

Analysis of Machine Learning Methods for Real-Time Recognition of Facial Expressions from Video

(MPLab TR 2003.05)

Gwen Littlewort, Marian Stewart Bartlett, Ian Fasel, Joel Chenu, Javier R. Movellan
Machine Perception Laboratory Institute for Neural Computation
University of California, San Diego
San Diego, CA 92093-0523

Abstract

We present a systematic comparison of machine learning methods applied to the problem of fully automatic recognition of facial expressions, including AdaBoost, support vector machines, and linear discriminant analysis. Each video-frame is first scanned in real-time to detect upright-frontal faces. The faces found are scaled into image patches of equal size and sent downstream for further processing. Gabor energy filters are applied at the scaled image patches followed by a recognition engine that codes facial expressions into 7 dimensions in real time: neutral, anger, disgust, fear, joy, sadness, surprise. We report results on a series of experiments comparing spatial frequency ranges, feature selection techniques, and recognition engines. Best results were obtained by selecting a subset of Gabor filters using AdaBoost and then training Support Vector Machines on the outputs of the filters selected by AdaBoost. The generalization performance to new subjects for a 7-way forced choice was 93% and 97% correct on two publicly available datasets, the best performance reported so far on these datasets. Surprisingly, registration of internal facial features was not necessary, even though the face detector does not provide precisely registered images. The outputs of the classifier change smoothly as a function of time and thus can be used for unobtrusive motion capture. We developed an end-to-end system that provides facial expression codes at 24 frames per second and animates a computer generated character in real time.

1. Introduction

In this paper we present results on a user independent fully automatic system for real time recognition of basic emotional expressions from video. The system automatically detects frontal faces in the video stream and codes each frame with respect to 7 dimensions: Neutral, anger, disgust, fear, joy, sadness, surprise. We report results on a series of experiments comparing spatial frequency ranges, feature

selection techniques, and recognition engines. Best results were obtained by selecting a subset of Gabor filters using AdaBoost and then training Support Vector Machines on the outputs of the filters selected by AdaBoost. The combination of AdaBoost and SVM's enhanced both speed and accuracy of the system. The system presented here also differs from previous work in that it is fully automatic and operates in real-time at a high level of accuracy (93% generalization to new subjects on a 7-alternative forced choice). Another distinction is that the preprocessing does not include explicit detection and alignment of internal facial features. This provides a savings in processing time which is important for real-time applications.

2. Facial Expression Data

The facial expression system was trained and tested on Cohn and Kanade's DFAT-504 dataset [5]. This dataset consists of 100 university students ranging in age from 18 to 30 years. 65% were female, 15% were African-American, and 3% were Asian or Latino. Videos were recoded in analog S-video using a camera located directly in front of the subject. Subjects were instructed by an experimenter to perform a series of 23 facial expressions. Subjects began and ended each display with a neutral face. Before performing each display, an experimenter described and modeled the desired display. Image sequences from neutral to target display were digitized into 640 by 480 pixel arrays with 8-bit precision for grayscale values.

For our study, we selected 313 sequences from the dataset. The only selection criterion was that a sequence be labeled as one of the 6 basic emotions. The sequences came from 90 subjects, with 1 to 6 emotions per subject. The first and last frames (neutral and peak) were used as training images and for testing generalization to new subjects, for a total of 625 examples. The trained classifiers were later applied to the entire sequence.

2.1 Real-time Face Detection

We developed a real-time face detection system based on [14] that consists of a cascade of classifiers trained by boosting techniques. The complete system, including enhancements to [14] such as employing Gentleboost instead of Adaboost, smart feature search, and a novel cascade training procedure, are described in [8]. Source code for the face detector is freely available at <http://kolmogorov.sourceforge.net>. The face detector was trained on 5000 faces and millions of non-face patches from about 8000 images collected from the web by Compaq Research Laboratories. Accuracy on the CMU-MIT dataset, a standard public data set for benchmarking frontal face detection systems, is 90% detections and 1/million false alarms, which is state-of-the-art accuracy. The CMU test set has unconstrained lighting and background. With controlled lighting and background, such as the facial expression data employed here, detection accuracy is much higher. The system presently operates at 24 frames/second on a 3ghz Pentium IV for 320x240 images.

