

Pose-Invariant Facial Expression Recognition Using Variable-Intensity Templates

Shiro Kumano · Kazuhiro Otsuka · Junji Yamato ·
Eisaku Maeda · Yoichi Sato

Received: 7 February 2008 / Accepted: 7 October 2008
© Springer Science+Business Media, LLC 2008

Abstract In this paper, we propose a method for pose-invariant facial expression recognition from monocular video sequences. The advantage of our method is that, unlike existing methods, our method uses a simple model, called the variable-intensity template, for describing different facial expressions. This makes it possible to prepare a model for each person with very little time and effort. Variable-intensity templates describe how the intensities of multiple points, defined in the vicinity of facial parts, vary with different facial expressions. By using this model in the framework of a particle filter, our method is capable of estimating facial poses and expressions simultaneously. Experiments demonstrate the effectiveness of our method. A recognition rate of over 90% is achieved for all facial orientations, horizontal, vertical, and in-plane, in the range of ± 40 degrees, ± 20 degrees, and ± 40 degrees from the frontal view, respectively.

Keywords Facial expression recognition · Variable-intensity templates · Intensity distribution models · Particle filter

Electronic supplementary material The online version of this article (<http://dx.doi.org/10.1007/s11263-008-0185-x>) contains supplementary material, which is available to authorized users.

S. Kumano (✉) · Y. Sato
Institute of Industrial Science, The University of Tokyo,
4-6-1 Komaba, Meguro-ku, Tokyo, 153-8505 Japan
e-mail: kumano@iis.u-tokyo.ac.jp

K. Otsuka · J. Yamato · E. Maeda
NTT Communication Science Laboratories, NTT,
3-1 Morinosato-Wakamiya, Atsugi-shi, Kanagawa,
243-0198 Japan


1 Introduction

Facial expression recognition is attracting a great deal of attention because of its usefulness in many applications such as human-computer interaction and the analysis of conversation structure (Otsuka et al. 2007). Most existing methods for facial expression recognition assume that the person in the video sequence does not rotate the head strongly and that the image shows a nearly frontal view of the face (e.g., Laniotis et al. 1997; Huang and Huang 1997; Tian et al. 2001; Fasel et al. 2004; Chang et al. 2006; Bartlett et al. 2006; Kotsia and Pitas 2007; Koelstra and Pantic 2008). However, in situations such as multi-party conversations (e.g. meetings), people will often turn their faces to look at other participants. Hence, unless a stressful head-mounted camera is used, e.g. Pantic and Rothkrantz (2000a), we must simultaneously handle the variations in head pose as well as facial expression changes.

The major approach to correctly recognizing facial expressions in a face image containing head movements is to prepare a three-dimensional facial shape model of the user's neutral expression (rigid model) and its deformation model for other facial expressions (non-rigid model). The shape model and the facial expression model are together referred to as the face model in this paper. In this approach, facial pose variations are described by globally translating and rotating the shape model in three-dimensional space, and facial expression changes are described by locally deforming the shape model according to the facial expression model. This is often called the geometry-based approach.

Existing methods require an accurate face model, because image variations caused by facial expression change are often smaller than those caused by facial pose change. Accordingly, the use of inaccurate face models degrades the accuracy of both the facial pose and the expression estimates

Fig. 1 Our method absorbs errors in shape models and recognizes facial expressions by treating the changes in intensity of multiple points defined around facial parts



because those two components cannot be differentiated reliably. One method generates a face model for each user, a person-specific model, by using stereo cameras (Gokturk et al. 2002). Accordingly, this approach cannot be applied to monocular video sequences. Other methods utilize a general face model which can be applied to arbitrary users (Zhu and Ji 2006; Dornaika and Davoine 2008). However, generating an accurate general deformable model is not so easy in practice. It has been reported that general models cannot cover the large interpersonal variations of face shape and facial expression expected with sufficient accuracy (Gross et al. 2005).

Another approach, the region-based approach, has been reported to be relatively robust for small out-of-plane head rotations. In this approach, a two-dimensional (plane) (Black and Yacoob 1997; Tong et al. 2007) or three-dimensional rigid shape model, e.g. a cylinder in Liao and Cohen (2006), is prepared. Facial expressions are recognized using regional features such as optical flow (Black and Yacoob 1997; Liao and Cohen 2006), and Gabor-wavelet coefficients (Tong et al. 2007). They first estimate the head pose, and then calculate regional features for the aligned face. However, the use of rough shape models such as plane or cylinder degrades the accuracy of the facial expression estimates, because large out-of-plane head rotations often cause image appearance variations that are more significant than those yielded by facial expression changes. In addition, the optical flow estimation is easily violated by illumination changes and non-rigid motion. The failure of optical flow estimation directly disturbs facial expression recognition.

Motivated by these problems, we propose a novel *point intensity-based* approach for facial expression recognition. Specifically, we propose *variable-intensity templates* for the following reasons:

1. Monocular video capture systems are supported.
2. A face model for each person can be easily prepared.
3. Facial expressions can be estimated even with a large change in facial pose.

Our idea is to describe the change in facial expression not as the deformation of a shape model or as optical flow in facial part regions but as just a change in the intensities of multiple points fixed on a rigid shape model. The variable-intensity

template makes it unnecessary to estimate non-rigid motion caused by facial expressions.

The variable-intensity template consists of three components: a rigid shape model, a set of interest points, and an intensity distribution model. As the shape model, we use an average face shape generated from the shapes of many people. The interest points are sparsely defined away from the edges of facial parts such as eye corners, in a frontal and neutral expression face image. The intensity distribution model is the facial expression model that describes how interest point intensity varies for different facial expressions. The set of interest points and the intensity distribution model are person-specific.

The mechanism and effect of the use of the variable-intensity template are as follows. The intensities of the interest points vary significantly due to the movements of the facial parts created when showing facial expressions (see Fig. 1). Focusing on this characteristic, we recognize facial expressions by matching the model intensities prepared in advance against the observed intensities. We want a system that is robust against head pose variations even if the shape model used produces some approximation error. To achieve this, we must handle the problem that the interest points will be shifted from their actual positions due to the shape model error as the out-of-plane head rotation angles increase. Hence, we site the interest points away from the edges of facial parts to keep the change in intensity small when the face is rotated (see Fig. 1).

Our contribution¹ is that we propose a facial expression recognition method for varying facial poses based on the key idea that facial expressions can be correctly recognized by knowing how the interest point intensities vary with the facial expressions. The main advantage of our method is that a face model for each person is created simply by capturing frontal face images of the person making the target facial expressions. Thus it is more efficient than existing

¹ A part of this work appeared in Asian Conference on Computer Vision (ACCV) (Kumano et al. 2007). The present work extends our previous work in several important aspects. First, to improve the performance, the face model is changed from the cylinder previously used to an average face shape model. Second, an intensity adjustment is applied to handle changes in intensity of the face caused by such as illumination variations and vertical head rotations.

methods that need measurement instruments (e.g. Gokturk et al. 2002), time-consuming manual intervention processes (e.g. Dornaika and Davoine 2008), large amounts of training data (e.g. Tong et al. 2007) and so on. Furthermore, we implement a particle filter that utilizes the variable-intensity template as a face model, to simultaneously estimate facial poses and expressions.

The remainder of this paper is organized as follows. First, a brief overview of related work is presented in Sect. 2. Section 3 describes our proposed method. Then, in Sect. 4, experimental results are given. Finally, a summary and future work are given in Sect. 5.

2 Related Work

While most existing facial expression recognition systems assume that the person in the video sequence does not rotate the head significantly and that the image shows a nearly frontal view of the face (e.g. Lanitis et al. 1997; Huang and Huang 1997; Tian et al. 2001; Chang et al. 2006; Bartlett et al. 2006; Fasel et al. 2004; Kotsia and Pitas 2007), some methods consider head pose variations. Facial pose and expression vary independently, and, image variations caused by facial expression change are often smaller than those caused by head movement.

Hence, the precise separation of these two components is needed to achieve correct recognition of facial expressions. Unfortunately, the highly non-linear effect on face appearance in the image due to large out-of-plane head rotations is not easy to handle. In addition, only a few video databases containing non-frontal faces have been published, e.g. Face Video Database,² so it is difficult to compare the recognition ability of existing methods quantitatively. Given these issues, pose-invariant facial expression recognition can be said to be an underdeveloped research area.

In this section, we briefly overview the facial expression recognition methods that allow some out-of-plane head rotations with regard to their face shape models. Excellent reviews of a number of recent studies assuming near-frontal faces can be found in Pantic and Rothkrantz (2000b); Fasel and Luetten (2003); Tian et al. (2005); Pantic and Bartlett (2007).

In what follows, we divide the existing methods into the following five groups with regards to the shape model used: (1) face shapes directly measured, (2) face shapes recovered from images, (3) general face shapes, (4) parametric shapes, and (5) shape-free approach. These shape models, except for group (2) and (5), are detailed below in descending order of their approximation accuracy. Group (1–4) are

divided into two classes by the following three kinds of classifications: measured models (1) versus non-measured models (2–4), person-specific models (1, 2) versus person-independent models (3, 4), and deformable models (1–3) versus rigid models (4).

