Assessing Feature Importance in Cardiovascular Variability for the Classification of Physiological Stress

Marta Iovino
Department of Engineering, University of Palermo, Palermo, Italy
marta.iovino@unipa.it

Marlis Ontivero-Ortega
Department of Physics, University of Bari Aldo Moro and Istituto Nazionale di Fisica Nucleare, Bari, Italy
marlis.ontiveroortega@uniba.it

Chiara Barà
Department of Engineering, University of Palermo, Palermo, Italy
chiara.bara@unipa.it

Sebastiano Stramaglia
Department of Physics, University of Bari Aldo Moro and Istituto Nazionale di Fisica Nucleare, Bari, Italy
sebastiano.stramaglia@ba.infn.it

Michal Javorka
Department of Physiology, Jessenius Faculty of Medicine, Comenius University, Martin, Slovakia
michal.javorka@uniba.sk

Riccardo Pernice
Department of Engineering, University of Palermo, Palermo, Italy
riccardo.pernice@unipa.it

Luca Faes
Department of Engineering, University of Palermo, Palermo, Italy
luca.faes@unipa.it

***Abstract*—Explainable artificial intelligence improves the interpretability of machine learning models, which is essential for reliable decision-making in healthcare. Integrating information-theoretic (IT) approaches into feature selection (FS) facilitates a more rational assessment of feature importance (FI) by capturing redundant and synergistic contributions. This study proposes a novel combined FI-FS method that extends the High-order Interactions Feature Importance (Hi-Fi) framework by replacing variance-based with IT metrics for improved quantification of high-order feature interactions. The method adapts the Leave One Covariate Out metric to identify feature subsets that maximise or minimise Conditional Mutual Information (CMI), prioritising features that increase synergy and reduce redundancy. Feature interpretation is further supported by analysing how selected variables interact, revealing both individual and joint predictive relevance. The proposed framework is applied to cardiovascular time series from 127 young, healthy individuals recorded at rest and under stress conditions (postural and mental stress). The analysed features are extracted from beat-to-beat electrocardiographic RR intervals, pulse-pulse intervals (PP), and systolic and diastolic blood pressure (SBP, DBP) time series and are computed in the time, frequency, and information domains. FI results show that features reflecting variability, spectral content, and entropy, especially from RR, DBP, and PP, are more informative than static measures. RR features consistently show the highest unique information, while PP contributes mainly through synergy. The FS process reduces feature count by ∼20%, and a Support Vector Machine trained on the selected features achieves ∼80% accuracy in multi-class stress classification. Overall, this IT-based Hi-Fi framework captures individual informativeness and complex signal interactions, improving dependency detection, interpretability, and physiological insight.**

***Keywords*—Feature Importance, Feature Selection, Information Theory, Cardiovascular Time Series, Stress Classification**

I. Introduction

Explainable Artificial Intelligence (XAI) has become a cornerstone of modern machine learning (ML), especially in domains such as healthcare, where interpretability is essential to ensure transparency, trust, and accountability in automated decision-making [1]. In contrast to black-box models, XAI methods aim to reveal the internal logic of algorithms by highlighting the contribution of input features to predictions. Feature Importance (FI) and Feature Selection (FS) play a key role in this context, as they allow for the identification of the most informative variables, thereby improving model interpretability, reducing overfitting, and enhancing computational efficiency. While FS focuses on selecting a subset of relevant features for model training, FI metrics provide a quantitative measure of the contribution of each feature to the prediction task, enabling a deeper understanding of the model's decision-making process [2].

Traditional FI and FS methods often fall short in capturing complex interactions among features and in distinguishing between redundant and complementary information. Information-theoretic (IT) approaches address these limitations by quantifying the dependencies among variables through entropy and mutual information [3], [4], offering a principled framework for analysing feature relevance [5].

