

Descriptive Models of Emotion: Learning Useful Abstractions from Physiological Responses during Affective Interactions

Rui Henriques¹ and Ana Paiva²

¹*KDBIO, Inesc-ID, Instituto Superior Técnico, University of Lisbon, Lisbon, Portugal*

²*GAIPS, Inesc-ID, Instituto Superior Técnico, University of Lisbon, Lisbon, Portugal*
{*rmch,ana.s.paiva*}@*tecnico.ulisboa.pt*

Keywords: Descriptive Models, Mining Physiological Signals, Measuring Affective Interactions

Abstract: Supervised recognition of emotions from physiological signals has been widely accomplished to measure affective interactions. Less attention is, however, placed upon learning descriptive models to characterize physiological responses. In this work we delve on why and how to learn discriminative, complete and usable descriptive models based on physiological signals from emotion-evocative stimuli. By satisfying these three properties, we guarantee that the target descriptors can be expressively adopted to understand the physiological behavior underlying multiple emotions. In particular, we explain why classification and unsupervised learning models do not address these properties, and point new directions on how to adapt existing learners to met them based on theoretical and empirical evidence.

1 Introduction

Monitoring the physiological responses to multiple emotion-evocative stimuli has been widely accomplished in order to understand and recognize emotions. The use of physiological signals to measure, describe and affect human-robot interactions is critical since they track subtle affective changes that are hard to perceive, and are neither prone to social masking nor have the heightened context-sensitivity of image, audio and survey-based analysis. A large stream of literature has been dedicated either to study how to learn classification models to recognize emotions from labeled signals (Jerritta et al., 2011; Wagner et al., 2005) or to convey emotion-dependent physiological aspects from scientific experiments (Cacioppo et al., 2007; Andreassi, 2007). However, there is a clear research gap on how to learn descriptive models from labeled signals. In fact, learning descriptive models is increasingly relevant to dynamically derive informative and usable abstractions from experiments, to monitor sensor-based data, and to take the study of emotion-evocative stimuli up to a wider range of affective states.

This paper makes two major contributions to the field. First, we see why the existing models learned from physiological responses are poorly descriptive.

Second, we go further on how to adapt them in order to guarantee that they are flexible, discriminative, complete and usable. We focus our contributions along three major sets of descriptors: feature-based learners, generative sequential learners, and pattern-centric learners. For each set, we discuss and present critical strategies for the learning of robust descriptive models from emotion-centered physiological data.

This paper is structured as follows. *Section 2* defines the major requirements that guarantee the utility of the target descriptors, and covers major contributions and limitations from existing work. *Section 3* relies on theoretical and empirical evidence to propose three major types of descriptive models that are able to address the surveyed limitations of existing models. An integrative view on the complementarity of the proposals is also discussed. Finally, the major implications of our work are synthesized.

2 Background

Physiological responses are increasingly measured to derive accurate analysis from affective interactions. Although there are numerous principles on how to recognize affective states from (streaming) signals, less attention is being paid to the task

of characterizing affective states. This task is referred as emotion description from physiological responses. Although we use throughout this work the term emotion, our contributions are generalizable to non-affective responses to stimuli, such as subject responses to non-empathic robot motion and speech.

Illustrative applications of emotion description include: measuring human interaction with *artificial agents*, assisting *clinical research* (emotion-centered understanding of addiction, affect dysregulation, alcoholism, anxiety, autism, attention deficit, depression, drug reaction, epilepsy, menopause, locked-in syndrome, pain management, phobias and desensitization therapy, psychiatric counseling, schizophrenia, sleep disorders, and sociopathy), studying the effect of body posture and exercises in *well-being*, disclosing responses to *marketing* and suggestive interfaces, reducing *conflict* in schools and prisons through the early detection of hampering behavior, fostering *education* by relying on emotion-centered feedback to escalate behavior and increase motivation, development of (pedagogic) *games*, and *self-awareness* enhancement.

Def.1 Consider a set of annotated signals $D=(x_1, \dots, x_m)$, where each instance is a tuple $x_i=(\vec{y}, a_1, \dots, a_n, c)$ where \vec{y} is a (multivariate) signal, a_i is an annotation related with the subject or experimental setting, and c is the labeled emotion, stimulus, task or environmental condition. Given D , *emotion description* task aims to learn a model M that characterizes the discriminative properties of \vec{y} for each emotion c in a complete and usable way.