(1) *Face shapes directly measured* The most conclusive way to make the estimation robust against head rotations is to directly measure the actual three-dimensional shape of each user with an instrument other than monocular camera. Gokturk et al. (2002) use a stereo cameras to accurately measure the three-dimensional positions of 19 facial feature points defined on the user's face and their deformations due to change in facial expression. In the test stage, they classify facial expressions using deformation parameters estimated by using an optical-flow-like technique. Their experimental results suggest that their method can recognize facial expressions in a variety of head poses. In the field of head tracking, there is a similar approach to handle facial expression variations, e.g. Oka and Sato (2005). However, these methods cannot be applied to monocular video sequences, because they require stereo cameras to measure the three-dimensional positions of feature points in the initial frame of test video sequences. Dispensing with images, the systems of Wang et al. (2006); Tang and Huang (2008) correctly recognize facial expressions regardless of head pose from face range data acquired by a three-dimensional digitizer.

(2) *Face shapes recovered from images* Some methods try to recover the actual three-dimensional shape of the user from a monocular video sequence by using a structure-from-motion technique. Xiao et al. (2004) propose 2D + 3D Active Appearance Models (AAM) where first the positions of the set of feature points in the video sequence are tracked by using a common two-dimensional AAM and then the three-dimensional face shape is recovered by using a non-rigid structure-from-motion technique. Lucey et al. (2006) apply 2D + 3D AAM to recognize facial expressions. However, in their experiments, the recovered three-dimensional shapes were inaccurate which rather degraded the recognition rates.

(3) *General face shapes* Some methods utilize models that can be applied to arbitrary users, i.e. person-independent models. Cohen et al. (2003) utilize the Piecewise Bezier Volume Deformation (PBVD) tracker which tracks feature points defined on a generic three-dimensional wireframe model of the face. They classify facial expressions using its motion parameters as the temporal cues. Dornaika and Davoine (2008) describe a system that uses a existing deformable face mesh model; its deformation parameters are related to action units. They estimate facial expressions in the framework of a particle filter (Isard and Blake 1998),

²Face Video Database of the Max Planck Institute for Biological Cybernetics in Tuebingen: <http://vdb.kyb.tuebingen.mpg.de/>.

after estimating the head pose by a gradient-based method. However, generating an accurate general deformable model is not so easy in practice. Gross et al. (2005) suggest that such person-independent models cannot cover the large interpersonal variations in face shape and facial expression expected with sufficient accuracy.

(4) *Parametric shapes* Some methods utilize simple rigid shapes rather than complicated deformable shapes. Black and Yacoob (1997), Tong et al. (2007), and Liao and Cohen (2006) approximate the human face as plane (the first two) and cylinder (the last), respectively. As the feature, Black and Yacoob (1997) and Liao and Cohen (2006) utilize the optical flow in the face region, while Tong et al. (2007) use Gabor wavelet coefficients. However, using the simple parametric shape causes large error in the alignment of the face. Hence, facial expressions at large head rotation angles tend to be misrecognized in the absence of a mechanism for handling the face misalignment. Tong et al. introduced Dynamic Bayesian Network (DBN) inference, where the relationships between action units are modeled. Their experiments suggest that this framework can compensate some misclassification of action units. However, it is reasonable to expect that it cannot handle heavy face misalignments caused by large out-of-plane head rotations. Furthermore, note that optical flow estimation is easily corrupted by illumination changes and non-rigid motion.

(5) *Shape-free approach* Hu and his colleagues developed a facial expression recognizer that does not utilize any face shape model (Hu et al. 2008). By utilizing Support Vector Machines (SVMs), they discriminated horizontal head orientations (discretized at 15 degree or larger angle interval) and facial expressions sequentially or simultaneously. Although their system can handle faces in the range of frontal-to profile-view, the discretization error for head pose in real situations, where the head pose is continuous, is expected to degrade the facial expression recognition. Note that preparing a large number of classifiers (each of which covers a small angle range) is undesirable in terms of processing time

and memory resources. Furthermore, they assume that the face regions in input data are precisely extracted. This, however, remains difficult for current computer vision technologies.

Overall, to the best of our knowledge, there is no work that matches all of our goals:

- Monocular video capture systems are supported. The approach that directly measures face shapes (1) does not have this advantage.
- A face model for each person can be easily prepared. The general model-based approach (3) and shape-free approach (5) fail to achieve this advantage because of model complexity.
- Facial expressions can be estimated even with a large change in facial pose. The recovered face shape model-based approach (2) and the parametric model-based approach (4) have not achieved this end.

The method proposed herein is a novel *point intensity-based* approach that can realize all of the above goals. We describe a change in facial expression not as the deformation of the shape model or as optical flow in facial part regions, but just as a change in the intensities of local face regions. In addition, none of the existing papers, other than Wang et al. (2006); Hu et al. (2008), describe how the recognition rate of facial expressions varies with out-of-plane head rotation angle. Hence, pose-invariability of their methods cannot be quantified while the problem of the difference in test data remains. In this paper, we also describe experiments that assess the impact of head pose on the recognition rate.

3 Proposed Method

Our method consists of two stages (see Fig. 2). First, we prepare a variable-intensity template for each person from just one frontal face image for each facial expression (referred to as the training images in this paper). Second, we estimate facial pose and expression simultaneously within the framework of a particle filter.

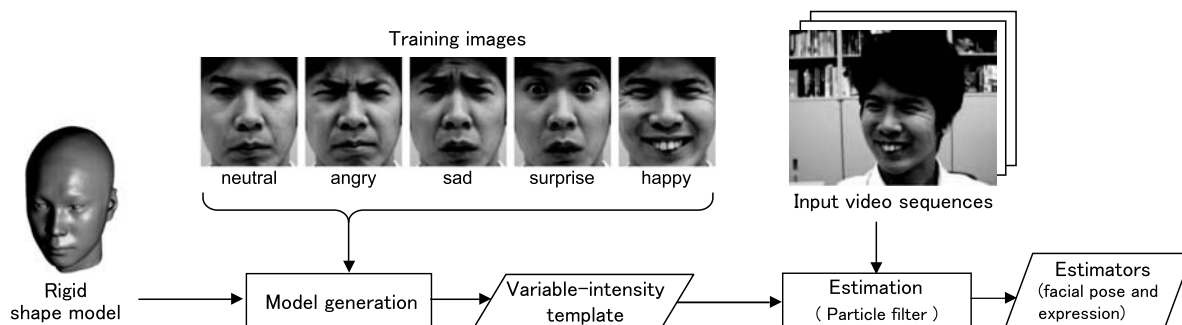


Fig. 2 System flow chart

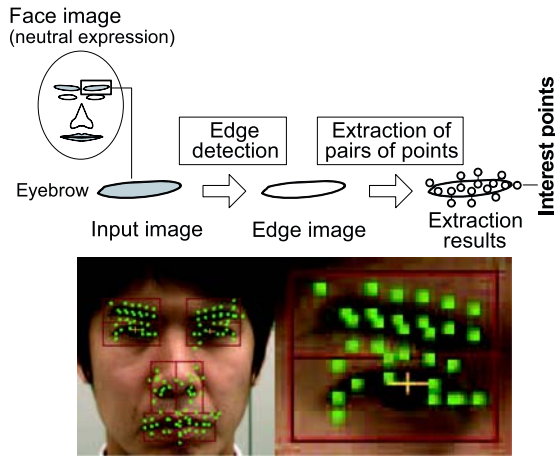


Fig. 3 Upper: The method used to extract interest points \mathcal{P} . Lower: Example of a set of paired interest points \mathcal{P} , shown as 128 green dots. The large rectangles represent the boundaries of facial parts, detected by the method described in Sect. 3.3. The right side shows the expanded right eyebrow and eye regions

3.1 Variable-Intensity Template

The variable-intensity template \mathcal{M} consists of the following three components:

$$\mathcal{M} = \{\mathcal{S}, \mathcal{P}, \mathcal{I}\} \tag{1}$$

where \mathcal{S} , \mathcal{P} , and \mathcal{I} denote a rigid face shape model, a set of interest points, and an intensity distribution model, respectively. The intensity distribution model describes the intensity distribution of each interest point for different facial expressions.

The variable-intensity template is generated from a set of training images and the rigid face shape model. The training image set consists of one face image for each facial expression, $\{\mathbf{g}_{e=1}, \dots, \mathbf{g}_{e=N_e}\}$, as shown in Fig. 2, where \mathbf{g}_e denotes the image labeled with facial expression e ($\in \{1, \dots, N_e\}$), and N_e is the number of target expressions. In particular, the training image for neutral expression is denoted by \mathbf{g}_{NEU} . The face in each image \mathbf{g} is frontal, and virtually fixed. The way to generate such training images is described in Sect. 3.3.

As the face shape model \mathcal{S} , we use the average face shape model³ shown in Fig. 2 in this paper. The approach used to fit it to each user is described in Sect. 3.4.