All faces in the DFAT-504 dataset were successfully detected. The automatically located faces were rescaled to 48x48 pixels. The typical distance between the centers of the eyes was roughly 24 pixels. A comparison was also made at double resolution (96x96). No further registration was performed. Many other approaches to automatic facial expression recognition include explicit detection and alignment of internal facial features. The recognition system presented here performs well without that step, providing a considerable savings in processing time. The images were converted into a Gabor magnitude representation, using a bank of Gabor filters at 8 orientations and 5 spatial frequencies (4:16 pixels per cycle at 1/2 octave steps) [7].

3. Facial Expression Classification

Facial expression classification was based on support vector machines (SVM's). SVM's are well suited to this task because the high dimensionality of the Gabor representation $O(10^5)$ does not affect training time, which depends only on the number of training examples $O(10^2)$. The system performed a 7-way forced choice between the following emotion categories: Happiness, sadness, surprise, disgust, fear, anger, neutral.

3.1. Strategies for multiclass decisions with SVM's

Support vector machines make binary decisions. There are a number of methods for making multiclass decisions with a set of binary classifiers. (See [4] for a review). Here, the seven-way forced choice for six emotions plus neutral was trained in two stages. In stage I, support vector ma-

chines performed binary decision tasks. We explored three approaches to training binary decisions: one-versus-one, one-versus-all, and all possible partitions. Stage II converts the representation produced by the first stage into a probability distribution over the seven expression categories. To this effect, we have implemented and evaluated several approaches: K-nearest neighbor, a simple voting scheme, and multinomial logistic ridge regression.

Partitioning into binary decisions. There are a number of strategies for partitioning the classification task into binary decisions. The simplest strategy is to train 1 versus all. Pairwise partitioning strategies have been advocated by [6] and [11], whereas others (e.g. [1]) advocate exploring the space of all possible partitions.

For pairwise partitioning (1:1), SVM's were trained to discriminate all pairs of emotions. For seven categories that makes 21 SVM's. In 1:1 partitioning, the number of training samples for each SVM may be relatively small. If some subjects performed some expressions and not others, as in this dataset, identity signals can interfere with the learning of expression. To avoid this, we trained on identity-matched pairs, where for example, the happy vs. surprise SVM is trained on only those subjects who gave samples of both happiness and surprise. An alternative to training SVM's to discriminate each pair of emotions was to train SVM's to discriminate one emotion from everything else (1:all). This strategy employed a larger number of training examples, 626, which diluted identity effects. An extension of the 1:all approach was to consider all possible non-trivial binary partitions of the classes. With 7 classes, there are seven 1:6 classifiers, twenty one 2:5 classifiers and thirty five 3:4 classifiers.

Combining outputs of multiple binary classifiers. In the system presented here, the SVM outputs were combined to make a 7 alternative forced choice. The most common way to combine SVM outputs for multiclass decisions is by voting. This procedure counts the number of stage 1 classifiers aligned with each emotion. For example, if one SVM indicates happiness and not surprise, happiness gets +1 and surprise gets -1. These votes are summed over all of the SVM's. Softmax ensures each class is allocated a number between 0 and 1, with unit sum over classes. We also explored a variation on voting which uses the sum of the classifier margins, which are typically clustered around +1 or -1, instead of the binary outputs. This variation made little difference, and the voting results presented here use binary outputs. We compared voting to nearest neighbor, and to a learned mapping based on multinomial logistic ridge regression (MLR). In nearest neighbor, the continuous SVM output (the margin) for each of the n SVM's gives an n-dimensional pattern vector. The test image is as-

signed the class of the training image with the shortest Euclidean distance between their pattern vectors. MLR learns the weight matrix that maps the outputs of Stage one classifiers onto the 7 emotions. MLR is a maximum likelihood approach, which is equivalent to a single layer perceptron with weight decay and with SoftMax competition between the outputs. The regression was implemented using the Newton-Raphson method and a ridge term coefficient of 0.001. The advantage of this data-dependent second stage is that it could learn common confusions and biases which lead to errors in a direct voting situation.

