NeuroLibre
Reproducible Preprints

NeuroMOSAICS: A collection of neurostimulation datasets - Multi-scale Open-Source Across Interfaces Conditions & Species

DOI: 10.55458/neurolibre.00033

Interactive Version

- Living Preprint

Code

- Technical Screening
- Submitted Repository

Reproducibility Assets

- Repository
- Dataset
- Living Preprint
- Container

Moderator: Mathieu Boudreau
Screener(s):

- @agahkarakuzu

Submitted: 19 March 2025
Published: 19 June 2025

Lison Kardassevitch[1,2,3], **Davide Burchielli**[2,3,5], **Numa Dancause**[1,3,4], **and Marco Bonizzato**[1,2,3,4,5]

1 Université de Montréal, Montreal, QC, Canada **2** Mila - Quebec Artificial Intelligence Institute, Montreal, QC, Canada **3** CIRCA - Centre Interdisciplinaire de Recherche sur le Cerveau et l'Apprentissage, Montreal, QC, Canada **4** CIB - Biomedical Innovation Centre, Montreal, QC, Canada **5** Electrical Engineering, Polytechnique Montreal, Montreal, QC, Canada

Disclaimer

This pdf is intended for content registration purposes only! For full access and interactive reading of this publication, please visit the live version.

Introduction

Neurostimulation can be applied on a large panel of nervous system disorders and injuries by modulating motor, sensory and cognitive functions through electrical stimulation. This technology offers, amongst others, a promising method to restore motor functions in patients living with paralysis due to neurotrauma, with strong potential for enhanced efficacy through improved spatiotemporal precision (James et al., 2018; Ting et al., 2021). Neuromodulation systems are becoming increasingly complex (Anderson et al., 2018; Schüpbach et al., 2017), and this trend is expected to continue and scale further in the future (De Ridder et al., 2021). Advances in manufacturing, such as the miniaturization of electrodes enabling higher spatial selectivity, the development of biocompatible designs for targeting areas closer to the nervous system, and the integration of wireless systems providing greater freedom of movement, are key drivers of this evolution (Jun et al., 2017; Tsai et al., 2017; Won et al., 2020). Whether in large channel-count or smaller devices such as current clinical deep brain stimulation (DBS) leads, the effectiveness of neuromodulation depends on system identification and on fine-tuning numerous control parameters to meet individual needs, problems that advanced optimization algorithms have the potential to solve more efficiently than traditional approaches.

A.I. In neurostimulation optimization

A wide range of pre-clinical studies, including on epilepsy (Pineau et al., 2009) and Parkinson's disease (Gao et al., 2023) are actively exploring how to enhance neuromodulation results using diverse machine learning (ML) methods such as decision trees (De Andres et al., 2021; Gebodh & Bikson, 2021), random forests (Gebodh & Bikson, 2021; Shukla et al., 2014), neural networks (NN) (Gebodh & Bikson, 2021; Shukla et al., 2012) support vector machines (SVM) (De Andres et al., 2021; Gebodh & Bikson, 2021; Medaglia et al., 2020) and Reinforcement Learning (RL) (Gao et al., 2023; Pineau et al., 2009; Yu et al., 2021). However, those ML methods rely on pre-training and require a huge number of iterations to converge to a solution, making them currently impractical to use in an online context. Online neurostimulation

optimization means in vivo trial-and-error tuning of a neural interface, where decisions involve selecting from a large number of parameter options, and the constantly evolving environment induces both noise and non-stationarity in the response.

To date, real-time neurostimulation optimization in clinical studies is performed by adaptive stimulation (Arlotti et al., 2018; Rosa et al., 2015; Swann et al., 2018). This closed-loop method traditionally relies on turning on and off DBS and varying the amplitude based on neurophysiological feedback, but this approach is not tailored for a complex multivariate problem such as selecting what specific stimulation pattern to activate, among an immense number of options for varying contact, frequency, current amplitude and pulse width. To offer a solution adapted to optimization of multi parameters, some studies focus on a class of algorithms that does not require training and that relies on parametric models of stimulation properties. Bayesian Optimization based on Gaussian Process (GPBO) can offer an effective solution (Aiello et al., 2023; Bonizzato et al., 2023; Choinière et al., 2024; Laferriere et al., 2020; Sarikhani et al., 2022). GPBO tests single-input parameter combinations ("queries"), leveraging responses to build an evidence-based approximation ("surrogate" function) that models how stimulus choices affect the desired output. This approach allows to efficiently explore high dimensional spaces in a small number of iterations.

Optimizing stimulation for movement disorders

Restoring motor control after neurotrauma or in disease is a very complex neuromodulation challenge since the movement is controlled by distributed brain and spinal networks with variable dynamic states and outputs. Promising studies already show how advanced optimization methods can enhance mobility (Bonizzato et al., 2023; Grahn et al., 2014), (Losanno et al., 2021), but the field of motor neuroprosthetics is still young. Novel exciting demonstrations of motor control restoration feature complex neuroprosthetic systems, with a high degree of control complexity (Rowald et al., 2022; Wagner et al., 2018). Research to develop optimization agents to fine-tune such systems once they are implanted largely benefits from offline modeling and simulation (Capogrosso et al., 2013; Rowald et al., 2022).

