

Feature Dimensionality Reduction for Video Affect Classification: A Comparative Study

Chenfeng Guo

*School of Printing and Packaging
Wuhan University
Wuhan, China*

E-mail: guochenfeng@whu.edu.cn

Dongrui Wu, *Senior Member, IEEE*

*School of Automation
Huazhong University of Science and Technology
Wuhan, China*

E-mail: drwu@hust.edu.cn

Abstract—Affective computing has become a very important research area in human-machine interaction. However, affects are subjective, subtle, and uncertain. So, it is very difficult to obtain a large number of labeled training samples, compared with the number of possible features we could extract. Thus, dimensionality reduction is critical in affective computing. This paper presents our preliminary study on dimensionality reduction for affect classification. Five popular dimensionality reduction approaches are introduced and compared. Experiments on the DEAP dataset showed that no approach can universally outperform others, and performing classification using the raw features directly may not always be a bad choice.

Index Terms—Affective computing, affect recognition, dimensionality reduction, feature extraction, feature selection

I. INTRODUCTION

Affective computing [31] is “*computing that relates to, arises from, or influences emotions.*” It is very important in human-machine interaction, as humans cannot have long-lasting intimate relationships with machines if they cannot understand our affects and respond appropriately.

Both affect classification and regression have been extensively studied in the literature [24], [43], [45], [46], [48]. For affect classification, the most commonly used categories are the six basic emotions (anger, disgust, fear, happiness, sadness, and surprise) proposed by Ekman et al. [5]. For regression, affects are usually represented as numbers in the 2D space of arousal and valence [35], or in the 3D space of arousal, valence, and dominance [25]. Recently, Yannakakis et al. [50] also argued that the nature of emotions is ordinal, and hence preference learning [51] should also play an important role in affective computing.

Various input signals could be used in affective computing, e.g., speech [21], [47], facial expressions [8], [29], physiological signals [7], [43], and multimodal combination [26], [53]. Numerous features could be extracted from each modality. For example, 6,373 acoustic features were extracted by OpenSMILE [6] in the InterSpeech 2013 Computational Paralinguistics Challenge. 465 Riemannian tangent space features were extracted from 30-channel EEG signals in [44]. The number would increase to 2,080 for 64-channel EEG signals, and 8,256 for 128-channel. And, 22,881 features were

extracted from 64-channel EEG signals in [15] for emotion recognition.

On the contrary, affects are very subjective, subtle, and uncertain. So, usually multiple human assessors are needed to obtain the groundtruth affect label for each video, audio, or facial expression. So, generally it is not easy to obtain a large number of labeled training samples in affective computing. As a result, the curse of dimensionality [14] becomes very significant in affect recognition, which implies high computational cost and poor generalization performance. Thus, it is critical to perform dimensionality reduction in affective computing. Though lots of dimensionality reduction approaches have been proposed in the literature [2], [23], to the authors’ knowledge, they have not been extensively studied specifically for affect recognition.

This paper compares five representative dimensionality reduction approaches in video affect classification. It represents our preliminary study of a comprehensive investigation on dimensionality reduction for affective computing.

The remainder of this paper is organized as follows: Section II introduces five representative dimensionality reduction approaches used in our study. Section III describes the DEAP dataset, the raw visual and audio features, and the experimental results. Section IV draws conclusions and points out several future research directions.

II. DIMENSIONALITY REDUCTION

Dimensionality reduction approaches can be categorized into two main classes [23]: feature extraction and feature selection.

Feature extraction projects a high-dimensional feature space to a low-dimensional one, which is usually a linear or nonlinear combination of the original feature space. Typical feature extraction approaches include Principal Component Analysis (PCA) [16], Linear Discriminant Analysis (LDA) [27], Canonical Correlation Analysis (CCA) [13], Singular Value Decomposition [12], Locally Linear Embedding (LLE) [34], etc.

Feature selection directly selects a subset of relevant features to be used in machine learning. According to how feature

selection is integrated with the machine learning model, feature selection approaches can be categorized into three groups [2]:

- 1) *Filter methods*, which select the features independent of the machine learning model. Typical criteria used in filter methods include correlation, mutual information [39], Relief [18], etc.
- 2) *Wrapper methods*, which wrap the machine learning model into a search algorithm to find the optimal feature set that gives the highest learning performance. Typical wrapper methods include sequential forward/backward selection [32], and heuristic search algorithms (e.g., genetic algorithms [11], particle swarm optimization [17], etc) for feature subset selection.
- 3) *Embedded methods*, which include feature selection as part of the machine learning model training process. Typical embedded methods include LASSO [38], 1-norm support vector machines (SVMs) [54], etc.

