# Energy-Aware Embedded Classifier Design for Real-Time Emotion Analysis*

Manoj Padmanabhan[1], Srinivasan Murali[2], Francisco Rincón[2] and David Atienza[1]

*Abstract*— Detection and classification of human emotions from multiple bio-signals has a wide variety of applications. Though electronic devices are available in the market today that acquire multiple body signals, the classification of human emotions in real-time, adapted to the tight energy budgets of wearable embedded systems is a big challenge. In this paper we present an embedded classifier for real-time emotion classification. We propose a system that operates at different energy budgeted modes, depending on the available energy, where each mode is constrained by an operating energy bound. The classifier has an offline training phase where feature selection is performed for each operating mode, with an energy-budget aware algorithm that we propose. Across the different operating modes, the classification accuracy ranges from 95% - 75% and 89% - 70% for arousal and valence respectively. The accuracy is traded off for less power consumption, which results in an increased battery life of up to 7.7 times (from 146.1 to 1126.9 hours).

## I. INTRODUCTION

Emotion classification holds key significance in the building of affective computing [1] systems. The motivation for emotion classification arises from its manifold applications ranging from human-computer interaction(HCI) to determining the emotional well-being of individuals. Recognition of stress levels in automobile drivers [2], continuous monitoring for epileptic seizures [3], occupational therapy for Autism and ADHD [4] are a few other examples. Numerous studies report the effects of emotions on the autonomic nervous system (ANS) and the specificity of ANS activity over distinct emotions [5]. ANS responses can be observed in physiological signals, thereby opening up the possibility of assessing emotions by physiological signal analysis. Relevant information is extracted from multiple bio-signals such as electrocardiogram (EKG or ECG), respiration and galvanic skin response (GSR), and multi-sensor data fusion is performed for classification. The MIT 8-emotion dataset is an early work in this direction [6], and another dataset from Augsburg University has physiological data from 4 emotions [7]. Usually external stimuli such as music [7] or videos [8] are employed in such studies to elicit the appropriate emotions. The results presented in most of the above works show classification accuracies ranging from 70% to 90%, suggesting that bio-signals can be used for emotion classification.

The two-dimensional valence-arousal model [9] for emotions is widely used by psychologists. Classifying emotions in the arousal and valence directions has been explored in many works [10], [7]. Emotions are classified according to attractiveness (positive) or averseness (negative) in the horizontal valence dimension. The emotion-induced excitement level is represented by high to low positions in the vertical arousal dimension of the valence-arousal space.

The selection of a suitable feature set for classification is an important problem. Classical methods have focused mainly on selecting a feature subset that yields a high classification accuracy. Meanwhile, choosing relevant features by also taking into account the energy spent in feature computation is still a challenge. A wide range of wearable devices are available today which can process bio-signals, such as the Jawbone UP3 activity tracker [11] and the Smartcardia INYU [12]. For such embedded systems with multiple physiological sensors and signal processing units, the acquisition of each signal and computation of every feature has an associated energy expenditure. Even though the classification accuracy is important, choosing relevant features by also taking into account the energy costs is crucial for an increased battery life. The desired target accuracy depends on the application (e.g., an epileptic seizure predicting device should have high accuracy, while an emotion classifier for an immersive game can limit its accuracy if energy savings are required). Our target of having high accuracy and low energy expenditure are inherently conflicting and therefore trade-offs need to be established. Feature selection by co-optimizing accuracy and energy cost is a major challenge.

In this paper, we present a real-time embedded emotion classifier operating under energy budget constraints for next-generation wearable sensors. The main contributions of this paper are the following:

- We propose an embedded emotion classification system that includes different operating modes, and the system transits between them to enable different trade-offs between emotion analysis and available energy. We have developed an energy-aware greedy algorithm for feature selection that is evaluated with real-life traces. Based on selected features, a classifier is trained for each mode.
- We measure the signal processing energy expenditure by implementing the proposed emotion classifier and feature calculations on a wearable embedded platform (the INYU device [12]) and present estimates of battery lifetime for the system executing this classifier.