		Nnbr	Voting	MLR
linear SVM's	1:1	82.7	81.6	85.8
	1:all	81.6	86.2	87.5
	all poss.	83.0	87.2	89.4
nonlinear SVM's	1:1	83.2	82.9	86.1
	1:all	81.4	88.0	89.8
	all poss.	85.1	89.9	90.4

Table 1: Comparison of strategies for multiclass decisions using SVM's.

Results Generalization to novel subjects was tested using leave-one-subject-out cross-validation. Results are given in Table 1. Linear, polynomial, and RBF kernels with Laplacian, and Gaussian basis functions were explored. Linear and RBF kernels employing a unit-width Gaussian performed best, and are presented here. The soft margin approach, allowing some training examples to lie within the margin, did not enhance performance.

For Stage I, partitioning the classification task into binary decisions, 1:all usually outperformed 1:1 partitioning, and all possible partitions gave the best performance. Of the Stage II strategies for combining the outputs of multiple SVM's into a 7-way forced choice, MLR was substantially better than nearest neighbor (5.3 percentage points). Voting was slightly but consistently less effective than MLR, typically 1.3 percent for 1:all and all partitions.

For the comparisons in the subsequent sections, 1:all partitioning followed by voting was employed due to training speed. The optimal strategies determined in this section (all possible partitions and MLR) will be reintroduced in the final system.

3.2. SVM's and Adaboost

SVM performance was next compared to Adaboost for emotion classification. The features employed for the Adaboost emotion classifier were the individual Gabor fil-

ters. The comparison was performed using 48x48 pixel images at 5 spatial scales (4:16 pixels per cycle). This gave $5 \times 8 \times 48 \times 48 = 92,160$ possible features. A subset of these features was chosen using Adaboost. On each training round, the Gabor feature with the best expression classification performance for the current boosting distribution was chosen. The performance measure was a weighted sum of errors on a binary classification task, where the weighting distribution (boosting) was updated at every step to reflect how well each training vector was classified.

Adaboost training continued until the classifier output distributions for the positive and negative samples were completely separated by a gap proportional to the widths of the two distributions (see Figure 1). The union of all features selected for each of the 7 emotion classifiers resulted in a total of 538 features.

Classification results are given in Table 2. The generalization performance with Adaboost was comparable to linear SVM performance. Adaboost had a substantial speed advantage, as shown in Table 3. There was a 170-fold reduction in the number of Gabor filters used. The convolutions were calculated in pixel space, rather than Fourier space which reduced the advantage of feature selection, but it nevertheless resulted in a substantial speed benefit.

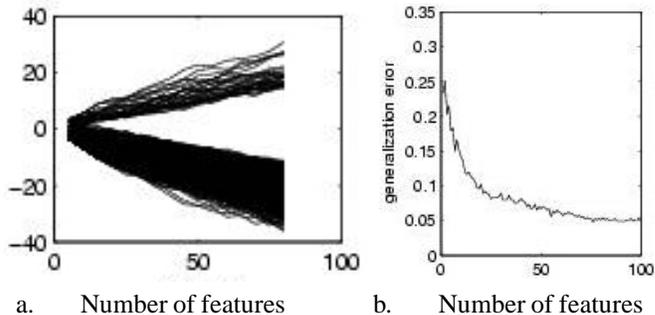


Figure 1: Stopping criteria for Adaboost training. a. Output of one expression classifier during Adaboost training. The response for each of the training examples is shown as a function of number of features as the classifier grows. b. Generalization error as a function of the number of features chosen by Adaboost. Generalization error does not increase with 'overtraining'.