Five representative dimensionality reduction approaches are introduced in more details next.

A. Principal Component Analysis (PCA)

PCA [16] uses an orthogonal transformation to convert a set of observations into a set of values on the principal components, which are linearly uncorrelated and ordered so that the first few retain most of the variation present in all of the original variables.

Let $\mathbf{X} \in \mathbb{R}^{N \times d}$ be the data matrix with N observations of d dimensions. Assume \mathbf{X} has been pre-processed such that each column has mean zero. Then, its principal components decomposition is:

$$\mathbf{X}' = \mathbf{X}\mathbf{W} \quad (1)$$

where the columns of $\mathbf{W} \in \mathbb{R}^{d \times d}$ are the eigenvectors of $\mathbf{X}^T \mathbf{X}$, sorted in descending order according to the corresponding eigenvalues. \mathbf{X}' can then be used to represent \mathbf{X} in the new space spanned by the columns of \mathbf{W} . Usually the most variation in \mathbf{X} is distributed along the first m columns of \mathbf{W} , and hence only the first m ($m < d$) columns of \mathbf{X}' are enough to represent \mathbf{X} , i.e., the dimensionality can be effectively reduced from d to m .

There can be different approaches to determine m . One is to find the minimum number of eigenvalues of $\mathbf{X}^T \mathbf{X}$ such that their sum exceeds a pre-defined threshold, e.g., 95% of the sum of all eigenvalues. In this paper we used 5-fold cross-validation to find the m that gave the highest classification accuracy on the training data.

B. Sequential Forward Selection (SFS)

Sequential forward selection (SFS) [32] is a very common and intuitive feature selection approach, in which features are sequentially added until the addition of further features does not improve the cross-validation performance. Starting from an empty feature set, SFS creates candidate feature subsets by successively adding each of the features not yet selected. For each candidate feature subset, SFS performs cross-validation to determine the optimal one.

C. ReliefF

The original Relief algorithm was proposed by Kira and Rendell in 1992 [18]. ReliefF [20], [33] is its improved version.

Relief is an iterative procedure that ranks the importance of the features according to how well their values distinguish between their nearest neighbors in different classes. For binary classification, in each iteration Relief first randomly selects a training sample \mathbf{x}_i , and identifies its two nearest neighbors, one from each class. Denote the neighbor from the same class (called *nearest hit*) as \mathbf{h} , and the one from the other class (called *nearest miss*) as \mathbf{m} . Then, Relief updates the feature weight vector $\mathbf{w} = [w_1, \dots, w_d]^T$ as:

$$w_j = w_j - \frac{\text{diff}(\mathbf{x}_{ij}, \mathbf{h}_j)}{M} + \frac{\text{diff}(\mathbf{x}_{ij}, \mathbf{m}_j)}{M}, \quad j = 1, \dots, d \quad (2)$$

where M is the pre-defined number of iterations, and $\text{diff}(\mathbf{x}_{ij}, \mathbf{h}_j)$ is the difference between the j th feature for \mathbf{x}_{ij} and \mathbf{h}_j . When the feature is discrete/categorical,

$$\text{diff}(\mathbf{x}_{ij}, \mathbf{h}_j) = \begin{cases} 1, & \mathbf{x}_{ij} \neq \mathbf{h}_j \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

When the feature is continuous,

$$\text{diff}(\mathbf{x}_{ij}, \mathbf{h}_j) = \mathbf{x}_{ij} - \mathbf{h}_j \quad (4)$$

provided that the continuous feature has been normalized to $[0, 1]$. $\text{diff}(\mathbf{x}_{ij}, \mathbf{m}_j)$ is computed similarly.

In summary, the rationale of (3) is to penalize features that have different values for neighbors from the same class (i.e., features that may lead to wrong classification), and reward features that have different values for neighbors from different classes (i.e., features with good distinguishability).

The ReliefF algorithm [20], [33] improved Relief from the following three perspectives:

- 1) It is more robust, by using k nearest neighbors from each class, instead of only one.
- 2) It can handle multi-class classification instead of binary classification only.
- 3) It can deal with incomplete and noisy data.

In this paper we used $k = 10$ in ReliefF, as suggested in [20].