This definition implies that the learning of descriptive models of emotions from labeled signals should satisfy four major *requirements*:

- *flexibility*: descriptive models cope with the complex and variable physiological expression of emotions within and among individuals;
- *discriminative power*: descriptive models capture and enhance the different physiological responses among emotions. Discriminatory ability can be seen at the level of a single emotion or at the level of a group of emotions from the target set (e.g. isolation of emotions with positive valence);
- *completeness*: descriptive models contain all of the discriminative properties and, when the reconstitution of the signal behavior is relevant, of flexible sequential abstractions;
- *usability*: descriptive models are compact and the abstractions of physiological responses are easily interpretable.

Multiple physiological modalities have been adopted to monitor emotions, including electrodermal activity to identify engagement and excitement states, respiratory volume and rate to recognize negative-valenced emotions, and heat contractile activity to separate positive-valenced emotions (Wu et al., 2011; Henriques et al., 2012; Cacioppo et al., 2007). Other modalities with studied emotion-driven behavior include multiple forms of brain activity, cardiovascular activity and muscular activity (Cacioppo et al., 2007). Descriptive models can either capture one or multiple physiological modalities under Def.1 by composing multivariate signals from the (either univariate or multivariate) physiological signals from each modality. When this is the case, we assume that proper dedicated pre-processing techniques are applied over each modality, such as smoothing, low-pass filtering and neutralization of cyclic behavior for respiratory and cardiac signals (Lessard, 2006).

Emotion description can be applied for experiments with different types of stimuli (discrete vs. continuous, high-agreement vs. self-report) and multiplicity of users (user-dependent vs. user-independent studies). Recovering Def.1, subjective stimuli commonly recurs to the optional a_j annotations to infer c , and the x_i instances can either be obtained for one or for multiple subjects. Additionally, not only discrete models of an high-agreement set emotions can be target (Ekman and W., 1988), but also more flexible models, such as recent work focused on recognizing states that are a complex mix of emotions (“the state of finding annoying usability problems”) (Picard, 2003). Additionally, a_i annotations can be combined to capture dimensional valence-arousal axes (Lang, 1995), Weiner’s attributions, and Ellsworth’s dimensions and agency (Oatley et al., 2006).

Emotion description as it is defined has been seen as an optional byproduct of emotion recognition from physiological signals. In particular, when measuring affective interactions with humans, solid contributions on recognizing emotions have been provided for interactions in social contexts (Wagner et al., 2005), with robots (Kulic and Croft, 2007; Leite et al., 2013), with computer interfaces (Picard et al., 2001), and with multi-modal adaptive virtual scenarios (Rani et al., 2006). Despite the large attention dedicated to classification and unsupervised learning models, they are not able to answer the previously introduced four requirements as they were developed for a different goal. Below we describe and enumerate the major limitations associated with the two major groups of classification models to perform emotion recognition (feature-based and generative models) and with unsupervised models.

Limitations of Feature-based Learning Models.

Feature-based models are a function of features extracted from the observed signals for each class (Povinelli et al., 2004; Nanopoulos et al., 2001). They can either be deterministic or probabilistic and are the typical choice for emotion recognition. Some of the most common learning functions include: random forests, k-nearest neighbors, Bayesian networks, support vector machines, canonical correlation analysis, neural networks, linear discriminant analysis, and Marquardt-back propagation (Maaoui et al., 2008; Jerritta et al., 2011; Mitsa, 2009). These extracted features are statistical (mean, deviation), temporal (rise and recovery time), or more complex (multi-scale sample entropy, sub-band spectra) (Haag et al., 2004). Methods for their extraction include rectangular tonic-phasic and moving-sliding windows; transformations (Fourier, wavelet, Hilbert, singular-spectrum); component analysis; projection pursuit; auto-associative networks; multidimensional scaling; and self-organizing maps (Haag et al., 2004; Lessard, 2006; Jerritta et al., 2011). The features that might not have significant correlation with the emotion under assessment are removed using sequential forward/backward/floating selection, branch-and-bound search, principal component analysis, Fisher projection, Davies-Bouldin index, and analysis of variance methods (Jerritta et al., 2011; Bos, 2006). This improves the simplicity and the discriminative power of feature-based models.