3.1.1 Set of Interest Points \mathcal{P}

The multiple interest points are sparsely defined away from the edges of facial parts such as eye corners in the training

image in neutral expression \mathbf{g}_{NEU} , as shown in the lower part of Fig. 3. All interest points are fixed on the rigid shape model \mathcal{S} . The set of interest points \mathcal{P} is described as

$$\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_{N_p}\} \tag{2}$$

where \mathbf{p}_i denotes the image coordinates of the i -th interest point in the training image \mathbf{g}_{NEU} , and N_p denotes the number of interest points. The set of interest points \mathcal{P} is different for each subject.

An interest point constitutes a pair of points (herein called *paired interest points* or *paired points*) (Matsubara and Shakunaga 2005) that straddle and are centered on the edge, to detect the bidirectional motions of the facial parts. The paired points are selected in the four facial part regions (eyebrows, eyes, nose, and mouth) in the training image, \mathbf{g}_{NEU} . These facial part regions are detected by the method described in Sect. 3.3. Examples of these pairs are shown in the lower part of Fig. 3: The total number of interest points for each subject is set to be 128 (eyebrows: 20×2 , eyes: 12×2 , nose: 24, and mouth: 40). This number was decided after conducting preliminary evaluations of the performance metrics of accuracy and processing speed.

The paired points are extracted in each facial part as follows (see also the upper part of Fig. 3).

- (1) Edges are detected in each of the four facial parts as zero-cross boundaries of the Laplacian-Gaussian filtered image.
- (2) All paired point candidates that straddle and are centered on the edge are extracted. These candidates are, then, grouped into eight according to their direction of intensity gradient (up, down, left, right, or one of four oblique directions). The distance between the points in each pair was empirically set to be $1/60$ and $1/45$ of the face width for eye and other regions, respectively. The short distance set in eye regions is based on the small movements of eyes.
- (3) The pairs are selected from the candidates one by one in ascending order of the interpair difference in intensity. The current target candidate is selected, if the center of the target candidate is separated from all of the centers of pre-selected pairs by at least a threshold distance. After one candidate is selected in the current target group or no candidate can be selected, the next target group is selected. The initial group is the upmost group, and the initial threshold distance is set to one-fifth of the width of the target facial part region. They were empirically decided.
- (4) The selection is terminated, as soon as the number of pairs reaches the limit. If the total number of selected pairs does not reach the limit, the current threshold of the interpair distance is decreased and the selection process is reentered at (3).

³Average head dummy of Japanese young males contains over 100,000 polygons without texture. It is published by Digital Human Research Center, Advanced Industrial Science and Technology, <http://www.dh.aist.go.jp/research/centered/facedummy/>.




Fig. 4 Intensity distribution model \mathcal{I} : The intensity distributions of interest points, described as normal distributions, indicate the change in facial expressions. The colors in the *right part* correspond to the interest points in the *left part*

Note that the resulting locations of the interest points are different for subjects due to the differences in positions of the edges.

3.1.2 Intensity Distribution Model \mathcal{I}

The intensity distribution model describes how the interest point intensity varies for different facial expressions. As shown in Fig. 4, the interest point intensity changes strongly due to the shift of its associated facial part. Focusing on this property, we recognize facial expressions from the changes in observed interest point intensities.

The observed intensity of each interest point is varied by small position shifts of the interest point. This is caused by error in the shape model as well as the errors in the intensity adjustment, described in Sect. 3.2.2. Such variation is represented as a normal distribution in our method. The intensity distribution model \mathcal{I} is described as:

$$\mathcal{I} = \{\mathcal{N}_1, \dots, \mathcal{N}_{N_p}\}, \quad \mathcal{N}_i = \mathcal{N}(\mu_i(e), \sigma_i(e)), \quad (3)$$

$$\sigma_i(e) = k\mu_i(e) \quad (4)$$

where $\mathcal{N}(\mu, \sigma)$ denotes a normal distribution with mean μ and standard deviation σ , and $\mu_i(e)$ and $\sigma_i(e)$ denote the mean and standard deviation of the intensity of point i for expression e , respectively. We set the intensity mean $\mu_i(e)$ to be the value recorded from the training image labeled with expression e , \mathbf{g}_e , at image coordinate \mathbf{p}_i . Furthermore, we assume that standard deviation σ_i is proportional to mean μ_i , and set the proportionality factor, k , to be 0.3.

3.2 Simultaneous Estimation of Facial Pose and Expression by Using a Particle Filter

Our method simultaneously estimates the facial pose and expression by calculating their likelihood given the observed intensity of interest points with the intensity distribution model. The joint distribution of facial pose and expression




Fig. 5 Dynamic Bayesian Network for head poses, facial expressions, and face images

at time t given all face images up to that time ($\mathbf{z}_{1:t}$) is recursively represented as follows:

$$p(\mathbf{h}_t, e_t | \mathbf{z}_{1:t}) = \alpha p(\mathbf{z}_t | \mathbf{h}_t, e_t) \int p(\mathbf{h}_t | \mathbf{h}_{t-1}) \times \sum_{e_{t-1}} P(e_t | e_{t-1}) p(\mathbf{h}_{t-1}, e_{t-1} | \mathbf{z}_{1:t-1}) d\mathbf{h}_{t-1} \quad (5)$$

where facial pose state \mathbf{h}_t and expression state e_t follow first order Markov processes (see Fig. 5); \mathbf{h}_t and e_t are assumed to be conditionally dependent given image \mathbf{z}_t ; Bayes' rule and conditional dependence are used along with marginalization (e.g. Russell and Norvig 2003); α is a normalization constant.

The facial pose state \mathbf{h}_t consists of the following six continuous variables: the coordinate of the center of the template on the image plane, three-dimensional rotation angles (yaw, pitch, and roll), and scale. For the head motion model, $p(\mathbf{h}_t | \mathbf{h}_{t-1})$, we utilize a random walk model where the components of head pose state are considered to be independent each other. The system noise for each component is created by a zero-mean Gaussian process. The variances of the system noises were empirically decided.

With regards to facial expression, we set $P(e_t | e_{t-1})$ to be equal for all expression combinations in this paper. Any expression transition matrix can be utilized in our framework, but it is difficult to obtain the actual transition. Accordingly, in this paper, we adopt the simple distribution model, where no prior knowledge of facial expression transitions is assumed.

Equation (5), unfortunately, cannot be calculated exactly, because parameters of facial pose \mathbf{h}_t are continuous, and their distributions are complex due to occlusion, etc. Hence, we use a particle filter (Isard and Blake 1998), which calculates (5) by approximating the posterior density as a set of weighted samples called particles. Each particle expresses a state and its weight. In our method, the state and weight of the l -th particle are expressed as $[\mathbf{h}_t^{(l)}, e_t^{(l)}]$ and $\omega_t^{(l)}$, where $\omega_t^{(l)}$ is proportional to $p(\mathbf{z}_t | \mathbf{h}_t^{(l)}, e_t^{(l)})$ calculated using (6) and $\sum_l \omega_t^{(l)} = 1$ is satisfied.

3.2.1 Likelihood of Facial Poses and Expressions

The likelihood of facial pose \mathbf{h}_t and expression e_t for face image \mathbf{z}_t is expressed as $p(\mathbf{z}_t | \mathbf{h}_t, e_t)$. Assuming that the in-




Fig. 6 Shape of robust function $\rho(x)$, weight function $w(x) = (d\rho(x)/dx)/x$, and $\exp(-\rho(x))$, similar to the likelihood function

tensities of the interest points are independent, we can transform likelihood $p(\mathbf{z}_t|\mathbf{h}_t, e_t)$ as follows:

$$p(\mathbf{z}_t|\mathbf{h}_t, e_t) = \prod_{i \in \mathcal{P}'} p(\tilde{z}_{i,t}|\mathbf{h}_t, e_t) \tag{6}$$

where $\tilde{z}_{i,t}$ denotes the illumination-adjusted intensity (see Sect. 3.2.2) of interest point i at time t , and \mathcal{P}' denotes the set of non-occluded interest points. Here, we consider that the interest point is not occluded if the surface normal of its corresponding point on the face shape model is pointing toward the camera.

We define the likelihood of facial pose \mathbf{h}_t and expression e_t for the illumination-adjusted intensity $\tilde{z}_{i,t}$, $p(\tilde{z}_{i,t}|\mathbf{h}_t, e_t)$, by adopting a robust estimation as:

$$p(\tilde{z}_{i,t}|\mathbf{h}_t, e_t) = \frac{1}{\sqrt{2\pi}\sigma_i(e_t)} \exp\left[-\frac{1}{2}\rho(d_i)\right], \tag{7}$$

$$d_i = \frac{\tilde{z}_{i,t} - \mu_i(e_t)}{\sigma_i(e_t)} \tag{8}$$

where function $\rho(\cdot)$ denotes a robust function. In this paper, we use the Geman McClure function (Geman and McClure 1987) with scaling factor $c(=9)$ which regulates an infinite input (see Fig. 6):

$$\rho(x) = c \cdot \frac{x^2}{1+x^2}. \tag{9}$$

This robust function $\rho(\cdot)$ makes the estimation more robust against such noise as imaging noise, intensity adjustment error, and large position shifts due to shape model error.

3.2.2 Intensity Adjustment

We should adjust the intensity of interest points observed in the input image to make the estimation robust against changes in intensity of face itself. Such changes in intensity caused by changes in illumination as well as head movements, especially up-down head rotations, which changes the orientation of the light source to the face.

