Deep Residual Neural Networks For Robust Denoising In Raman Spectroscopy

Matteo Matera[1], Luca Polenta[1] and Christian Napoli[1]

[1]Department of Artificial Intelligence, Czestochowa University of Technology, Czestochowa, Poland

Abstract

Raman spectroscopy is a key tool for material analysis, but its accuracy is often hindered by noise and baseline distortions. This paper presents a robust denoising method using a parallel deep residual neural network architecture based on DnCNN, designed for one-dimensional spectral data. The model learns noise patterns through multiple convolutional branches, enabling effective denoising without assumptions about the signal origin. We evaluate several pre-processing techniques, with minimum-shift normalization proving most effective in preserving spectral features. Trained on datasets with varying noise levels, the network achieves high peak detection accuracy and low error rates, outperforming traditional and recent methods. This approach enhances the reliability of Raman analysis and demonstrates the potential of AI-driven models in spectroscopy and time-series signal processing.

Keywords

Residual Learning Neural Network, One-Dimensional Signal Denoising, Raman Spectrum, Time-Series Signal

1. Introduction

Denoising methods are fundamental across numerous scientific disciplines. Their purpose is to eliminate noise introduced during the acquisition of a signal, image, or other data. In this project, the focus is on one-dimensional signals obtained from Raman spectroscopy, a non-destructive material analysis technique based on the scattering of monochromatic electromagnetic radiation by a sample.

Raman spectroscopy is crucial for studying materials in solid, liquid, or gaseous states, particularly carbon-based materials such as graphite and graphene. It also finds applications in geological research, industrial process control, planetary exploration, internal security, and even medical diagnostics when combined with artificial intelligence to analyze cancer cells and melanomas.

The key strength of Raman spectroscopy lies in its ability to detect subtle molecular vibrations that are specific to the chemical bonds within the sample. This makes it extremely valuable for both qualitative and quantitative analyses. However, the full potential of Raman spectroscopy can only be realized if the acquired signals are clear, stable, and free of noise. Unfortunately, this is rarely the case in practical applications. Raman signals are often weak, and their acquisition is sensitive to many external factors such as temperature, laser fluctuations, photobleaching, sample heterogeneity, or even instrument drift.

Obtaining a clean, accurate signal is critical for reliable material characterization. However, Raman signals often suffer from noise that obscures peaks and non-uniform baselines, complicating interpretation. Among the various types of signal corruption [1], baseline distortion and high peak noise are the most problematic [2], as the analysis heavily relies on peak characteristics [3]. These peak characteristics, such as position, intensity, width, and shape, are used to identify substances and measure their concentrations. When noise is too strong, even expert analysts may misinterpret spectra, leading to incorrect conclusions.

To address these issues, the scientific community has developed numerous denoising techniques. Traditional signal processing methods such as Empirical Mode Decomposition (EMD) [4] and wavelet analysis have been widely applied. EMD is attractive because it does not require any prior knowledge of the signal or noise structure, but it is often unstable and sensitive to noise itself. Wavelet transforms, on the other hand, allow localized analysis in both time and frequency domains, but they require careful selection of wavelet families and thresholds.

Wiener filtering is another classical method that performs well under certain assumptions about signal stationarity and noise properties. It can produce satisfactory results, especially in laboratory settings, but its performance decreases significantly in more dynamic or uncontrolled environments [5, 6, 7, 8, 9, 10].

In recent years, attention has turned toward machine learning and, in particular, deep learning, for denoising tasks. These approaches do not rely on handcrafted rules but instead learn directly from examples. This ability to capture complex, non-linear patterns from data has led to performance levels that surpass those of classical methods in many domains. In particular, convolutional neural

ICYRIME 2025: 10th International Conference of Yearly Reports on Informatics, Mathematics, and Engineering. Czestochowa, January 14-16, 2025

christian.napoli@pcz.pl (C. Napoli)

© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR-WS.org/Vol-3984/p04.pdf

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

networks (CNNs) have proven effective for processing structured data like time series and spectra.