To effectively train and test complex ML models or simulate online experimental neurostimulation sessions, a diverse range of neurostimulation datasets is essential (Bonizzato et al., 2023; Wang et al., 2023). Such datasets must account for variability across subjects, species, and neural interface targets. For example, rodents provide the advantage of generating large datasets due to relatively restrained costs, while NHPs offer a closer approximation to human brain complexity, making both species invaluable for research. The characteristics of any single dataset can vary widely due to environmental factors and experimental design, leading to significant differences not only between rodent and NHP data but also across neural targets such as the spinal cord and cortex; even within subjects of the same species.

Although synthetic data can serve as an initial benchmark for algorithm performance, it falls short in addressing the complexities inherent to in vivo data. Many neurophysiological mechanisms, such as the biological variability of motor responses, remain poorly modeled. This variability causes, among other things, inconsistencies in neural output to identical stimulation parameters and is compounded by noise introduced by experimental setups. These challenges make in vivo validation important, not only for assessing algorithm functionality but also for uncovering patterns of responses to neural stimuli and deepening our understanding of neural circuits response dynamics and non-stationarities. However, expanding the collection of neurostimulation datasets faces significant barriers, including high costs, technical complexity, and ethical considerations.

The recent surge in optimization algorithms within this field opens new possibilities for neurostimulation applications. To further push the boundaries of AI in neurostimulation, it is essential to broaden the dataset pool, thereby supporting a wider range of innovative use cases.

Our goal is to develop NeuroMOSAICS (Neurostimulation Multi-scale Open-Source Across

Interface, Conditions & Species), a collection of open-access datasets comprising a wide range of different neurostimulation scenarios: in rats and non-human primates (NHP), in sedated and awake animals, in intact and injured subjects, in small and large input spaces optimization, and in both cortical and spinal neurostimulation interventions. NeuroMOSAICS aims to support the development of neurostimulation optimization algorithms. For each dataset, we highlight common challenges that any optimization algorithm must address when applied in this context. Our goal is not only to provide a series of datasets, but also to shed light on these recurring challenges for anyone interested in advancing this field and accelerating the development of related technologies. This collection is an attempt to regroup and valorize all neurostimulation characterization datasets opened by the community. Beyond enabling the training and evaluation of data-demanding algorithms, this collaborative initiative also aims to make valuable knowledge more accessible, thereby encouraging the ML community to design solutions tailored to neurostimulation problems.

This living white paper intends to evolve in time with new incoming open-source datasets. New datasets will be added as we target additional unmet needs in neurostimulation optimization. Moreover, we aim to incorporate complex datasets that establish benchmarks and open avenues for novel algorithmic solutions. Datasets may also be introduced in response to community demand or as invited contributions. We welcome opportunities to review contributions from groups exploring high-dimensional optimization, to collaboratively grow this evolving body of work.

Materials and methods

In this section we describe how each featured dataset has been collected. We wrote this section with a focus on clarity and a level of readership accessible to those without a background in neuroscience.

Dataset 1 – Cortically evoked muscle responses: NHP forelimb and rat hindlimb

Important

Input: coordinates of an electrode in a 96-electrode intracortical array for NHP and 32-electrode intracortical array for rats.Output: Motor Evoked Potentials (MEPs) from 4 to 8 muscles, forearm and hand muscles for NHP and leg muscles for rat. Replicates: Data were collected from N=4 NHP and N=6 rats (biological replicate). For each independent muscle, technical replicates ranged from 10 to 28 repetitions. Depth: The same neurostimulation is repeated 10 to 28 times per site (muscle) per subject, resulting in an equivalent number of MEP measurements being collected. Each MEP represents the muscle response within a time window of -100 to 200 ms relative to the stimulation delivery (0 ms). Rationale: To develop optimization algorithms to efficiently optimize neurostimulation within a set of electrode location options.

This dataset contains limb EMG responses to cortical stimulation in 4 male NHPs (2 macaques, 2 capuchins) and 6 female rats. The dataset, available on OSF, was collected to develop a Gaussian Process Bayesian Optimization (GPBO) framework, an algorithm designed to optimize neurostimulation parameters for maximizing MEPs. Cross-species data can challenge the robustness and adaptability of optimization algorithms, testing their ability to generalize across scales and from simpler to more complex cortical organizations, that approximate better the human brain.

