Frequency-Based Filtering for Automated Spike Detection in Human Microneurography

Dominika Darabos, Alina Troglio, Andrea Fiebig, Anna Maxion,
Ekaterina Kutafina, Barbara Namer, Gergő Bognár, Péter Kovács

***Abstract*—Although microneurography is a powerful tool for recording activity from individual human peripheral nerve fibers, its broader adaptation is still limited due to the characteristics of the data, which prevent most existing spike detection algorithms from being used effectively. Since these are extracellular recordings collected with microelectrodes, they contain a significant amount of noise, resulting in a low signal-to-noise ratio. This noise can easily mask the spike-shaped responses of nerve fibers, whose shape can also vary over time, making it difficult to develop detectors that generalize reliably across different spike shapes. In this study, we analyze the recordings in the frequency domain and identify frequency components that can be removed with minimal compromise to the underlying neural activity. Using the weighted Hermite Variable Projection Neural Network applied to sliding-window segments of the signal, we compare classification performance on the original, low-pass filtered, and band-pass filtered data. The results identify the frequency ranges that carry the most informative components of the neural signal and demonstrate that appropriate filtering can reduce the number of false alarms to 22% of the level obtained with unfiltered data, while only modestly affecting true detections.**

***Index Terms*—microneurography, spike detection, band-pass filtering, low-pass filtering, noise reduction, variable projection, sliding-window classification, time-series analysis, VPNet**

I. Introduction

C-type fibers [1], [2] are responsible for transmitting neural information such as nociception or pruriception within the peripheral nervous system. Understanding the patterns of these neural responses could contribute significantly to the development of treatments for neuropathic pain. When a nociceptive stimulus is detected by sensory receptors, the resulting neural activity travels through the nerve fibers toward the central nervous system. This activity can be recorded with microelectrodes inserted into the nerve bundles. The number of detectable fibers depends on on the position of the single electrode, and these fibers are referred to as the target fibers. The technique enabling such measurements is microneurography [3], [4].

Previous studies [5], [6] have shown that artificial, constant-frequency electrical stimulation evokes nearly constant response times in C-type fibers, either immediately or—depending on the stimulation frequency—after an initial period of slowing, while the introduction of additional stimuli leads to nonlinear changes in response latency. If spike-like neural responses were always cleanly observable in the signal, this temporal variability would cause less of a challenge.

In practice, microneurography recordings are extracellular and contain activity from nearby non-target fibers as well as substantial in vivo, environment-related noise, while spike shapes may vary over time. These factors further challenges the ability of detectors to generalize to a single canonical shape. Detection algorithms specifically designed for microneurography remain extremely limited. Currently, a semi-automated approach is commonly used in which experts identify evoked spikes based on expected latencies using a vertical waterfall-style visualization [7]. While spikes may be partially masked by noise, vertical alignment allows reliable identification with stable conduction velocity. However, this approach provides only partial detection, as spontaneous spikes cannot be captured. Once spikes are detected, spike-sorting methods are applied to assign events to individual fibers, for which algorithms already exist [8].

In our previous work [9], we presented, to our knowledge, the first neural network evaluated on microneurography data for spike detection. The sequential recordings were converted into sliding-window segments, and a knowledge-augmented classification model—the weighted Hermite Variable Projections Neural Network (WHVPNet) [10]—was used to predict whether a given window contained spike-like activity. That work focused primarily on handling extreme class imbalance and extracting shape-based features for spike characterization.

In the present study, we shift attention to the signal structure itself. We analyze the raw recordings in the frequency domain to identify potential patterns related to the underlying neural activity and to determine which frequency bands are dominated by noise. These findings are then used to create low-pass and band-pass filters designed to remove uninformative frequency components with minimal distortion of the neural signal. The filtered recordings are integrated into the same WHVPNet framework, allowing us to directly compare model performance on raw versus frequency-filtered data.