Recent contributions in this area have applied deep residual learning to denoising problems with encouraging results. Residual networks learn to estimate the noise component instead of reconstructing the entire signal, simplifying the learning task and improving convergence. One of the most successful architectures in this field is the DnCNN model, originally developed for image denoising. Adaptations of this architecture to one-dimensional signals have been proposed, including in biomedical and physical measurement contexts [11, 12, 13, 14, 15, 16].

In this project, a discriminative learning model based on the DnCNN structure is implemented. It uses parallel convolutional branches to better learn noise characteristics without assuming any specific origin or structure of the signal. This makes the model suitable for a wide range of applications, including Raman spectroscopy. Special emphasis is placed on ensuring that the model remains lightweight, which is important for future integration into embedded systems or portable spectrometers used in field operations.

By focusing on generalization, efficiency, and minimal assumptions, this study aims to contribute to the growing field of AI-assisted spectroscopy, where deep learning is rapidly becoming a standard tool for signal interpretation and preprocessing.

2. Related Work

Earlier attempts to denoise Raman spectra primarily relied on wavelet-based methods [17, 18], but these were quickly surpassed by techniques better suited to nonlinear, non-stationary signals, such as EMD [19]. Although EMD decomposes signals into intrinsic mode functions (IMFs) without prior signal characterization, its performance is limited by requirements such as symmetry in the upper and lower signal envelopes, which are not always met in Raman data.

Raman spectroscopy, due to its sensitivity and specificity, has long been used as a key method for material identification. However, the presence of noise and baseline fluctuations strongly limits its reliability. The scientific community has proposed many denoising techniques to overcome this challenge, starting from classical signal processing methods and gradually moving towards machine learning and deep learning approaches.

One of the most traditional approaches to Raman signal denoising has been based on wavelet transforms. These methods decompose the signal into multiple frequency components, allowing selective suppression of noise at various scales. For instance, Kumar et al. [17] and Chen et al. [18] used discrete wavelet transforms to smooth Raman signals while preserving spectral peaks. However, these methods depend on the choice of the wavelet basis and thresholding strategy, which often needs manual tuning.

Other researchers adopted Empirical Mode Decomposition (EMD), which is a fully data-driven technique suitable for non-linear and non-stationary data [19]. The main advantage of EMD is its capacity to adaptively separate noise from signal through decomposition into Intrinsic Mode Functions (IMFs). Nevertheless, its results can suffer from mode mixing, and the quality of denoising is not always guaranteed, especially in cases where noise is not additive or where the signal does not meet the envelope symmetry conditions.

Wiener filtering has also been considered an effective solution for Raman signal denoising. For example, Bai et al. [20] applied a modified Wiener filter to improve Raman signal quality in conditions with low signal-to-noise ratio. This technique does not require prior experimental data, which is an advantage, but its performance degrades in the presence of non-Gaussian noise, which is common in real-world spectroscopic data.

In the last few years, the research community has shifted toward deep learning techniques due to their ability to automatically learn representations from raw data. The most basic models used convolutional neural networks (CNNs) trained on labeled spectral datasets. For example, Pan et al. [21] proposed a dual-path CNN that processes the signal in parallel branches, each specialized in detecting different noise features. This architecture improves robustness against baseline shifts and peak distortion, but it increases the computational cost.

More advanced models such as UHRED (Unsupervised Hyperspectral Residual Encoder-Decoder) introduced by Abdolghader et al. [22] combined denoising and segmentation in an unsupervised pipeline, demonstrating promising results for hyperspectral data. However, these models are often tailored for imaging applications and may not directly generalize to one-dimensional Raman spectra without significant adaptation.

A separate line of work has explored residual learning for signal restoration. One of the most well-known architectures is the DnCNN [23], originally developed for image denoising but recently adapted for one-dimensional signals as well [24]. The residual learning principle allows the model to focus on learning only the noise component, which simplifies training and improves convergence. This approach has been particularly successful in biomedical signals and time-series applications where precise peak localization is required.

In addition, some authors have explored hybrid methods, combining classical signal processing techniques with deep learning to balance interpretability and performance. For example, Zhou et al. [25] designed a model that applies adaptive baseline correction before feeding the signal into a neural network. These hybrid systems

often achieve good results, but they can be sensitive to errors in the preprocessing stage.