3.3 Combining feature selection by Adaboost with classification by SVM's

Adaboost is not only a fast classifier, it is also a feature selection technique. An advantage of feature selection by Adaboost is that features are selected contingent on the features that have already been selected. In feature selection by Adaboost, each Gabor filter is treated as a weak classifier. Adaboost picks the best of those classifiers, and then boosts the weights on the examples to weight the errors more. The

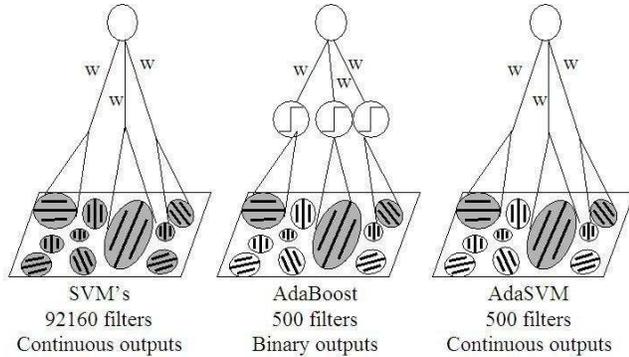


Figure 2: SVM’s learn weights for the continuous outputs of all 92160 Gabor filters. AdaBoost selects a subset of features and learns weights for the thresholded outputs of those filters. AdaSVM’s learn weights for the continuous outputs of the selected filters.

next filter is selected as the one that gives the best performance on the errors of the previous filter. At each step, the chosen filter can be shown to be uncorrelated with the output of the previous filters [3, 12].

We explored training SVM classifiers on the features selected by Adaboost. When the SVM’s were trained on the thresholded outputs of the selected Gabor features, they performed no better than Adaboost. However, we trained SVM’s on the continuous outputs of the selected filters. We informally call these combined classifiers AdaSVM in abbreviation of Adaptive Boosting Selected Feature representations in Support Vector Machines. AdaSVM’s outperformed straight Adaboost by 3.8 percent points, a difference that was statistically significant ($z=1.99$, $p=0.02$). AdaSVM’s outperformed SVM’s by an average of 2.7 percent points, an improvement that was marginally significant ($z = 1.55$, $p = 0.06$).

ω	kernel	Adaboost	SVM	AdaSVM
4:16	Linear	87.2	86.2	88.8
4:16	RBF		88.0	90.7
2:32	Linear	90.1	88.0	93.3
2:32	RBF		89.1	93.3

Table 2: Leave-one-out generalization performance of Adaboost, SVM’s and AdaSVM’s (48x48 images). ω : Gabor wavelength range, sampled at 0.5 octave intervals.

Distribution of spatial frequencies selected by Adaboost. The features selected by Adaboost showed no preference for orientation, but the highest frequencies were chosen

	SVM		Adaboost	AdaSVM	
	Lin	RBF		Lin	RBF
Time t	t	$90t$	$0.01t$	$0.01t$	$0.0125t$
Time t'	t	$90t$	$0.16t$	$0.16t$	$0.2t$
Memory	m	$90m$	$3m$	$3m$	$3.3m$

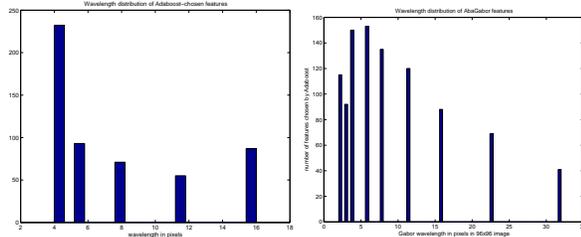
Table 3: Processing time and memory considerations. Time t' includes the extra time to calculate the outputs of the 538 Gabors in pixel space for Adaboost and AdaSVM, rather than the full FFT employed by the SVM’s.