The detailed protocol for surgery and experiment is available on M. Bonizzato et al. (Bonizzato et al., 2023; Choinière et al., 2024). The rats were implanted with 32-microelectrode arrays in the hindlimb sensorimotor cortex and the NHPs were implanted with 96-electrode arrays

implanted in the hand area of the primary motor cortex (M1). For rats, the EMGs were implanted in leg muscles. For NHPs, the EMGs were implanted in hand and finger muscles. When the stimulation protocol was running, the two capuchin NHP were under sedation and the two macaques were awake. In the next sections of this preprint, we will refer to **Dataset 1** as the NHP results; however, the rat outputs follow the same underlying concept.

Dataset 2 – Forelimb muscle responses to spinal stimulation in rats

> **Important**
>
> Dataset 2 * Input: coordinates of an electrode in an epidural 64-electrode array. * Output: motor evoked potentials from 8 to 10 arm muscles. * Replicates: Data were collected from N = 9 biological replicates, defined as the number of distinct physical implants performed across 4 rats. Each implant corresponds to a separate array placement on the spinal cord between levels C3 and C6. For each array position, 10 to 15 repetitions of single-pulse stimulation were delivered to each electrode in random order using Tucker-Davis Technologies (TDT) equipment. * Depth: The same neurostimulation is repeated 10 to 15 times per site (muscle) per subject, resulting in an equivalent number of MEP measurements being collected. Each MEP represents the muscle response within a time window of -35 to 40 ms relative to the stimulation delivery (0 ms). * Rationale: To develop optimization algorithms to efficiently optimize neurostimulation within a set of electrode location options and multiple concurrent output muscle objectives (8-10 bilateral and varied agonist/antagonist muscles).

This dataset contains EMG responses to spinal cord stimulation in the forelimbs of 4 female Long Evans rats. EMG electrodes were implanted bilaterally in the Biceps, Triceps, Flexor Carpi Radialis, Deltoid, and Dorso muscles. Spinal cord stimulation was delivered via a 64-electrode array placed epidurally.

- Experiment Description

A 64-electrode array was placed epidurally on the cervical spinal cord (C3–C6) in 9 distinct placements across 4 sedated rats. In each animal, up to three array positions were tested within this spinal segment. EMG activity was recorded from muscles in both forelimbs in response to single-pulse stimulation delivered through the array.

- Surgical Procedure

The experiments were terminal. Rats were anesthetized with urethane (1 g/kg) prior to surgery. After 30 minutes, they were transitioned to isoflurane anesthesia (2%). Isoflurane was withdrawn prior to stimulation, and the rat maintained under the initial dose of urethane. Vital signs, temperature, and hydration were continuously monitored.

EMG Implants EMG microwires were implanted in the Biceps, Triceps, Flexor Carpi Radialis, Deltoid, and Dorso muscles on both sides. Correct placement was confirmed by electrical stimulation through the wires and observation of evoked movements. Wires were inserted transcutaneously.

Spinal Cord Implants The dura mater was exposed from C3 to C6. Initially, a single electrode was used to determine muscle response thresholds. A 64-electrode array was then placed epidurally at up to three different locations.

- Data Collection

For each placement, 10 repetitions of single-pulse stimulation were delivered to each electrode in random order using Tucker-Davis Technologies (TDT) equipment. Stimulation amplitude was set to the threshold value identified during single-electrode testing.

Processing

The raw signals in both datasets have been rectified.

Important

In addition to the raw data, a filtered version of the signal is provided in this dataset. The following processing steps were applied in order: - Blanking from stimulation artefact; - Band-pass filtering (low cut at 70Hz, high cut at 1000Hz); - Notch filtering at 60Hz; - Full-wave rectification.

Admonition

Sometimes, the EMG baseline signal may appear elevated above zero due to noise or low-frequency drift. While NeuroMOSAICS open-source datasets include such baseline variations, we applied an additional preprocessing step to the signals shown in the figure to normalize the baseline to zero. To do so, the mean value of a 30 ms pre-stimulation window was computed for each trial and subtracted from the corresponding EMG signal. This normalization ensures consistent baseline levels across channels and reduces the risk of misinterpretation, for instance, mistaking baseline offsets (e.g., 2 mV vs. 0.5 mV) as differences in signal quality or responsiveness.

Results

In this section, we frequently refer to the heatmaps in each Figure. Each heatmap corresponds to a specific muscle from a subject and displays the average quantification of motor evoked potentials (MEPs) across different electrodes. Those electrode arrays allow us to map the topographic organization of the implanted motor region. Typically, the best-performing electrodes cluster spatially. Each query refers to an individual stimulation.

Dataset 1

Dataset 1 was collected during in vivo sessions that involved repeated stimulations of electrode locations in M1. These stimulation sessions were initially conducted on sedated capuchins and on awake rats and macaques. The average MEP responses across spatial locations show a topographic organization of M1 such that nearby cortical locations have similar MEP patterns (Figure A.). While identifying the optimal electrode in this condition may seem straightforward, it presents significant challenges. The heatmaps represent averages of MEP, but individual responses are influenced by excitability, noise factors, and possibly non-stationarity, leading to considerable variability. This variability can obscure important details at the single-trial level. In Figure B., the upper-left panel displays the average response across multiple stimulations delivered through a specific electrode, providing a global overview of the evoked activity. In contrast, the two other panels offer a more detailed perspective : the distribution of peak amplitudes reveals the consistency of responses across trials, while the stack of individual traces allows visualization of how responses evolve over time. Given the goal of minimizing the number of queries, the task becomes finding the right balance between using enough queries for a robust solution and limiting the search to minimize stimulation.