To summarize, although many strategies have been proposed, there is still no universally accepted solution for Raman signal denoising. Classical methods often fail to generalize across different noise types and signal conditions. Deep learning models show great promise but sometimes require large labeled datasets and careful tuning. In this context, our work proposes a lightweight parallel residual CNN, trained on synthetically noised spectra, capable of generalizing across varying noise patterns and maintaining high peak fidelity even in distorted baselines. This contributes to improving the practical usability of Raman spectroscopy in both research and applied contexts.

3. Dataset and Pre-Processing

The performance of any deep learning model for signal denoising strongly depends on the quality and diversity of the training data. In this project, we worked with four synthetic datasets, each composed of one-dimensional Raman-like spectra. Each sample has 801 time steps, corresponding to a realistic spectral resolution often used in practical Raman analysis. The datasets were designed to cover a wide range of noise conditions and allow the model to learn robust denoising strategies that generalize well across various scenarios.

The first three datasets were created to represent distinct noise environments: low-noise, high-noise, and mixed-noise settings. Each of these contains 5000 examples. The low-noise dataset includes signals with minor Gaussian noise, mimicking laboratory conditions where samples are well-prepared, and the spectrometer operates in ideal calibration. The high-noise dataset simulates challenging measurement conditions, such as field experiments or low-light environments, where the signal-to-noise ratio is heavily degraded. The mixed-noise dataset contains a random combination of both low- and high-noise patterns, and helps the model learn noise characteristics in a more varied context.

In addition to these, a fourth and larger dataset was created, containing 15000 samples with mixed noise levels. This extended dataset was specifically used for training and testing the final version of our deep learning model. The larger sample size allows the network to generalize better, while the mixed noise ensures that it learns to handle diverse real-world signal conditions. Each signal in the datasets is paired with a clean ground truth version, allowing supervised training through direct comparison.

Pre-processing plays a crucial role in any machine learning pipeline, especially when dealing with spectral data. Different strategies were tested to prepare the data before feeding it into the network. We experimented with several normalization and baseline correction techniques, including Z-score normalization, max scaling, baseline subtraction using polynomial fitting, and wavelet-based background removal. While some of these methods produced visually appealing results, they often introduced small distortions in peak shapes or amplitudes, which could negatively affect the learning process.

Among all tested methods, the most effective and robust was a simple transformation: shifting the minimum value of each spectrum to zero. This method ensures that the entire signal lies in the positive domain and removes negative values, which are not expected in Raman intensity data. More importantly, this approach preserves the relative intensity of peaks, avoids rescaling artifacts, and maintains the original structure of the noise. This makes it easier for the neural network to distinguish noise from real spectral features. For this reason, minimum-shift normalization was selected as the default pre-processing step for all experiments.

Finally, in order to make the data compatible with convolutional neural network layers, each input signal was reshaped into a three-dimensional tensor with the shape `(samples, timesteps, 1)`. This format treats the Raman signal as a one-dimensional image with a single channel, allowing the model to apply 1D convolutions and learn local noise patterns effectively. This reshaping is a standard procedure when working with convolutional architectures in time-series or spectral signal processing tasks.

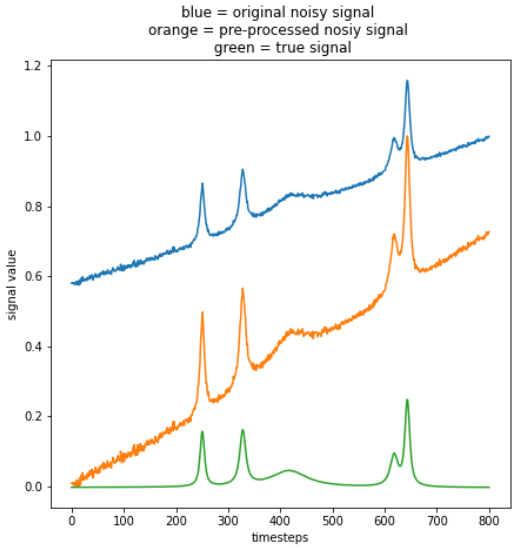

Figure 1: Signal normalization between 0 and 1

Three pre-processing methods were evaluated:

- **Normalization:** Scaling signals between 0 and 1

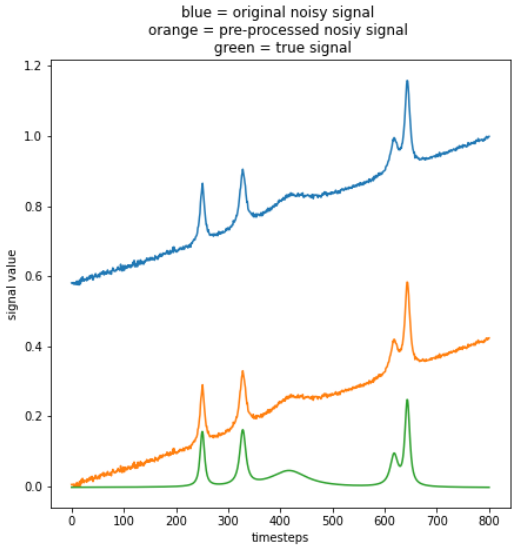

Figure 2: Shift the noisy signal so that the minimum value is zero

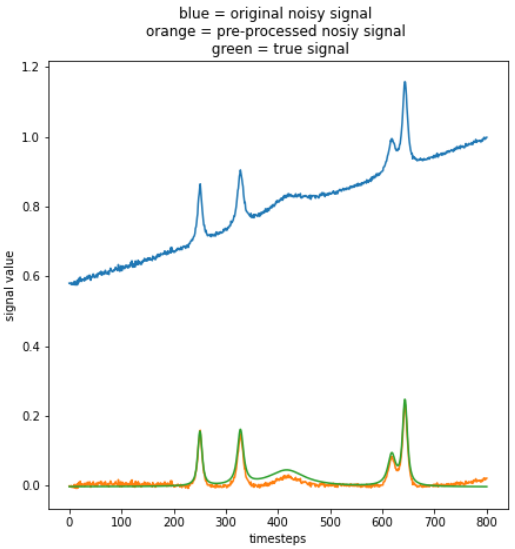

Figure 3: Application of airPLS algorithm on noisy signal

(Figure 1).

- **Minimum Shift:** Shifting each signal so its minimum value is zero (Figure 2), which preserves the signal's shape.
- **airPLS Algorithm:** Adaptive baseline correction (Figure 3) [26, 27].

4. Model

The proposed denoising model is a parallel deep residual neural network architecture based on the DnCNN framework, specifically adapted for one-dimensional Raman spectral data. The core idea is to employ multiple independent branches, each learning to capture distinct noise patterns present in the input signal. This parallel configuration improves the model's robustness and generalization, especially across datasets with different noise characteristics.

Each branch of the network processes the same input signal in parallel, applying a series of 1D convolutional layers with varying kernel sizes and dilation rates. This allows each branch to extract features at different temporal scales. The layers use LeakyReLU activation functions to maintain gradient flow and allow for learning non-linear transformations. Batch Normalization is included as an optional component to stabilize and speed up training, although experiments indicate that omitting it can sometimes yield better performance in this context.

To encourage generalization and reduce overfitting, L2 regularization is applied to the convolutional layers, with regularization strengths fine-tuned in the range of $1e^{-5}$ to $1e^{-6}$. Each branch outputs an intermediate estimate of the noise, and these outputs are averaged to form a combined noise prediction.

The final denoised signal is obtained by subtracting this aggregated noise estimate from the original input. This residual learning formulation focuses the model's capacity on learning the noise component rather than reconstructing the full signal, which simplifies the learning task and improves convergence.

The model is compiled using the Adam optimizer, chosen for its efficiency and adaptability. The loss function is the Residual Sum of Squares (RSS), which emphasizes penalizing large errors in noise estimation. Additionally, Mean Squared Error (MSE) and Mean Absolute Error (MAE) are used as evaluation metrics during training and validation.

The training process is supported by a suite of callbacks, including:

- **ReduceLROnPlateau:** Dynamically reduces the learning rate when performance plateaus.
- **EarlyStopping:** Prevents overfitting by halting training when validation performance stops improving.
- **ModelCheckpoint:** Saves the best-performing model during training.
- **TensorBoard:** Provides real-time visualization of training metrics.