Because ReliefF ranks the features instead of selecting a subset of them, in this paper we applied PCA to the $d/2$ most important raw features, and then used 5-fold cross-validation on the training data to find the optimal number of PCA features.

D. Minimal-Redundancy-Maximal-Relevance (mRMR)

For discrete/categorical variables X and Y , their mutual information $I(X; Y)$ is defined as:

$$I(X; Y) = \sum_{y \in Y} \sum_{x \in X} p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \quad (5)$$

where $p(x, y)$ is the joint probabilistic distribution, and $p(x)$ and $p(y)$ are the marginal probabilities.

Consider a C -class classification problem, and we want to select a feature subset S with m discrete features $\{x_i\}_{i=1}^m$. The minimal-redundancy-maximal-relevance (mRMR) [4], [30] feature selection approach optimizes the following mutual information difference criterion:

$$\max D - R \quad (6)$$

or the mutual information quotient criterion:

$$\max D/R \quad (7)$$

where

$$D = \frac{1}{|S|} \sum_{x_j \in S} I(x_j; c) \quad (8)$$

is the mean of all mutual information values between individual feature x_j and the class label, and

$$R = \frac{1}{|S|^2} \sum_{x_j, x_r \in S} I(x_j; x_r) \quad (9)$$

is the redundancy among all features in S . In practice incremental search is used to find the near-optimal feature subset S [4], [30].

Note that we only introduce the mRMR approach for discrete features. mRMR approaches for continuous features have also been proposed [4], but it was found that the discrete versions usually work better, so the discrete mRMR is preferred. In this paper we converted each continuous feature x_j into three discrete levels, by thresholds $\bar{x}_j \pm std(x_j)$ (i.e., values larger than $\bar{x}_j + std(x_j)$ were mapped to 1, smaller than $\bar{x}_j - std(x_j)$ to -1, and the rest to 0), where \bar{x}_j is the mean of x_j , and $std(x_j)$ is the standard deviation. The objective function (7) was used in our study, because it gave slightly better performance than (6), as demonstrated in [4]. Finally, the optimal number of features m was determined by 5-fold cross-validation on the training dataset, from the value set of $\{1, \dots, d/2\}$.

E. Neighborhood Component Analysis (NCA)

Neighborhood component analysis (NCA) [49] is a non-parametric feature selection scheme that can be used for both classification and regression. Next we introduce its formulation for classification.

Let $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ be N training samples, where $\mathbf{x}_i \in \mathbb{R}^d$ is the feature vector and $y_i \in \{1, \dots, C\}$ is the corresponding class label. In classification, NCA tries to find a feature weighting vector $\mathbf{w} = [w_1, \dots, w_d]^T$ that maximizes the average leave-one-out cross-validation (LOOCV) accuracy.

The weighted distance between \mathbf{x}_i and \mathbf{x}_l is:

$$d(\mathbf{x}_i, \mathbf{x}_l) = \sum_{j=1}^d w_j^2 |x_{ij} - x_{lj}| \quad (10)$$

Consider a nearest neighbor classifier, which randomly selects the neighbor \mathbf{x}_l for \mathbf{x}_i according to the following probabilities:

$$p_{il} = \begin{cases} \frac{\kappa(d(\mathbf{x}_i, \mathbf{x}_l))}{\sum_{n \neq l} \kappa(d(\mathbf{x}_i, \mathbf{x}_n))}, & i \neq l \\ 0, & i = l \end{cases} \quad (11)$$

where $\kappa(d(\mathbf{x}_i, \mathbf{x}_l)) = \exp(-d(\mathbf{x}_i, \mathbf{x}_l)/\sigma)$ is a kernel function, in which σ is the kernel width. Then, the LOOCV classification accuracy for \mathbf{x}_i is:

$$p_i = \sum_{l \neq i} p_{il} y_{il} \quad (12)$$

where

$$y_{il} = \begin{cases} 1, & y_i = y_l \\ 0, & y_i \neq y_l \end{cases} \quad (13)$$

NCA for classification then maximizes the following regularized objective function:

$$f(\mathbf{w}) = \sum_{i=1}^N p_i - \lambda \sum_{j=1}^d w_j^2 \quad (14)$$

where λ is a regularization parameter. $\sigma = 1$ and $\lambda = 1$ were used in this paper. After finding \mathbf{w} , we sorted w_j in descending order, identified the first a few such that their sum accounts for at least 95% of $\sum_{j=1}^d w_j$, and selected the corresponding features in classification.

