from a coarse representation to an almost photo-realistic one by simply varying the number of used wavelets.

We will discuss each single point in section 2.

The use of Gabor filters implies a model for the actual representation of the object information. In fact, as we will see, the GWN represents object information as a set of local image features, which leads to a higher level of abstraction and to a considerable data reduction. Both, textural and geometrical information is encoded at the same time, but can be split to some degree.

The variability in precision and the data reduction are the most important advantage in this context, that has several consequences:

1. Because the parameters of the Gabor wavelets and the weights of the network are directly related to the structure of the training image and the Gabor filter responses, a GWN can be seen as a task oriented optimal filter bank: given the number of filters, a GWN defines *that* set of filters that extracts the maximal possible image information.
2. For real-time applications one wants to keep the number of filtrations low to save computational resources and it makes sense in this context to relate the number of filtrations to the amount of image information really needed for a specific task: In this sense, it is possible to relate the representation precision to the specific task and to increment the number of filters if more information is needed. This, we call *progressive attention* [26].
3. The training speed of neural networks, that correlates with the dimensionality of the input vector.

The *progressive attention* is related to the *incremental focus of attention (IFA)* for tracking [19] or the attentive processing strategy (*GAZE*) for face feature detection [8]. Both works are inspired by [20] and relate features to scales by using a coarse-to-fine image resolution strategy. In contrary, the *progressive attention* should not relate features to scale but to the object itself that is described by these features. In this sense, the object is considered as a collection of image features and the more information about the object is needed to fulfill a task the more features are extracted from the image.

In the following section we will give a short introduction to GWNs. Also, we will discuss each single point mentioned above, including the invariance properties, the abstraction properties and specificity of the wavelet parameters for the object representation and a task oriented image filtration.

In section 3 we will present the results of our pose estimation experiment where we exploited the optimality of the filter bank and the *progressive attention* property to speed up the response time of the system and to optimize the training of the neural network.

In the last section we will conclude with some final remarks.



Figure 1: The very right image shows the original face image I , the other images show the image I , represented with 16, 52, 116 and 216 Gabor wavelets (left to right)

2 Introduction to Gabor Wavelet Networks

The basic idea of the wavelet networks is first stated by [27], and the use of Gabor functions is inspired by the fact that they are recognized to be good feature detectors [13].

To define a GWN, we start out, generally speaking, by taking a family of N odd Gabor wavelet functions $\Psi = \{\psi_{\mathbf{n}_1}, \dots, \psi_{\mathbf{n}_N}\}$ of the form $\psi_{\mathbf{n}}(x, y) =$

$$= \exp\left(-\frac{1}{2}\left[s_x\left((x-c_x)\cos\theta - (y-c_y)\sin\theta\right)\right]^2 + \left[s_y\left((x-c_x)\sin\theta + (y-c_y)\cos\theta\right)\right]^2\right) \times$$

$$\sin\left(s_x\left((x-c_x)\cos\theta - (y-c_y)\sin\theta\right)\right), \text{ with } \mathbf{n} = (c_x, c_y, \theta, s_x, s_y)^T. \text{ Here, } c_x, c_y \text{ de-}$$

note the translation of the Gabor wavelet, s_x, s_y denote the dilation and θ denotes the orientation. The choice of N is arbitrary and is related to the maximal representation precision of the network. The parameter vector \mathbf{n} (translation, orientation and dilation) of the wavelets may be chosen arbitrarily at this point. In order to find the GWN for image I , the energy functional for wavelet networks: $E = \min_{\mathbf{n}_i, w_i \text{ for all } i} \|I - \sum_i w_i \psi_{\mathbf{n}_i}\|_2^2$ is minimized with respect to the weights w_i and the wavelet parameter vector \mathbf{n}_i . We therefore define a Gabor wavelet network as follows:

Definition: Let $\psi_{\mathbf{n}_i}, i = 1, \dots, N$ be a set of Gabor wavelets, I a DC-free image and w_i and \mathbf{n}_i chosen according to the above energy functional. The two vectors $\Psi = (\psi_{\mathbf{n}_1}, \dots, \psi_{\mathbf{n}_N})^T$ and $\mathbf{w} = (w_1, \dots, w_N)^T$ define then the *Gabor wavelet network* (Ψ, \mathbf{w}) for image f .