Overall, this architecture is designed to be lightweight, modular, and efficient, making it suitable not only for

Table 1
Quantitative results of model trainings with pre-processing shift to zero

Type of dataset	Num. of real peak	Peak correctly predicted	Peak not predicted	Peak incorrectly predicted (added)	Accuracy peak intensity THR 10%-20%	Accuracy peak intensity THR 15%-30%	RMSE	MAPE
Low Noise	3244	2754 (84.90%)	490 (15.10%)	89	83.51%	92.05%	1.15e-2	4.14
Mixed Noise	3278	2775 (84.66%)	503 (15.34%)	100	78.99%	89.23%	1.32e-2	9.13
High Noise	3222	2627 (81.53%)	595 (18.47%)	125	78.72%	90.67%	1.30e-2	17.82
Mixed Noise (17 Aug)	9219	7719 (83.72%)	1500 (16.27%)	1070	83.61%	92.54%	1.55e-2	4.72

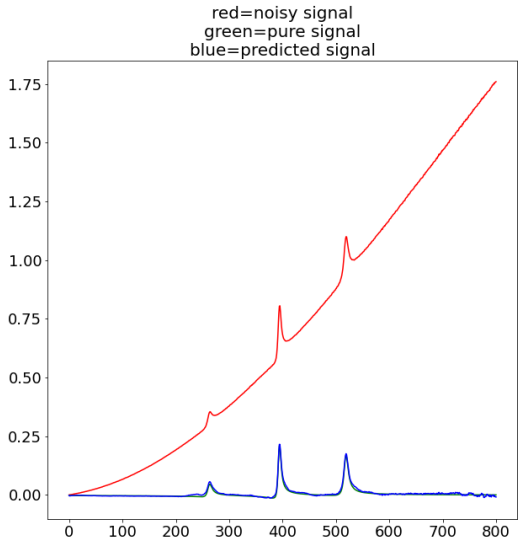

Figure 4: Minimum of signal to zero on sample 1

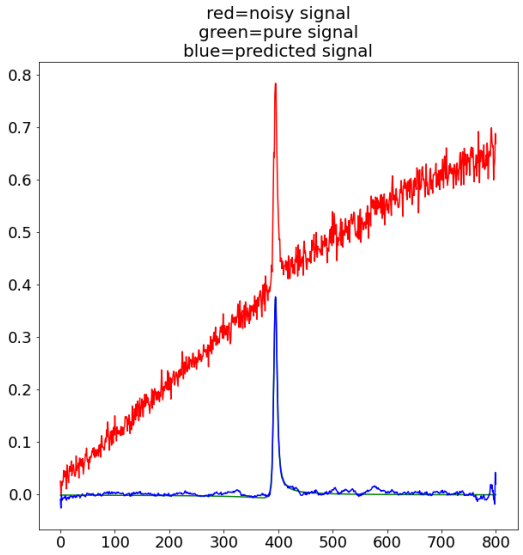

Figure 5: Minimum of signal to zero on sample 2

Raman spectroscopy but also for other one-dimensional denoising tasks in time-series analysis.

5. Experiments and Results

The model was optimized through extensive testing. The best-performing configuration includes depth settings of [17, 12, 7, 3], corresponding filter counts of [96, 64, 32, 96], kernel sizes of [5, 15, 30, 7], dilation rates of 5, no Batch Normalization, and L2 regularization between $1e^{-5}$ and $1e^{-6}$.

Among pre-processing strategies, shifting the minimum to zero led to the best results, achieving excellent denoising without signal distortion (Figures 4 and 5).

Quantitative evaluations involved peak detection accuracy and intensity prediction, as shown in Table 1. Peak localization accuracy consistently exceeded 81%, and peak intensity accuracy ranged from 78% to 92% depending on noise level.

Comparative results with prior work [21] (Table 2) demonstrate that our model achieves substantially lower RMSE and MAPE values, outperforming state-of-the-art methods even under high-noise conditions.