[1]M. Padmanabhan and D. Atienza are with Embedded Systems Lab. (ESL), EPFL, Switzerland {manoj.padmanabhan, david.atienza}@epfl.ch

[2]S. Murali and F. Rincón are with SmartCardia Inc., Ltd, Switzerland {Srinivasan.murali, Francisco.rincon}@smartcardia.com
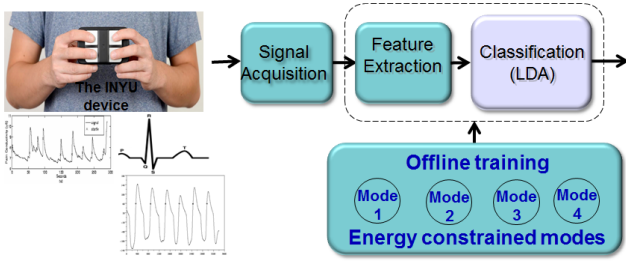
Fig. 1. Mode-based emotion classifier using multi-sensor data fusion

To the best of our knowledge, this is the first work in the context of emotion classification where the energy cost of features from a large feature set is measured on an embedded device. We consider this work to be a first step towards energy cost aware multi-sensor data fusion in target embedded platforms for emotion classification.

The paper is organized as follows: Section II describes the mode-based design of the embedded emotion classifier. The offline training algorithm used for energy cost aware feature selection is explained in Section III. The experimental setup and results are discussed in Section IV, and finally Section V provides the main conclusions of this workd.

## II. DESIGN METHODOLOGY

We propose a system as shown in Fig. 1. The embedded emotion classifier operates in one of the modes at a time. Each mode is defined by certain energy budget constraints. The first phase of the design - a relatively simple step of finding out which mode to operate on, is based on the energy budget value. In the second phase, for every operating mode energy budget, an offline algorithm selects the appropriate features such that the overall energy consumption will be less than the budget. The classifier is trained to operate in the different modes, and while operating in real-time it can transit from one mode to another.

### A. Energy-Aware Operating Modes Design

Modes of operation are defined for the embedded classifier, each governed by an energy budget. The number of operating modes depends on design convenience, and the requirements of the target application. For illustration, we discuss the case of having four operating modes and perform experiments for this case with the INYU device. We assume that the energy budget information is always known to the embedded classifier, so that it can operate in a corresponding mode. When there is surplus power, the classifier can be made to operate in one of the higher energy modes, and vice-versa when the battery is running low on energy. Based on the operating mode, the classification performance is traded off for lower energy consumption.

The highest energy mode is Mode 1, which is not constrained by any energy budget limit. The classification accuracy is expected to be the highest here, since there is a liberty to choose all features that provide the highest accuracy. Mode 2 is an energy budgeted mode designed such that from the considered initial feature set, the energy-expensive features are not chosen. The energy cost considered is the overall cost, which includes the pre-processing cost (eg. signal acquisition, low pass filtering, R-R peak detection

for ECG etc.) and the individual feature computation cost. The pre-processing cost is a one-time cost, i.e. if a feature from a particular signal is chosen, this cost appears once and does not recur for another feature from the same signal. The feature computation cost is unique, depending solely on the computational complexity of the feature. For example, the computational cost of a mean is significantly less than that of a median or an FFT.

Mode 3 and Mode 4 consume less energy because they use less features, Mode 4 being the most energy conservative. Once in an energy constrained mode, depending on the selected features, only the corresponding signal processing circuitry are activated. Thus, in the low energy modes, energy savings occurs from thrifty feature selection, signal acquisition and processing. As the freedom in choosing features is expected to reduce due to energy constraints, trade-offs are explored for classification accuracy.

### B. Training Phase

The training for the embedded classifier and the parameter selection for each mode is determined by an offline algorithm. The energy costs for computing various features and the overall energy budget for the operating modes are given to the algorithm. The algorithm selects a feature set and trains an LDA (linear discriminant) model for valence and arousal classification. The feature set and the trained LDA model for the various modes, are programmed onto the embedded platform for classification. The energy-budget aware offline training algorithm is detailed in Section III.

## III. ENERGY AWARE FEATURE SELECTION

In this section, we present the algorithm for feature selection under energy budget constraints. The proposed method is a greedy algorithm that co-optimizes classification accuracy and energy expenditure.

### A. Mathematical Notation

A vector $\mathbf{g}$ is a row vector $\mathbf{g} = [g_1, g_2, \ldots, g_n]$, where the $k^{\text{th}}$ element of $\mathbf{g}$ is written as $g_k$. Upper case bold letters denote matrices, and the dimension is written as a subscript, as in $\mathbf{M}_{m \times n}$. The notation $\mathbf{M}[:, \mathbf{f}]$ denotes the number of columns of $\mathbf{M}$, where $\mathbf{f}$ is a vector of column indices. The value of $\mathbf{g}$ at iteration $i$ is written as $\mathbf{g}^{(i)}$. An all-zeros vector is denoted by $\mathbf{0}$. All non-bold alphabets denote scalars.