more often. Figure 3a shows the number of chosen features at each of the 5 wavelengths used (4:16 pixels per cycle on 48x48 pixel images). The distribution was skewed towards the high spatial frequencies. Examination of this distribution suggested that including higher spatial frequency Gabors may improve recognition performance. Indeed, doubling the resolution to 96x96 and adding 4 more wavelengths in the shorter end of the distribution improved performance of the nonlinear AdaSVM to 93.1% correct. The new frequency range was 2:32 pixels per cycle sampled at 0.5 octave steps for a total of 9 wavelengths. At this resolution, the performance advantage of AdaSVM’s over SVM’s was statistically significant. At higher resolution, the speed benefit of AdaSVM’s becomes even more apparent. The full Gabor representation was 7 times larger than before, whereas the number of Gabors selected by Adaboost only increased by a factor of 1.7 (900 from 538).

It turned out that using a larger range of Gabor wavelengths in the 48x48 images was sufficient to achieve the performance improvement. The bottom rows of Table 2 show performance for 48x48 images sampled at 9 Gabor wavelengths, from 2 to 32 pixels per cycle. The result of 93% accuracy for a user-independent 7-alternative forced choice was encouraging, given that previously published results on this database were 81-83% accuracy.

We then reintroduced the approaches to multiclass SVM’s found to be optimal in Section 3.1, and applied them to the AdaSVM system. Results for using all possible class partitions and training an MLR matrix instead of voting are shown in Table 4. The performance enhancement with these approaches is small, if any. Optimal performance with the AdaSVM was obtained with the simpler paradigm of 1:all partitions and voting, which is a considerable savings in training time over all possible partitions and MLR.

Number of Support Vectors We next examined the effect of feature selection by Adaboost on the number of support vectors. Smaller numbers of support vectors proffer two advantages: (1) the classification procedure is faster, and (2) the expected generalization error decreases as the



a. 48x48 pixel images b. 96x96 pixel images

Figure 3: Wavelength distribution of features selected by Adaboost. With 48x48 images, the distribution of selected features was skewed to the shorter wavelengths. Doubling the resolution and including more Gabors in the shorter wavelengths made the distribution more balanced. For comparison, 4 to 16 pixels per cycle in 48x48 images is equivalent to 8 to 32 pixels per cycle in the 96x96 images.

Partitioning Combining	SVM	AdaSVM	AdaSVM	AdaSVM
	1:all vote	1:all vote	all poss. vote	all poss. MLR
	89.8	93.1	93.8	93.5

Table 4: Performance of all possible partitions and MLR for AdaSVM’s. Performance is shown for nonlinear SVM’s and AdaSVM’s (with 900 features) for 96x96 images and 9 Gabor wavelengths (2:32).

number of support vectors decreases [13]. The number of support vectors for the linear SVM ranged from 10 to 33 percent of the total number of training vectors. Nonlinear SVM’s employed 14 to 43 percent, despite better generalization performance. Feature selection by Adaboost reduced the number of support vectors employed by the nonlinear SVM in the AdaSVM system, to 12 to 26 percent.

4 Comparison to Linear Discriminant Analysis

A previous successful approach to basic emotion recognition used Linear Discriminant Analysis (LDA) to classify Gabor representations of images [9]. While LDA may be optimal when the class distributions are Gaussian, SVM’s may be more effective when the class distributions are not Gaussian. Table 5 compares LDA with linear SVM’s. The classifiers were tested on 48x48 images using the nine wavelength Gabor representation (2:32 pix/cyc). A small ridge term was used in LDA.

The performance results for LDA were dramatically lower than SVMs. Performance with LDA improved by adjusting the decision threshold for each emotion so as to bal-

ance the number of false detects and false negatives. This approach is labeled LDA- θ in Table 5. This form of threshold adjustment is commonly employed with LDA classifiers, but it uses post-hoc information, whereas the SVM performance was without post-hoc information. Even with the threshold adjustment, the linear SVM performed significantly better.