Let us see why feature-based models as-is are not prone for emotion description. First, features hardly capture *flexible* behavior (e.g. motifs underlying complex rising and decaying responses) and are strongly dependent on directive thresholds (e.g. peak amplitude to compute frequency-based features). Additionally, the multiplicity of physiological expression per emotion is hardly modeled since the majority of methods rely on prototype features for each emotion. Second, although a wide range of complementary features can be extracted to provide a *complete* description of each emotion, they cannot be adopted to retrieve abstractions of sequential behavior of the signal. Third, *discriminative power* is highly dependent on the chosen methods and mainly consider at emotion's individual level. Discriminative power is achieved through: *i*) feature selection, *ii*) weighting of features, *iii*) entropy-based functions, among others. Note, however, that these strategies are only prepared to assess the ability to differentiate individual emotions. If a particular feature is able to differentiate two sets of emotions but does not isolate a single emotion, it is commonly discarded. Finally, *usability* is highly dependent on the chosen classifiers.

Limitations of Generative Learning Models.

Generally, generative sequential models learn lattices from (multivariate) physiological signals. Although they are the common option for speech and video recognition, only recently became more prominently adopted for emotion recognition (Henriques et al., 2013; Kulic and Croft, 2007; Henriques et al., 2012). Common sequential models include dynamic Bayesian networks, such as hidden Markov models (HMMs) (Murphy, 2002), time-sensitive neural networks (NNs), such as time-sensitive NNs (Guimarães, 2000) or time-delay NNs (Berthold and Hand, 1999), (temporal) support vector machines (SVMs) (Burges, 1998), and logistic regressions. In particular, we use HMMs as the illustrative model due to their maturity, expressive power, inherent simplicity and flexible parameter-control (Rabiner and Juang, 2003). Lattices are commonly defined by the underlying automaton (characterized by transition and value emission probabilities) according to a specific architecture.

Similarly to feature-based models, these models have properties that deteriorate their ability to be used as descriptive models. First, although the use of large interconnected lattices can capture multiplicity of physiological expression due to the large number of paths, it becomes hard to abstract such multiplicity of expressions per emotion from the analysis of the lattices. This hampers their *flexibility* as a descriptive model. Second, although the most probable sequences characterizing the physiological can be retrieved from the lattices, specific behavior can be lost during the learning process (e.g. frequency features), deteriorating the *completeness* of the learned models. Third, commonly one lattice is learned independently for each emotion (classification is then performed by evaluating the generative probability of a new signal on each one of the learned lattices), and, thus, the models do not accentuate *discriminative* behavior. In fact, the differences among lattices can be very subtle as they are typically observed for a small subset of transition and emission probabilities. Finally, although generative models offer a compact view of physiological responses per emotion, they tend to be highly complex and, therefore, hardly *usable*. This is particularly problematic if there is not a clear convergence for a specific subsets of transitions and emissions.

Limitations of Unsupervised Learning Models.

The properties of unsupervised learners, such as (bi)clustering models (Madeira and Oliveira, 2004) and collections of temporal patterns (Mörchen, 2006; Han et al., 2007), deserve a closer attention in the context of emotion description from physiological re-

sponses. Although we refer to these models as unsupervised, we assume that the target local regularities are learned in the context of a specific emotion. For instance, biclusters can disclose strong correlations between a specific subset of features that are only observed for a subset of signals. Similarly, temporal pattern, such as sequential patterns or motifs extracted directly from the physiological signal, can be discovered and used to enrich the target descriptive models.

These local models are flexible, as they can isolate multiple responses per emotion either within and across individuals. However, they do not meet the remaining desirable descriptor's properties. First, similarly to generative models, unsupervised local models typically are not prepared to discover *discriminative* patterns or biclusters. Second, mining local regularities do not guarantee *completeness* since unsupervised methods are not exhaustive, i.e., they easily fail to consider specific features or sequential aspects of the signal of interest. Finally, although listing local regularities promotes simplicity, it is necessary to guarantee that this set is compact and navigable. In the absence of an organized structure for the presentation of these regularities, these models are hardly *usable*.

3 Descriptive Models

In the previous section, we explored the limitations of relying on widely-adopted classification models and on unsupervised models to perform emotion description from physiological data. Although these models show multiple properties of interest that can be seized within descriptive models, they fail to satisfy the introduced four requirements. In this section, we propose a set of principles for each class of models in order to guarantee their compliance with these requirements.