### B. Algorithm Description

We define a single objective function of two parameters - the classification accuracy and an energy based reward. The objective function $\phi$ is calculated as follows:

$$\phi = \alpha A + (1 - \alpha) E \quad (1)$$

where $\alpha$ is an adaptive weight, $A$ is the classification accuracy obtained by augmenting a new feature to the set of selected features and $E$ is the fraction of energy. The available energy fraction $E$ is calculated as follows:

$$E = \frac{(E_{tot} - (\sum \mathbf{e}[\mathbf{f}]) + e_k)}{E_{tot}} \quad (2)$$

where $\mathbf{f}$ is the index set of features already selected, $e_k$ is the energy of the feature under consideration and $E_{\text{tot}}$ is the energy budget of the operating mode.

The objective function balances the parameters $A$ and $E$, and the value of the variable $\alpha$ determines the prominence of one parameter over the other. The accuracy $A$ and the energy reward $E$ are desired to be high, and therefore we propose to select the feature that maximizes the objective function $\phi$. A greedy search is performed in each iteration to select the features that maximize the objective function. The algorithm is outlined in Algorithm 1.

---

**Algorithm 1:** Energy aware feature selection

**Input**: Feature matrix $\mathbf{M}_{m \times n}$ ; Accuracy threshold $T$;
Feature energy cost vector $\mathbf{e}_{1 \times n}$ ; Energy budget $E_{\text{tot}}$;
**Output**: Set of selected feature indices $\mathbf{f}$;
Total energy expenditure of selected features $E_{\text{eff}}$

1   $\mathbf{f} \leftarrow \emptyset$ ; $\mathbf{g}^{(1)} \leftarrow \{1, 2, \ldots, n\}$ ; $\alpha^{(1)} \leftarrow 1$ ; $A^{(0)} \leftarrow 0$ ;
2   **while do**
3     $\phi^{(i)} \leftarrow \mathbf{0}_{1 \times \text{length}(\mathbf{g}^{(i)})}$;
4     **for** $k = 1$ *to* $length(\mathbf{g}^{(i)})$ **do**
5       **if** $\left( \mathbf{e}[g_k^{(i)}] + \sum \mathbf{e}[\mathbf{f}] \right) < E_{tot}$ **then**
6         $p \leftarrow \text{LDA\_LOSO}(\mathbf{M}[:, \mathbf{f}], \mathbf{M}[:, g_k^{(i)}])$;
7         **if** $p \geq A^{(i-1)}$ **then**
8           Calculate $\phi_k^{(i)}$ ; $A^{(i)} \leftarrow p$;
9         **else**
10           $\phi_k^{(i)} \leftarrow 0$;
11         **end**
12       **end**
13     **end**
14     **if** $\phi^{(i)} = \mathbf{0}$ **then**
15       $E_{\text{eff}} \leftarrow \sum \mathbf{e}[\mathbf{f}]$ ; Exit;
16     **end**
17     Augment index set $\mathbf{f}$ with $\mathbf{g}^{(i)}[\arg\max_k \phi^{(i)}]$;
18     Remove element $\mathbf{g}^{(i)}[\arg\max_k \phi^{(i)}]$ ;
19     Update $\mathbf{e}$, pre-processing costs counted once;
20     **if** $A^{(i)} \geq T$ **then**
21       $\alpha^{(i+1)} \leftarrow \alpha^{(i)} - c$;
22     **end**
23   **end**

---

The feature matrix $\mathbf{M}_{m \times n}$ contains the data from all the $n$ features over $m$ recordings. The energy costs for the $n$ features are put together into a vector $\mathbf{e}$ of dimension $1 \times n$. An energy budget upper bound, based on the mode of operation, is denoted by $E_{\text{tot}}$. These parameters, along with an accuracy threshold $T$ are the inputs to the algorithm. The algorithm outputs an index set of selected features $\mathbf{f}$, the total energy expenditure for the selected feature set, $E_{\text{eff}}$ and the parameters for an LDA model trained for separating the classes using the selected feature set. The algorithm starts with $\mathbf{f}$ being an empty set and $\mathbf{g}$ having all feature indices. The initial value of $\alpha$ is set to 1 and the initial classification accuracy value is zero.