We obtain the illumination-adjusted intensity $\tilde{z}_{i,t}$ as:

$$\tilde{z}_{i,t} = \gamma_{b,t} z_{i,t} \tag{10}$$

where $z_{i,t}$ denotes the intensity of interest point i , and $\gamma_{b,t}$ represents the intensity adjustment factor for facial block b

at time t . That is, we divide the face into four blocks, {left eyebrow and eye, right eyebrow and eye, left parts of nose and mouth, and right parts of nose and mouth}, and assume that the rate of change in interest point intensity is uniform in each block. Although this assumption is not strictly valid, the small error present in intensity adjustment estimation does not severely disturb facial expression recognition. The reason is that the interest points defined in the vicinity of facial parts yield significant differences in intensity between facial expressions.

Intensity $z_{i,t}$ is obtained as follows. The image coordinate of the i -th interest point under head pose \mathbf{h}_t is obtained via three processes: (1) orthogonal projection from the training image plane onto the shape model \mathcal{S} , (2) translation and rotation of \mathcal{S} according to pose \mathbf{h}_t , and (3) projection of the interest point i on the shape model \mathcal{S} onto the target image plane.

We define the intensity adjustment factor $\gamma_{b,t}$ as the maximum likelihood estimator (MLE) of $p(\gamma_b z_{i,t}|\mathbf{h}_t, e_t)$. This estimation of $\gamma_{b,t}$ can be simplified to the following equation by considering the log likelihood:

$$\hat{\gamma}_{t,b} = \arg \max_{\gamma_b} \prod_{i \in \mathcal{P}'_b} p(z_{i,t}|\gamma_b, \mathbf{h}_t, e_t) \tag{11}$$

$$= \arg \min_{\gamma_b} \sum_{i \in \mathcal{P}'_b} \rho\left(\frac{\gamma_b z_{i,t} - \mu_i(e_t)}{\sigma_i(e_t)}\right) \tag{12}$$

where \mathcal{P}'_b denotes the set of non-occluded interest points in facial block b . It is noted that this robust estimation problem can be effectively solved by using an iterative method such as iteratively reweighted least squares (Beaton and Tukey 1974). The estimation of the intensity adjustment factor $\hat{\gamma}_{b,t}$ by the iteratively reweighted least squares technique proceeds as follows:

1. Initialize $N_b \times N_b$ weight matrix $\mathbf{W} = \text{diag}[1, \dots, 1]^T$, where N_b denotes the number of elements in set \mathcal{P}'_b .
2. Estimate current intensity adjustment factor $\gamma_b^{(m)}$ using the previous weight matrix $\mathbf{W}^{(m-1)}$:

$$\gamma_b^{(m)} = (\mathbf{z}_b^T \mathbf{W}^{(m-1)} \mathbf{z}_b)^{-1} \mathbf{z}_b^T \mathbf{W}^{(m-1)} \boldsymbol{\mu}_b \tag{13}$$

where m denotes the iteration step number, and \mathbf{z}_b and $\boldsymbol{\mu}_b$ denote $N_b \times 1$ vectors of observed intensity and mean of intensity distribution model divided by standard deviation, $\mathbf{z}_b = [\dots, z_i/\sigma_i, \dots]^T$ and $\boldsymbol{\mu}_b = [\dots, \mu_i/\sigma_i, \dots]^T, i \in \mathcal{P}'_b$.

3. Update weight matrix \mathbf{W} :

$$W_{ii}^{(m)} = w\left(\frac{\gamma_b^{(m)} z_i - \mu_i}{\sigma_i}\right) \tag{14}$$

where W_{ii} denotes the i -th diagonal element of weight matrix \mathbf{W} , and $w(\cdot)$ denotes the weight function related

to the robust function $\rho(\cdot)$ as $w(x) = (d\rho(x)/dx)/x$ (see Fig. 6).

4. Steps 2 and 3 are repeated until convergence.

3.2.3 Estimators of Facial Pose and Expression

Estimators of facial pose and expression at time t , \hat{h}_t and \hat{e}_t , are calculated as:

$$\hat{h}_t = \sum_l \omega_t^{(l)} \mathbf{h}_t^{(l)}, \quad (15)$$

$$\hat{e}_t = \arg \max_e \sum_l \omega_t^{(l)} \delta_e(e_t^{(l)}) \quad (16)$$

where the recognized facial expression \hat{e}_t is defined to be the expression that maximizes the marginal probability $p(e_t | z_{1:t})$. The function $\delta_e(e')$ is the indication function such that $\delta_e(e') = 1$ if $e = e'$, and $\delta_e(e') = 0$ otherwise.

3.3 Generating Training Images and Detecting Facial Parts

Generating training images We prepare training images, \mathbf{g} , where the iris centers are laid on a horizontal line. Such training images are generated from source face images (referred to as training source images) that contain a human face in one of the target facial expressions. We assume that the face in the training source images is stationary between neutral expression and non-neutral expressions for the same subject. The position and size of the faces, however, may be different for each person. The training images are the training source images that are rotated in-plane until iris centers fall on a horizontal line. The iris centers are detected in the eye regions by utilizing Fast Radial Symmetry (Loy and Zelinsky 2003; Zhang et al. 2006). This rotation is calculated with each training source image only for neutral expression. The non-neutral expression images are rotated by the same rotation matrix for the neutral expression image of the same person.

Detection of face and facial parts Before selecting the interest points (described in Sect. 3.1.1), the face and the four facial parts (eyebrows, eyes, nose and mouth) are detected in the neutral training image \mathbf{g}_{NEU} . The face region is roughly detected by using a cascaded AdaBoost detector based on Haar-like features (Viola and Jones 2001). The four facial parts are detected as rectangular boundaries in the face region, as shown in Fig. 3.

Eye and mouth regions are detected by using cascaded AdaBoost detectors (Castrillon et al. 2007), respectively. These detectors often return multiple candidates and we select the most likely candidate. We empirically defined the likelihoods with rough previous knowledge about their positions and sizes: $L(\mathbf{y}) = \prod_k \mathcal{N}(y_k; \mu_k, \sigma_k)$, where $L(\mathbf{y})$



Fig. 7 Example of the shape model fitted to a person

is the likelihood of the feature vector of target candidate $\mathbf{y} (= [X, Y, W, H]^T)$. Vector $[X, Y]^T$ denotes the position of the center of the candidate region, and W and H are its width and height, respectively. Variables μ_k and σ_k describe the empirically defined mean and standard deviation of the k -th component of \mathbf{y} , respectively.

Nose region is given by the position of nostrils extracted as the valley in the horizontal and vertical intensity profiles between eye and mouth regions. Eyebrow regions are simply set as the regions that contact (above) the eye regions with the same size.

3.4 Shape Model Fitting

We fitted the average face shape to each user as follows: (1) fit the center of the shape model to the center of face region in the training image in neutral expression, \mathbf{g}_{NEU} , detected by the method of Viola and Jones (2001), (2) stretch the shape model in the horizontal and vertical directions to match both face width and height; stretching in the depth direction used the scaling factor given as the square root of the product of vertical and horizontal scaling factors. An example of the fitting results is given in Fig. 7.

4 Experimental Results

To evaluate the estimation performance of the proposed method, we performed three experiments: The objective of the first test was to evaluate the accuracy of our system in recognizing facial expressions when the face does not face the camera. The second test evaluates the performance of our system with person-independent strategy. The third test evaluates the performance of the system in head pose tracking.

At the same time, to evaluate the effectiveness of our interest points (paired points described in Sect. 3.1.1), we prepared two other kinds of variable-intensity templates: those with randomly allocated points and those with points defined on the edges (these points are hereinafter referred to as random points and edge points, respectively). These points




Fig. 8 Example of two kinds of set of interest points: randomly allocated (*left*) and defined on edges (*right*). These points are used for the comparison with our interest points shown in Fig. 3

are also defined in each of the four facial part regions as shown in Fig. 8.

In the following experiments, the recognition rates were calculated as the ratio between the number of frames wherein the estimated expression matched the ground truth to the total number of target frames. In this paper, we ran our system offline to evaluate the recognition rates of facial expression. The number of particles was set to 1,500, and the processing time was about 80 ms/frame on a Pentium D processor at 3.73 GHz with 3.0 GB RAM. Currently, the identity of a person is given a priori.

4.1 Performance Evaluation for Non-Frontal-View Faces

Unless the true face shape model is used, the interest points are misaligned due to error in the shape model. Such shift increases as the head is more rotated in out-of-plane. Accordingly, it is important to evaluate the robustness of the system against head rotation. Unfortunately, to the best of our knowledge, there is no facial expression database that contains video sequences where the face moves together with various changes in facial expression. Hence, we generated such video datasets and used them to evaluate our system.

4.1.1 Our Original Facial Expression Datasets

Our datasets contain the following two kinds of video sequences: *fixed-pose dataset* and *free-pose dataset*. Both datasets contain facial expression categories of neutral, angry, sad, surprise and happy. Both sets include training images and test video sequences for each subject. The training images were captured immediately prior to the capture of the test video sequences. That is, each couple of training images and test sequences was captured in the same situation on the same day. These video sequences were captured by the same IEEE 1394 XGA (1024 × 768 pixel) color camera at 15 fps. In this paper, we utilized these video sequences after converting them into grayscale 512 × 384 pixel sequences.