Alongside the growing interest in high-order interactions [6]–[8], there has been increasing attention devoted to the emergent properties of complex systems. These phenomena manifest through high-order behaviours in observed data,

Research supported by PRIN 2022 project "HONEST" (funded by MUR, code 2022YMHNPY, CUP B53D23003020006 and by European Union-NextGenerationEU - funds from Italian MUR, D.M. 737/2021 - EUROSTART 2021 research project "Novel Computational tools for Patient Stratification in cardiovascular diseases and brain disorders".

979-8-3315-2037-3/25/$31.00 ©2025 IEEE

which can be detected by moving beyond traditional dyadic descriptions that consider only pairwise dependencies. A key framework in this literature is the Partial Information Decomposition [9] and its subsequent developments, which utilise information-theoretic tools to reveal high-order dependencies among groups of random variables and to describe their synergistic and redundant nature. Redundancy refers to information that is retrievable from multiple sources. In contrast, synergy refers to statistical relationships that emerge only at the group level and cannot be captured by observing the variables in isolation [8], [10]. These tools enable the decomposition of feature contributions into unique, redundant, and synergistic components, supporting a more nuanced interpretation of model behaviour.

Recently, a method for quantifying high-order effects in feature importance (Hi-Fi) has been proposed, based on regression models and predictability decomposition [10]. This approach leverages the Leave One Covariate Out (LOCO) strategy, which evaluates the importance of each feature by measuring the drop in model performance, typically in terms of variance or prediction error, when that feature is omitted. A notable implementation of this strategy applies LOCO to a linear regression model, assessing the contribution of individual features to the explained variance [11]. Although linear regression can capture some non-linear interactions, such as quadratic effects, the standard model estimated via ordinary least squares may not be well-suited for multi-class classification tasks. Moreover, this approach is inherently model-dependent and assumes linear additive relationships, which may limit its generalizability in scenarios involving more complex, non-linear feature interactions.

In recent years, these IT approaches have been increasingly applied to the classification of physio-pathologic conditions using physiological signals. Various strategies have been proposed to classify different types of stress (e.g., cognitive, postural) based on cardiovascular dynamics. Stress is defined as the organism's physiological and psychological response to disruptive stimuli [12], and it affects the autonomic nervous system (ANS), altering heart rate dynamics and vascular tone [13]. Heart rate variability (HRV), derived from electrocardiographic (ECG) signals, is a widely accepted marker of autonomic function and is commonly used for stress assessment [14]. An alternative signal, the photoplethysmogram (PPG), allows for the estimation of pulse rate variability (PRV), offering a wearable-compatible proxy to HRV, though with known limitations due to vascular and motion-related artefacts [15]. In addition to HRV and PRV indices, blood pressure measures such as systolic (SBP) and diastolic (DBP) blood pressure provide complementary information on cardiovascular reactivity and autonomic control. These parameters reflect vascular resistance and arterial compliance, and their modulation under stress conditions can reveal distinct patterns associated with sympathetic activation and baroreflex sensitivity.

Herein, we introduce a novel approach that integrates FI and FS through a Conditional Mutual Information (CMI)-based framework based on the Hi-Fi methodology to quantify the information shared between a target variable, the class, and a candidate feature, conditioned on a subset of already selected features. While this approach assumes a Gaussian distribution of the features, it remains model-agnostic and well-suited to multi-class classification problems. By leveraging CMI, the method allows for a principled decomposition of feature contributions into unique, redundant, and synergistic components, yielding a more accurate and interpretable estimate of each feature's importance. This CMI-based FI and FS strategy overcomes the limitations of traditional regression-based explainability tools, as it does not rely on restrictive model assumptions and captures complex, high-order interactions between variables. The proposed framework also unifies FI and FS into a single process that prioritises features not only for their individual relevance, but also for their synergistic value in the presence of others. This enables the detection of subtle yet physiologically meaningful dependencies among cardiovascular indices. Herein, the framework is applied to stress classification using features extracted from cardiovascular signals, demonstrating that the selected subset improves both classification accuracy and physiological interpretability.