III. EXPERIMENTS

This section describes the DEAP dataset used in our study, the raw visual and audio features, and the experimental results

A. The DEAP Dataset

The DEAP dataset [19] was used in our study. It consists of 40 1-minute music video clips, each of which had been evaluated by 14-16 assessors online. Each assessor watched the music videos and rated them on a discrete 9-point scale for valence, arousal, and dominance. Among the 40 videos, 10 had high arousal and high valence, 10 high arousal and low valence, 10 low arousal and high valence, and 10 low arousal and low valence. We would like to classify valence, arousal and dominance independently into two levels (high and low), from the visual, audio, and video signals.

B. Visual Features

The visuals were first converted to MPEG format files at 25 FPS. The 16 features extracted, shown in Table I, consisted of the following valence-related frame-based static features and arousal-related motion features:

- 1) **Static features:** We used the lighting key [52], lightness, and color variance to describe the brightness and color information of frames. The lighting key was defined as the product of the mean and variance of the V-channel in the HSV color space. The lightness was defined as the maximum, minimum, and mean value of the V-channel in the HSV color space. To calculate the color variance, the key-frames were first identified by comparing the histogram distances of two adjacent frames, and then the mean of the determinant of the covariance matrix of the L, U, and V components in the CIELUV color space of the key-frames were computed.
- 2) **Motion features:** Motion features show the changes between frames and the movements of shot, including

the shot change rate, shot length, visual excitement, and motion component. The shot change rate was simply defined as the numbers of key-frames. The shot length features consisted of the longest, shortest, and mean shot lengths. The visual excitement, which measures the degree of video arousal, was calculated from the amount of local pixel changes according to the definition in [40]. The motion components were calculated by accumulating the absolute values of the x and y of the motion vectors, and their sum of squares.

TABLE I
THE 16 VISUAL FEATURES.

Feature category	Number	Value
Lighting key	3	Mean
Lightness	3	Median
Color variance	1	Mean
Shot change rate	1	Mean
Shot length	3	Mean
Visual excitement	2	Mean, variance
Motion component	3	Mean

C. Audio Features

Mono MP3 format audio was first extracted from each video at a sampling rate of 44.1 kHz. Each audio was split into frames and frame-level features were extracted, as follows:

- 1) **Low-level features**, which describe the basic properties of audio in time- and frequency- domains, including the spectral centroid, band energy ratio, delta spectrum magnitude, zero crossing rate, short-time average energy, and pitch. More details about these low-level features can be found in [22].
- 2) **Silence ratio**, which is the ratio of the amount of silence frames to the time window [3]. A frame is considered as a silence frame when its root mean square is less than 50% of the mean root mean square of the fixed-length audio fragments.
- 3) **MFCCs and LPCCs**. In order to combine the static and dynamic characteristics of audio signals, 12 Mel Frequency Cepstral Coefficients (MFCCs), 11 Linear Predictive Cepstral Coefficients (LPCCs), and 12 first-order differential MFCC coefficients were calculated.
- 4) **Formant**, which reflects the resonant frequencies of the vocal tract. Formant frequencies F1-F5 in each frame were extracted.

We then computed the mean and/or variance of these frame-level features, resulting in a total of 76 audio features, as shown in Table II.

D. Experimental Results

We compared the performances of the five dimensionality reduction approaches introduced in Section II, plus a baseline approach (denoted as *Raw*) that does not use any dimensionality reduction, i.e., all the extracted raw features were used. The performance measure was the LOOCV accuracy on the 40 videos.

TABLE II
THE 76 AUDIO FEATURES.

Feature category	Number	Value
Spectral centroid, Band energy ratio, Delta spectrum magnitude, Zero crossing rate, Pitch, Short-time average energy	12	Mean, variance
Silence ratio	1	Mean
MFCC coefficients,	24	Mean, variance
Delta MFCC,	12	Mean
LPCC	22	Mean, variance
Formant	5	Mean

We assume that the features for the 40 videos are all available, the labels for 39 videos are known, and we would like to estimate the label for the remaining video. We first normalized each dimension of the features to [0,1], and then applied different dimensionality reduction approaches. Finally, a radial basis function (RBF) SVM was used as the classifier, where the best SVM parameters were found through 5-fold cross-validation on the training data (39 videos). The final LOOCV classification accuracies are shown in Table III (the highest ones are in bold), and also illustrated in Fig. 1, for different feature sets. Observe that:

- 1) Generally better classification accuracies were obtained from audio than from visual. This may be because more audio features were extracted.
- 2) Interestingly, combining visual and audio features and then performing dimensionality reduction did not necessarily improve the classification performance. In fact, most of the time the performance was actually decreased. These results suggested that the feature selection approaches were not always able to select the global optimal features: otherwise the classification accuracies on the video features would not be lower than those on the visual or audio features alone.
- 3) The highest performances on different affect dimensions and different modalities were achieved by different dimensionality reduction approaches (and sometimes the raw features), and there was not a single approach that was always better than others. This is consistent with the well-known *no free lunch theorems for optimization* [42], which state that “*for any algorithm, any elevated performance over one class of problems is offset by performance over another class.*”
- 4) Surprisingly, on average the raw features achieved the best overall performance in our experiments. This may be because the dimensionality of our features was not high enough (the video features had 76+16=92 dimensions). In the future we will increase the dimensionality of the features, and also take the computational cost into consideration.
- 5) Our preliminary study showed that ReliefF and PCA were two of the better dimensionality reduction approaches among the five. More extensive comparisons

will be performed in the future to verify this.

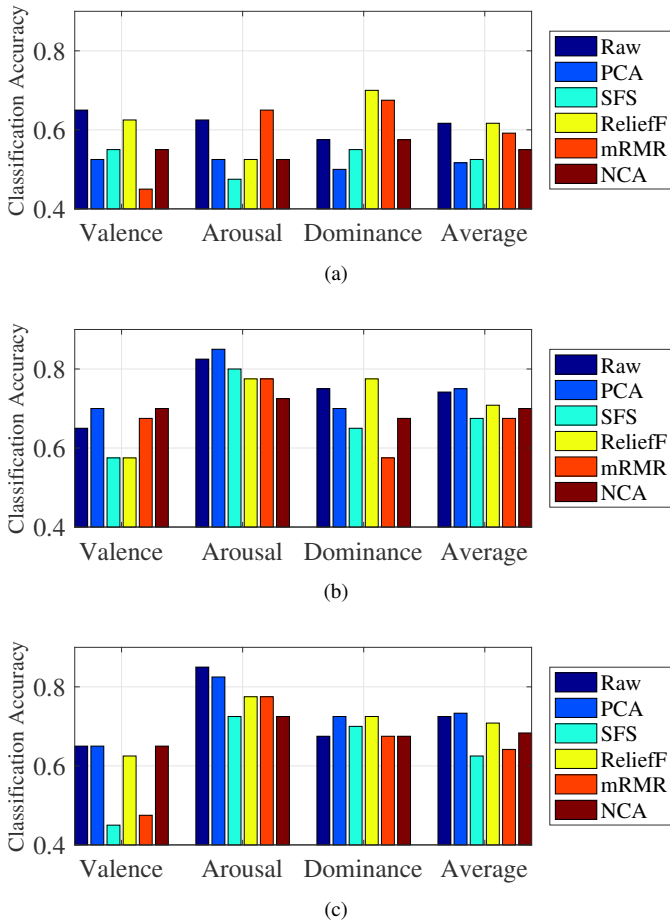


Fig. 1. LOOCV classification accuracies on the DEAP dataset. (a) visual only; (b) audio only; (c) video (visual + audio).

TABLE III
LOOCV CLASSIFICATION ACCURACIES OF DIFFERENT FEATURE SELECTION APPROACHES.

Modality	Affect	Raw	PCA	SFS	ReliefF	mRMR	NCA
Visual	Valence	0.65	0.52	0.55	0.62	0.45	0.55
	Arousal	0.62	0.52	0.47	0.52	0.65	0.52
	Dominance	0.57	0.50	0.55	0.70	0.67	0.57
	Average	0.61	0.51	0.52	0.61	0.59	0.55
Audio	Valence	0.65	0.70	0.57	0.57	0.67	0.70
	Arousal	0.82	0.85	0.80	0.77	0.77	0.72
	Dominance	0.75	0.70	0.65	0.77	0.57	0.67
	Average	0.74	0.75	0.67	0.70	0.67	0.70
Video	Valence	0.65	0.65	0.45	0.62	0.47	0.65
	Arousal	0.85	0.82	0.75	0.77	0.75	0.72
	Dominance	0.67	0.72	0.70	0.72	0.67	0.67
	Average	0.72	0.73	0.62	0.70	0.64	0.68
Overall Average		0.694	0.667	0.608	0.678	0.636	0.644