In the fixed-pose dataset, subjects exhibited multiple facial expressions with the head fixed in one of three directions relative to the camera: $0/\pm 20/\pm 40$ degree in yaw (horizontal), $0/\pm 20$ degree in pitch (vertical), or $0/\pm 20/\pm 40$ degree in roll (in-plane) (see Figs. 9–11). Yaw, pitch, and roll sequences contain nine subjects (seven males and two females in their 20 s to 40 s), four subjects (four males of the

nine subjects), and one subject (one male of the four males), respectively. For each direction, each subject provided one sequence.

All sequences start with neutral expression face looking toward the camera, then move to one direction relative to the camera. Finally, the subjects create five facial expressions one by one without moving their head, each for a duration of 60 frames followed by a 60 frame interval with relaxed face as indicated by the instructions displayed on a monitor, which are used as truth of the facial expression.

The free-pose dataset, on the other hand, contains one sequence of the subject, who is also contained in the fixed-pose dataset. The subject freely shows five facial expressions one by one while shaking the head left and right.

To capture the training images, we gave the subjects the following instructions: (1) face the head directly at the camera during the capture process, (2) show facial expressions one by one as requested by the instructions with category name on the monitor, (3) press a key on a keyboard, when the target expression was being expressed, to have a facial image captured as the training image.

4.1.2 Evaluation with Fixed-Pose Dataset

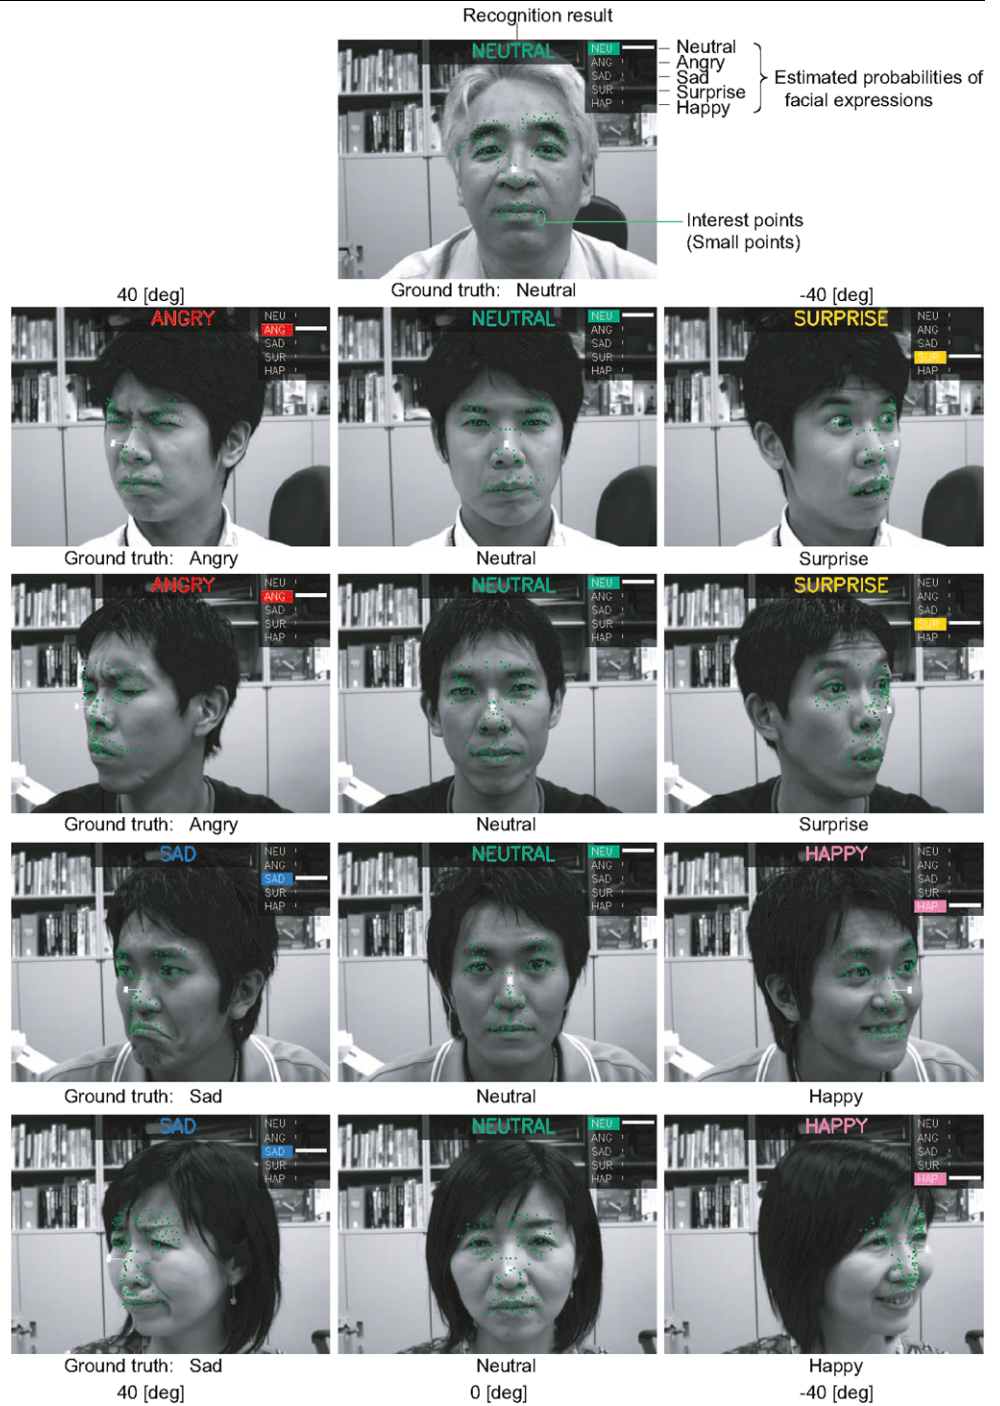
First, by using the fixed-pose dataset, we evaluate how accurately our system can recognize facial expressions for various head orientations. In calculating the recognition rates, we excluded the first 20 frames of each expression just after the instruction was displayed, because of the time lag between the instruction and the creation of the facial expression.

Figures 9, 10, and 11 show some estimation results of facial poses and expressions by using our paired points for horizontal, vertical and in-plane head rotations, respectively. They show the facial poses and expressions were correctly estimated for all subjects.

Table 1 shows a comparison of our paired points and random points with the average facial expression recognition rates for each target head rotation angle. For out-of-plane (yaw and pitch) rotations, the recognition rate decreased as the head rotation angle increased. Nevertheless, our paired interest points are more robust against the out-of-plane rotations, than the random points. For in-plane direction, our method recognized all facial expressions perfectly, because the misalignment of interest points due to error in the shape model does not occur with in-plane rotation. Accordingly, we evaluated it just for one subject. Our paired points achieved at least 90(%) in average for all rotations, yaw, pitch and roll.

Table 2 shows the confusion matrix for the average recognition rates with the paired interest points. These results suggest that, for many subjects, the difference in face appearance between sad expression and both neutral and angry ex-

Fig. 9 Some estimation results of facial poses and expressions with fixed-pose dataset (*Horizontal*): The expression in the upper part of each image denotes the recognized result. The width of each bar in the upper right part of each image denotes the estimated probability of each facial expression, $P(e_t | z_{1:t})$. *Small points* on each face denote interest points



pressions is smaller than for other expression combinations, accordingly, they were sometimes confused with each other.

On the other hand, with the edge interest points, head tracking was lost for many video sequences. The reason for this seems to be that such interest points are too sensitive to misallocation of interest points. The misallocation is caused by errors in the shape model and slight differences in facial expression between the test sequences and training images.

4.1.3 Evaluation with Free-Pose Dataset

We next evaluated our system with the free-pose dataset, where the subject changes both the facial expression and the head's horizontal orientation.

Five frames of the video sequence of the free-pose dataset and the estimated results of facial expression and pose in

Fig. 10 Some estimation results of facial poses and expressions with the fixed-pose dataset (*Vertical*): The expression in the upper part of each image denotes the recognized result. The width of each bar in the upper right part of each image denotes the estimated probability of each facial expression, $P(e_t|z_{1:t})$. *Small points* on each face denote interest points




Fig. 11 Some estimation results of facial poses and expressions with the fixed-pose dataset (*In-plane rotation*): The expression in the upper part of each image denotes the recognized result

each frame⁴ are shown in (a), (b) and (c) of Fig. 12, respectively. The ground truth of the facial expression in every frame was hand-labeled by the subject. Figure 12(b) shows that facial expressions were recognized correctly in almost all frames. In addition, the correct expressions were assigned significantly higher probabilities than other expressions. A quantitative evaluation for head pose is given in Sect. 4.3.

4.2 Performance Evaluation with Person-Independent Strategy

The main advantage of our system is that our system makes it possible to generate the variable-intensity template for each person with very little time and effort. As demonstrated in Sect. 4.1.2, by using our person-specific model, we can

⁴A part of video sequences for the result with both the fixed-pose dataset and the free-pose dataset are available from <http://www.hci.iis.u-tokyo.ac.jp/~kumano/papers/IJCV2008/>.

