

Decision-Level Fusion for Robust Wearable Affect Recognition

Anonymous Author(s)

Submission Id: «submission id»

ABSTRACT

Automatic recognition of affective state from wearable physiology has clear societal impact for public health, preventive care, and stress-aware interventions, but real deployments require robustness to non-stationary dynamics, artefacts, and missing sensors. We study this problem on WESAD (baseline, stress, amusement), where common fixed-basis spectral features (e.g., FFT bandpower and Welch PSD) can oversmooth short-lived discriminative patterns. We propose a non-stationary pipeline that combines Fourier-Bessel Series Expansion (FBSE) with EWT data driven spectral segmentation to extract mode wise transient descriptors. For multimodal integration, we adopt decision-level aggregation over per-modality predictors and weight each modality by predictive uncertainty and modality reliability. On WESAD (15 subjects; ECG, EDA, BVP, EMG, ACC three classes) indicate that decision-level aggregation is $\approx 84\%$ times at least as good as feature level aggregation, and $\approx 48\%$ strictly better, suggesting improved robustness under heterogeneous and partially reliable sensing.

KEYWORDS

Wearable affect recognition, non-stationary physiological signals, Fourier-Bessel series expansion, empirical wavelet transform, uncertainty-aware fusion, decision-level aggregation.

ACM Reference Format:

Anonymous Author(s). 2026. Decision-Level Fusion for Robust Wearable Affect Recognition. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 5 pages.

1 INTRODUCTION

Automatic affective state detection (a person’s current subjective, psychological, and physiological state) from physiological signals is a key problem in AI for health, with applications such as continuous mental-health monitoring, stress-aware interventions, and adaptive human-computer interaction [9, 11]. Among affective states, stress and amusement capture complementary high-arousal conditions and enable continuous assessment of emotional well-being using wearable sensing in naturalistic settings [13]. Recent advances in wearable technology make it possible to collect multimodal signals such as electrodermal activity (EDA), electrocardiography (ECG), electromyography (EMG), blood volume pulse (BVP), and motion (ACC), which reflect autonomic and somatic correlates of affect. Benchmark datasets such as WESAD facilitate systematic evaluation of stress and affect recognition under standard protocols,

accelerating progress toward deployable affect-aware health systems [14]. Despite this progress, accurate detection remains challenging because physiological time series are noisy, non-stationary, and subject-dependent, and affect-related dynamics often appear as transient, context-dependent changes rather than stable periodic patterns.

A major limitation of many pipelines is feature extraction. Common approaches rely on statistical summaries or Fourier-based representations computed over fixed windows (e.g., FFT bandpower or Welch PSD). These methods implicitly assume approximate local stationarity and can discard time-varying structure that is critical for distinguishing affective state [10]. Time-frequency methods such as STFT and wavelets partially mitigate this limitation, but they retain trade-offs between time resolution, frequency resolution, and adaptability to signal-specific structure [16]. Motivated by these limitations, we focus on non-stationarity decomposition methods and, in particular, we propose the FBSE-EWT signal analysis method, which combines Fourier Bessel-series expansion with Empirical Wavelet Transform like data-driven segmentation to better preserve transient spectral structure in physiological windows [7].

Beyond limitations in physiological signal analysis, wearable affect monitoring highly relies on multimodal sensing to leverage complementary physiological mechanisms. Combining various physiological signals (such as ECG, EMG, EDA, etc.) using machine learning has proven to improve the affective state detection [3]. However, this introduces a second major challenge: the robust integration under sensor noise, signal failure, and missing modalities, common issues in real-life deployments [8]. Early fusion [5] approaches and models that assume complete sensor availability can therefore fail under realistic conditions. To address this challenge, we adopt a modular approach in which each modality computes an independent prediction, followed by decision-level aggregation that accounts for predictive uncertainty. This design is suitable for agent-based decision support because it outputs calibrated belief states that can be used to trigger interventions conservatively under uncertainty.