II. Methods

All training, testing, and evaluation were performed on a single in vivo human microneurography recording containing two C-type fibers [6]. The recording was sampled at 10 kHz, and the total duration of the recording is approximately 13 minutes. The study was approved by the Ethics Board of the University Hospital RWTH Aachen (Vo-Nr. EK141-19), and written informed consent was obtained. Ground-truth spike labels were generated by a single expert using

D.D., G.B., and P.K. are with the Eötvös Loránd University, Budapest, Hungary. A.T., A.F., and A.M are with the RWTH Aachen University, Aachen, Germany. E.K. is with the University of Cologne, Cologne, Germany. B.N. is with the University of Würzburg, Würzburg, Germany. The project was supported by the University Excellence Fund of Eötvös Loránd University, Budapest, Hungary (ELTE). P. Kovács was supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences.

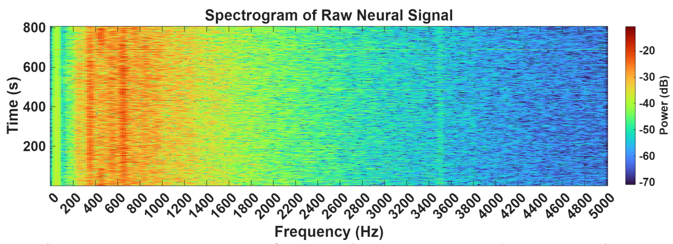

Fig. 1: Spectrogram of the microneurography recording, showing the time–frequency representation.

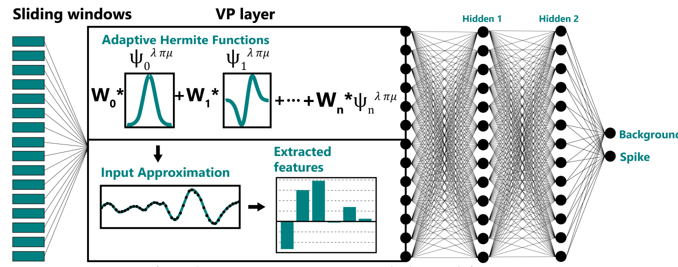

Fig. 3: WHVPNet model architecture.

the vertical-alignment procedure described in [6]. The signal underwent fixed hardware filtering during recording, however, as the hardware predates detailed documentation, precise filter specifications are unavailable, and the data were treated as raw signals. Due to high inter-recording variability in microneurography, generalization to other recordings is left to future work involving calibration and transfer-learning strategies.

A. Frequency-domain analysis

To characterize the spectral structure of the recordings, short-time Fourier transforms were computed using a 256-sample window, 200-sample overlap, and a 512-point fast Fourier transformation at a 10 kHz sampling rate. The resulting spectrogram (Fig. 1) indicated that most neural energy resides in lower-frequency bands, which motivated the subsequent filtering strategy.

B. Filtering

Guided by the spectral analysis, 4th-order Butterworth low-pass and band-pass filters were applied (Fig. 2) using zero-phase forward–reverse filtering. Several band-pass configurations with varying lower and upper cutoff frequencies were evaluated. Because the relevant neural bands are tightly clustered, a single band-pass setup was used rather than a multiband combination.

C. Integration with VPNet

Filtered data were processed using the same sliding-window WHVPNet framework (Fig. 3) as in the previous work. The model adapts Hermite basis functions through learnable parameters to capture spike-shape variability. Their combination provides an input approximation from which features are extracted, followed by fully connected layers that perform binary, window-level spike classification, deciding between spike and non-spike. The number of extracted features and Hermite-system weights are model parameters, selected through prior hyperparameter optimization. Due to the model's compact size (64 parameters), two evaluation strategies were considered:

- retraining VPNet on each filtered dataset, and
- applying filters only during validation/testing while keeping the original model unchanged.

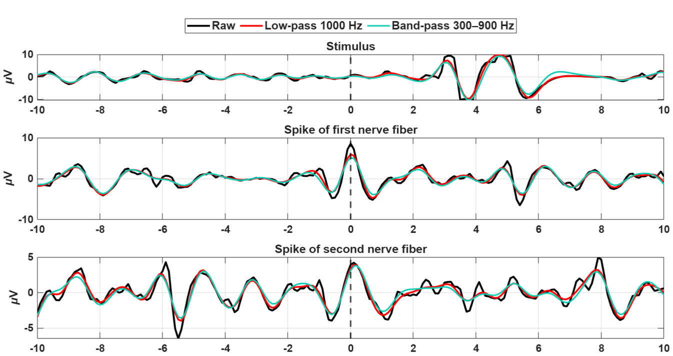

Fig. 2: Raw, low-pass and band-pass filtered signal around the a stimulus, and two types of nerve fiber responses.