3.1 Feature-based Classification Models

To guarantee that feature-based models are *flexible*, it is either necessary to choose a classification model that is able to group distinctive physiological responses per emotion. This need is derived from the observation that single emotion-evocative stimulus can elicit small-to-large groups of significantly different physiological responses. Lazy learners and random forests implicitly are able to deal with this aspect. However, many of remaining classification models are not prepared to deal with such flexible paths of expression per class. In order to adapt them,

it is crucial to understand how to refine these models. This can be done by analyzing the variances of features per emotion or by clustering responses per emotion with a non-fixed number of clusters as proposed in (Henriques and Paiva, 2014). For these cases, new labels associated with the major physiological responses per emotion are considered during the learning process, so the learned models properly capture these ramifications of expression per emotion.

To guarantee the *completeness* of feature-based models two strategies are suggested. First option is to combine these models with the output of sequential generative models as they provide different but complementary views (Henriques et al., 2013). However, this option does not solve the fact that disclosing prototype values for the features from each emotion can be misleading. Exemplifying, mean or median values are inconclusive when variance is high. Nevertheless, outputting simultaneously the feature and a deviation metric can still be misleading as often the observed values can hardly be approximated by Gaussian distributions (Lessard, 2006). Therefore, we propose a second option, that is to rely on an approximated distribution for the most significant and discriminative sets of features. At least, Gaussian and Poisson tests should be considered for a more correct interpretation of how a particular feature characterize a particular emotion.

To guarantee the *discriminative power* of the selected features, not only feature selection should be considered, but also feature weighting methods should be adopted in order to rank the features according to their ability to separate emotions (Liu and Motoda, 1998). However, the majority of these methods rely on metrics, such as entropy ratios, that are only prepared to deal with differentiation at an individual level. Thus, we advise the adoption of feature weighting methods that are additionally able to separate sets of emotions (e.g. isolate intense emotions or separate positive from negative valence), such as some methods for the analysis of the variance of statistical models (Surendiran and Vadivel, 2011).

Finally, substantial research has been dedicated to guarantee the *usability* of feature-based models, with particular incidence on its understandability, reproducibility and ability to retrieve embedded knowledge. Rule-based models, as decision trees, and Bayesian networks are illustrative usable models as is. In order to foster the usability of more complex models, such as support vector machines in high-dimensional spaces and neural networks, visualization techniques and rule-extraction methods have been largely proposed (Mörchen, 2006).

3.2 Generative Classification Models

For generative models, the physiological responses for a specific emotion are characterized by the learned lattice. However, lattices are commonly complex, and, therefore, providing them as-is as descriptors is not a good option. Specific states can have a large number of equi-probable emissions and there is an heightened complexity associated with the identification of the most probable paths (Murphy, 2002). Therefore, we propose two strategies in order to guarantee a more delineated convergence and the retrieval of the most probable prototype patterns. The combined application of these strategies guarantee the *flexibility* of the target generative models.

The first strategy relies on an adapted learning setting that guarantees acceptable levels of convergence in combination with an algorithmic search that is able to decode the most probable sequences from the lattices. In the context of HMMs, Viterbi is a learning setting that guarantees acceptable levels of convergence for both the transition and emission probabilities (Murphy, 2002). Other more restrictive learning settings can be adopted either recurring to minimum entropy prior (Brand, 1999) or mixtures of Dirichlets (Brown et al., 1993). The goal here is to guarantee that a only a compact set of symbols can be emitted per state. Retrieving the most probable sequential responses per emotion is simply a matter of defining efficient methods to explore the most probable transitions and emissions (Henriques and Antunes, 2014). For this goal, the specification of the minimum probability threshold for the path and for individual emissions and transitions are the adopted criteria to decode the abstractions of the signal for a particular emotion. This step is illustrated in Fig.1.




Figure 1: Decoding of most probable signal responses for an emotion-centered fully-interconnected Markov model

The second strategy enhances flexibility through the use of more expressive architectures for the lattices. The commonly adopted fully-interconnected lattices are complex. Additionally, the more restrictive left-to-right architectures are only prepared to abstract a major physiological response per emotion (Murphy, 2002), degrading the flexibility of the

model. In a left-to-right architecture, lattices have a set of main states from which abstracted sequential responses can be derived. To tackle the problems of these commonly adopted architectures, we propose the use of multi-path architectures (Henriques and Antunes, 2014). A multi-path architecture is the sound parallel composition of left-to-right architectures. The number of paths most exceed the expected number of distinct physiological responses per emotion. When such knowledge is not available, the number of paths can be increased until a specific convergence criterion is satisfied. An illustrative multi-path architecture is presented in Fig.2. Under this architecture, the retrieval of sequential prototype behavior is just a matter of retrieving the most probable sequences from the set of main states from each one of these paths.