Figure 1: Interactive content placeholder

Please see the living preprint to interact with this figure.

A. Heatmap of motor responses across a 96-electrode array in M1 of an anesthetized Cebus: This panel displays the mean motor responses for a selected muscle, averaged across repetitions, for each electrode in a 96-electrode array (10x10 grid with 4 unused electrodes). To view data for a specific muscle, select the desired muscle from the list on the right. The available muscles are: Flexor Carpi Ulnaris (FCU), Extensor Digitorum Communis (EDC), Extensor Carpi Radialis (ECR), Opponens Pollicis (OP), Flexor Pollicis Brevis (FPB), and Adductor Pollicis (AP). B. Detailed Motor Evoked Potential (MEP) response for a single electrode : This panel provides detailed visualizations for the MEP response corresponding to a selected electrode from the heatmap in A (click on the corresponding cell to select it). - Top Left: The mean MEP waveform across all repetitions. - Top Right: The distribution of peak amplitudes for the selected electrode. - Bottom Left: A stack plot showing the rectified raw MEP waveforms for individual repetitions. Red color hilights outliered trials.

The clear topographic organization of the responses observed in sedated NHPs become considerably noisier and more selective in awake NHPs (see Figure A.). These NHP implants were older, and the quality of the electrical interface at the EMG wire-muscle contact can degrade over time, potentially contributing to variations in noise levels in the recorded muscle responses across subjects. We can also identify two ranges of electrodes (column 8 and 9 in Figure A.) that consistently exhibit very low signal amplitudes across all muscles. This suggests that these portions of the electrode matrix may not have been in full contact with the tissue. Additionally, the NHPs were awake, which may have further contaminated MEPs with spontaneous muscle activations, which are possibly present even when the condition of no muscle activity pre-stimulation is enforced. Consequently, this dataset provides an opportunity to test optimization approaches on data that differ significantly from the previous dataset.

Figure 2: Interactive content placeholder

Please see the living preprint to interact with this figure.

A. Heatmap of motor responses across a 96-electrode array in M1 of an awaken Macaque : This panel displays the mean motor responses for a selected muscle, averaged across repetitions, for each electrode in a 96-electrode array (10x10 grid with 4 unused electrodes). To view data for a specific muscle, select the desired muscle from the list on the right. The available muscles are: First Dorsal Interosseous (FDI), Flexor Digitorum Superficialis (FDS), Extensor Carpi Radialis (ECR), Extensor Digitorum Communis (EDC). The EDC muscle exhibits only noisy signals and does not contain clear evoked responses. However, it serves as a useful example of how a single stimulation contaminated by an artefact can significantly distort the average response. This highlights the importance of analyzing multiple repetitions of the same stimulation to ensure that the averaged signal accurately reflects true underlying activity rather than artefactual noise. B. Detailed Motor Evoked Potential (MEP) response for a single electrode : This panel provides detailed visualizations for the MEP response corresponding to a selected electrode from the heatmap in A (click on the corresponding cell to select it). - Top Left: The mean MEP waveform across all repetitions. - Top Right: The distribution of peak amplitudes for the selected electrode. - Bottom Left: A stack plot showing the rectified raw MEP waveforms for individual repetitions. Red color hilights outliered trials.

High selectivity becomes particularly problematic when the search space is expanded (such as by adding additional parameters like stimulation amplitude and frequency) and even more so

in noisy conditions. If it makes it harder to identify the best electrode, the global maximum of the function to optimize, and increases the risk of getting stuck in local maxima during optimization. For example, while searching for the electrode that grants maximal response in the First Dorsal Interosseous muscle (see Figure A.), one could easily get stuck at coordinates `(7,7)` on the heatmap. If the algorithm hasn't yet queried the optimal electrode at coordinates `(4,1)`, and instead queries `(7,7)` at a given query q, it will produce the best result thus far. The algorithm would then continue to exploit the local maximum around `(7,7)` without knowing that a better option exists unless it explores very specific electrodes `(4,1)` or `(5,1)` (see Figure B.). With around 50 electrodes to query in a limited number of stimulations, the algorithm's exploration strategy becomes critical to avoid local maxima. This issue of local maxima and high selectivity is also observed in awake rats, as seen in the Left Medial Gastrocnemius available in **Dataset 1** on OSF. However, the smaller array (32 electrodes, compared to 96 for NHPs) reduces the complexity of the search.

The temporal aspect of the responses can play an important role in the nature of the response. In any scenario from this dataset, the stronger MEP are peaking at the same time as we can observe in the stacked MEP where deeper blue lines are vertically crossing all the trials.