The parameter vectors \mathbf{n}_i are chosen from *continuous* phase space \mathbb{R}^5 [5] and the Gabor wavelets are positioned with sub-pixel accuracy. This is precisely the main advantage over the discrete approach [5; 12]. While in case of a discrete phase space local image structure has to be approximated by a combination of wavelets, a *single* wavelet can be chosen selectively in the continuous case to reflect *precisely* the local image structure. This assures that a maximum of the image information is encoded.

Using the optimal wavelets Ψ and weights \mathbf{w} of the Gabor wavelet network of an image f , I can be (closely) reconstructed by a linear combination of the weighted wavelets: $\hat{I} = \sum_{i=1}^N w_i \psi_{\mathbf{n}_i} = \Psi^T \mathbf{w}$. Of course, the quality of the image representation and of the reconstruction depends on the number N of wavelets used and can be varied to reach almost any desired precision. In section 2.2 we will discuss the relation between I and \hat{I} in more detail. An example reconstruction can be seen in fig. 1: A family of 216 wavelets has been distributed over the inner face region of the very right image I by the energy functional. Different reconstructions \hat{I} with variable N are shown in the first four images.

A further example can be seen in fig. 2: The left image shows a reconstruction with 16 wavelets and the right image indicates the corresponding wavelet positions. It should be pointed out that at each indicated wavelet position, just *one* single wavelet is located.



Figure 2: The images show a Gabor wavelet network with $N = 16$ wavelets after optimization (left) and the indicated positions of each single wavelet(right).



Figure 3: The figure shows images of a wooden toy block on which a GWN was trained. The black line segments sketch the positions, sizes and orientations of all the wavelets of the GWN (left), and of some automatically selected wavelets (right).

2.1 Feature Representation with Gabor Wavelets

It was mentioned in the introduction that the Gabor wavelets are recognized to be good feature [13] detectors, that are directly related to the local image features by the energy functional. This means that an optimized wavelet has e.g. ideally the exact position and orientation of a local image feature. An example can be seen in fig. 3. The figure shows the image of a little wooden toy block, on which a Gabor wavelet network was trained. The left image shows the positions, scales and orientations of the wavelets as little black line segments. By thresholding the weights, the more “important” wavelets may be selected, which leads to the right image. Ideally, each Gabor wavelet should be positioned *exactly* on the image line after optimization. Furthermore, since large weights indicate that the corresponding wavelets represents an edge segment (see sec. 2.2), these wavelets encode local geometrical object information. In reality, however, interactions with other wavelets of the network have to be considered so that most wavelet parameters reflect the position, scale, and orientation of the image line closely, but not precisely. This fact is clearly visible in fig. 3. As it can be seen in fig. 1 an object can be represented almost perfectly with a relatively small set of wavelets. The considerable data reduction is achieved by the introduction of the model for local image primitives, i.e. the introduction of Gabor wavelets.

The use of Gabor filters as a model for local object primitives leads to a higher level of abstraction where object knowledge is represented by a set of local image primitives. The Gabor wavelets in a network that represent edge segments can be easily identified. How to identify wavelets, however, that encode specific textures is not really clear, yet, and subject to future investigation.

2.2 Direct Calculation of Weights and Distances

As mentioned earlier, the weights w_i of a GWN are directly related to the filter responses of the Gabor filters $\psi_{\mathbf{n}_i}$ on the training image.