In summary, we propose a three-stage pipeline for affective state detection from wearable physiology. First, we extract meaningful representations using feature extraction techniques tailored to non-stationary time series. Second, we train single-sensor predictors to classify sensor-specific features into affective state, enabling each modality to independently learn affect-discriminative patterns. Finally, we aggregate the multiple sensor classifiers to reach one final prediction. We do so by weighting each sensor’s contribution according to its predictive uncertainty in order to enhance robustness to sensor variability and partial failure. We evaluate the proposed framework on WESAD for three-class affect recognition (baseline / stress/ amusement). Our results suggest that non-stationary feature extraction and decision-level aggregation improve reliability in multimodal settings. The key contributions of this paper are as follows:

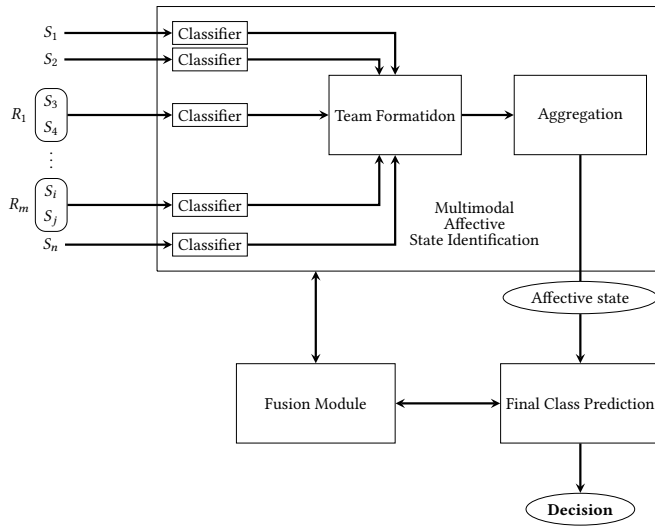


Figure 1: Wearable affect recognition flow: windowing, FBSE–EWT features, per-modality prediction, and uncertainty-weighted decision fusion.

- (1) A non-stationary representation for wearable physiology based on FBSE–EWT that preserves transient spectral structure via data-driven mode isolation.
- (2) A modular per-modality prediction design and an uncertainty-aware, entropy-weighted decision fusion rule that remains stable under noisy or missing modalities.
- (3) An empirical evaluation on WESAD demonstrating the benefit of decision-level aggregation across sensors.

2 RELATED WORK

2.1 Wearable affect detection

Wearable affect recognition is commonly framed as supervised classification from multimodal physiological signals, with WESAD dataset becoming a standard benchmark because it provides synchronized chest and wrist modalities and three affective conditions (baseline, stress, amusement) [14]. Recent work explores both classical Machine learning and deep learning approaches. CNN/LSTM-attention style models are frequently reported as strong baselines for multimodal WESAD classification [15], reflecting the value of temporal modelling when affect clues over time. However, cross-subject generalization remains difficult due to subject-dependent physiology and non-stationary responses. Many methods on WESAD still depend on representations that are either overly aggregated (losing transient nature) or not designed to be useful across modalities and subjects, leaving space for representation learning that better matches the signal properties of wearable physiology.

2.2 Frequency and time-frequency representation

A major line of work relies on frequency domain coefficients, Fourier-domain descriptors and bandpower features are widely used due to simplicity and interpretability. Welch’s method is a standard PSD

estimator that reduces variance relative to raw FFT periodograms through averaging of modified periodogram [17]. These estimators are effective when signals are approximately stationary within short windows, but physiological affect responses can be non-stationary and event-like, where averaging may smooth discriminative structure. Time-frequency methods partially address this limitation. STFT provides a localised spectrum but requires a fixed window and thus faces inherent time–frequency resolution trade-offs [2]. Wavelet decompositions (e.g., DWT) provide multi-resolution analysis and are widely used for physiological signals, but still rely on pre-defined basis structure rather than fully data-adaptive segmentation [12]. Fixed basic spectral and time frequency methods impose analysis choices (band, windows, wavelet families) that can underfit the transient and signal specific structure present in wearable affect data, motivating more adaptive decomposition methods that followed the observed spectrum.

