A Feature-Generalizable Technique for Neural Conditioning

Jack D. Carson
Booker T. Washington High School, Tulsa, Oklahoma

Abstract

Advances in brain-computer interface technology seek to enable nuanced control of neural activity patterns. This study demonstrates a framework combining artificial intelligence and high-resolution optical sensing to identify and modulate intricate signaling dynamics. Genetically-encoded calcium indicators provide fluorescent readouts of neuronal firing. An autoencoder neural network compresses these optical data into a compact latent space, extracting interpretable features. The novel software *Rasa* coordinates data flow, deploying the models to translate neural sequences into vector representations. By comparing latent vectors of target and observed activity, *Rasa* identifies desired patterns and administers neurofeedback accordingly. Initial *in vivo* validation demonstrates increased activity of simple correlated firing patterns under this novel stimulation paradigm versus unmodulated recordings in rat motor cortices. While limited by experimental constraints, these preliminary results highlight the potential of integrated machine learning techniques and fine-grained optical sensing to reinforce complex behavior. Looking forward, such AI systems could unlock new therapeutic abilities to remedy dysfunctional neurological signaling underlying disease states.

1 Introduction

As scientific understanding of functional neural activity increases, researchers have begun to take an interest in the treatment of the human brain as an input-output (I/O) computational machine [1]. This interest arises from the neurological capacity for conditioning that the brain uses in adapting to its physical environment. A major question in the field of neuroengineering is to what extent neural plasticity can be controlled via an external electrical stimulus [2] [3]. This question informs methods in closed-loop neurofeedback, a technique relying on electrical stimulation of dopaminergic pathways in the brain to reinforce synaptic potentiation [4]. It is believed that this method can be used to modulate activity of neuron groups [4]. The psychological implications of this have made the modulation of neural circuits a candidate of great interest for therapeutic treatments [5]. However, the ability of closed-loop neurofeedback to be demonstrated on a cellular scale, where it could have the most powerful impact for therapeutic treatment, has been inhibited previously by the difficulty of accurately providing feedback to a target and then training that target neuron group on signals of any non-trivial informational value.

Current neuroimaging and neurophysiological sensors, such as electroencephalography (EEG) and electrocorticography (ECoG), expand our ability to model brain activity, but both technologies fall short in capturing specific activity with high spatiotemporal accuracy, instead measuring the summation of synchronized postsynaptic potentials over a large electrode application region [6] [7]. Furthermore, both ECoG and EEG confound haemodynamic methods of analysis such as fMRI, a valuable tool for mapping the brain activity on a global, albeit lower-resolution, scale. To address this, calcium sensors based on genetically embedded fluorescent proteins such as GCaMP have emerged as a potential alternative for capturing neurological activity at finer scales [8]. The data generated from calcium sensors can form the foundation of a neurofeedback output stream.

In the wake of recent advancements in neuroscience, there is also a growing consensus that electrical deep brain stimulation, apart from acting as a data receiver from neural sensors, can potentially serve as a robust input channel for neurofeedback and thus form an integral component of a Brain-Machine Interface (BMI). This process commences with an observation of neural activity, which is subjected to online analysis to provide a requisite stimulation [9] [10] [11]. However, the current neurofeedback techniques are bound by both the complexity of the conditionable signals and also the experimental difficulty of conditioning, which together conspire to severely inhibit the encoding of explicit brain signals and confine conditioning to simplistic binary-enforceable signals (i.e., Do we observe a spike?: Yes/No) [12] [13].

This study seeks to challenge these limitations, targeting the complexity of encodable signals and the rate at which these signals can be trained. I propose to optimize the training of neural signals by defining a continuous similarity function that can be used to robustly evaluate the closeness of two arbitrary signals. I use an autoencoder neural model to define a vector representation of a particular neural signal as the values of latent variables at the bottleneck of an autoencoder model. Through comparing the distance of a target and active vector in the latent space of an autoencoder model, we were able to define an optimizable point to approach for conditioning. This technique aims to amplify the diversity and complexity of encodable signals within the existing feedback framework, while also enhancing the precision and speed of training for even simple signals. In tandem with this, I introduce a new software package, developed in the Rust programming language and specifically crafted to address the glaring issues of data latency, integrity, and performance encountered with existing data logging and real-time visualization software. By integrating the new software with the innovative optimization technique, this study aims to introduce a refined, comprehensive solution to the existing limitations in neurofeedback and conditionable signal complexity.

2 Ethics

The study of how active electric feedback relates to neural potentiation requires the use of living brain tissue. As the purpose of our research relates to the possible therapeutic implications of closed-loop neurofeedback, laboratory rats (*Rattus norvegicus domestica*) were obtained from Braintree Scientific. These rats serve as ideal candidates for feedback due to a neural anatomy similar to humans, as well as detailed existing maps of rodent neural potentiation. Due to the surgical and sensor limitations attendant to small animals and invertebrates, there is extensive precedent for neurofeedback on this species [14]. Indeed, *Rattus norvegicus domestica* has been used as a model organism for investigating neurofeedback since the 1960s [14]. Four female rats were kept in accordance with guidelines set by the Institutional Animal Care and Use Committee of the Massachusetts Institute of Technology. For the purposes of this study, feedback was administered to a single subject, six months in age, 250 grams in weight, and approximately 18 cm in length. After all procedures, adequate pain relief was granted to the subject, and sufficient rest time was allowed between operations and testing. This study did not require the use of any additional subjects or feedback experimentation than was already approved for a large-scale feedback study on rat motor cortex modulation that is unrelated to the computational model set forth in this paper.