4.1 Feature selection using PCA

Many approaches to LDA also employ PCA to perform feature selection prior to classification. For each classifier we searched for the number of PCA components which gave maximum LDA performance, which was typically 40 to 70 components. The PCA step resulted in a substantial improvement. The combination of PCA and threshold adjustment gave performance accuracy of 80.7% for the 7-alternative forced choice, which was comparable to other LDA results in the literature [9]. Nevertheless, the linear SVM outperformed LDA even with the combination of PCA and threshold adjustment. SVM performance on the PCA representation was significantly reduced, indicating an incompatibility between PCA and SVM’s for the problem.

4.2 Feature selection using Adaboost

We next examined whether feature selection by Adaboost gave better performance with LDA than feature selection by PCA. Adaboost was used to select 900 features from $9 \times 8 \times 48 \times 48 = 165888$ possible Gabor features which were then classified by LDA (Table 5). Feature selection with Adaboost gave better performance with the LDA classifier than feature selection by PCA. Using Adaboost for feature selection reduced the difference in performance between LDA and SVM’s. Nevertheless, SVM’s continued to outperform LDA.

Feature selection	LDA- θ	SVM (linear)
None	44.4	88.0
PCA	80.7	75.5
Adaboost	88.2	93.3

Table 5: Top Row: Comparing SVM performance to LDA on 48x48 pixel images. The two classifiers are compared with no feature selection, with feature selection by PCA, and feature selection by Adaboost.

4.3 Image alignment

Another difference from previous implementations of LDA for expression recognition was image alignment. Was LDA

more sensitive to alignment noise than SVM’s? Expression recognition performance using the automatically detected face images was compared to performance using images that were aligned using hand-labeling of internal feature points. Six points on each face image were manually located with a mouse (the corners of each eye, the nose tip, and the mouth center). Eye centers were defined as the mean of the eye corners. Images were then rotated in the plane so that the eyes were horizontal and scaled to align the eye centers as well as the midpoint between the mouth and nose tip.

As shown in Table 6, the hand alignment offered no improvement in performance over the automatically aligned face images for either LDA or SVM’s.

	PCA-LDA- θ	SVM	AdaSVM
Face Finder	80.7	88.0	93.3
Hand Aligned	76.8	86.2	91.3

Table 6: Comparison of performance with automatically located faces (top row) and hand aligned faces (lower row).

5. Generalization to other datasets

We tested the system on a second publicly available data set, Pictures of Facial Affect (POFA) [2]. POFA contains 110 images from 14 subjects posing facial expressions. The facial displays were guided by Ekman’s observations of the facial expressions of basic emotion. Previous recognition results reported for this dataset ranged from 85-88%.

Results are shown in Table 7. AdaSVM’s trained and tested on this dataset using leave one subject out cross-validation obtained 97.3% accuracy with a linear kernel, and 95.5% with an RBF kernel. Feature selection by Adaboost had a significant impact on performance for this dataset. SVM’s trained on the full set of Gabors obtained only 79.1% correct. Feature selection may be particularly important for small datasets such as this.

Training and testing on a combined dataset consisting of both DFAT-504 and POFA also gave strong recognition results. Generalization performance was again tested using leave-one-subject-out cross-validation.

Generalization across datasets was substantially lower. A nonlinear AdaSVM trained on DFAT-504 and tested on POFA obtained 60% correct. This highlights the need for large training datasets of facial expressions with variations in image conditions in order to generalize across image collection environments. While the Face Finder was trained on a large number of faces (5000 positive and millions of negative examples) with many lighting conditions and other irregularities, the only condition being roughly frontal pose,

	AdaSVM linear	AdaSVM RBF
POFA	97.3	95.5
DFAT-504+POFA	91.4	93.1
Train: DFAT-504 Test: POFA	56.0	60.0

Table 7: Generalization performance using leave-one-out cross-validation on the POFA dataset alone and on the combined DFAT-504 and POFA datasets. The bottom row gives performance for training on DFAT-504 and testing on POFA. The AdaSVMs were tested for 96x96 images, 9 frequencies, and 953 Adaboost features.

the expression coder was trained on a single dataset with a uniformly controlled environment. The result is that the face finder is robust to real-world application, while the expression coder performs well only within a given dataset or combination of datasets.