II. Information-Theoretic Decomposition of Feature Importance

Let us consider a set of n stochastic variables $Z = \{z_\alpha\}_{\alpha=1,\dots,n}$, a driver variable X, and a target variable Y. In this study, we aim to quantify the contribution of the variable X to the prediction of Y, by leveraging an information-theoretic framework based on CMI. Specifically, the relevance of X is assessed through the CMI $I(Y;X|Z)$, which measures the amount of information that X provides about Y, given the conditioning set Z. This approach aligns with the interpretation of feature importance in terms of predictability, similar to the LOCO measure [11], but is framed within an entropy-based perspective. The CMI is interpreted as the reduction in uncertainty about Y when X is observed, beyond what is already explained by Z, and it has a direct analogue with feature importance in regression-based settings. Within a linear-Gaussian framework, we assess the relevance of X to the prediction of Y employing the CMI, defined as [16]:

$$I(Y;X \mid Z) = I(Y;X,Z) - I(Y;Z). \tag{1}$$

In this setting, both MI terms are computed in closed form by assuming joint Gaussianity over features and class. To disentangle the distinct contributions of X to the prediction of Y, we define the unique information U as:

$$U = \min_{z \subseteq Z} I(X;Y \mid z), \tag{2}$$

which quantifies the non-redundant predictive contribution of X that cannot be attributed to any subset of Z. Conversely, the redundant information R and the synergistic information S are defined as:

$$R = I(X;Y) - \min_{z \subseteq Z} I(X;Y \mid z), \tag{3}$$

$$S = \max_{z \subseteq Z} I(X;Y \mid z) - I(X;Y), \tag{4}$$

yielding the decomposition of the maximal predictive power as:

$$\max_{z} I(X;Y \mid z) = U + R + S. \tag{5}$$

Due to the combinatorial complexity of exhaustively searching all 2^n subsets of Z, the Hi-Fi framework employs a greedy approximation strategy. Starting from an empty set, variables are sequentially added to maximise or minimise the $I(X;Y \mid z)$. At each step, the statistical relevance of the selected variable is assessed through a permutation-based significance test, comparing the observed gain against a null distribution obtained by randomly permuting the candidate variable [10]. A variable is added to the subset only if the resulting increase or decrease in CMI is statistically significant, using a corrected p-value according to Bonferroni correction for multiple comparisons. Subsequently, FS was based on a composite importance score defined as $I = U + S - R > 0$, prioritising features that provide both unique and cooperative information while penalising redundancy.

III. Application to Stress Classification

In the context of stress classification, the method presented in Sect. II is applied to cardiovascular indices to identify the most informative features.

A. Subjects and experimental protocol

Data were analysed from 127 young, healthy volunteers (75 females and 52 males, average age: 18.63 ± 3.27 years), who were all normotensive and maintained a normal body mass index (BMI: 21.42 ± 2.20 kg/m²). The study was conducted in accordance with the approval granted by the Ethics Committee of the Jessenius Faculty of Medicine at Comenius University, Martin, Slovakia. Each participant provided signed informed consent. For individuals under 18, consent was obtained from a parent or legal guardian to authorise their participation in the study. The physiological signals examined included horizontal bipolar thoracic lead ECG recordings and arterial blood pressure measured using the volume-clamp method at a sampling rate of 1 kHz. The experimental protocol encompassed three conditions: a resting phase (REST), a head-up tilt that induces orthostatic stress (HUT), and a mental stress condition involving a mental arithmetic test (MA). For more details on the dataset, the experimental protocol and the ethical approval, we refer the readers to [17], [18].