2.3 Adaptive decompositions for non-stationary physiological signals

To better match non-stationary physiology, research has increasingly used adaptive decompositions such as empirical mode decomposition (EMD/HHT), Variational mode decomposition (VMD). EMD decomposes a signal into intrinsic mode functions that capture time varying oscillatory modes [10]. VMD formulates decomposition as a variational optimisation problem with band-limited modes, offering stronger control over mode bandwidth and noise sensitivity [6]. EWT learns wavelet bands from the signal’s spectrum via data driven segmentation, providing an adaptive alternative to fixed wavelet banks [7]. These methods motivate our focus to affect related cues in physiology often seen as transient spectral shifts, where adaptive segmentation can improve separability. In short noisy wearable windows, mode separation can still be unstable, and many decompositions remain sensitive to boundary selection and closely spaced components. Moreover, their output are not always converted into fixed dimension coefficients in a way that consistently benefits machine learning models across modalities.

2.4 Multi model fusion

Multimodal fusion can improve affect recognition by combining complementary physiological cues, as supported by benchmark settings such as WESAD. However, wearable data in practice is affected by artefacts and occasional modality failure, so fusion strategies that operate at the decision level are often preferred over approaches that assume all modalities are always available [4, 8]. In this context, uncertainty aware weighting (e.g., using predictive entropy) provides a principled way to down weight unreliable modality predictions and achieve graceful degradation when signals are missing or noisy [1].

Motivated by these gaps, our approach targets two limitations in wearable affect detection. Representation quality under non-stationary, transient physiological dynamics, and robustness of multimodal integration under missing or unreliable modalities. We introduce an adaptive transient-preserving representation (FBSE–EWT) and couple it with uncertainty and reliability weighted decision-level fusion for multimodal inference in realistic sensing conditions.

3 AFFECTIVE STATE DETECTION: A THREE-STAGE PIPELINE

3.1 Feature Extraction

3.1.1 Problem formulation and windowing. Let subjects be indexed by $s \in \mathcal{S}$ and modalities by $i \in \{1, \dots, M\}$. Each modality provides a discrete time physiological signal

$$x_{s,i}[n], \quad n = 0, \dots, N_{s,i} - 1. \quad (1)$$

Signals are segmented into overlapping windows of fixed duration L with overlap ratio $\alpha = 0.75$. With sampling rate f_i , the window length is $W = Lf_i$ samples and the hop size is $H = (1 - \alpha)W$. The k -th window is defined as

$$x_{s,i}^{(k)}[n] = x_{s,i}[kH + n], \quad n = 0, \dots, W - 1. \quad (2)$$

The WESAD label stream is provided at 700 Hz. Each window receives a label by majority voting with purity threshold ρ :

$$y_s^{(k)} = \arg \max_{c \in \mathcal{C}} \sum_{t \in \mathcal{I}_k} \mathbf{1}\{\ell_s[t] = c\}, \quad (3)$$

$$\max_{c \in \mathcal{C}} \frac{\sum_{t \in \mathcal{I}_k} \mathbf{1}\{\ell_s[t] = c\}}{|\mathcal{I}_k|} \geq \rho, \quad (4)$$

where $\mathcal{C} = \{\text{Baseline, Stress, Amusement}\}$ and \mathcal{I}_k denotes the corresponding label index range.

3.1.2 Fourier Bessel representation. Given a windowed signal $y[n]$ of length $U = W$, the zero-order FBSE represents y using Bessel bases:

$$y[n] = \sum_{m=1}^U C_m J_0\left(\beta_m \frac{n}{U}\right), \quad n = 0, \dots, U - 1, \quad (5)$$

with coefficients

$$C_m = \frac{2}{U^2 (J_1(\beta_m))^2} \sum_{n=0}^{U-1} n y[n] J_0\left(\beta_m \frac{n}{U}\right), \quad (6)$$

where β_m is the m -th positive root of $J_0(\beta) = 0$. Orders m map to physical frequencies via the approximation

$$\beta_m \approx \frac{2\pi f_m U}{f_s}, \quad \Rightarrow \quad m \approx \frac{2f_m U}{f_s}, \quad (7)$$

so the FBSE spectrum is $|C_m|$ as a function of f_m .

This basis is advantageous for wide-band, non-stationary signals because Bessel bases exhibit AM-like behaviour and can yield compact spectral representations for AFM like components.