6 Real-time expression recognition from video

We combined the face detection and expression recognition into a system that operates on live digital video in real time. Face detection operates at 24 frames/second in 320x240 images on a 3 ghz Pentium IV. The expression recognition step operates in less than 10 msec. Figure 4 shows the output of the expression recognizer for a test video in which the subject posed a series of facial expressions. The traces show outputs of each of the seven emotion detectors. The output of the sadness detector increases as he poses a sad expression, and anger increases as he poses anger. The output for neutral increases as the subject passes through neutral in between.

Although each individual image is separately processed and classified, the outputs change smoothly as a function of time, particularly under illumination and background conditions that are favorable for alignment. (See Figure 5). This enables applications for measuring the magnitude and dynamics of facial expressions.

To demonstrate the potential of this system we developed a real time ‘emotion mirror’ which renders a 3D character in real time that mimics the emotional expression of a person. (See Figure 6). The emotion mirror is a prototype system that recognizes the emotion of the user and responds in an engaging way.

In the emotion mirror, the face-finder captures a face image which is sent to the emotion classifier. For speed, the emotion classifier employed the linear AdaSVM. The outputs of the 7-emotion classifier constitutes a 7-D emotion

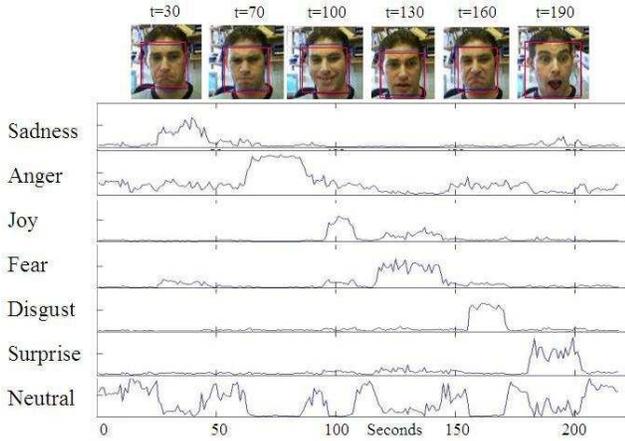


Figure 4: Examples of real-time emotion code traces from a test video sequence. The top row shows frames from the sequence. Continuous outputs of each of the 7 expression detectors is given below.

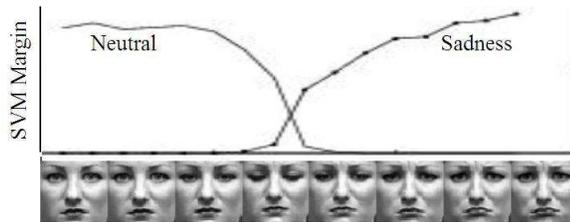


Figure 5: Outputs of the SVM's trained for neutral and sadness for a full test image sequence of a subject performing sadness from the DFAT-504 database. The SVM output is the distance to the separating hyperplane (the margin).

code. This code was sent to CU Animate, a set of software tools for rendering 3D computer animated characters in real time, developed at the Center for Spoken Language Research at CU Boulder. The 7-D emotion code gave a weighted combination of morph targets for each emotion.

We are presently exploring applications of this system including automatic evaluation of human-robot interaction [8], and deployment in automatic tutoring systems [10] and social robots. We are also exploring clinical applications, including psychiatric diagnosis and measuring response to treatment.