B. Time series processing and feature extraction

The analysed cardiovascular time series comprised beat-to-beat RR intervals (RR) and pulse-pulse intervals (PP), respectively derived as the temporal distance between successive ECG QRS complexes and the blood pressure peaks. This allowed for the extraction of HRV and PRV time series. Additionally, systolic and diastolic blood pressure time series were generated for each subject under each condition by identifying the maximum and minimum values of the pressure signal within each detected heartbeat. Fifty-three features were computed across three domains (time, frequency, and information) on short-term cardiovascular variability time series (approximately 300 beats or around 5 minutes). In the time domain, metrics such as mean (MEAN), standard deviation (STD), and variance (VAR) were calculated for each time series. In the frequency domain, parametric spectral analysis was performed by fitting the preprocessed time series with an autoregressive model of fixed order $p = 10$, using ordinary least squares. This choice ensured the representation of key oscillatory components, as detailed in [14]. The frequency domain involved calculating absolute spectral power in the Low Frequency (LF, 0.04-0.15 Hz), High Frequency (HF, 0.15-0.4 Hz), and Total Power (TP, 0-0.4 Hz) bands [13], [17], [19]. Normalised power values in the LF and HF bands (LF_n and HF_n) were also computed by dividing for the total absolute power, alongside the respiratory peak frequency in the HF band (f_{HF}) [13], [17], [19]. In the information domain, metrics such as entropy (H), conditional entropy (CE), and self-entropy (SE) were calculated using the k-nearest neighbour estimation approach, $k = 10$, as in [17], [19]. For the RR time series, additional metrics including the root mean square of successive differences (RMSSD), the percentage of successive RR intervals (pNN50), the coefficient of variation (CV) as the STD to MEAN ratio, the absolute spectral power in the Very Low Frequency (VLF, 0-0.04 Hz) band, and the sympathovagal balance index (SVB) were also assessed [19].

C. Feature Importance, Feature Selection and Stress Classification

To evaluate the predictive utility of cardiovascular features under different stress conditions, we applied the information-theoretic feature analysis to a set of cardiovascular-derived indices. The proposed method relies on the decomposition of CMI, computed between continuous features and target, into unique, redundant and synergistic components, enabling a nuanced characterisation of each variable's contribution to allow a more thorough characterisation of the contribution of each variable to the prediction task. As a preliminary step, pairwise correlations among features were computed to identify and remove highly collinear variables. Specifically, features exhibiting a correlation coefficient of exactly ± 1 were considered redundant and excluded from further analysis. In particular, the MEAN PP was highly correlated with the MEAN RR, leading to the removal of the first one and reducing the feature set to fifty-two. This step ensures numerical stability and avoids multicollinearity artefacts that could bias the information-theoretic estimates. The main analysis was conducted within a 10-fold cross-validation framework, ensuring subject-wise independence between training and test sets to prevent data leakage. Within each fold, features were standardised via z-score normalisation based on the training data and applied consistently to the corresponding test set. The Hi-Fi decomposition was performed, and the resulting U, R, and S scores were computed for each fold. For each fold, the top-ranked features were evaluated using a repeated hold-out validation (100 repetitions) on the training data to estimate their classification performance. At each step, a Support Vector Machine (SVM) classifier was trained with an increasing number of features, and the subset yielding the highest average accuracy was retained. The integrated "Statistics and Machine Learning Toolbox" of MATLAB 2022b (The MathWorks, Inc., Natick, MA, USA) was used to apply the classifier, leaving the hyperparameters at their default MATLAB 2022b values. The final model was trained using the features selected according to the selection criteria reported in II on the full training data and evaluated on the corresponding test set. Classification performance was measured in each fold, and expressed by performance metrics and a confusion matrix. Finally, the selection frequency across folds was analysed to identify the most consistently informative features. This synergy-aware selection and evaluation pipeline enables the identification of cardiovascular features that are both individually relevant and collectively complementary, enhancing model interpretability and performance in stress detection tasks.

The analyses were entirely conducted on MATLAB 2022b, and the "Information Theory for the Analysis of Physiological Time Series - ITS Toolbox v 2.1" was used to compute information-theoretic indices.

IV. Results and Discussion

Figure 1 summarises the results of the FI and FS analyses. Panel (a) presents the average (bar column) and standard deviation (error bar) of the unique information, redundancy, synergy, and total CMI for each cardiovascular feature across the 10-fold cross-validation. Panel (b) illustrates the number of folds (out of 10) in which each feature was selected based on the selection criteria. These results provide insights into the contribution and stability of each feature in stress classification across different folds.