Considering noise (such as muscle activations unrelated to stimulation), some responses may convey misleading information about the overall response pattern. A clear example is provided by the EDC muscle in Figure . The response map indicates a high activity level for electrode `(0,6)`. However, upon closer examination of the stacked EMG traces, it becomes evident that this apparent response is driven almost entirely by a single trial, the 18th stimulation, which shows an isolated increase in MEP amplitude. Notably, this sharp deflection is also visible in other channels, such as electrode`(7,2)`, but occurs at a time (~170 ms post-stimulation) that is unrelated to the actual stimulation event. This suggests the presence of a transient artefact, underscoring the importance of examining individual trials alongside average responses to avoid misinterpretation.

Identifying these outliers in real time is challenging, but it is valuable to do so afterward as it raises awareness of their potential impact on optimization processes. The parameter `is_valid` classifies outliers in two categories: those that can be excluded due to excessive EMG activity prior to stimulation and those that cannot be excluded, where no prior EMG activity was observed, but the response amplitude is unusually high (such as the 18th trial of electrode `(0,6)` in ECD muscle, Figure). Accurate online processing can easily reject the first of the two classes, while the second is unlikely to be rejected in a real use case. The first category is highlighted in red in the stack MEP figures. The mean and standard deviation of responses in **Dataset 1** still include these outliers. A data scientist wishing to remove these spurious responses from the optimization "ground truth", or objective value, must though recalculate the mean responses.

Admonition

The collection of Dataset 1 was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC; RGPIN-2022-04210).

Dataset 2

This dataset was collected from four anesthetized rats, resulting in smooth topographic organization of the spinal cord with high variance for the average responses. Similar to capuchin in **Dataset 1**, the relaxed muscle state in sedated animals, with no intention of motor control or reflex interference, contributes to the smooth responses. However, unlike **Dataset 1**, this dataset involves spinal cord (spinal cord) stimulation, leading to some key differences in the results. One notable distinction is that the high-response electrodes are more widely distributed spatially, which may be attributed to the different spatial recruitment properties of spinal cord stimulation compared to motor cortex stimulation. Because dorsal spinal cord

stimulation is, from a circuit perspective, closer to the muscle target than M1 stimulation, there is less variability in the responses. Indeed, less synapses (and thus less pulses) lead to a higher repeatability of the responses. This very clear repeatability (a straight deep blue line across each trial in stack MEP figures) makes the identification of outliers easier, and they are also listed in a parameter is_valid and highlighted in red in the stack MEP figures. In this case, the mean and the standard deviation were computed without those outliers.

Figure 3: Interactive content placeholder

Please see the living preprint to interact with this figure.

A. Heatmap of motor responses across a 64-electrode array placed on C5 in an anesthetized rat : This panel displays the mean motor responses for a selected muscle, averaged across repetitions, for each electrode in a 64-electrode array (8x8 grid with 4 unused electrodes). To view data for a specific muscle, select the desired muscle from the list on the right. The available muscles are: Right Triceps (R Tr), Right Biceps (R Bi), Right Torso (R Trs), Right Wrist flexor (R WrFlex), Right Spinal Deltoid (R SpDel), Left Triceps (L Tr), Left Biceps (L Bi), Left Tors (L Trs), Left Wrist flexor (L WrFlex), Left Spinal Deltoid (L SpDel). B. Detailed Motor Evoked Potential (MEP) response for a single electrode : This panel provides detailed visualizations for the MEP response corresponding to a selected electrode from the heatmap in A (click on the corresponding cell to select it). - Top Left: A stack plot showing the rectified raw MEP waveforms for individual repetitions. Red color hilights outliered trials. - Top Right: A stack plot showing the filtered MEP waveforms for individual repetitions. - Bottom Left: The mean MEP waveform across all repetitions. - Bottom Right: The distribution of peak amplitudes for the selected electrode.

Whereas the clear repeatability from this dataset has a positive impact on the optimization process, the excitability and noise can be, in some cases, particularly challenging. Indeed, the excitability can increase and decrease drastically in a few queries (like for electrode `(2,4)` in Left Biceps (see Figure B.)), and some electrodes are losing contact on half of the process (like in electrode `(4,3)` in Left Biceps), making it harder to identify a best electrode and to take a decision on the definition of what is the best electrode. This phenomenon suggests potential future use of this dataset to take into account the non-stationarity of the responses.