Figure 2: Multi-path architecture: composition of left-to-right architectures to model distinct responses per emotion

To promote *completeness*, the most probable physiological responses can be synthesized according to regular expressions. Additionally, two further options can be adopted. Hierarchical disclosure of less probable behaviors can be done by discovering probable sequences in lattices under more relaxed thresholds. Second, these regular expressions can be integrated with the complementary feature-based classification models.

Since each lattice is learned independently for each emotion, differences among lattices can be very subtle, which deteriorates the ability to retrieve discriminative probable sequences. In order to guarantee the *discriminative power* of generative models, two strategies can be adopted. First, after learning the lattices, a method based on graph differences can be applied in order to accentuate the points of divergence among a set of lattices by adapting their transition and emission probabilities (Murphy, 2002). Second, methods for the mining of discriminative frequent sequences can be adopted in order to retrieve sequential behavior that is highly specific for one or few of the overall lattices (emotions) (Han et al., 2007).

The *usability* of generative models can be easily achieved recurring to two major strategies. First,

by listing the most probable sequential behavior for each emotion by decoding the lattices. Similarly to feature-based models, weights or data structures can be used to prioritize and organize the most probable sequential behaviors per emotion, promoting usability. Second, multi-path architectures can be visually displayed due to the highly constrained transition paths. In this setting, only the main states per path need to be displayed, and only the most probable emissions per states should be included.

3.3 Unsupervised Local Models

Feature-based and generative classification models can be extended and combined for the delivery of adequate descriptive models. In this section, we see how unsupervised models can be adopted to enhance the properties of the target descriptive models. In particular, we focus our discussion on two major types of local models: (bi)clustering models and temporal pattern mining models.

Biclusters can be either used over feature-based datasets to find subsets of instances with strong correlation among a subset of features (Madeira and Oliveira, 2004) or over signal-based datasets to find subsets of instances correlated with local properties of the signal (Madeira et al., 2010). In particular, extensions are available in literature to allow for scaling and shifting biclusters for both options. Although one bicluster can be seen as a simple feature, it is of interest to capture it within descriptive models not as a simple boolean variable (meaning that they can either be or not be discovered for a specific emotion) but to further disclose its properties. And, therefore, it is important to guarantee that this disclosure is compliant with the target requirements. Models of biclusters are flexible and, by nature, not complete. To guarantee the discriminative power of these models, meaning that only discriminative biclusters are discovered, several strategies have been proposed (Wang et al., 2010). These same strategies can be used to compose models where biclusters are ranked by discriminatory relevance, which fosters their usability.

Multiple temporal patterns have been proposed, such as sequential patterns, calendric rules, temporal association rules, motifs, episodes, containers and partially-ordered tones (Mueen et al., 2009; Mörchen, 2006). Each one holds different properties that have been largely considered to be of interest for the analysis of signal-based data (Geurts, 2001). For instance, in (Leite et al., 2013), sequential patterns have been retrieved to visually display the differences of multiple affective states. In fact, and similarly to biclusters, temporal patterns can be captured as a boolean feature

by feature extraction methods. However, in order to zoom in their characteristics for emotion description, more detailed models can be targeted. In particular, when discriminative temporal patterns are considered (Exarchos et al., 2008; Tseng and Lee, 2009), these models are commonly a rank of temporal patterns according to their confidence in relation to a particular emotion (percentage of instances supporting this pattern for a particular emotion from the overall dataset) (Li et al., 2001). Again, although completeness cannot be achieved (unless feature-based or generative models are also present), the remaining properties can be satisfied recurring to principles from methods prone to find structured models of compact discriminative patterns (Tseng and Lee, 2009).

4 Conclusion

In this work we introduce the task of emotion description from physiological signals and motivate its relevance for measuring affective interactions. Four structural properties for the definition of useful descriptive models are synthesized. They are: flexibility, discriminative power, completeness and usability. The limitations from considering emotion description as a byproduct of emotion recognition are covered. In particular, we explore why feature-based and generative classification models are not able to satisfy the introduced properties.