Among the analysed cardiovascular features, the SVB RR consistently exhibited the highest unique information across folds, indicating its strong individual contribution to the classification of stress conditions. This suggests that RR alone provides relevant and non-redundant information about the output variable. Interestingly, SBP and DBP presented a balanced profile, with moderate levels of unique and synergistic information, suggesting that their features contribute both individually and in interaction with other features. PP, while exhibiting relatively low unique information, displayed the highest synergy among all features (e.g., VAR), indicating that its relevance emerges predominantly through joint interactions with the rest of the feature set. This pattern suggests that PP features may act as modulatory or complementary variables in the presence of others rather than providing standalone predictive power. Overall, these results highlight heterogeneous contributions among cardiovascular features, with RR playing a dominant and independent role and PP contributing primarily in a synergistic fashion.

The highly unique information observed in entropy-based features (e.g., SE, CE for both RR and PP) and frequency components (e.g., HF, LF) can be physiologically explained by their strong link to ANS control. In this context, SE and CE provide complementary insights into the complexity and regularity of heart rate dynamics, which are known to vary

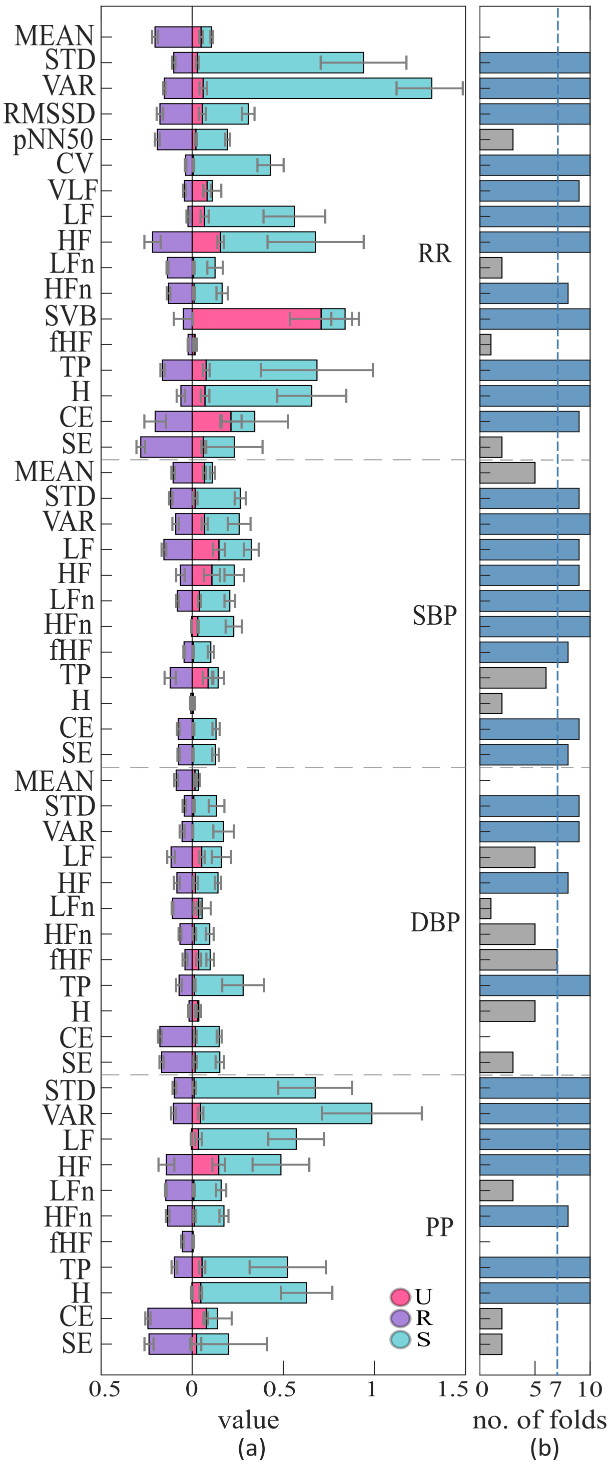

Fig. 1: Feature Importance (FI) and Feature Selection (FS) results. (a) The bar plot shows the average values (column) and standard deviation (error bar) of the unique information (U), redundancy (R), and synergy (S) for each cardiovascular feature across the 10 folds. R values are reported using a secondary axis to highlight their role in the selection criterion, specifically against the sum of U and S. (b) The number of folds (out of 10) in which each feature was selected is displayed based on the selection criterion ($I = U + S - R > 0$). In blue, features selected in more than 7 folds out of 10.

under different stress conditions [20]. Similarly, frequency domain indices reflect both sympathetic and parasympathetic modulation, which changes significantly between rest and stress [21].