Gabor wavelet functions are not orthogonal. For a given family Ψ of Gabor wavelets it is therefore not possible to calculate a weight w_i directly by a simple projection of the Gabor wavelet $\psi_{\mathbf{n}_i}$ onto the image. Instead one has to consider the family of dual wavelets $\tilde{\Psi} = \{\tilde{\psi}_{\mathbf{n}_1} \dots \tilde{\psi}_{\mathbf{n}_N}\}$. The wavelet $\tilde{\psi}_{\mathbf{n}_j}$ is the dual wavelet to the wavelet $\psi_{\mathbf{n}_i}$ iff $\langle \psi_{\mathbf{n}_i}, \tilde{\psi}_{\mathbf{n}_j} \rangle = \delta_{i,j}$. With $\tilde{\Psi} = (\tilde{\psi}_{\mathbf{n}_1}, \dots, \tilde{\psi}_{\mathbf{n}_N})^T$, we can write $[\langle \Psi, \tilde{\Psi} \rangle] = \mathbb{I}$. In other words: $w_i = \langle I, \tilde{\psi}_{\mathbf{n}_i} \rangle$. We find $\tilde{\psi}_{\mathbf{n}_i}$ to be $\tilde{\psi}_{\mathbf{n}_i} = \sum_j (\Psi^{-1})_{i,j} \psi_{\mathbf{n}_j}$, where $\Psi_{i,j} = \langle \psi_{\mathbf{n}_i}, \psi_{\mathbf{n}_j} \rangle$.

The equation $w_i = \langle I, \tilde{\psi}_{\mathbf{n}_i} \rangle$ allows us to define the operator $\mathcal{T}_{\Psi} : \mathbb{L}^2(\mathbb{R}^2) \mapsto \langle (\psi_{\mathbf{n}_1}, \dots, \psi_{\mathbf{n}_N}) \rangle$ as follows: Given a set Ψ of optimal wavelets of a GWN, the operator

\mathcal{T}_Ψ realizes an orthogonal projection of a function J onto the closed linear span of Ψ i.e. $\hat{J} = \mathcal{T}_\Psi(J) = J\tilde{\Psi}\Psi = \sum_{i=1}^N w_i\psi_{\mathbf{n}_i}$, with $\mathbf{w} = J\tilde{\Psi}$.

2.3 Reparameterization of Gabor Wavelet Networks

The “reverse” task of finding the position, the scale and the orientation of a GWN in a new image is most important because otherwise the filter responses are without any sense. Here, PCA, bunch graphs and GWN have similar properties: In case of the PCA and bunch graph it is important to ensure that corresponding pixels are aligned into a common coordinate system, in case of the GWN, local image primitives have to be aligned. For example, consider an image J that shows the person of fig. 1, left, possibly distorted affinely. Given a corresponding GWN we are interested in finding the correct position, orientation and scaling of the GWN so that the wavelets are positioned on the same facial features as in the original image, or, in other words, how should the GWN be deformed (warped) so that it is aligned with the coordinate system of the new object. An example for a successful warping can be seen in fig. 2, where in the right image the wavelet positions of the *original* wavelet network are marked and in fig. 4, where in new images the wavelet positions of the *reparameterized* Gabor wavelet network are marked. Parameterization of




Figure 4: The images show the positions of each of the 16 wavelets after reparameterizing the wavelet net and the corresponding reconstruction. The reconstructed faces show the same orientation, position and size as the ones they were reparameterized on.

a GWN is established by using a *superwavelet* [18]:

Definition: Let (Ψ, \mathbf{w}) be a Gabor wavelet network with $\Psi = (\psi_{\mathbf{n}_1}, \dots, \psi_{\mathbf{n}_N})^T$, $\mathbf{w} = (w_1, \dots, w_N)^T$. A *superwavelet* $\Psi_{\mathbf{n}}$ is defined to be a linear combination of the wavelets $\psi_{\mathbf{n}_i}$, such that $\Psi_{\mathbf{n}}(\mathbf{x}) = \sum_i w_i \psi_{\mathbf{n}_i}(\mathbf{S}\mathbf{R}(\mathbf{x} - \mathbf{c}))$, where the parameters of vector \mathbf{n} of superwavelet Ψ define the dilation matrix $\mathbf{S} = \text{diag}(s_x, s_y)$, the rotation matrix \mathbf{R} , and the translation vector $\mathbf{c} = (c_x, c_y)^T$.