D. Metrics

Using overlapping-window preprocessing, metrics defined in [9] were applied. Each window was initially treated as an independent prediction in a *window-wise* manner. Consecutive positive windows corresponding to the same spike were then *merged*, with the *proximity* metric removing merged predictions of only one window. The *latency* metric further ignored predictions inconsistent with expected response timing.

Given the extreme class imbalance, recall was used to evaluate detection performance, while false positives were quantified as the ratio per true positive (FP/TP).

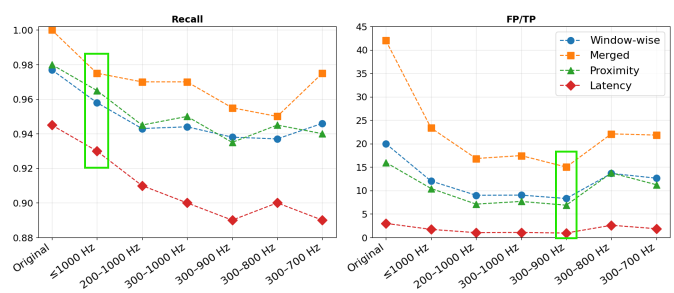

Fig. 4: Recall and FP/TP for different filtered signals using the introduced metrics, evaluated on the validation set after retraining. Green rectangles mark the best setting per subplot.

III. Results

The spectrogram (Fig. 1) shows that most neural energy is concentrated below 1000 Hz, with prominent bands between 300–700 Hz. A higher-frequency band around 3500 Hz was also visible, but including it in a multiband filter had no measurable impact on performance. Based on the 0–1000 Hz range, Fig. 4 compares model performance across different filtered signals using the metrics described in Sec. II-D. Among the filtered signals, the low-pass filter achieved the highest recall across all metrics, while the 300–900 Hz band-pass filter minimized the FP/TP ratio.

TABLE I: Performance metrics for the datasets using the original signal, the low-pass filtered (1000 Hz), and a band-pass filtered signal (300–900 Hz). Results are shown for retrained models (R) and for the baseline model trained only on the original data but evaluated on the filtered signals (E). An 80% decision threshold is used.

Dataset	Condition	Window-wise	Merged Windows	Proximity	Proximity + Latency
			RECALL		
Validation	**Original (R)**	**0.977**	**1.000**	**0.980**	**0.945**
Validation	Low-pass (R)	0.958	0.975	0.965	0.930
	Low-pass (E)	0.96	0.98	0.96	0.925
Validation	Band-pass (R)	0.938	0.955	0.935	0.890
	Band-pass (E)	0.945	0.97	0.95	0.9
Test	**Original (R)**	**0.957**	**0.983**	**0.967**	**0.943**
Test	Low-pass (R)	0.916	0.94	0.934	0.918
	Low-pass (E)	0.912	0.935	0.918	0.89
Test	Band-pass (R)	0.884	0.92	0.894	0.853
	Band-pass (E)	0.888	0.92	0.89	0.85
			FP / TP		
Validation	Original (R)	20	42	16	3
Validation	Low-pass (R)	12.04	23.37	10.45	1.72
	Low-pass (E)	11.3	22.0	8.76	1.46
Validation	Band-pass (R)	8.31	15.02	6.88	0.94
	Band-pass (E)	**8.15**	**14.94**	**6.37**	**0.81**
Test	**Original (R)**	**24**	**51**	**18**	**3**
Test	Low-pass (R)	13.17	25.44	11.08	1.42
	Low-pass (E)	11.78	23.37	8.66	1.21
Test	Band-pass (R)	9.47	**16.33**	7.51	0.84
	Band-pass (E)	**9.21**	16.57	**6.96**	**0.68**

This illustrates the deep-rooted trade-off between recall and FP/TP: improving one metric tends to compromise the other. Therefore, both filters were selected for further evaluation.