Redundant information is higher among features that capture similar physiological mechanisms. High redundancy is often seen in time- and information-domain features that covary due to their shared physiological underpinnings [22].

Features showing high synergy, such as VAR and some entropy indices, likely capture interaction effects between different regulatory components (e.g., sympathetic modulation combined with non-linear heart rate dynamics) [23]. Synergistic information highlights combinations of features that together reveal patterns not evident from individual variables alone. This suggests that complex physiological responses to stress are better captured when considering the joint behaviour of features.

Based on the adopted selection criterion, several features were consistently selected across the 10 folds of cross-validation, indicating a robust and stable contribution to the classification task. The proposed FS method yields approximately a 20% reduction in the number of features, highlighting the physiological relevance of the retained variables. RR-derived features STD, VAR, LF, and CE are selected with high frequency, while features such as pNN50, fHF, SE, and MEAN are rarely or never selected. This suggests that the most consistently selected features capture dynamic aspects of heart rate variability, such as temporal variability (e.g., pNN50), low-frequency oscillations (e.g., VLF, LF) [13], and signal complexity (e.g., CE) [20]. These results indicate that stress-related autonomic regulation is more effectively characterised by fluctuating and irregular patterns in cardiovascular signals rather than by static or aggregate measures (e.g., MEAN).

A similar trend is observed for the SBP derived features. STD, VAR, LF, HF, and SE are selected with high frequency, while features such as MEAN, H, and TP are rarely or never selected, indicating that variability and frequency-domain information are more informative than total power or high-frequency components in the SBP signal.

Within the DBP features, a broad set of features, MEAN, LF, HFn, fHF, H, CE, and SE, are rarely or never selected across folds. This reinforces the relevance of both spectral and entropy-based descriptors in diastolic pressure, underlining the complex physiological modulation of this signal during stress.

PP-derived features also show high consistency in selection, particularly LF, HF, STD, VAR, TP and H, again emphasising the role of both time and frequency content and signal variability. In contrast, the fHF and some high-level descriptors such as CE and SE appear to be less informative.

This pattern supports the idea that stress-related changes in autonomic function are better captured through dynamic, oscillatory, and irregular signal components. For instance, the consistent selection of LF and HF features across DBP and PP points to the role of relative power in distinguishing sympathetic/parasympathetic balance. Similarly, the consistent selection of SE and CE suggests that complexity-related descriptors are essential for capturing subtle modulations in cardiovascular control [18], particularly during mental and postural stress.

Overall, the most robustly selected features reflect variability (e.g., STD, VAR, H) and frequency components rather than simple mean or total power. This reinforces the hypothesis that dynamic aspects of cardiovascular signals, such as oscillatory behaviour and irregularity, carry more discriminative information for stress classification than static features. The results highlight not only the direct influence of individual features but also the complex interactions between physiological markers, offering a more interpretable assessment of ANS dynamics.

	TP	FP	FN	TN	Acc	Pre	Spe	F1
REST	334	68	47	186	81.9	83.1	87.7	85.3
HUT	103	28	24	480	91.9	78.6	81.1	79.9
MA	73	29	54	479	86.9	71.6	57.5	63.7
Tot					80.3	77.8	75.4	76.3

TABLE I: Confusion Matrix and Classification performance, expressed in terms of accuracy (Acc), specificity (Spe), precision (Pre), and the F1-score for each condition (REST, HUT, and MA) and overall (values in percentage).