7 Conclusions

We presented a systematic comparison of machine learning methods applied to the problem of fully automatic recognition of facial expressions, including AdaBoost, support vector machines, and linear discriminant analysis. We reported results on a series of experiments comparing methods for multiclass decisions, spatial frequency ranges, feature se-

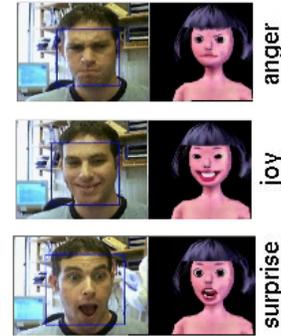


Figure 6: Examples of the emotion mirror. The animated character mirrors the facial expression of the user.

lection methods, and recognition engines. Best results were obtained by selecting a subset of Gabor filters using AdaBoost and then training Support Vector Machines on the outputs of the filters selected by AdaBoost. The combination of Adaboost and SVM's enhanced both speed and accuracy of the system. The generalization performance to new subjects for a 7-way forced choice was 93.3% and 97% correct on two publicly available datasets, the best performance reported so far on these datasets. Surprisingly, registration of internal facial features was not necessary, even though the face detector does not provide precisely registered images. The outputs of the classifier change smoothly as a function of time and thus can be used for unobtrusive motion capture.

Our results suggest that user independent fully automatic real time coding of basic expressions is an achievable goal with present computer power, at least for applications in which frontal views can be assumed. The problem of classification into 7 basic expressions can be solved with high accuracy by a simple linear system, after the images are pre-processed by a bank of Gabor filters.

References

- [1] O. Dekel and Y. Singer. Multiclass learning by probabilistic embedding. In S. Becker and K. Obermayer, editors, *Advances in Neural Information Processing Systems*, volume 15. MIT Press, 2003.
- [2] P. Ekman and W. Friesen. Pictures of facial affect. Photographs, 1976. Available from Human Interaction Laboratory, UCSF, HIL-0984, San Francisco, CA 94143.
- [3] J. Friedman, T. Hastie, and R. Tibshirani. Additive logistic regression: a statistical view of boosting, 1998.
- [4] S. Har-Peled, D. Roth, and D. Zimak. Constraint classification for multiclass classification and ranking. In S. Becker and K. Obermayer, editors, *Advances in Neural Information Processing Systems*, volume 15. MIT Press, 2003.

- [5] T. Kanade, J.F. Cohn, and Y. Tian. Comprehensive database for facial expression analysis. In *Proceedings of the fourth IEEE International conference on automatic face and gesture recognition (FG'00)*, pages 46–53, Grenoble, France, 2000.
- [6] U. Kressel. Pairwise classification and support vector machines. In B. Scholkopf, C. J. C. Burges, and A. J. Smola, editors, *Advances in Kernel Methods: Support Vector Learning*, pages 255–268. MIT Press, Cambridge, MA, 1999.
- [7] M. Lades, J. Vorbrüggen, J. Buhmann, J. Lange, W. Konen, C. von der Malsburg, and R. Würtz. Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions on Computers*, 42(3):300–311, 1993.
- [8] G. Littlewort, M.S. Bartlett, Chenu J, I. Fasel, T. Kanda, H. Ishiguro, and J.R. Movellan. Towards social robots: Automatic evaluation of human-robot interaction by face detection and expression classification. In *Advances in neural information processing systems*, volume 16, Cambridge, MA, in press. MIT Press.
- [9] M. Lyons, J. Budynek, A. Plante, and S. Akamatsu. Classifying facial attributes using a 2-d gabor wavelet representation and discriminant analysis. In *Proceedings of the 4th international conference on automatic face and gesture recognition*, pages 202–207, 2000.
- [10] Jiyong Ma, Jie Yan, Ron Cole, and CU Animate. Cu animate: Tools for enabling conversations with animated characters. In *Proceedings of ICSLP-2002*, Denver, USA, 2002.
- [11] J.C. Platt, N. Cristianini, and J. Shawe-Taylor. Large margin dags for multiclass classification. In S.A. Solla, T.K. Leen, and K.-R. Muller, editors, *advances in neural information processing systems*, volume 12, pages 547–553, Cambridge, MA, 2000. MIT Press.
- [12] Robert E. Schapire. A brief introduction to boosting. In *IJCAI*, pages 1401–1406, 1999.
- [13] V. N. Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, 1995.
- [14] Paul Viola and Michael Jones. Robust real-time object detection. Technical Report CRL 20001/01, Cambridge Research-Laboratory, 2001.