As summarized in Table I, retraining the model on filtered data (R) generally produced similar outcomes to evaluating the original model on filtered data only (E), with some cases performing better in the latter. While filtering caused a moderate decrease in recall, the FP/TP ratio dropped dramatically—by roughly 45% to 80%—demonstrating that the reduction in false positives far outweighs the loss in sensitivity.

Performance on the test set was slightly lower than on the validation set, which is expected given natural variability and the limited size of in vivo microneurography recordings. Despite this, the overall trends and improvements observed with filtering remain consistent.

IV. Conclusion

In this work, the frequency content of noisy microneurography recordings was analyzed, revealing that the majority of neural-related information is located below 1000 Hz, despite power being present up to 5000 Hz. Applying a low-pass filter at 1000 Hz removed approximately 80% of the spectral power while preserving the relevant neural signal. A selected band-pass filter, targeting 300–900 Hz, excluded around 88% of the total frequency content, further isolating the frequency range most informative for spike detection. The primary benefit of these filtering steps was a substantial reduction in false positives. Although to some extent, recall decreased as well, this reduction was relatively small compared to the improvement in false positives. In some cases, the FP/TP ratio dropped to just 22% of the value obtained with the unfiltered data. As the data is inherently filtered during recording, supposedly, re-annotating the signal would not change the spike occurrences. Future work could extend this approach by incorporating learnable filter parameters, enabling the pipeline to automatically identify cutoff frequencies that optimize model performance.

References

[1] C. K. Arcilla and P. Tadi, *Neuroanatomy, Unmyelinated Nerve Fibers*. Treasure Island, FL: StatPearls Publishing, 2023.

[2] B. Namer and A. Lampert, “Functional signatures of human somatosensory c fibers by microneurography,” *Pain*, vol. 166, no. 10, pp. 2220–2235, 10 2025.

[3] A. B. Vallbo and K.-E. Hagbarth, “Activity from skin mechanoreceptors recorded percutaneously in awake human subjects,” *Experimental Neurology*, vol. 21, no. 3, pp. 270–289, 1968.

[4] H. E. Torebjörk and R. G. Hallin, “Identification of afferent C units in intact human skin nerves,” *Brain Research*, vol. 67, no. 3, pp. 387–403, 1974.

[5] E. Kutafina, S. Becker, and B. Namer, “Measuring pain and nociception: Through the glasses of a computational scientist. transdisciplinary overview of methods,” *Frontiers in Network Physiology*, vol. 3, 02 2023.

[6] E. Kutafina, A. Troglio, R. de Col, R. Röhrig, P. Rossmanith, and B. Namer, “Decoding neuropathic pain: Can we predict fluctuations of propagation speed in stimulated peripheral nerve?” *Frontiers in Computational Neuroscience*, vol. 16, 2022.

[7] M. Schmelz, C. Forster, R. Schmidt, M. Ringkamp, H. O. Handwerker, and H. E. Torebjörk, “Delayed responses to electrical stimuli reflect C-fiber responsiveness in human microneurography,” *Experimental Brain Research*, vol. 104, no. 2, pp. 331–336, 1995.

[8] A. Troglio, P. Konradi, A. Fiebig, A. Pérez Garriga, R. Röhrig, J. Dunham, E. Kutafina, and B. Namer, “Supervised spike sorting feasibility of noisy single-electrode extracellular recordings: Systematic study of human c-nociceptors recorded via microneurography,” *PLOS ONE*, vol. 20, no. 9, pp. 1–28, 09 2025.

[9] D. Darabos, A. Troglio, A. Fiebig, A. Maxion, E. Kutafina, B. Namer, G. Bognár, and P. Kovács, “Explainable spike detection algorithm on c-fiber microneurography data using weighted hermite variable projection neural networks,” in *Proc. 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, July 2025, accepted.

[10] T. Dozsa, C. Bock, J. Meier, and P. Kovács, “Weighted Hermite variable projection networks for classifying visually evoked potentials,” *IEEE Transactions on Neural Networks and Learning Systems*, October 2024.