The results of the classification performance further validate these findings and are summarised in Table I. The classification performance is reported in terms of accuracy (Acc), precision (PRE), specificity (Spec), and the F1-score for each condition (REST, HUT, and MA). The results highlight the strong performance of the classification model in distinguishing between stress and non-stress states across different physiological conditions, demonstrating the importance of dynamic cardiovascular features for accurate stress detection. For the REST condition, the model exhibited good accuracy, successfully distinguishing between stress and non-stress states, as found in [17]. In the HUT condition, where participants underwent a head-up tilt test, the model performed even better, especially in terms of specificity, suggesting its ability to identify non-stress states correctly. However, in the MA condition, which involved mental arithmetic, the performance was lower, particularly in specificity, confirming that cognitive load posed a greater challenge for accurate stress detection [17].

Taken together, these results highlight the direct influence of individual features and the complex interactions between physiological markers, offering a more interpretable assessment of ANS dynamics. The selection of features emphasising variability, spectral content, and entropy, combined with the classification results, reinforces the importance of dynamic cardiovascular regulation metrics for stress classification. The classification performance across the different conditions supports the utility of advanced feature decomposition methods to isolate informative components in physiological data and suggests that our approach offers a reliable and context-sensitive way to assess stress based on ANS activity.

V. Conclusion

This study presents a comprehensive framework that integrates information decomposition and classification to identify stress-related physiological markers. The proposed method combines CMI analysis with a principled selection criterion to isolate features conveying unique and synergistic information while discarding redundancy. The results demonstrate that cardiovascular features reflecting variability, spectral content, and entropy, particularly those derived from RR, DBP, and PP signals, are more consistently informative than static measures like mean or total power values.

Classification outcomes confirm the effectiveness of the selected features, especially in distinguishing REST from orthostatic conditions. Although MA stress detection was more challenging, overall performance supports the validity of the decomposition strategy. Features with high unique information, such as RR-based entropy and spectral markers, contributed independently, whereas redundancy among features like CE and SE reflected shared physiological content. The consistency of selected features across folds highlights the relevance of dynamic, non-linear traits, such as variability, complexity, and oscillatory patterns, as key indicators of autonomic regulation under stress.

These findings emphasise not only the informativeness of individual features but also the importance of capturing complex interactions among physiological signals. By leveraging information-theoretic tools, the proposed approach enables a richer and more interpretable assessment of ANS dynamics, enhancing both detection and understanding of stress responses. The framework also offers a systematic path for Hi-Fi quantification, improving feature interpretability and revealing hidden dependencies in physiological data.

Future work will extend this methodology to broader datasets and explore alternative CMI estimators, moving beyond the current linear-Gaussian assumptions to better capture nonlinear dependencies and further improve the generalizability and interpretability of the proposed framework. Systematic comparisons with alternative feature importance methods, such as SHAP and ReliefF [24], [25], are also envisaged to further validate the advantages of the proposed approach.

References

[1] L. H. Gilpin, D. Bau, B. Z. Yuan, A. Bajwa, M. Specter, and L. Kagal, "Explaining explanations: An overview of interpretability of machine learning," in *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*. IEEE, 2018, pp. 80–89.

[2] C. Molnar, G. Casalicchio, and B. Bischl, "Interpretable machine learning–a brief history, state-of-the-art and challenges," in *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 2020, pp. 417–431.

[3] G. Mijatovic *et al.*, "Measuring the rate of information exchange in point-process data with application to cardiovascular variability," *Frontiers in Network Physiology*, vol. 1, p. 765332, 2022.

[4] M. T. Valderas, J. Bolea, P. Laguna, R. Bailón, M. Vallverdú *et al.*, "Mutual information between heart rate variability and respiration for emotion characterization," *Physiological measurement*, vol. 40, no. 8, p. 084001, 2019.

[5] J. R. Vergara and P. A. Estévez, "A review of feature selection methods based on mutual information," *Neural computing and applications*, vol. 24, pp. 175–186, 2014.

[6] F. Battiston *et al.*, "The physics of higher-order interactions in complex systems," *Nature Physics*, vol. 17, no. 10, pp. 1093–1098, 2021.