IV. CONCLUSIONS AND FUTURE WORKS

Affective computing problems typically have a small number of training samples, compared with the number of possible features we could extract. Thus, dimensionality reduction becomes a necessity. This paper reports our preliminary

results on dimensionality reduction for affect classification. Five popular dimensionality reduction approaches have been introduced and compared. Experiments on the DEAP dataset showed that no approach can universally outperform others, and performing classification using the raw features directly (without dimensionality reduction) may sometimes result in even better performance.

Our current study has some limitations, e.g., we only considered one dataset, and the dimensionality of the features was not high enough. We will deal with them in our future research, by considering more affective computing datasets, e.g., MAHNOB-HCI [36], MSP-IMPROV [1], and AMIGOS [28], and by extracting more features, e.g., through OpenSMILE [6]. Additionally, we will:

- 1) Optimize the parameters in the feature selection approaches, e.g., k (the number of nearest neighbors) in ReliefF, the thresholds for discretization in mRMR, and σ (the kernel width) and λ (the regularization parameter) in NCA.
- 2) Investigate multi-view feature selection approaches. Because visual and audio, and sometimes also physiological signals, represent different facets of the same affect, it is more intuitive to perform feature selection in a multi-view setting, instead of combining features from different modalities directly and then performing an overall feature selection. Potential multi-view feature selection approaches include sparse group LASSO [10], adaptive unsupervised multi-view feature selection [9], unsupervised multi-view feature selection [37], and multi-view clustering and feature learning via structured sparsity [41].
- 3) Compare also the computational cost of different dimensionality reduction approaches.
- 4) Study also dimensionality reduction in affect regression [47] and ranking [50].

REFERENCES

- [1] C. Busso, S. Parthasarathy, A. Burmanian, M. AbdelWahab, N. Sadoughi, and E. M. Provost, "MSP-IMPROV: An acted corpus of dyadic interactions to study emotion perception," *IEEE Trans. on Affective Computing*, vol. 8, no. 1, pp. 67–80, 2017.
- [2] G. Chandrashekar and F. Sahin, "A survey on feature selection methods," *Computers and Electrical Engineering*, vol. 40, pp. 16–28, 2014.
- [3] L. Chen, S. Gunduz, and M. T. Ozsu, "Mixed type audio classification with support vector machine," in *Proc. IEEE Int'l Conf. on Multimedia and Expo*, Toronto, ON, Canada, July 2006, pp. 781–784.
- [4] C. Ding and H. Peng, "Minimum redundancy feature selection from microarray gene expression data," in *Proc. IEEE Bioinformatics Conf.*, Stanford, CA, August 2003, pp. 523–528.
- [5] P. Ekman, W. Friesen, M. O'Sullivan, A. Chan, I. Diacoyanni-Tarlatzis, K. Heider, R. Krause, W. LeCompte, T. Pitcairn, and P. Ricci-Bitti, "Universals and cultural differences in the judgments of facial expressions of emotion," *Journal of Personality and Social Psychology*, vol. 53, no. 4, pp. 712–717, 1987.
- [6] F. Eyben, M. Wöllmer, and B. Schuller, "OpenSMILE: the Munich versatile and fast open-source audio feature extractor," in *Proc. 18th ACM Int'l Conf. on Multimedia*. Firenze, Italy: ACM, October 2010, pp. 1459–1462.
- [7] S. H. Fairclough, "Fundamentals of physiological computing," *Interacting with Computers*, vol. 21, pp. 133–145, 2009.
- [8] B. Fasel and J. Luetttin, "Automatic facial expression analysis: A survey," *Pattern Recognition*, vol. 36, no. 1, pp. 259–275, 2003.