EWT constructs empirical wavelet filters whose supports depend on the signal's spectral content. Boundaries are obtained by locating meaningful minima in a spectrum derived histogram using a scale space persistence criterion, and an automatic threshold (e.g., Otsu) selects which minima are retained. Let the resulting ordered boundaries be

$$0 = \omega_0 < \omega_1 < \dots < \omega_N = \pi. \quad (8)$$

Given a transition parameter ξ , empirical scaling and wavelet functions in the Fourier domain are defined piecewise. These filters form a tight frame under suitable ξ , and EWT coefficients are computed by inner products with the empirical wavelets/scaling function, reconstruction follows from the same frame.

In FBSE-EWT, boundary detection and filter-bank construction operate on the FBSE spectrum rather than the conventional FFT

spectrum, improving mode separation in challenging cases (e.g., closely spaced components, short-duration activity).

3.1.3 Feature construction from modes. Let $K = N$ be the number of resulting modes for window $x_{s,i}^{(k)}$, and let $m_{s,i,k}^{(r)}[n]$ denote the r -th reconstructed mode. A fixed-dimensional feature vector $\phi_i(x_{s,i}^{(k)})$ is formed by aggregating per-mode descriptors. In particular, the mode energy and log-energy are

$$E_r = \frac{1}{W} \sum_{n=0}^{W-1} (m_{s,i,k}^{(r)}[n])^2, \quad (9)$$

$$\log(E_r + \epsilon), \quad (10)$$

and the mode entropy is computed from the normalised energy distribution

$$p_r[n] = \frac{(m_{s,i,k}^{(r)}[n])^2}{\sum_{u=0}^{W-1} (m_{s,i,k}^{(r)}[u])^2}, \quad (11)$$

as

$$H_r = - \sum_{n=0}^{W-1} p_r[n] \log p_r[n]. \quad (12)$$

This representation preserves transient spectral structure while remaining compatible with standard supervised learning.

3.2 Single-Sensor Prediction

The physiological features extracted from each modality are fed into a single-sensor predictor (SSP) to compute the probability of the predicted classes P_i along with the predictor's confidence score $F1_i$. In this study, we adopt a Multi-Layer Perceptron (MLP) as the predictor, denoted as $SSP_{MLP}(S_i)$, to model non-linear relationships in the feature data collected from a single sensor S_i as presented in Eq. 13. The choice of MLP is motivated by the nature of the input features, which provide fixed-dimensional representations that already encode transient spectral dynamics. This enables effective learning without the need for explicit temporal models such as CNNs or LSTMs, while maintaining computational efficiency and robustness for relatively small datasets.

In our designed classifier, the MLP consists of multiple fully connected layers with ReLU activation, followed by a softmax output layer that produces class probabilities. Dropout and L2 regularisation are applied to mitigate overfitting. The model is trained using categorical cross-entropy loss and optimised with Adam, with early stopping based on validation performance. We also explore variations in the number of hidden layers and select the configuration that achieves the highest prediction accuracy. The output P_i is interpreted as a probabilistic estimate of the affective state, enabling uncertainty quantification via entropy, while the confidence score $F1_i$ represents the overall reliability of the sensor-specific predictor. These two outputs are subsequently used in the multi-sensor aggregation phase discussed in the following section.

$$P_i, F1_i \leftarrow SSP_{MLP}(S_i) \quad \forall i \in \{1, 2, \dots, n\} \quad (13)$$

3.3 Multi-Sensor Aggregation

Wearable affect recognition benefits from combining complementary physiological cues, but in realistic settings different modalities

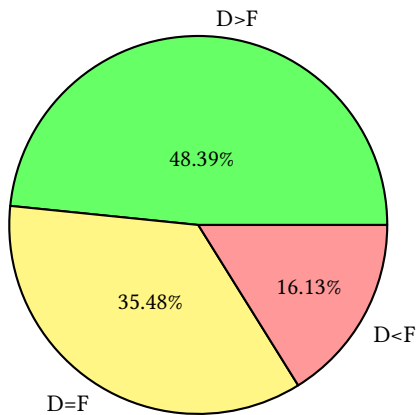


Figure 2: Overall Accuracy

can vary in quality across time due to motion artefacts, skin contact variation, or temporary sensor failure. We therefore perform aggregation at the decision level using a weighted average of per-modality probability vectors. The information can be aggregated at the data, feature, or decision level [8], and each approach has its challenges and benefits. In this work, we propose to perform the aggregation on decision-level, and use a weighted average of the individual sensors’ decisions. Let $T \subseteq \{1, 2, \dots, n\}$ be a team of sensors, and each sensor feeds data to a single-sensor predictor. The decision P_i yielded by the SSP operating on sensor $i \in T$, contributes to the final aggregated decision according to its *entropy* and to the *confidence score* of the SSP. Specifically, we obtain the aggregated decision as:

$$P_T(e) = \frac{1}{\gamma} \sum_{i \in T} (1 - H(P_i))^{F1_i} \cdot P_i(e) \quad \forall e \in \text{classes} \quad (14)$$

where T is a team of sensors, P_i is the probability distribution over the classes obtained by sensor $i \in T$; $H(\cdot)$ denotes the information entropy function; $F1_i$ is the F1 score of the single-sensor predictor of sensor $i \in T$; and $\gamma = \sum_{i \in T} H(P_i)^{F1_i}$ is a normalisation factor. Intuitively, the more certain an SSP is (i.e., the lower the entropy $H(P_i)$), the higher is its contribution to the final decision. Similarly, the more confident the SSP is (i.e., higher confidence score $F1_i$), the more the decision is valued. Notably, the entropy is predicted-instance specific, while the confidence score is descriptive of the SSP in general. As such, using the entropy as the base and the confidence score as the power, we regulate the influence of each individual decision, prioritising primarily according to the entropy, and secondly according to the confidence score.

4 IMPLEMENTATION AND RESULTS

We performed some preliminary evaluation using the WESAD dataset [14] for three-class affect recognition (baseline, stress, amusement). We use five sensors, electrocardiogram (ECG), electrodermal activity (EDA), blood volume pulse (BVP), electromyogram (EMG) and three-axis acceleration, for 15 subjects. In these experiments, we compare the decision-level aggregation against the feature-level one. Specifically, Figure 2 shows the overall accuracy in correctly predicting an class when the aggregation was performed at the

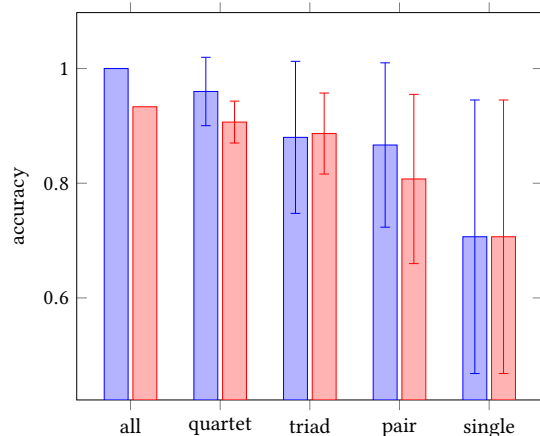


Figure 3: Team level accuracy

feature (F) and decision (D) levels. We note that in $\approx 84\%$ cases, decision aggregation is as good as or better than feature aggregation. Next, we explore the impact of the *team size* on the accuracy. That is, instead of considering all five sensors, we test the accuracy of a team consisting of varying team-size combinations of sensors (including teams of a single sensor as well). Figure 3 shows the comparison of feature-level and decision-level aggregation of all the team sizes. Notably, regardless of the team size, decision-level aggregation is better. Additionally, we observe that when more sensors participate in the aggregation results in better accuracy.

To conclude, our preliminary results indicate that the proposed decision-level aggregation achieves better accuracy than the typical feature-level one. Decision-level aggregation is more generic and can be used to aggregate a diverse set of sensors without expert knowledge regarding specific signal features. It is also robust, since if one of the sensors fail, the decisions of the other sensors can still be aggregated.

5 CONCLUSIONS AND FUTURE WORK

This paper presented a three-stage pipeline for wearable affect recognition that combines a non-stationary representation (FBSE – EWT) with uncertainty and reliability weighted decision-level fusion. The proposed approach is motivated by deployment realities in societally important health and care settings, where physiological sensing is noisy, non-stationary, and occasionally incomplete, and where downstream systems should act on calibrated beliefs rather than overconfident predictions. Preliminary experiments on WESAD (baseline, stress, amusement) indicate that decision-level fusion is competitive with, and often superior to, feature-level fusion across sensor-team sizes, suggesting improved robustness under heterogeneous sensing conditions.