The higher variance in this dataset tends to blur the distinction between the best average solution and its neighboring electrodes. This makes it easy to identify the best location area, but also more challenging to point out the precise optimal electrode. As the variability in the responses increases, it becomes more likely that, at query q, the best electrode yields a lower response than nearby ones. For instance, in the Right Biceps muscle (see Figure A.), the peak amplitude distribution from the best electrode (coordinates `(2,5)`) is centered at `13.1 mV`, but there is an outlier stimulation yielding a response of `11.2 mV`. Meanwhile, the electrode at coordinates `(2,4)` has stimulations with outputs around `13.1 mV`. If an algorithm repeatedly queries electrode `(2,4)` and obtains consecutively good motor responses, while a single query of electrode `(2,5)` produces an outlier of `11.2 mV`, the algorithm might incorrectly infer that electrode `(2,4)` is superior. Although `(2,4)` is not the global maximum, it remains a strong candidate because it provides stable, high-quality responses and is situated within the highest region of the heatmap.

Another key insight from this dataset is the spatial organization of the spinal cord for motor control. There is a clear lateralization, with the right side of the spinal cord predominantly evoking stronger responses in muscles on the right side of the body, and the left side evoking responses in left-side muscles. Additionally, there appears to be spatial clustering in spinal cord sites evoking motor potential in specific muscles. The spatial organization is non-trivial given that the dorsal aspect of the spinal cord, which we stimulated, is associated with afferent pathways, and motor output depends on reflex responses (Capogrosso et al., 2013). This

insight has significant implications for optimization in neuromodulation. The function linking stimulation parameters to motor responses is inherently complex and unknown, posing a "black-box" optimization problem. In order to facilitate the optimization task, incorporating prior knowledge (Guay-Hottin et al., 2025) about the spatial organization of the spinal cord can significantly enhance the efficiency of optimization algorithms, such as BO, that is designed to leverage prior information to reduce the complexity of the search space and accelerate convergence to optimal solutions (Greenhill et al., 2020). For example, integrating the knowledge that stimulating the left (or right) dorsal spinal cord enhances motor responses on the corresponding side of the body allows the algorithm to narrow its spatial exploration. This targeted approach can substantially decrease the number of iterations required to identify the optimal stimulation parameters, improving both the speed and accuracy of the optimization process.

Admonition

Funding The collection of Dataset 2 was supported by New Frontiers in Research Fund Exploration (NFRFE-2022-00394).

Conclusion

In this living white paper, we introduced NeuroMOSAICS, an open-access, reproducible and diverse dataset collection covering neurostimulation across different species, neural interfaces and search space dimensions. The two scenarios presented in this first version support optimization research in diverse tasks and objectives. The first dataset provides EMG responses from cortical stimulation in both awake and sedated NHPs and rats, while the second dataset focuses on bilateral EMG responses from spinal cord stimulation in sedated rats.

This collection is an attempt to bring together open neurostimulation characterization datasets shared by the community. These datasets, often difficult to collect due to cost, ethical, and technical constraints, are typically shared only on a limited scale. By making them more widely available, we aim to support researchers in developing and validating complex optimization and data-intensive algorithms tailored to neurostimulation challenges. We aim to foster collaboration between neurostimulation researchers and the ML community, encouraging the creation of algorithms better suited to concrete neurostimulation problems.

Availability

The datasets presented in this living white paper are openly available on the platform OSF: **Dataset 1** (M1 stimulations on NHPs) : https://osf.io/54vhx/, reference number : `54vhx`, **Dataset 1 bis** (M1 stimulations on rats) : A similar dataset that **Dataset 1** but conducted on rats is also available at https://osf.io/54vhx/, reference number : `54vhx`, **Dataset 2** (spinal cord stimulations on rats) : https://osf.io/kpb7x/, reference number : `kpb7x`. Additional datasets will be progressively integrated into this living white paper. For each dataset already included, this paper is augmented: a supporting text is provided to clarify its relevance and potential applications, along with a direct link to access the data.

CRediT authorship contribution statement

- **Lison Kardassevitch:** Writing – review & editing, Writing – original draft, Visualization, Project administration, Methodology, Formal analysis, Conceptualization.
- **Davide Burchielli:** Writing – review & editing.

- **Numa Dancause:** Writing – review & editing, Supervision.
- **Marco Bonizzato:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

Funding sources

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC; RGPIN-2023-04370) and New Frontiers in Research Fund Exploration (NFRFE-2022-00394).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors extend their gratitude to all scholars who have openly shared datasets for neurostimulation optimization, fostering the growth of the community. This includes the contributors who released **Dataset 1** in 2021. NeuroMOSAICS owes its existence to a steadfast community commitment to open science.