A superwavelet $\Psi_{\mathbf{n}}$ is again a wavelet (because of the linearity of the sum) and in particular a continuous function that has the wavelet parameters dilation, translation and rotation. Therefore, we can handle it in the same way as we handled each single wavelet in the previous section. For a new image J we may arbitrarily deform the superwavelet by optimizing its parameters \mathbf{n} with respect to the energy functional E : $E = \min_{\mathbf{n}} \|J - \Psi_{\mathbf{n}}\|_2^2$. The above energy functional defines the operator $\mathcal{P}_\Psi : \mathbb{L}^2(\mathbb{R}^2) \mapsto \mathbb{R}^5$, $g \mapsto \mathbf{n} = (c_x, c_y, \theta, s_x, s_y)$, where \mathbf{n} minimizes the energy functional E of the above equation. In the eq. of the operator \mathcal{P} Ψ is defined to be a superwavelet. For optimization of the superwavelet parameters, the same optimization procedure as for the optimization of the GWNs may be used. An example of an optimization process can be seen in fig. 5: Shown are the initial values of \mathbf{n} , the values after 2 and 4 optimization cycles of the gradient decent method and the final values after 8 cycles, each marked with the white square. The square marks the inner face region and its center position marks the center position of the

corresponding superwavelet. The superwavelet used in fig. 5 is the one of fig. 2, i.e. it is derived from the person in fig. 1.



Figure 5: The images show the 1st, the 2th, the 4th and the 8th (final) step of the gradient descent method optimizing the parameters of a superwavelet. The top left image shows the initial values with 10 px. off from the true position, rotated by 10° and scaled by 20%. The bottom right image shows the final result. As superwavelet, the GWN of figure 1 was used.

The image distortions of a planar object that is viewed under orthographic projection is described by six parameters: translation c_x, c_y , rotation θ , and dilation s_x, s_y and s_{xy} .

The reparameterization (warping) works quite robust: Using the superwavelet of fig. 1 we have found in several experiments on the various subjects with ≈ 60 pixels in width that the initialization of \mathbf{n}_0 may vary from the correct parameters by approx. ± 10 px. in x and y direction, by approx. 20% in scale and by approx. $\pm 10^\circ$ in rotation (see fig. 5). Compared to the AAM, these findings indicate a much better robustness [4]. Furthermore, we found that the warping algorithm converged in 100% of the cases to the correct values when applied on the *same* individual, independently of pose and gesture. The tests were done on the images of the Yale face database [22] and on our own images. The poses were varied within the range of $\approx \pm 20^\circ$ in pan and tilt where all face features were still visible. The various gestures included *normal, happy, sad, surprised, sleepy, glasses, wink*. The warping on other faces depended certainly on the similarity between the training person and the test person and on the number of used wavelets. We found that the warping algorithm always converged correctly on $\approx 80\%$ of the test persons (including the training person) of the Yale face database. The warping algorithm has also been successfully applied for a wavelet based affine real-time face tracking application [11].

2.4 Related Work

There are other models for image interpretation and object representation. Most of them are based on PCA [9], such as the eigenface approach [21]. The eigenface approach has shown its advantages especially in the context of face recognition. Its major drawbacks are its sensitivity to perspective deformations and to illumination changes. PCA encodes textural information only, while geometrical information is discarded. Furthermore, the alignment of face images into a common coordinate system is still a problem.

Another PCA based approach is the active appearance model (AAM)[4]. This approach enhances the eigenface approach considerably by including geometrical information. This allows an alignment of image data into a common coordinate system while the formulation of the alignment technique can be elegantly done with techniques of the AAM framework. Also, recognition and tracking applications are presented within this framework [6]. An advantage of this approach was demonstrated in [4]: they showed the ability of the AAM to model, in a photo-realistic way, almost any face gesture and




Figure 6: The left image shows the original doll face image I , the right image shows its reconstruction $\hat{I}_{4,6}$ using the reconstruction formula with an optimal wavelet net Ψ of just $N = 52$ odd Gabor wavelets, distributed over the inner face region. For optimization, the scheme that was introduced in section 2 was applied.

gender. However, this is undoubtedly an expensive task and one might ask for which task such a precision is really needed. In fact, a variation to different precision levels in order to spare computational resources and to restrict considerations to the data actually needed for a certain application seems not easily possible.