In order to overcome the problems associated with existing models, this work proposes multiple strategies to extend existing models in order to guarantee the delivery of robust descriptive models. This set of strategies, derived from theoretical and empirical evidence, is the central contribution of our work. In particular, we show how to improve the flexibility, discriminative power, completeness and usability of supervised models (either deterministic or generative) and unsupervised models recurring to state-of-the-art methods from the fields of data mining, signal processing and multivariate analysis.

Promoting the quality of descriptive physiological models of emotions opens a new door for the psychophysiological study and real-time monitoring of affective interactions. Therefore, in a context where the use of non-intrusive wearable sensors is rapidly increasing, this paper establishes solid foundations for upcoming contributions on this novel and critical field of research.

Acknowledgment

This work was supported by Fundação para a Ciência e a Tecnologia under the project PEst-OE/EEI/LA0021/2013 and PhD grant SFRH/BD/75924/2011, and by the project EMOTE from the EU 7th Framework Program (FP7/2007-2013) under grant agreement n.317923.

REFERENCES

- Andreassi, J. (2007). *Psychophysiology: Human Behavior And Physiological Response*. Lawrence Erlbaum.
- Berthold, M. and Hand, D. J., editors (1999). *Intelligent Data Analysis: An Introduction*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1st edition.
- Bos, D. O. (2006). Eeg-based emotion recognition the influence of visual and auditory stimuli. *Emotion*, 57(7):1798–806.
- Brand, M. (1999). Structure learning in conditional probability models via an entropic prior and parameter extinction. *Neural Comput.*, 11(5):1155–1182.
- Brown, M., Hughey, R., Krogh, A., Mian, I. S., Sjölander, K., and Haussler, D. (1993). Using dirichlet mixture priors to derive hidden markov models for protein families. In *Int. Sys. for Molecular Bio.*, pages 47–55. AAAI Press.
- Burges, C. J. C. (1998). A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2:121–167.
- Cacioppo, J., Tassinary, L., and Berntson, G. (2007). *Handbook of psychophysiology*. Cambridge University Press.
- Ekman, P. and W. F. (1988). Universals and cultural differences in the judgments of facial expressions of emotion. *Personality and Social Psychology*, 53:712–717.
- Exarchos, T., Tspouras, M., Papaloukas, C., and Fotiadis, D. (2008). A two-stage methodology for sequence classification based on sequential pattern mining and optimization. *Data Knowl. Eng.*, 66(3):467–487.
- Geurts, P. (2001). Pattern extraction for time series classification. In *PKDD*, pages 115–127, London, UK. Springer-Verlag.
- Guimarães, G. (2000). The induction of temporal grammatical rules from multivariate time series. In *Colloquium on Grammatical Inference: Alg. and App.*, pages 127–140, London, UK. Springer-Verlag.
- Haag, A., Goronzy, S., Schaich, P., and Williams, J. (2004). Emotion recognition using bio-sensors: First steps towards an automatic system. In *Affective D.Sys.*, volume 3068 of *LNCIS*, pages 36–48. Springer.
- Han, J., Cheng, H., Xin, D., and Yan, X. (2007). Frequent pattern mining: current status and future directions. *Data Min. Knowl. Discov.*, 15(1):55–86.
- Henriques, R. and Antunes, C. (2014). Learning predictive models from integrated healthcare data: Capturing temporal and cross-attribute dependencies. In *HICSS*. IEEE.
- Henriques, R. and Paiva, A. (2014). Seven principles to mine flexible behavior from physiological signals for effective emotion recognition and description in affective interactions. In *PhyCS*.
- Henriques, R., Paiva, A., and Antunes, C. (2012). On the need of new methods to mine electrodermal activity in emotion-centered studies. In *AAMAS'12, 8th IW on ADMI*. Springer-Verlag LNAI series.
- Henriques, R., Paiva, A., and Antunes, C. (2013). Accessing emotion patterns from affective interactions using electrodermal activity. In *Affective Comp. and Intel. Interaction (ACII)*. IEEE Computer Society.
- Jerritta, S., Murugappan, M., Nagarajan, R., and Wan, K. (2011). Physiological signals based human emotion recognition: a review. In *CSPA, 2011 IEEE 7th International Colloquium on*, pages 410–415.
- Kulic, D. and Croft, E. A. (2007). Affective state estimation for human-robot interaction. *Trans. Rob.*, 23(5):991–1000.
- Lang, P. (1995). The emotion probe: Studies of motivation and attention. *American psychologist*, 50:372–372.
- Leite, I., Henriques, R., Martinho, C., and Paiva, A. (2013). Sensors in the wild: Exploring electrodermal activity in child-robot interaction. In *HRI*, pages 41–48. ACM/IEEE.
- Lessard, C. S. (2006). *Signal Processing of Random Physiological Signals*. S.Lectures on Biomedical Eng. Morgan and Claypool Publishers.
- Li, W., Han, J., and Pei, J. (2001). Cmar: Accurate and efficient classification based on multiple class-association rules. In *ICDM*, pages 369–376. IEEE CS.
- Liu, H. and Motoda, H. (1998). *Feature Selection for Knowledge Discovery and Data Mining*. Kluwer Academic Pub.
- Maaoui, C., Pruski, A., and Abdat, F. (2008). Emotion recognition for human-machine communication. In *IROS*, pages 1210–1215. IEEE/RSJ.
- Madeira, S., Teixeira, M. N. P. C., Sá-Correia, I., and Oliveira, A. (2010). Identification of regulatory modules in time series gene expression data using a linear time biclustering algorithm. *IEEE/ACM TCBB*, 1:153–165.
- Madeira, S. C. and Oliveira, A. L. (2004). Biclustering algorithms for biological data analysis: A survey. *IEEE/ACM Trans. Comput. Biol. Bioinformatics*, 1(1):24–45.
- Mitsa, T. (2009). *Temporal Data Mining*. DMKD. Chapman & Hall/CRC.
- Mörchen, F. (2006). *Time series knowledge mining*. Wissenschaft in Dissertationen. Görlich & Weiershäuser.
- Mueen, A., Keogh, E. J., Zhu, Q., Cash, S., and Westover, M. B. (2009). Exact discovery of time series motifs. In *SDM*, pages 473–484.
- Murphy, K. (2002). *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, UC Berkeley, CS Division.
- Nanopoulos, A., Alcock, R., and Manolopoulos, Y. (2001). Information processing and technology. chapter Feature-based classification of time-series data, pages 49–61. Nova Science Pub.