Table 1 Average recognition rates of facial expressions for a variety of head orientations with fixed-pose dataset by comparing two types of interest points: Pair: our paired interest points. Random: points randomly defined in each facial part region

Point type	Total	Angle [deg]				
		-40	-20	0	20	40
Yaw (Horizontal): nine sequences						
Pair	92.3	83.3	94.3	95.4	95.9	92.5
Random	90.1	90.2	92.1	95.4	96.1	76.7
Pitch (Vertical): four sequences						
Pair	94.0	N/A	87.0	97.5	97.6	N/A
Random	86.4	N/A	74.1	98.1	87.0	N/A
Roll (In-plane): one sequence						
Pair	100.0	100.0	100.0	100.0	100.0	100.0
Random	100.0	100.0	100.0	100.0	100.0	100.0

Unit is (%)

Table 2 Average confusion matrix of facial expressions with fixed-pose dataset for horizontal and vertical directions: GT and RCG denote ground truth and recognition, respectively. The overall recognition rate is 93.1[%]

GT\RCG	Neutral	Angry	Sad	Surprise	Happy
Neutral	88.9	1.5	7.9	1.6	0.1
Angry	0.4	97.6	0.9	1.1	0.0
Sad	3.2	8.0	85.4	3.3	0.1
Surprise	0.0	0.0	4.2	95.8	0.0
Happy	0.2	1.1	0.0	1.1	97.7

Unit is (%)

achieve the high facial expression recognition performance. Such a person-specific strategy is useful for applications where the number of users is small. However, some applications require that it can be applied to arbitrary users without such training for each user. Accordingly, we evaluate the potential of our system to be extended to the person-independent strategy.

To compare our system (with person-independent strategy) against existing methods, we utilized a widely used public facial expression database, the Cohn-Kanade DFAT-504 database (Kanade et al. 2000). The Cohn-Kanade database consists of FACS-coded (Ekman and Friesen 1978) video sequences of 104 subjects, starting from a neutral expression and ending in the peak of the facial expression. For each person, there are on average 8 frames for each expression. The desired facial displays were described and modeled before the material was recorded.

Although there are a variety of facial expression sequences coded with Action Units (Ekman and Friesen 1978), or components of facial expression, in the Cohn-

Kanade database, they are not labeled with facial expression categories. Hence, we hand-labeled the last frame of them with one of the non-neutral expressions, following the FACS rule (Ekman et al. 2002). We excluded about one third of sequences that were not labeled with six basic facial expressions (Ekman and Friesen 1975): angry, sad, surprise, happy, fear, and disgust. Next, we excluded about one half of the remaining sequences where the stationary head assumption (described in Sect. 3.3) was violated, that is, where the head moved while the facial expressions were being demonstrated. 129 video sequences were finally selected (13, 20, 14, 43, 12, and 27 for angry, sad, surprise, happy, fear, and disgust expressions, respectively) from the 53 subjects. All initial frames of the selected 129 sequences were inevitably labeled as neutral. All labeled frames, i.e. the first and last frames in each selected video sequence, were used as training image candidates. That is, our target facial expressions in the person-independent strategy are seven categories: neutral and six basic facial expressions.

The average facial expression recognition rates were calculated as follows. First, we tracked (estimated facial pose and expression) over all video sequences for each subject in the database. Next, we picked up the recognition results of the first frames and the last frames for all tracked sequences. Thus, the total number of the target frames for neutral expression is the sum of those for the non-neutral expression (each of the six categories started with a neutral expression). Finally, we calculated the recognition rate separately for each of the seven expressions.

Learning stage We developed a fully automatic system for learning a person-independent, or general, variable-intensity template. Given training images, it proceeds as follows.

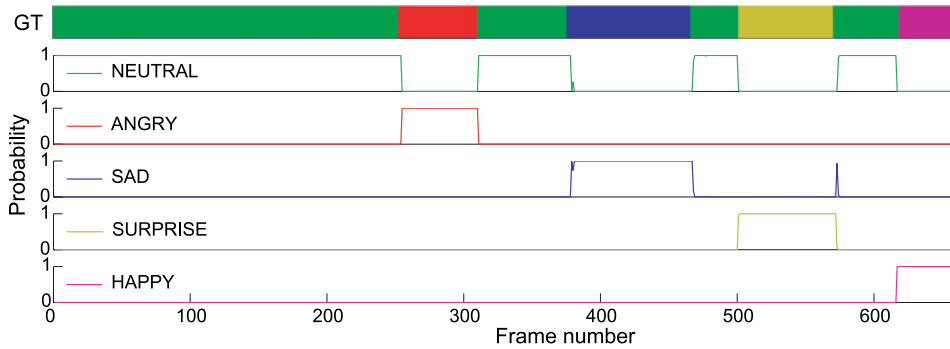
First, facial parts are detected in the rotationally aligned training image (Sect. 3.3) in neutral expression of each subject. The face part regions in non-neutral expressions are set to be the same as those of the same subject in neutral expression. Hereinafter, these detected face images are just referred to as training images, and those in expression e for the j -th subject are denoted as $\mathbf{g}_{j,e}$, especially those in neutral expression also as $\mathbf{g}_{j,NEU}$.

Next, all training images of each subject are normalized to remove the effect of global lighting, skin color, etc. The normalization is given by $\tilde{\mathbf{g}}_{j,e} = (\mathbf{g}_{j,e} - \alpha^{(j)}\mathbf{1})/\beta^{(j)}$. The normalization factors for the j -th subject, $\alpha^{(j)}$ and $\beta^{(j)}$, are calculated as the mean and variance of intensity, respectively, in the rectangle in $\mathbf{g}_{j,NEU}$ that encloses eyebrow and mouth regions.


Then, average face images of each facial expression e , $\bar{\mathbf{g}}_e$, are generated individually for every facial part p with the normalized images $\tilde{\mathbf{g}}_{j,e}$ for all subject j . This average images of each facial part p are hereinafter referred to as average facial part images, and denoted as $\bar{\mathbf{g}}_{e,p}$. An example of the average facial part image $\bar{\mathbf{g}}_{e,p}$ is shown in Fig. 13.



(a) Input video sequence (from left to right, frame number 100, 290, 400, 560, 660)



(b) Ground truth (*top*) and recognition results (*others*) of facial expression: The probability of correct expression is remarkably higher than that of other expressions



(c) Estimation results of facial pose (horizontal axis equals that of (b)): Facial poses are estimated with enough accuracy to detect three cycles of head shake movement (*red solid line*)

Fig. 12 Images and estimation results in the free-pose dataset



Fig. 13 The average face image for each facial expression, \bar{g}_e , learned from the Cohn-Kanade DFAT-504 database. Each average face image consists of the average facial part images, $\bar{g}_{e,p}$

Finally, a set of interest points is selected in the average facial part images $\bar{g}_{e,p}$ as described in Sect. 3.1.1. A sample of the resulting set of interest points is shown in the left side of Fig. 14.

Testing stage In the test stage, the variable-intensity template for the person is prepared by using the average face images \bar{g}_e .

First, the facial parts are detected in the initial frame of the test video sequence, after rotating the frame until the iris centers fall on a horizontal line using the approach described in Sect. 3.3 (hereafter this rotated image is just referred to as the initial frame).

Next, the coordinates of the i -th interest point in the initial frame, p_i , is calculated by scaling and translating its coordinates in the average facial part images $\bar{g}_{e,p}$. This mapping adapts the average facial part image to the same facial part region in the initial frame both in size and position. The right part of Fig. 14 shows the mapping results.

Finally, the mean intensity in the intensity distribution in each facial expression $\mu(e)$ is restored via denormalization of the average facial part images: $\hat{g}_{e,p} = \beta^{(j)} \bar{g}_{e,p} + \alpha^{(j)} \mathbf{1}$, where $\hat{g}_{e,p}$ is the restored (training) image, $\alpha^{(j)}$ and $\beta^{(j)}$ are calculated as described in the above test stage process, where the face region in the initial frame is used as $g_{j,NEU}$.

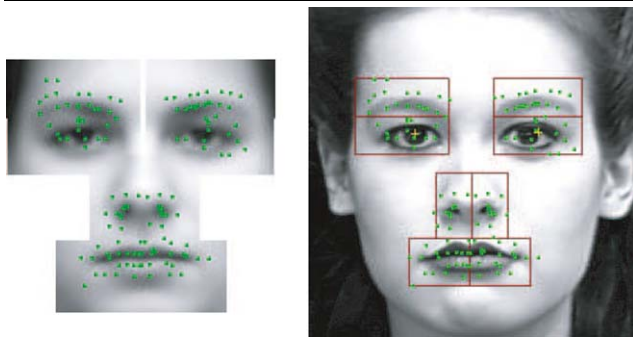


Fig. 14 *Left*: Example of set of interest points defined in the average facial part images $\bar{g}_{e,p}$ that are learned with the Cohn-Kanade DFAT-504 database. *Right*: The set of interest points \mathcal{P} generated by projecting the average facial part image into the initial frame, or the training image for neutral expression $g_{j,NEU}$, of a target subject



Fig. 15 Sample recognition results with the set of interest points \mathcal{P} shown in the right part of Fig. 14

Finally, the mean intensity $\mu_i(e)$ is set to be the intensity of the restored training image, $\hat{g}_{e,p}$, at coordinates p_i .