Aiello, G., Valle, G., & Raspopovic, S. (2023). Recalibration of neuromodulation parameters in neural implants with adaptive bayesian optimization. *Journal of Neural Engineering*, *20*(2), 026037. https://doi.org/10.1088/1741-2552/acc975

Anderson, D. N., Osting, B., Vorwerk, J., Dorval, A. D., & Butson, C. R. (2018). Optimized programming algorithm for cylindrical and directional deep brain stimulation electrodes. *Journal of Neural Engineering*, *15*(2), 026005. https://doi.org/10.1088/1741-2552/aaa14b

Arlotti, M., Marceglia, S., Foffani, G., Volkmann, J., Lozano, A. M., Moro, E., Cogiamanian, F., Prenassi, M., Bocci, T., Cortese, F., & others. (2018). Eight-hours adaptive deep brain stimulation in patients with parkinson disease. *Neurology*, *90*(11), e971–e976. https://doi.org/10.1212/wnl.0000000000005121

Bonizzato, M., Hottin, R. G., Côté, S. L., Massai, E., Choinière, L., Macar, U., Laferrière, S., Sirpal, P., Quessy, S., Lajoie, G., & others. (2023). Autonomous optimization of neuroprosthetic stimulation parameters that drive the motor cortex and spinal cord outputs in rats and monkeys. *Cell Reports Medicine*, *4*(4). https://doi.org/10.1016/j.xcrm.2023.101008

Capogrosso, M., Wenger, N., Raspopovic, S., Musienko, P., Beauparlant, J., Luciani, L. B., Courtine, G., & Micera, S. (2013). A computational model for epidural electrical stimulation of spinal sensorimotor circuits. *Journal of Neuroscience*, *33*(49), 19326–19340. https://doi.org/10.1523/jneurosci.1688-13.2013

Choinière, L., Guay-Hottin, R., Picard, R., Lajoie, G., Bonizzato, M., & Dancause, N. (2024). Gaussian-process-based bayesian optimization for neurostimulation interventions in rats. *STAR Protocols*, *5*(1), 102885. https://doi.org/10.1016/j.xpro.2024.102885

De Andres, J., Ten-Esteve, A., Harutyunyan, A., Romero-Garcia, C. S., Fabregat-Cid, G., Asensio-Samper, J. M., Alberich-Bayarri, A., Marti-Bonmati, L., & others. (2021). Predictive clinical decision support system using machine learning and imaging biomarkers in patients with neurostimulation therapy: A pilot study. *Pain Physician*, *24*(8), E1279.

De Ridder, D., Maciaczyk, J., & Vanneste, S. (2021). The future of neuromodulation: Smart neuromodulation. *Expert Review of Medical Devices*, *18*(4), 307–317. https://doi.org/10.1080/17434440.2021.1909470

Gao, Q., Schmidt, S. L., Chowdhury, A., Feng, G., Peters, J. J., Genty, K., Grill, W. M., Turner, D. A., & Pajic, M. (2023). Offline learning of closed-loop deep brain stimulation controllers for parkinson disease treatment. *Proceedings of the ACM/IEEE 14th International Conference on Cyber-Physical Systems (with CPS-IoT Week 2023)*, 44–55. https://doi.org/10.1145/3576841.3585925

Gebodh, N., & Bikson, M. (2021). A novel approach to closed-loop neuromodulation with machine learning. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *14*(6), 1662. https://doi.org/10.1016/j.brs.2021.10.235

Grahn, P. J., Mallory, G. W., Berry, B. M., Hachmann, J. T., Lobel, D. A., & Lujan, J. L. (2014). Restoration of motor function following spinal cord injury via optimal control of intraspinal microstimulation: Toward a next generation closed-loop neural prosthesis. *Frontiers in Neuroscience*, *8*, 296. https://doi.org/10.3389/fnins.2014.00296

Greenhill, S., Rana, S., Gupta, S., Vellanki, P., & Venkatesh, S. (2020). Bayesian optimization for adaptive experimental design: A review. *IEEE Access*, *8*, 13937–13948. https://doi.org/10.1109/access.2020.2966228

Guay-Hottin, R., Kardassevitch, L., Pham, H., Lajoie, G., & Bonizzato, M. (2025). Robust prior-biased acquisition function for human-in-the-loop bayesian optimization. *Knowledge-Based Systems*, 113039. https://doi.org/10.1016/j.knosys.2025.113039

James, N. D., McMahon, S. B., Field-Fote, E. C., & Bradbury, E. J. (2018). Neuromodulation in the restoration of function after spinal cord injury. *The Lancet Neurology*, *17*(10), 905–917. https://doi.org/10.1016/s1474-4422(18)30287-4

Jun, J. J., Steinmetz, N. A., Siegle, J. H., Denman, D. J., Bauza, M., Barbarits, B., Lee, A. K., Anastassiou, C. A., Andrei, A., Aydın, Ç., & others. (2017). Fully integrated silicon probes for high-density recording of neural activity. *Nature*, *551*(7679), 232–236. https://doi.org/10.1038/nature24636

Laferriere, S., Bonizzato, M., Côté, S. L., Dancause, N., & Lajoie, G. (2020). Hierarchical bayesian optimization of spatiotemporal neurostimulations for targeted motor outputs. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *28*(6), 1452–1460. https://doi.org/10.1109/tnsre.2020.2987001

Losanno, E., Badi, M., Wurth, S., Borgognon, S., Courtine, G., Capogrosso, M., Rouiller, E., & Micera, S. (2021). Bayesian optimization of peripheral intraneural stimulation protocols to evoke distal limb movements. *Journal of Neural Engineering*, *18*(6), 066046. https://doi.org/10.1088/1741-2552/ac3f6c