In future, we plan to work and extend this line of work in some of the following ways: (i) use multi-modal input, like physiological sensors, video, audio, images and natural language input, (ii) augment the data by using data generators for different sensors, (iii) detect a variety of relevant classes, (iv) deploy the model in different types of care homes. (v) process temporal information, (vi) communication and coordination among diverse robots from

different manufacturers and having a different set of sensors. We also plan to validate on additional datasets and integrate the resulting uncertainty-aware affect estimates into agent-based decision support pipelines for stress-aware interventions and continuous wellbeing monitoring.

REFERENCES

- [1] Jose Ignacio Aizpurua, Victoria M Catterson, Brian G Stewart, Stephen DJ McArthur, Brandon Lambert, and James G Cross. 2018. Uncertainty-aware fusion of probabilistic classifiers for improved transformer diagnostics. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 51, 1 (2018), 621–633.
- [2] Jont B Allen and Lawrence R Rabiner. 2005. A unified approach to short-time Fourier analysis and synthesis. *Proc. IEEE* 65, 11 (2005), 1558–1564.
- [3] Değer Ayata, Yusuf Yaslan, and Mustafa E Kamasak. 2020. Emotion recognition from multimodal physiological signals for emotion aware healthcare systems. *Journal of Medical and Biological Engineering* 40, 2 (2020), 149–157.
- [4] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence* 41, 2 (2018), 423–443.
- [5] Jing Chen, Bin Hu, Lixin Xu, Philip Moore, and Yun Su. 2015. Feature-level fusion of multimodal physiological signals for emotion recognition. In *2015 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 395–399.
- [6] Konstantin Dragomiretskiy and Dominique Zosso. 2013. Variational mode decomposition. *IEEE transactions on signal processing* 62, 3 (2013), 531–544.
- [7] Jerome Gilles. 2013. Empirical wavelet transform. *IEEE transactions on signal processing* 61, 16 (2013), 3999–4010.
- [8] Raffaele Gravina, Parastoo Alinia, Hassan Ghasemzadeh, and Giancarlo Fortino. 2017. Multi-sensor fusion in body sensor networks: State-of-the-art and research challenges. *Information Fusion* 35 (2017), 68–80.
- [9] Jennifer A Healey and Rosalind W Picard. 2005. Detecting stress during real-world driving tasks using physiological sensors. *IEEE Transactions on intelligent transportation systems* 6, 2 (2005), 156–166.
- [10] Norden E Huang, Zheng Shen, Steven R Long, Manli C Wu, Hsing H Shih, Quanan Zheng, Nai-Chyuan Yen, Chi Chao Tung, and Henry H Liu. 1998. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society of London. Series A: mathematical, physical and engineering sciences* 454, 1971 (1998), 903–995.
- [11] Sylvia D Kreibig. 2010. Autonomic nervous system activity in emotion: A review. *Biological psychology* 84, 3 (2010), 394–421.
- [12] Stephane G Mallat. 1989. Multifrequency channel decompositions of images and wavelet models. *IEEE Transactions on Acoustics, speech, and signal processing* 37, 12 (1989), 2091–2110.
- [13] Rosalind W Picard. 2000. *Affective computing*. MIT press.
- [14] Philip Schmidt, Attila Reiss, Robert Duerichen, Claus Marberger, and Kristof Van Laerhoven. 2018. Introducing wesad, a multimodal dataset for wearable stress and affect detection. In *Proceedings of the 20th ACM international conference on multimodal interaction*. 400–408.
- [15] Ritu Tanwar, Orchid Chetia Phukan, Ghanapriya Singh, Pankaj Kumar Pal, and Sanju Tiwari. 2024. Attention based hybrid deep learning model for wearable based stress recognition. *Engineering Applications of Artificial Intelligence* 127 (2024), 107391.
- [16] Zhiguang Wang and Tim Oates. 2015. Imaging time-series to improve classification and imputation. *arXiv preprint arXiv:1506.00327* (2015).
- [17] Peter Welch. 2003. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on audio and electroacoustics* 15, 2 (2003), 70–73.