Recognition performance We evaluated the recognition performance achieved by this person-independent strategy, based on the leave-one-subject-out cross-validation. The results with our paired points and those with random points are shown in Table 3 (also in Fig. 15) and Table 4, respectively. The overall recognition rate was about 60% with our paired points, and about 70% with the random points. In both results, fear expression recorded significantly lower recognition rates than the other expressions. The average recognition rate excluding fear expression with random points is about 80%.

We were encouraged by these results with the person-independent strategy, even though these recognition rates are lower than those with our person-specific strategy in Sect. 4.1.2 and those of state-of-the-art methods such as Littlewort et al. (2006); Sebe et al. (2007); Zhao and Pietikainen (2007);⁵ Kotsia and Pitas (2007); Yang et al. (2008) (over 90% for all methods). Note that, our original target is the person-specific strategy, and the objective of this evaluation was to assess if our approach has the potential to be extended to the person-independent version.

⁵The neutral expression was not one of their targets.

Table 3 Confusion matrix with the Cohn-Kanade DFAT-504 database by using our paired interest points with person-independent strategy. The overall recognition rate is 59.8[%]

GT\RCG	N	A	Sd	Sp	H	F	D
Neutral (N)	82.9	5.7	1.9	1.6	0.3	5.7	1.9
Angry (A)	23.1	30.7	7.7	7.7	7.7	7.7	15.4
Sad (Sd)	20.0	15.0	35.0	5.0	5.0	10.0	10.0
Surprise (Sp)	0.0	0.0	0.0	100.0	0.0	0.0	0.0
Happy (H)	5.0	3.7	0.0	2.5	73.8	2.5	12.5
Fear (F)	16.7	0.0	8.3	8.3	33.4	25.0	8.3
Disgust (D)	3.8	19.2	3.8	0.0	0.0	1.9	71.3

Unit is (%)

Table 4 Confusion matrix with the Cohn-Kanade DFAT-504 database by using randomly defined points with person-independent strategy. The overall recognition rate is 70.2[%]

GT\RCG	N	A	Sd	Sp	H	F	D
Neutral (N)	82.1	4.1	4.6	3.8	0.0	1.6	3.8
Angry (A)	7.7	69.2	0.0	0.0	0.0	15.4	7.7
Sad (Sd)	15.0	5.0	70.0	5.0	0.0	5.0	0.0
Surprise (Sp)	0.0	0.0	0.0	100.0	0.0	0.0	0.0
Happy (H)	0.0	3.7	0.0	2.5	83.8	5.0	5.0
Fear (F)	25.0	0.0	8.3	8.3	33.4	25.0	0.0
Disgust (D)	9.6	3.8	3.8	3.8	0.0	17.3	61.7

Unit is (%)

The main difference between the person-specific strategy and the person-independent strategy is whether the interest points are correctly allocated. With the person-specific strategy, the paired points are correctly allocated directly to the face image of the target subject, as shown in Fig. 3. On the other hand, with the person-independent strategy, we locate the paired points in the average faces learned with the database. Hence, without precise facial part detection, these interest points are likely to be misallocated in the actual face image for each subject, as shown in Fig. 16. Consequently, the recognition rates with person-independent strategy seem to be significantly degraded. By improving the facial part detection, the performance of our paired points with the person-independent strategy is expected to outperform those with the random points.

The interest point misallocation affects more severely for the facial expressions that produce similar appearances each other. For example, as shown in Fig. 13, angry, sad and disgust expressions resemble each other, especially in the eyebrow regions. Inevitably, the eyebrow regions are likely to yield large localization errors, especially with our rectangle-based detector, due to their large interpersonal variation in position, size, and shape. Furthermore, misrecognition of




Fig. 16 Example of misallocation of interest points in eyebrow region. *Left*: interest points defined in the average face. *Right*: those allocated in the actual face image. Both images shows neutral expression. The interest points on the eyebrow in the average image do not so in the actual image




Fig. 17 Comparison between the estimated poses and the ground truth for one sequence. *Dashed*: estimations by our method; *Solid*: the ground truth. *Black*: yaw; *Red*: pitch; *Blue*: roll

some neutral expressions, in both Tables 3 and 4, also seems to be caused by the facial part mislocalization. The neutral eyebrow image in the right part of Fig. 16 associates angry and disgust expressions (lowered eyebrow).

We suggest that the low recognition rates for fear expression are caused by the large interpersonal difference in the display of fear expression in the database, especially in the eyebrow regions (the mouth region is chiefly determiner of the fear expression). Consequently, the eyebrow regions in average face image in Fig. 13 are significantly blurred compared to the other expressions. Such a variety of fear expressions are hardly to be recognized correctly with the degenerated intensity distribution model. Consequently, many fear expressions are mistakenly recognized as happy expression, where the mouth region resembles those in fear expression. This issue could be mitigated by modifying our method to divide the intensity distribution model into upper and lower face regions (cf. Action Units).

4.3 Performance Evaluation for Head Tracking

Finally, we evaluate the performance of the proposed method for head tracking, separated from those for facial expression recognition, by using the Boston University (BU) face tracking database⁶ (Cascia et al. 2000). The

⁶<http://www.cs.bu.edu/groups/ivc/HeadTracking/>



Fig. 18 Sample tracking results on BU face tracking database

Table 5 Comparison with mean absolute error in head pose estimation between the proposed method, its variations in interest points, and existing methods. The accuracies of these two existing methods are taken from Murphy-Chutorian and Trivedi (2008)

Methods	Errors [deg]		
	Yaw	Pitch	Roll
Pair	7.1	4.2	2.9
Edge	9.5	6.6	5.5
Random	7.1	4.9	3.0
Cascia et al. (2000)	3.3	6.1	9.8
Xiao et al. (2003)	3.8	3.2	1.4

BU face tracking database contains sets of 30[fps] QVGA (320 × 240 pixel) MPEG compressed video sequences. These sequences were captured under uniform illumination or varying illumination together with head pose recorded by a magnetic sensor. Subjects were asked to perform free head motion, which included translation and rotation. We use the set of video sequences captured under uniform illumination. Moreover, the BU database does not contain any frames with distinctive facial expression. Accordingly, we generated the intensity distribution model that describes only neutral expression.

The tracking results are shown in Figs. 17 and 18. The mean absolute angular error of our head pose estimation module for yaw, pitch, and roll is shown in Table 5 together with the results of two existing methods listed in Murphy-Chutorian and Trivedi (2008). Our system is ranked between them. Table 5 also shows that vertical head pose detection is more accurate than horizontal detection.

These results demonstrated that the performance of head pose estimation with the paired points almost matches that with random points and bests that with edge points. The cause of the low accuracy with edge points seems that they

are sensitive to the misalignment caused by shape model error and to the slight change in neutral expression. Compared to the random points, little effect of the use of the paired points is found in Table 5.

4.4 Discussion

This section evaluated the estimation performance of the proposed method in three aspects: facial expressions in non-frontal-view with person-specific strategy, person-independent strategy, and head tracking accuracy.

The results with person-specific strategy using our dataset demonstrate that our method (person-specific) can achieve the high performance for facial expression recognition even with a large head rotation in almost all orientations. This is because it avoids the problem of personal differences in terms of the geometric arrangement of facial parts and facial expressions. These results suggest that locating interest points away from the edges of facial parts can suppress the severe impact of interest point shift produced by out-of-plane rotation of the head with shape model error.

By enhancing interest point selection module, the performance of our system would be increased. First, the interest point selection module in this paper uses only a neutral expression image. However, selecting salient points for each facial expression is expected to increase the recognition performance. Second, the head tracking will become more stable by more accurately fitting the face shape model.

Our experiments also demonstrate the potential of our approach in supporting person-independent applications. The performance of the person-independent version is expected to be improved by enhancing the facial part detection module using techniques such as active contours (Kass et al. 1988) or profile-based matching (Huang and Huang 1997), i.e. shifting each paired point along the line connecting the interpair points to be centered on the edge.

Based on these results with both strategies, we believe that our person-independent version could handle large head rotations as well as the person-specific version. Compared to the existing deformation-based facial expression recognition methods, our approach has the large advantage that facial expression models can be learned just from monocular images, that is, many facial expression databases can be utilized. Furthermore, unlike the dense-feature-based approach (such as optical-flow-based one), our approach can introduce any three-dimensional face shape model to improve the robustness against large head rotations, without increasing the computational cost.

5 Summary and Future Work

In this paper, we presented a particle filter-based method for estimating facial pose and expression simultaneously; it uses

a novel face model called the variable-intensity template. Our method has the distinct advantage that a face model for each person can be prepared very easily with a simple step. Our method recognized the five facial expressions with overall accuracy of 95.4% for horizontal, vertical and in-plane facial orientations in the range of ± 40 , ± 20 , and ± 40 degrees from the frontal view, respectively.