Every complete iteration of the loop (step 2) is designed to select a feature and augment the feature index to the target set $\mathbf{f}$. To this end, an inner loop (step 4) runs through the unselected features' set $\mathbf{g}$. A feature is considered only when the total energy cost remains below the energy budget (step 5). For such a feature, the classification accuracy using an LDA model and a leave-one-subject-out (LOSO) cross validation scheme is computed. The data from the feature matrix $\mathbf{M}$ corresponding to the already selected feature indices $\mathbf{f}$ and the current feature are fed to an LDA_LOSO model to compute the accuracy (step 6). As long as the addition of the feature of interest does not bring about a decrease in classification accuracy, the algorithm proceeds to calculate the objective function as per (1) (steps 7 to 9). The elements of the objective function vector corresponding to the feature indices that are not considered, remain zero.

If none of the features pass the search criteria, the algorithm exits. When this is not the case, the feature set $\mathbf{f}$ is augmented with the element of $\mathbf{g}$ corresponding to the maximum value of the objective function. Since the input feature energy vector also contains the associated signal pre-processing costs, this has to be removed accordingly after a one-time accounting. For example, the cost of acquisition and processing (low-pass filtering and baseline removal) of GSR should be taken into account only once when a GSR feature is selected. For further GSR features, this cost does not appear and only the feature calculation cost is accounted. Therefore $\mathbf{e}$ is updated in step 19 to make sure that the pre-processing costs are accounted only once. Steps 20 through 22 define the procedure to decrease the variable $\alpha$. Only when the classification accuracy crosses a minimum threshold, the decrease of $\alpha$ is initiated. Till then, $\alpha$ remains 1, meaning that features are selected keeping solely accuracy in mind. Once this minimum threshold $T$ is reached, $\alpha$ can be decreased at every iteration. There are many methods for iterative parameter decrease that are employed in classical heuristic algorithms, namely, linear, quadratic, exponential etc. For our analysis, we choose to decrease $\alpha$ linearly (step 21). The decrease constant $c$ is heuristically set as 10% of the initial starting value. The algorithm is scalable, with a worst case time complexity of $O(n^2)$, where $n$ is the total number of considered features.

## IV. EXPERIMENTAL VALIDATION

### A. Experimental Setup

Our target embedded platform is the INYU device [12], which is capable of acquiring ECG, respiration and GSR signals. The INYU device measures ECG at a configurable frequency (250 Hz to 1 KHz) and 16-bit resolution, using the top two metal pads. Real-time respiration is measured by impedance pneumography - by injecting a minute current through the metal pads. The change in impedance due to chest movement and oxygen/blood change is captured using proprietary circuitry, resulting in real-time estimate of breathing rate and volume. The GSR is calculated by real-time voltage difference measurement between two contact metal pads. The INYU is equipped with the STM32L151RDT6, an
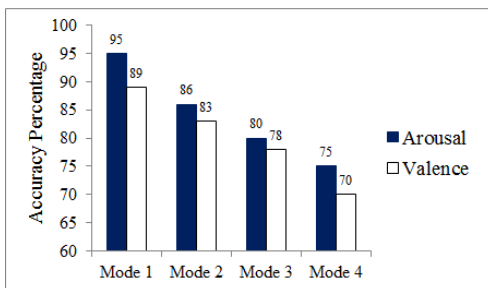
Fig. 2.  Classification accuracies in the different operating modes

| Mode | Average current (mA) | Lifetime (h) |
|------|---------------------|--------------|
| 1 | 4.86 | 146.1 |
| 2 | 1.69 | 420.1 |
| 3 | 1.09 | 651.3 |
| 4 | 0.63 | 1126.9 |

TABLE I

<small>CURRENT CONSUMPTION AND BATTERY LIFETIME FOR EACH MODE</small>

ultra-low-power 32-bit microcontroller which can operate at a maximum frequency of 32 MHz, has 48 KB RAM, 384 KB Flash and analog peripherals including a 12-bit ADC.

We consider the following sampling rates: 256 Hz for ECG, 32 Hz for respiration and GSR. The energy cost of feature computation is calculated by measuring the computation time taken by the CPU. For the active mode of the STM32, the current consumption is 10.5 mA at an operating voltage of 3.3 V, with the system clock running at a frequency of 32 MHz. The CPU energy consumption is calculated as follows:

$$e_k = (I_{\text{active}}t_1 + I_{\text{sleep}}t_2)(3.3V) \tag{3}$$

where $t_1$, $t_2$ are the durations, and $I_{\text{active}}$, $I_{\text{sleep}}$ are the currents of active and sleep modes respectively. The appropriate signal acquisition cost and other pre-processing costs (low-pass filtering, baseline removal, etc.) are added to $e_k$. The analog front ends for ECG and respiration have the current consumption of 136 $\mu$A and 970 $\mu$A respectively. The GSR measurement consumes 1.15 mA.