Table 2
Quantitative results of papers [21]

Models	RMSE	MAPE
CNN	183.75	13.79
parallelCNNs	129.96	9.50
Empirical Bayes	422.59	27.07
Block James-Stein	489.30	27.57
False Discovery Rate	467.36	27.79
Minimax Estimation	512.98	28.15
Stein's Unbiased Risk Estimate	416.64	34.80
Universal Threshold G	674.92	29.48

6. Conclusion

This paper proposes a parallel convolutional neural network model for denoising highly variable one-dimensional signals, such as those encountered in Raman spectroscopy. Each network branch independently learns noise characteristics, and the final model combines these learnings for superior results. Pre-processing via minimum shifting, which preserves signal integrity, further enhances performance.

Experimental results demonstrate excellent qualitative and quantitative denoising performance across datasets, establishing the proposed method as a significant advancement over previous techniques.

7. Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT, Grammarly in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] S. Russo, S. Ahmed, I. E. Tibermacine, C. Napoli, Enhancing eeg signal reconstruction in cross-domain adaptation using cyclegan, in: Proceedings - 2024 International Conference on Telecommunications and Intelligent Systems, ICTIS 2024, 2024. doi:10.1109/ICTIS62692.2024.10894543.

[2] C. Randieri, V. F. Puglisi, Design of a software defined radio using soc builder, in: CEUR Workshop Proceedings, volume 3398, 2022, p. 73 – 77.

[3] S. Russo, F. Fiani, C. Napoli, Remote eye movement desensitization and reprocessing treatment of long-covid- and post-covid-related traumatic disorders: An innovative approach, Brain Sciences 14 (2024). doi:10.3390/brainsci14121212.

[4] L. Max, E. Andrew, D. Matt, B. Ben, Empirical mode decomposition, https://www.clear.rice.edu/elec301/Projects02/empiricalMode/index.html, 2018.

[5] F. Bonanno, G. Capizzi, G. L. Sciuto, C. Napoli, G. Pappalardo, E. Tramontana, A cascade neural network architecture investigating surface plasmon polaritons propagation for thin metals in openmp, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), volume 8467 LNAI, 2014, p. 22 – 33. doi:10.1007/978-3-319-07173-2_3.

[6] F. Bonanno, G. Capizzi, S. Coco, C. Napoli, A. Laudani, G. L. Sciuto, Optimal thicknesses determination in a multilayer structure to improve the spp efficiency for photovoltaic devices by an hybrid fem - cascade neural network based approach, in: 2014 International Symposium on Power Electronics, Electrical Drives, Automation and Motion, SPEEDAM 2014, 2014, p. 355 – 362. doi:10.1109/SPEEDAM.2014.6872103.

[7] F. Bonanno, G. Capizzi, A. Gagliano, C. Napoli, Optimal management of various renewable energy sources by a new forecasting method, in: SPEEDAM 2012 - 21st International Symposium on Power Electronics, Electrical Drives, Automation and Motion, 2012, p. 934 – 940. doi:10.1109/SPEEDAM.2012.6264603.

[8] G. Capizzi, F. Bonanno, C. Napoli, A wavelet based prediction of wind and solar energy for long-term simulation of integrated generation systems, 2010, p. 586 – 592. doi:10.1109/SPEEDAM.2010.5542259.

[9] F. Bonanno, G. Capizzi, G. L. Sciuto, C. Napoli, Wavelet recurrent neural network with semi-parametric input data preprocessing for micro-wind power forecasting in integrated generation systems, in: 5th International Conference on Clean Electrical Power: Renewable Energy Resources Impact, ICCEP 2015, 2015, p. 602 – 609. doi:10.1109/ICCEP.2015.7177554.

[10] F. Bonanno, G. Capizzi, G. L. Sciuto, C. Napoli, G. Pappalardo, E. Tramontana, A novel cloud-distributed toolbox for optimal energy dispatch management from renewables in igss by using wrnn predictors and gpu parallel solutions, in: 2014 International Symposium on Power Electronics, Electrical Drives, Automation and Motion, SPEEDAM 2014, 2014, p. 1077 – 1084. doi:10.1109/SPEEDAM.2014.6872127.