- Oatley, K., Keltner, and Jenkins (2006). *Understanding Emotions*. Blackwell P.
- Picard, R. W. (2003). Affective computing: challenges. *International Journal of Human-Computer Studies*, 59(1-2):55–64.
- Picard, R. W., Vyzas, E., and Healey, J. (2001). Toward machine emotional intelligence: Analysis of affective physiological state. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(10):1175–1191.
- Povinelli, R. J., Johnson, M. T., Lindgren, A. C., and Ye, J. (2004). Time series classification using gaussian mixture models of reconstructed phase spaces. *IEEE Trans. on Knowl. and Data Eng.*, 16(6):779–783.
- Rabiner, L. and Juang, B. (2003). An introduction to hidden Markov models. *ASSP Magazine*, 3(1):4–16.
- Rani, P., Liu, C., Sarkar, N., and Vanman, E. (2006). An empirical study of machine learning techniques for affect recognition in human-robot interaction. *Pattern Anal. Appl.*, 9(1):58–69.
- Surendiran, B. and Vadivel, A. (2011). Feature selection using stepwise anova discriminant analysis for mammogram mass classification. *IJ on Signal Image Proc.*, 2(1):4.
- Tseng, V. and Lee, C.-H. (2009). Effective temporal data classification by integrating sequential pattern mining and probabilistic induction. *Expert Sys.App.*, 36(5):9524–9532.
- Wagner, J., Kim, J., and Andre, E. (2005). From physiological signals to emotions: Implementing and comparing selected methods for feature extraction and classification. In *ICME*, pages 940–943. IEEE.
- Wang, M., Shang, X., Zhang, S., and Li, Z. (2010). Fdcluster: Mining frequent closed discriminative bicluster without candidate maintenance in multiple microarray datasets. In *ICDM Workshops*, pages 779–786. IEEE CS.
- Wu, C.-K., Chung, P.-C., and Wang, C.-J. (2011). Extracting coherent emotion elicited segments from physiological signals. In *WACI*, pages 1–6. IEEE.