- [9] Y. Feng, J. Xiao, Y. Zhuang, and X. Liu, "Adaptive unsupervised multiview feature selection for visual concept recognition," in *Proc. 11th Asian Conf. on Computer Vision*, Daejeon, Korea, Nov. 2012, pp. 343–357.
- [10] J. Friedman, T. Hastie, and R. Tibshirani, "A note on the group lasso and a sparse group lasso," *arXiv preprint arXiv:1001.0736*, 2010.
- [11] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, 1989.
- [12] G. H. Golub and C. F. V. Loan, *Matrix computations*. JHU Press, 2012.
- [13] D. R. Hardoon, S. Szedmak, and J. Shawe-Taylor, "Canonical correlation analysis: An overview with application to learning methods," *Neural computation*, vol. 16, no. 21, p. 26392664, 2004.
- [14] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. Springer, 2009.
- [15] R. Jenke, A. Peer, and M. Buss, "Feature extraction and selection for emotion recognition from EEG," *IEEE Trans. on Affective Computing*, vol. 5, no. 3, pp. 327–339, 2014.
- [16] I. Jolliffe, *Principal component analysis*. Wiley Online Library, 2002.
- [17] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proc. IEEE Int'l. Conf. on Neural Networks*, vol. 4, Perth, Australia, November 1995, pp. 1942–1948.
- [18] K. Kira and L. A. Rendell, "A practical approach to feature selection," in *Proc. Int'l. Conf. on Machine Learning*, Aberdeen, Scotland, UK, Jul. 1992, pp. 249–256.
- [19] S. Koelstra, C. Muhl, M. Soleymani, J. S. Lee, A. Yazdani, T. Ebrahimi, T. Pun, A. Nijholt, and I. Patras, "DEAP: A database for emotion analysis using physiological signals," *IEEE Trans. on Affective Computing*, vol. 3, no. 1, pp. 18–31, 2012.
- [20] I. Kononenko, "Estimating attributes: Analysis and extensions of RELIEF," in *Proc. European Conf. on Machine Learning*, Catania, Italy, Apr. 1994, pp. 171–182.
- [21] C. M. Lee and S. S. Narayanan, "Toward detecting emotions in spoken dialogs," *IEEE Trans. on Speech and Audio Processing*, vol. 13, no. 2, pp. 293–303, 2005.
- [22] D. Li, I. Sethi, N. Dimitrova, and T. McGee, "Classification of general audio data for content-based retrieval," *Pattern Recognition Letters*, vol. 22, no. 5, pp. 533–544, 2001.
- [23] J. Li, K. Cheng, S. Wang, F. Morstatter, R. P. Trevino, J. Tang, and H. Liu, "Feature selection: A data perspective," *CoRR*, vol. abs/1601.07996, 2016. [Online]. Available: <http://arxiv.org/abs/1601.07996>
- [24] H. P. Martinez, Y. Bengio, and G. N. Yannakakis, "Learning deep physiological models of affect," *IEEE Computational Intelligence Magazine*, vol. 8, no. 2, pp. 20–33, 2013.
- [25] A. Mehrabian, *Basic Dimensions for a General Psychological Theory: Implications for Personality, Social, Environmental, and Developmental Studies*. Oelgeschlager, Gunn & Hain, 1980.
- [26] A. Metallinou, S. Lee, and S. S. Narayanan, "Decision level combination of multiple modalities for recognition and analysis of emotional expression," in *Proc. Int'l. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, Dallas, TX, March 2010, pp. 2462–2465.
- [27] S. Mika, G. Ratsch, J. Weston, B. Scholkopf, and K. R. Mullers, "Fisher discriminant analysis with kernels," in *Proc. IEEE Signal Processing Society Workshop on Neural Networks for Signal Processing*, August 1999, pp. 41–48.
- [28] J. A. Miranda-Correa, M. K. Abadi, N. Sebe, and I. Patras, "AMIGOS: A dataset for mood, personality and affect research on individuals and groups," *arXiv preprint arXiv:1702.02510*, 2017.
- [29] M. Pantic and L. Rothkrantz, "Automatic analysis of facial expressions: The state of the art," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 22, no. 12, pp. 1424–1445, 2000.
- [30] H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 27, no. 8, pp. 1226–1238, 2005.
- [31] R. Picard, *Affective Computing*. Cambridge, MA: The MIT Press, 1997.
- [32] P. Pudil, J. Novovicova, and J. Kittler, "Floating search methods in feature selection," *Pattern Recognition Letters*, vol. 15, no. 11, pp. 1119–1125, 1994.
- [33] M. Robnik-Sikonja and I. Kononenko, "Theoretical and empirical analysis of ReliefF and RReliefF," *Machine Learning*, vol. 53, pp. 23–69, 2003.