We intend to tackle the recognition of subtle spontaneous facial expressions. To this end, we would like to apply unsupervised learning with an online clustering technique, and to estimate the intensity of facial expressions from changes in interest point intensity by referring to optical flow estimation.

References

- Bartlett, M. S., Littlewort, G., Frank, M. G., Lainscsek, C., Fasel, I. R., & Movellan, J. R. (2006). Automatic recognition of facial actions in spontaneous expressions. *Journal of Multimedia*, 1(6), 22–35.
- Beaton, A. E., & Tukey, J. W. (1974). The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *Technometrics*, 16(2), 147–185.
- Black, M. J., & Yacoob, Y. (1997). Recognizing facial expressions in image sequences using local parameterized models of image motion. *International Journal of Computer Vision*, 25(1), 23–48.
- Cascia, M. L., Sclaroff, S., & Athitsos, V. (2000). Fast, reliable head tracking under varying illumination: An approach based on registration of texture-mapped 3D models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(4), 322–336.
- Castrillon, M., Deniz, O., Guerra, C., & Hernandez, M. (2007). Encara2: Real-time detection of multiple faces at different resolutions in video streams. *Journal of Visual Communication and Image Representation*, 18(2), 130–140.
- Chang, Y., Hu, C., Feris, R., & Turk, M. (2006). Manifold based analysis of facial expression. *Image and Vision Computing*, 24(6), 605–614.
- Cohen, I., Sebe, N., Garg, A., Chen, L. S., & Huang, T. S. (2003). Facial expression recognition from video sequences: temporal and static modeling. *Computer Vision and Image Understanding*, 91(1–2), 160–187.
- Dornaika, F., & Davoine, F. (2008). Simultaneous facial action tracking and expression recognition in the presence of head motion. *International Journal of Computer Vision*, 76(3), 257–281.
- Ekman, P., & Friesen, W. V. (1975). *Unmasking the face: a guide to recognizing emotions from facial expressions*. Englewood Cliffs: Prentice Hall.
- Ekman, P., & Friesen, W. V. (1978). *The facial action coding system: a technique for the measurement of facial movement*. Palo Alto: Consulting Psychologists Press.
- Ekman, P., Friesen, W. V., & Hager, J. C. (2002). *FACS investigator's guide. A human face*.
- Fasel, B., & Luetttin, J. (2003). Automatic facial expression analysis: Survey. *Pattern Recognition*, 36, 259–275.
- Fasel, B., Monay, F., & Gatica-Perez, D. (2004). Latent semantic analysis of facial action codes for automatic facial expression recognition. In *Proceedings of the ACM SIGMM international workshop on multimedia information retrieval* (pp. 181–188).
- Geman, S., & McClure, D. E. (1987). Statistical methods for tomographic image reconstruction. *Bulletin of the International Statistical Institute*, LII, 5–21.

- Gokturk, S. B., Tomasi, C., Girod, B., & Bouquet, J. (2002). Model-based face tracking for view-independent facial expression recognition. In *Proceedings of the IEEE international conference on automatic face and gesture recognition* (pp. 287–293).
- Gross, R., Matthews, I., & Baker, S. (2005). Generic vs. person specific active appearance models. *Image and Vision Computing*, 23(11), 1080–1093.
- Hu, Y., Zeng, Z., Yin, L., Wei, X., Zhou, X., & Huang, T. S. (2008). Multi-view facial expression recognition. In *Proceedings of the IEEE international conference on automatic face and gesture recognition*.
- Huang, C. L., & Huang, Y. M. (1997). Facial expression recognition using model-based feature extraction and action parameters classification. *Journal of Visual Communication and Image Representation*, 8(3), 278–290.
- Isard, M., & Blake, A. (1998). Condensation—conditional density propagation for visual tracking. *International Journal of Computer Vision*, 29(1), 5–28.
- Kanade, T., Cohn, J., & Tian, Y. L. (2000). Comprehensive database for facial expression analysis. In *Proceedings of the IEEE international conference on automatic face and gesture recognition* (pp. 46–53).
- Kass, M., Witkin, A., & Terzopoulos, D. (1988). Snakes: Active contour models. *International Journal of Computer Vision*, 1(4), 321–331.
- Koelstra, S., & Pantic, M. (2008). Non-rigid registration using free-form deformations for recognition of facial actions and their temporal dynamics. In *Proceedings of the IEEE international conference on automatic face and gesture recognition*.
- Kotsia, I., & Pitas, I. (2007). Facial expression recognition in image sequences using geometric deformation features and support vector machines. *IEEE Transactions on Image Processing*, 16(1), 172–187.
- Kumano, S., Otsuka, K., Yamato, J., Maeda, E., & Sato, Y. (2007). Pose-invariant facial expression recognition using variable-intensity templates. In *Proceedings of Asian conference on computer vision* (Vol. 1, pp. 324–334).
- Lanitis, A., Taylor, C. J., & Cootes, T. F. (1997). Automatic interpretation and coding of face images using flexible models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7), 743–756.
- Liao, W. K., & Cohen, I. (2006). Belief propagation driven method for facial gestures recognition in presence of occlusions. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition workshop* (pp. 158–163).
- Littlewort, G., Bartlett, M. S., Fasel, I. R., Susskind, J., & Movellan, J. R. (2006). Dynamics of facial expression extracted automatically from video. *Image and Vision Computing*, 24(6), 615–625.
- Loy, G., & Zelinsky, A. (2003). Fast radial symmetry for detecting points of interest. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(8), 959–973.
- Lucey, S., Matthews, I., Hu, C., Ambadar, Z., Torre, F., & Cohn, J. (2006). AAM derived face representations for robust facial action recognition. In *Proceedings of the IEEE international conference on automatic face and gesture recognition* (pp. 155–160).
- Matsubara, Y., & Shakunaga, T. (2005). Sparse template matching and its application to real-time object tracking. *IPSJ Transactions on Computer Vision and Image Media*, 46(9), 17–40 (in Japanese).
- Murphy-Chutorian, E., & Trivedi, M. M. (2008). Head pose estimation in computer vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (in press).
- Oka, K., & Sato, Y. (2005). Real-time modeling of face deformation for 3D head pose estimation. In *Proceedings of the IEEE international conference on automatic face and gesture recognition* (pp. 308–320).
- Otsuka, K., Sawada, H., & Yamato, J. (2007). Automatic inference of cross-modal nonverbal interactions in multiparty conversations: “who responds to whom, when, and how?” from gaze, head gestures, and utterances. In *Proceedings of the international conference on multimodal interfaces* (pp. 255–262).
- Pantic, M., & Bartlett, M. (2007). Machine analysis of facial expressions. In *I-Tech education and publishing* (pp. 377–416).
- Pantic, M., & Rothkrantz, L. (2000a). Expert system for automatic analysis of facial expression. *Image and Vision Computing*, 18, 881–905.
- Pantic, M., & Rothkrantz, L. J. M. (2000b). Automatic analysis of facial expressions: the state of the art. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22, 1424–1445.
- Russell, S., & Norvig, P. (2003). *Artificial intelligence—a modern approach*. Paris: Pearson Education.
- Sebe, N., Lew, M. S., Sun, Y., Cohen, I., Gevers, T., & Huang, T. S. (2007). Authentic facial expression analysis. *Image and Vision Computing*, 25(12), 1856–1863.
- Tang, H., & Huang, T. S. (2008). 3D facial expression recognition based on properties of line segments connecting facial feature points. In *Proceedings of the IEEE international conference on automatic face and gesture recognition*.
- Tian, Y. L., Kanade, T., & Cohn, J. F. (2001). Recognizing action units for facial expression analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(2), 97–115.
- Tian, Y. L., Kanade, T., & Cohn, J. (2005). *Facial expression analysis*. Berlin: Springer.
- Tong, Y., Liao, W., & Ji, Q. (2007). Facial action unit recognition by exploiting their dynamic and semantic relationships. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(10), 1683–1699.
- Viola, P., & Jones, M. (2001). Rapid object detection using a boosted cascade of simple features. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition* (pp. 511–518).
- Wang, J., Yin, L., Wei, X., & Sun, Y. (2006). 3D facial expression recognition based on primitive surface feature distribution. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition* (pp. 1399–1406).
- Xiao, J., Moriyama, T., Kanade, T., & Cohn, J. (2003). Robust full-motion recovery of head by dynamic templates and re-registration techniques. *International Journal of Imaging Systems and Technology*, 13, 85–94.
- Xiao, J., Baker, S., Matthews, I., & Kanade, T. (2004). Real-time combined 2D + 3D active appearance models. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition* (Vol. 2, pp. 535–542).
- Yang, P., Liu, Q., Cui, X., & Metaxas, D. N. (2008). Facial expression recognition based on dynamic binary patterns. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition*.
- Zhang, W., Chen, H., Yao, P., Li, B., & Zhuang, Z. (2006). Precise eye localization with AdaBoost and fast radial symmetry. In *Proceedings of the international conference on computational intelligence and security* (Vol. 1, pp. 725–730).
- Zhao, G., & Pietikainen, M. (2007). Dynamic texture recognition using local binary patterns with an application to facial expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(6), 915–928.
- Zhu, Z., & Ji, Q. (2006). Robust real-time face pose and facial expression recovery. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition* (pp. 681–688).