[7] F. E. Rosas *et al.*, "Disentangling high-order mechanisms and high-order behaviours in complex systems," *Nature Physics*, vol. 18, no. 5, pp. 476–477, 2022.

[8] S. Stramaglia, L. Faes, J. M. Cortes, and D. Marinazzo, "Disentangling high-order effects in the transfer entropy," *Phys. Rev. Res.*, vol. 6, p. L032007, Jul 2024.

[9] P. L. Williams and R. D. Beer, "Nonnegative decomposition of multivariate information," *CoRR*, vol. abs/1004.2515, 2010.

[10] M. Ontivero-Ortega, L. Faes, J. M. Cortes, D. Marinazzo, and S. Stramaglia, "Assessing high-order effects in feature importance via predictability decomposition," *Phys. Rev. E*, vol. 111, p. L033301, Mar 2025.

[11] J. Lei, M. G'Sell, A. Rinaldo, R. J. Tibshirani, and L. W. and, "Distribution-free predictive inference for regression," *Journal of the American Statistical Association*, vol. 113, no. 523, pp. 1094–1111, 2018.

[12] R. J. Gerrig and P. G. Zimbardo, *American Psychological Association: Glossary of Psychological Terms*. Pearson Education, Education, Incorporated (COR), 2002.

[13] F. Shaffer and J. P. Ginsberg, "An overview of heart rate variability metrics and norms," *Frontiers in public health*, p. 258, 2017.

[14] R. Pernice *et al.*, "Comparison of short-term heart rate variability indexes evaluated through electrocardiographic and continuous blood pressure monitoring," *Medical & biological engineering & computing*, vol. 57, no. 6, pp. 1247–1263, 2019.

[15] F. Scardulla *et al.*, "Photoplethysmograhic sensors, potential and limitations: Is it time for regulation? a comprehensive review," *Measurement*, p. 113150, 2023.

[16] L. Faes, A. Porta, and G. Nollo, "Information decomposition in bivariate systems: Theory and application to cardiorespiratory dynamics," *Entropy*, vol. 17, no. 1, pp. 277–303, 2015. [Online]. Available: https://www.mdpi.com/1099-4300/17/1/277

[17] M. Iovino, I. Lazic, T. Loncar-Turukalo, M. Javorka, R. Pernice, and L. Faes, "Comparison of automatic and physiologically-based feature selection methods for classifying physiological stress using heart rate and pulse rate variability indices," *Physiological Measurement*, vol. 45, no. 11, p. 115004, 2024.

[18] M. Valente *et al.*, "Univariate and multivariate conditional entropy measures for the characterization of short-term cardiovascular complexity under physiological stress," *Physiological measurement*, vol. 39, no. 1, p. 014002, 2018.

[19] M. Iovino, I. Lazic, C. Barà, L. Faes, and R. Pernice, "Conditional mutual information-based feature selection for sex differences characterization," in *2024 13th Conference of the European Study Group on Cardiovascular Oscillations (ESGCO)*, 2024, pp. 1–2.

[20] M. Javorka *et al.*, "Towards understanding the complexity of cardiovascular oscillations: Insights from information theory," *Computers in biology and medicine*, vol. 98, pp. 48–57, 2018.

[21] A. Malliani, M. Pagani, F. Lombardi, and S. Cerutti, "Cardiovascular neural regulation explored in the frequency domain." *Circulation*, vol. 84, no. 2, pp. 482–492, 1991.

[22] A. Porta *et al.*, "Quantifying net synergy/redundancy of spontaneous variability regulation via predictability and transfer entropy decomposition frameworks," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 11, pp. 2628–2638, 2017.

[23] T. Scagliarini, L. Sparacino, L. Faes, D. Marinazzo, and S. Stramaglia, "Gradients of o-information highlight synergy and redundancy in physiological applications," *Frontiers in Network Physiology*, vol. 3, p. 1335808, 2024.

[24] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)*, 2017, pp. 4765–4774.

[25] I. Kononenko, "Estimating attributes: Analysis and extensions of relief," *Machine Learning: ECML-94*, vol. 784, pp. 171–182, 1994.