During real-time operation, when the CPU is inactive, the microcontroller can go to a sleep mode or to a low-power sleep mode. A transition to the sleep mode or to the low-power sleep mode is dependent on the peripherals that are required to remain active. In sleep mode, the microcontroller current consumption is 2.1 mA and in the low power sleep mode, the current consumption is $18\mu$A. The energy source for the device is a battery of 710 mAh rating, at 3.7V. The battery life in each case is calculated based on the average current consumption for signal acquisition and processing.

*B. Experimental Results*

For the training algorithm, we consider the ECG, respiration and GSR data from the Augsburg dataset [7]. The dataset has recordings of four emotions (anger, joy, sadness and pleasure) over 20 days, thereby having 80 recordings ($m = 80$) in total. The 4-emotions are labeled in the 4 respective valence-arousal quadrants for training, as in [7].

The initial feature set we consider has 45 features ($n = 45$) in total. These include statistical features like mean, median, standard deviation etc., statistical difference features and frequency domain features, of the three bio-signals. From the diversity of the features it can be noticed that their computational costs cover a wide range. The time spent in calculating each feature is computed by counting the number of independent timer cycles running in parallel. The average current consumed in signal acquisition and feature calculation of selected features is presented for each operating mode. Based on the current consumption, we predict the battery life of the device for each of the modes. The results are tabulated in Table I. The classification accuracy obtained for each of the modes is shown in Fig. 2. In Mode 1, spectral and time domain features from all signals are present, while in Mode 2 features from GSR are avoided. Mode 3 includes features from respiration. In Mode 4, only time domain features of ECG (heart rate and SDNN) are extracted, as spectral features consume too much energy to be calculated.

## V. CONCLUSIONS

This work presented the design of an embedded emotion classifier based on energy-budgeted operating modes. The motivation for and the idea of a mode-based design was explained, and an offline training algorithm for energy budgeted feature selection was proposed. The energy-aware feature selection capability of the algorithm was demonstrated with the Augsburg emotion dataset. The accuracy trade-offs incurred in such energy-constrained scenarios show a range of 95% - 75% and 89% - 70% for arousal and valence classification accuracy respectively, depending on the operating mode. The calculated battery lifetime estimates for the different modes range from 146.1 to 1126.9 hours, thereby illustrating the significance of energy budgeted feature selection for embedded classifiers. Further analysis of the proposed algorithm with additional data from human subjects in a controlled setup is intended for future work.

## REFERENCES

[1] R. Picard, "Affective computing," *MIT Technical Report*, 1995.

[2] J. Healey and R. Picard, "Detecting stress during real-world driving tasks using physiological sensors," *IEEE Trans. ITS*, vol. 6, no. 2, pp. 156 – 166, 2005.

[3] M.-Z. Poh *et al.*, "Continuous monitoring of electrodermal activity during epileptic seizures using a wearable sensor," in *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, Sept. 2010, pp. 4415–4418.

[4] E. Hedman *et al.*, "Measuring autonomic arousal during therapy," in *Proceedings of 8th International Design and Emotion Conference*, Sept. 2012.

[5] S. Kreibig, "Autonomic nervous system activity in emotion: A review," *Biological Psychology*, vol. 84, no. 3, pp. 394 – 421, July 2010.

[6] J. A. Healey, "Wearable and automotive systems for affect recognition from physiology," *PhD thesis, MIT, Cambridge, MA*, May 2000.

[7] J. Wagner *et al.*, "From physiological signals to emotions: Implementing and comparing selected methods for feature extraction and classification," in *IEEE Intl. Conf. Multimedia and Expo (ICME)*, 2005, pp. 940 – 943.

[8] S. Koelstra *et al.*, "Deap: A database for emotion analysis ;using physiological signals," *IEEE Transactions on Affective Computing*, vol. 3, no. 1, pp. 18 – 31, June 2011.

[9] J. A. Russell, "A circumplex model of affect," *Journal of Personality and Social Psychology*, vol. 39, no. 6, pp. 1161 – 1178, 1980.

[10] G. Valenza *et al.*, "Revealing real-time emotional responses: a personalized assessment based on heartbeatdynamics," *Nature Scientific Reports*, vol. 4, no. 4998, May 2014.

[11] "JAWBONE." [Online]. Available: https://jawbone.com/store/buy/up3

[12] "SMARTCARDIA." [Online]. Available: http://www.smartcardia.com