The bunch graph approach [24] is based, on the other hand, on the discrete wavelet transform. A set of Gabor wavelets are applied at a set of hand selected prominent object points, so that each point is represented by a set of filter responses, called *jet*. An object is then represented by a set of jets, that encode each a single local texture patch of the object. The jet topology, the so-called *image graph*, encodes geometrical object information. A precise positioning of the image graph onto the test image is important for good matching results and the positioning is quite a slow process. The feature detection capabilities of the Gabor filters are not exploited since their parameters are fixed and a variation to different precision levels has not been considered so far.

3 Pose Estimation with GWN

In this section we will present the approach for the estimation of the pose of a head. There exist many different approaches for pose estimation, including pose estimation with color blobs [3; 17], pose estimation applying a geometrical approach [7], stereo information [25] or neural networks [1], to cite just a few. While in some approaches, such as in [17], only an approximate pose is estimated, other approaches have the goal to be very precise so that they could even be used as a basis for gaze detection such as in [23]. The precision of the geometrical approach [7] was extensively tested and verified in [15]. The minimal mean pan/tilt error that was reached was $> 1.6^\circ$. In comparison to this, the neural network approach in [1] reached a minimal pan/tilt error of $> 0.64^\circ$.

The good result in [1] was reached by first detecting the head using a color tracking approach. Within the detected color blob region, 4×4 sets of 4 complex Gabor filters with the different orientations of $0, \frac{\pi}{4}, \frac{\pi}{2}$ and $\frac{3}{4}\pi$ were evenly distributed. The 128 coefficients of these 64 complex projections of the Gabor filters were then fed into a neural LLM network.

At this point, it is reasonable to assume that a precise positioning of the Gabor filters would result into an even lower mean pan/tilt error. In our experiments we therefore trained a GWN on an image I showing a doll's head. For the training of the GWN we used again the optimization scheme introduced in section 2 with $N = 52$ Gabor wavelets (see fig. 6). In order to be comparable with the approach in [1] we used in our experiments *exactly* the same neural network and the same number of training examples as described in [1]. A subspace variant of the Local Linear Map (LLM) [16] was used for learning input - output mappings [2]. The LLM rests on a locally linear (first order) approximation of the unknown function $f : \mathbb{R}^n \mapsto \mathbb{R}^k$ and computes its output as (winner-take-all-variant) $y(x) = A_{bmu}(x - c_{bmu}) + o_{bmu}$. Here, $o_{bmu} \in \mathbb{R}^k$ is an output vector attached to the best

matching unit (zero order approximation) and $A_{bmu} \in \mathbb{R}^{k \times n}$ is a local estimate of the Jacobian matrix (first order term). Centers are distributed by a clustering algorithm. Due to the first order term, the method is very sensitive to noise in the input. With a noisy version $x' = x + \eta$ the output differs by $A_{bmu}\eta$, and the LLM largely benefits from projecting to the local subspace, canceling the noise component of η orthogonal to the input manifold M . As basis functions normalized Gaussians were used.

The doll's head was connected to a robot arm, so that the pan/tilt ground truth was known. During the training and testing, the doll's head was first tracked using our wavelet based face tracker [11]. For each frame we proceeded in two steps:

1. optimal reparameterization of the GWN by using the positioning operator \mathcal{P}
2. calculating the optimal weights for the optimally repositioned GWN by using the projection operator \mathcal{T} .

See fig. 7 for example images. The weight vector that was calculated with the operator




Figure 7: The images show different orientations of the doll's head. The head is connected to a robot arm so that the ground truth is known. The white square indicates the detected position, scale and orientation of the GWN.