- [34] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [35] J. Russell, "A circumplex model of affect," *Journal of Personality and Social Psychology*, vol. 39, no. 6, pp. 1161–1178, 1980.
- [36] M. Soleymani, J. Lichtenauer, T. Pun, and M. Pantic, "A multimodal database for affect recognition and implicit tagging," *IEEE Trans. on Affective Computing*, vol. 3, no. 1, pp. 42–55, 2012.
- [37] J. Tang, X. Hu, H. Gao, and H. Liu, "Unsupervised feature selection for multi-view data in social media," in *Proc. SIAM Int'l. Conf. on Data Mining*, Dallas, TX, Dec. 2013, pp. 270–278.
- [38] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society*, vol. 58, no. 1, pp. 267–288, 1996.
- [39] J. R. Vergara and P. A. Estévez, "A review of feature selection methods based on mutual information," *CoRR*, vol. abs/1509.07577, 2015. [Online]. Available: <http://arxiv.org/abs/1509.07577>
- [40] H. L. Wang and L. F. Cheong, "Affective understanding in film," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 16, no. 6, pp. 689–704, 2006.
- [41] H. Wang, F. Nie, and H. Huang, "Multi-view clustering and feature learning via structured sparsity," in *Proc. Int'l. Conf. on Machine Learning*, Atlanta, GA, June 2013, pp. 352–360.
- [42] D. H. Wolpert and W. G. Macready, "No free lunch theorems for optimization," *IEEE Trans. on Evolutionary Computation*, vol. 1, no. 1, pp. 67–82, 1997.
- [43] D. Wu, C. G. Courtney, B. J. Lance, S. S. Narayanan, M. E. Dawson, K. S. Oie, and T. D. Parsons, "Optimal arousal identification and classification for affective computing: Virtual Reality Stroop Task," *IEEE Trans. on Affective Computing*, vol. 1, no. 2, pp. 109–118, 2010.
- [44] D. Wu, V. J. Lawhern, B. J. Lance, S. Gordon, T.-P. Jung, and C.-T. Lin, "EEG-based user reaction time estimation using Riemannian geometry features," *IEEE Trans. on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 11, pp. 2157–2168, 2017.
- [45] D. Wu and T. D. Parsons, "Active class selection for arousal classification," in *Proc. 4th Int'l. Conf. on Affective Computing and Intelligent Interaction*, vol. 2, Memphis, TN, October 2011, pp. 132–141.
- [46] D. Wu and T. D. Parsons, "Inductive transfer learning for handling individual differences in affective computing," in *Proc. 4th Int'l. Conf. on Affective Computing and Intelligent Interaction*, vol. 2, Memphis, TN, October 2011, pp. 142–151.
- [47] D. Wu, T. D. Parsons, E. Mower, and S. S. Narayanan, "Speech emotion estimation in 3D space," in *Proc. IEEE Int'l. Conf. on Multimedia & Expo (ICME)*, Singapore, July 2010, pp. 737–742.
- [48] D. Wu, T. D. Parsons, and S. S. Narayanan, "Acoustic feature analysis in speech emotion primitives estimation," in *Proc. InterSpeech*, Makuhari, Japan, September 2010.
- [49] W. Yang, K. Wang, and W. Zuo, "Neighborhood component feature selection for high-dimensional data," *Journal of Computers*, vol. 7, no. 1, pp. 161–168, 2012.
- [50] G. N. Yannakakis, R. Cowie, and C. Busso, "The ordinal nature of emotions," in *Proc. Int'l. Conf. on Affective Computing and Intelligent Interaction*, San Antonio, TX, Oct. 2017.
- [51] G. Yannakakis, M. Maragoudakis, and J. Hallam, "Preference learning for cognitive modeling: A case study on entertainment preferences," *IEEE Trans. on Systems, Man and Cybernetics-A*, vol. 39, no. 6, pp. 1165–1175, 2009.
- [52] A. Yazdani, E. Skodras, and N. Fakotakis, "Multimedia content analysis for emotional characterization of music video clips," *EURASIP Journal on Image and Video Processing*, vol. 2013, no. 1, p. 26, 2013.
- [53] Z. Zeng, M. Pantic, G. I. Roisman, and T. S. Huang, "A survey of affect recognition methods: Audio, visual, and spontaneous expressions," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 31, no. 1, pp. 39–58, 2009.
- [54] J. Zhu, S. Rosset, R. Tibshirani, and T. J. Hastie, "1-norm support vector machines," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, Vancouver, BC, Dec. 2004, pp. 49–56.