[11] N. Brandizzi, A. Fanti, R. Gallotta, S. Russo, L. Iocchi, D. Nardi, C. Napoli, Unsupervised pose estimation by means of an innovative vision transformer, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial In-

telligence and Lecture Notes in Bioinformatics), volume 13589 LNAI, 2023, p. 3 – 20. doi:10.1007/978-3-031-23480-4_1.

[12] G. Zimatore, M. Cavagnaro, Recurrence analysis of otoacoustic emissions, Understanding Complex Systems (2015) 253 – 278. doi:10.1007/978-3-319-07155-8_8.

[13] M. C. Gallotta, V. Bonavolontà, G. Zimatore, S. Iazzoni, L. Guidetti, C. Baldari, Effects of open (racket) and closed (running) skill sports practice on children's attentional performance, Open Sports Sciences Journal 13 (2020) 105 – 113. doi:10.2174/1875399X02013010105.

[14] C. Randieri, A. Pollina, A. Puglisi, C. Napoli, Smart glove: A cost-effective and intuitive interface for advanced drone control, Drones 9 (2025). doi:10.3390/drones9020109.

[15] M. C. Gallotta, G. Zimatore, L. Falcioni, S. Migliaccio, M. Lanza, F. Schena, V. Biino, M. Giuriato, M. Bellafiore, A. Palma, et al., Influence of geographical area and living setting on children's weight status, motor coordination, and physical activity, Frontiers in pediatrics 9 (2022) 794284.

[16] N. Boutarfaia, S. Russo, A. Tibermacine, I. E. Tibermacine, Deep learning for eeg-based motor imagery classification: Towards enhanced human-machine interaction and assistive robotics, in: CEUR Workshop Proceedings, volume 3695, 2023, p. 68 – 74.

[17] N. Kumar, A. Siddiqi, K. Alam, Raman spectral data de-noising based on wavelet analysis, https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.662.4653&rep=rep1&type=pdf, 2014.

[18] C. Chen, F. Peng, Q. Cheng, D. Xu, Raman spectral data denoising based on wavelet analysis, 2009.

[19] F. León-Bejarano, M. Ramírez-Elías, M. O. Mendez, G. Dorantes-Méndez, M. d. C. Rodríguez-Aranda, A. Alba, Denoising of raman spectroscopy for biological samples based on empirical mode decomposition, https://www.worldscientific.com/doi/abs/10.1142/S0129183117501169, 2017.

[20] C. Bai, H. Liu, Denoising of low-snr raman spectra using wiener filter with improved noise estimation, Applied Spectroscopy 74 (2020) 317–324.

[21] L. Pan, P. Pipitsunthonsan, P. Zhang, C. Daengngam, A. Booranawong, M. Chongcheawchamnan, Noise reduction technique for raman spectrum using deep learning network, https://ieeexplore.ieee.org/abstract/document/9325766, 2020.

[22] P. Abdolghader, A. Ridsdale, T. Grammatikopoulos, G. Resch, F. Légaré, A. Stolow, A. F. Pegoraro, I. Tamblyn, Unsupervised hyperspectral stimulated raman microscopy image enhancement: denoising and segmentation via one-shot deep learning, https://opg.optica.org/oe/fulltext.cfm?uri=oe-29-21-34205&id=460202, 2021.

[23] K. Zhang, W. Zuo, Y. Chen, D. Meng, L. Zhang, Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising, IEEE Transactions on Image Processing 26 (2017) 3142–3155.

[24] F. Brandizzi, C. Napoli, G. Russo, Residual learning for time-series denoising using 1d dncnn: Applications in bio-signal processing, Computers in Biology and Medicine 157 (2023) 106810.

[25] L. Zhou, Z. Zhang, Hybrid baseline correction and deep learning for robust raman spectrum denoising, Journal of Raman Spectroscopy 52 (2021) 45–55.

[26] Z.-M. Zhang, S. Chen, Y.-Z. Liang, Baseline correction using adaptive iteratively reweighted penalized least squares, https://pubs.rsc.org/en/content/articlelanding/2010/an/b922045c/unauth, 2010.

[27] S. C. Renato Lombardo, Zhi-Min Zhang, Y.-Z. Liang, adaptive iteratively reweighted penalized least squares (airpls), https://github.com/zmzhang/airPLS, 2016. [Online; accessed November-2021].