\mathcal{T} was then fed into the same neural network that was used in [1]. The training was done exactly as it was described in [1]: We used 400 training images, evenly distributed within the range of $\pm 20^\circ$ in pan and tilt direction (this is the range where all face features appeared to be visible). With this, we reached a minimal mean pan/tilt error of 0.19° for a GWN with 52 wavelets and a minimal mean pan/tilt error of 0.29° for a GWN with 16 wavelets. The maximal errors were 0.46° for 52 wavelets and 0.81° for 16 wavelets, respectively. The experiments were carried out on an experimental setup, that has not yet been integrated into a complete, single system. A complete system should reach a speed on a 450 MHz Linux Pentium of $> \approx 5$ fps for the 52 wavelet network and $> \approx 10$ fps for the 16 wavelet network ¹.

In comparison, for the *gaze* detection in [23], 625 training images were used, with a 14-D input vector, to train an LLM-network. The user was advised to fixate a 5×5 grid on the computer screen. The minimal errors after training for pan and tilt were 1.5° and 2.5° , respectively, while the system speed was 1 Hz on a SGI (Indigo, High Impact). A direct comparison to geometrical approaches is difficult, because, by their very nature, the cited ones are less precise, less robust but much faster.

4 Conclusions

The contribution of this article is twofold: First, we introduced the concepts of the *Gabor wavelet network* and the *Gabor superwavelet* that allow a data reduction and the use of the *progressive attention* approach:

¹This is a conservative estimation, various optimizations should allow higher frame rates.

- The representation of an object with variable degree of precision, from a coarse representation to an almost photo-realistic one,
- the definition of an optimal set of filters for a selective filtering
- the representation of object information on a basis of local image primitives and
- the possibility for affine deformations to cope with perspective deformations.

In the second section we discussed these various properties in detail. In [10; 11], GWNs have already been used successfully for wavelet based affine real time face tracking and pose invariant face recognition. It is future work, to fully exploit the advantages of the data reduction by reducing considerations to the vector space over the set of Gabor wavelet networks. And second, we exploited all these advantages of the GWN for the estimation of the head pose. The experimental results show quite impressively that it is sensible for an object representation to reflect the specific individual properties of the object rather than being independent of the individual properties such as general representations are. This can especially be seen when comparing the presented approach with the one in [1]: While having used the same experimental setup and the same type of neural network, the precision of the presented approach is twice as good with only 16 coefficients (vs. 128), and three times as good with only about half the coefficients. Furthermore, the experiment shows, how the precision in pose estimation and the system speed change with an increasing number of filters. A controllable variability of precision and speed has a major advantage: The system is able to decide how precise the estimation should be in order to minimize the probability that the given task is not fulfilled satisfactorily. It is future work to incorporate the experimental setup into a complete system. An enhancement for the evaluation of the positions of the irises for a precise estimation of gaze is about to be tested.

Acknowledgment

The images used are derived from the Yale Face Database.

References

- [1] J. Bruske, E. Abraham-Mumm, J. Pauli, and G. Sommer. Head-pose estimation from facial images with subspace neural networks. In *Proc. of Int. Neural Network and Brain Conference*, pages 528–531, Beijing, China, 1998.
- [2] J. Bruske and G. Sommer. Intrinsic dimensionality estimation with optimally topology preserving maps. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 20(5):572–575, 1998.
- [3] Q. Chen, H. Wu, T. Fukumoto, and M. Yachida. 3d head pose estimation without feature tracking. In *Int. Conf. on Automatic Face- and Gesture-Recognition*, pages 88–93, Nara, Japan, April 14–16, 1998.
- [4] T.F. Cootes, G.J. Edwards, and C.J. Taylor. Active appearance models. In *Proc. Fifth European Conference on Computer Vision*, volume 2, pages 484–498, Freiburg, Germany, June 1–5, 1998.
- [5] I. Daubechies. The wavelet transform, time-frequency localization and signal analysis. *IEEE Trans. Informat. Theory*, 36, 1990.