Medaglia, J. D., Erickson, B., Zimmerman, J., & Kelkar, A. (2020). Personalizing neuromodulation. *International Journal of Psychophysiology*, *154*, 101–110. https://doi.org/10.1016/j.ijpsycho.2019.01.002

Pineau, J., Guez, A., Vincent, R., Panuccio, G., & Avoli, M. (2009). Treating epilepsy via adaptive neurostimulation: A reinforcement learning approach. *International Journal of Neural Systems*, *19*(04), 227–240. https://doi.org/10.1142/s0129065709001987

Rosa, M., Arlotti, M., Ardolino, G., Cogiamanian, F., Marceglia, S., Di Fonzo, A., Cortese, F., Rampini, P. M., & Priori, A. (2015). Adaptive deep brain stimulation in a freely moving parkinsonian patient. *Movement Disorders*, *30*(7), 1003. https://doi.org/10.1002/mds.26241

Rowald, A., Komi, S., Demesmaeker, R., Baaklini, E., Hernandez-Charpak, S. D., Paoles, E., Montanaro, H., Cassara, A., Becce, F., Lloyd, B., & others. (2022). Activity-dependent spinal cord neuromodulation rapidly restores trunk and leg motor functions after complete paralysis. *Nature Medicine*, *28*(2), 260–271. https://doi.org/10.1038/s41591-021-01663-5

Sarikhani, P., Ferleger, B., Mitchell, K., Ostrem, J., Herron, J., Mahmoudi, B., & Miocinovic, S. (2022). Automated deep brain stimulation programming with safety constraints for tremor suppression in patients with parkinson's disease and essential tremor. *Journal of Neural Engineering*, *19*(4), 046042. https://doi.org/10.1088/1741-2552/ac86a2

Schüpbach, M., Chabardes, S., Matthies, C., Pollo, C., Steigerwald, F., Timmermann, L., Visser Vandewalle, V., Volkmann, J., & Schuurman, P. R. (2017). Directional leads for deep brain stimulation: Opportunities and challenges. *Movement Disorders*, *32*(10), 1371–1375. https://doi.org/10.1002/mds.27096

Shukla, P., Basu, I., Graupe, D., Tuninetti, D., & Slavin, K. V. (2012). A neural network-based design of an on-off adaptive control for deep brain stimulation in movement disorders. *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 4140–4143. https://doi.org/10.1109/embc.2012.6346878

Shukla, P., Basu, I., & Tuninetti, D. (2014). Towards closed-loop deep brain stimulation: Decision tree-based essential tremor patient's state classifier and tremor reappearance predictor. *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2605–2608. https://doi.org/10.1109/embc.2014.6944156

Swann, N. C., De Hemptinne, C., Thompson, M. C., Miocinovic, S., Miller, A. M., Ostrem, J. L., Chizeck, H. J., Starr, P. A., & others. (2018). Adaptive deep brain stimulation for parkinson's disease using motor cortex sensing. *Journal of Neural Engineering*, *15*(4), 046006. https://doi.org/10.1088/1741-2552/aabc9b

Ting, W. K.-C., Fadul, F. A.-R., Fecteau, S., & Ethier, C. (2021). Neurostimulation for stroke rehabilitation. *Frontiers in Neuroscience*, *15*, 649459. https://doi.org/10.3389/fnins.2021.649459

Tsai, D., Sawyer, D., Bradd, A., Yuste, R., & Shepard, K. L. (2017). A very large-scale microelectrode array for cellular-resolution electrophysiology. *Nature Communications*, *8*(1), 1802. https://doi.org/10.1038/s41467-017-02009-x

Wagner, F. B., Mignardot, J.-B., Le Goff-Mignardot, C. G., Demesmaeker, R., Komi, S., Capogrosso, M., Rowald, A., Seáñez, I., Caban, M., Pirondini, E., & others. (2018). Targeted neurotechnology restores walking in humans with spinal cord injury. *Nature*, *563*(7729), 65–71. https://doi.org/10.1038/s41586-018-0649-2

Wang, D. X., Ng, N., Seger, S. E., Ekstrom, A. D., Kriegel, J. L., & Lega, B. C. (2023). Machine learning classifiers for electrode selection in the design of closed-loop neuromodulation devices for episodic memory improvement. *Cerebral Cortex*, *33*(13), 8150–8163. https://doi.org/10.1093/cercor/bhad105

Won, S. M., Song, E., Reeder, J. T., & Rogers, J. A. (2020). Emerging modalities and implantable technologies for neuromodulation. *Cell*, *181*(1), 115–135. https://doi.org/10.1016/j.cell.2020.02.054

Yu, C., Liu, J., Nemati, S., & Yin, G. (2021). Reinforcement learning in healthcare: A survey. *ACM Computing Surveys (CSUR)*, *55*(1), 1–36. https://doi.org/10.1145/3477600