- [6] G.J. Edwards, T.F. Cootes, and C.J. Taylor. Face recognition using active appearance models. In *Proc. Fifth European Conference on Computer Vision*, volume 2, pages 581–595, Freiburg, Germany, June 1-5, 1998.
- [7] A. Gee and R. Cipolla. Determining the gaze of faces in images. *Image and Vision Computing*, 12(10):639–647, 1994.
- [8] R. Herpers, H. Kattner, H. Rodax, and G. Sommer. Gaze: An attentive processing strategy to detect and analyze the prominent facial regions. In *Int. Workshop on Automatic Face- and Gesture-Recognition, FG*, pages 214–220, Zurich, Switzerland, June 26-28, 1995.
- [9] I. Jolliffe. *Principal Component Analysis*. Springer Verlag, New York, 1986.
- [10] V. Krüger and G. Sommer. Gabor wavelet networks for object representation. In *Proc. of the Int. Dagstuhl 2000 Workshop*, 2000. to be published.
- [11] V. Krüger and Gerald Sommer. Affine real-time face tracking using gabor wavelet networks. In *Proc. Int. Conf. on Pattern Recognition*, pages 141–150, Barcelona, Spain, Sept. 3-8, 1999.
- [12] T. S. Lee. Image representation using 2D Gabor wavelets. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 18(10):959–971, 1996.
- [13] B.S. Manjunath and R. Chellappa. A unified approach to boundary perception: edges, textures, and illusory contours. *IEEE Trans. Neural Networks*, 4(1):96–107, 1993.
- [14] B. Moghaddam and A. Pentland. Probabilistic visual learning for object detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 17(7):696–710, Juli 1997.
- [15] Eleni Petraki. Analyse der blickrichtung des menschen und er kopforientierung im raum mittels passiver bildanalyse. Master's thesis, Technical University of Hamburg-Harburg, 1996.
- [16] H. Ritter, T. Martinez, and K. Schulten. *Neuronale Netze*. Addison-Wesley, 1991.
- [17] B. Schiele and A. Waibel. Gaze tracking based on face-color. In *Int. Workshop on Automatic Face- and Gesture-Recognition, FG*, pages 344–349, Zurich, Switzerland, June 26-28, 1995.
- [18] H. Szu, B. Telfer, and S. Kadambe. Neural network adaptive wavelets for signal representation and classification. *Optical Engineering*, 31(9):1907–1961, 1992.
- [19] K. Toyama and G. Hager. Incremental focus of attention for robust visual tracking. In *IEEE Conf. Computer Vision and Pattern Recognition, CVPR*, pages 189–195, 1996.
- [20] J.K. Tsotsos. Analyzing vision at the complexity level. *Behavioral and Brain Sci.*, 13:423–469, 1990.
- [21] M. Turk and A. Pentland. Eigenfaces for recognition. *Int. Journal of Cognitive Neuroscience*, 3(1):71–89, 1991.
- [22] Yale University. Yale face database. <http://cvc.yale.edu/projects/yalefaces/yalefaces.html>.
- [23] A.C. Varchmin, R. Rae, and H. Ritter. Image based recognition of gaze direction using adaptive methods. In I. Wachsmuth, editor, *Proceedings of the International Gesture Workshop, Incs*, pages 245–257. Springer, 1997.
- [24] L. Wiskott, J. M. Fellous, N. Krüger, and C. v. d. Malsburg. Face recognition by elastic bunch graph matching. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 19(7):775–779, July 1997.
- [25] M. Xu and T. Akatsuka. Detecting head pose from stereo image sequences for active face recognition. In *Int. Conf. on Automatic Face- and Gesture-Recognition*, pages 82–87, Nara, Japan, April 14-16, 1998.
- [26] H. Zabrodsky and S. Pelec. Attentive transmission. *Journal of Visual Communication and Image Representation*, 1(2):189–198, 1990.
- [27] Q. Zhang and A. Benviste. Wavelet networks. *IEEE Trans. Neural Networks*, 3(6):889–898, Nov. 1992.