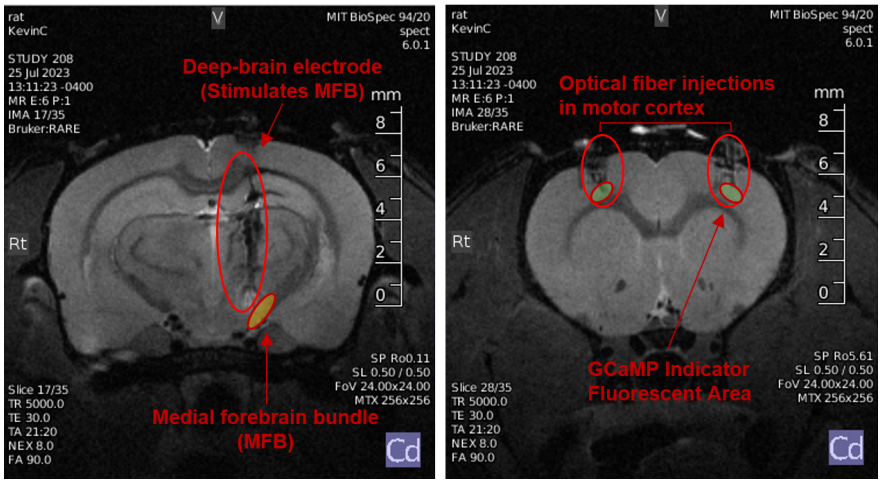

Figure 1. Left: Bipolar deep brain stimulation electrode: (+3.4mm AP, +1.7mm ML, 8.5mm DV); Right: Virus (AAV9 containing DNA encoding GCaMP8f cargo) and optic fiber: (-2mm AP, +/-3.5mm ML, 2.5mm DV). Medial forebrain bundle (central dopeminergic pathway) highlighted orange. Approximate GCaMP fluorescent area highlighted green.

3 Methods

3.1 Surgical Preparation

We performed an initial surgery to implant a 200 μm-diameter optic fiber canula in the rat's motor cortex, allowing for the reading of GCaMP8f-based photometry data [15] [16] [17] [18]. The animals were provided with isofluorane analgesic pain relief during and after the operation in addition to general anaesthetization [19]. During surgery, the genetically encoded calcium indicator GCaMP8f was introduced virally at stereotaxic coordinates of 2.0mm anterior of bregma, 3.5mm lateral, and 2.5mm ventral from bregma, which are part of a broad area of the motor cortex hypothesized to be under the control of volitional movement [20]. The clarity of their purpose and lack of confounding noise makes these regions an ideal target for feedback. Following the recovery period from this surgery, the rats were again anaesthetized for a craniotomy to insert a deep-brain electrode that stimulates the Medial Forebrain Bundle (MFB). The MFB is a dopaminergic pathway that plays a significant role in the brain's reward system. In neurofeedback, the ventral tegmental area's connection to the nucleus accumbens influences mood and behavior, allowing detectable brain signals to be correlated to emotional change in the deep brain [21]. This provides an incentive for behavior modulation in the testing rats via the closed-loop neurofeedback circuit.

3.2 Dataset

GCaMP8f fluorescence data was read from a highly sensitive photosensor integrated into an Arduino setup. The fluorescence intensity values of the GCaMP complex read from the fiber were passed through an amplification tube [17], and the intensity from the left and right hemisphere sensors was read concurrently by the Arduino, which serialized the two input streams and directed the data flow to software. Under this setup, 180 million sequential time points were saved for each hemisphere, representing approximately 30 hours of sensing under both dormant and scanning environments. This volume of data was made possible by a novel software suite *Rasa*, which, in addition to offering machine learning solutions for data processsing, was built as a low-level software system for interpreting the high volume of data attainable from GCaMP photometry. With this purpose-built software, the control dataset was captured for five rats who were subsequently undergoing fMRI experimentation. As such, the dataset represents a wide character of neural activity, including both sleep spindles during anaesthetization and high frequency activity in response to the environment of the EPI scan. The unified two-channel data stream from each was subdivided into discrete subsets of $2,560$ values by a sliding window with no self-overlap (i.e., points 0-2,560 are subset 1; points 2,561-5,120 are subset 2). The size of the window was chosen to be approximately equal to the number of data points processed per second in the new software, making each subset approximately one second in length. Within both channels, groups of 40 points were averaged to pool the 2560-value subsets into 64-value subsets. This may seem overly destructive; however, it serves the important purposes of (a) reducing noise in highly sensitive photometry data, (b) shrinking the data such that it can be analyzed on a millisecond timescale, and (c) increasing recognition of similar features for the signal comparison algorithm that forms the primary innovation of this study. Furthermore, the algorithm still perceives neural activity down to $1/64$th of a second in duration, and with improved confidence against noise, due to the vastly increased data volume.

Despite the pooling operation, the subsets are still two-channeled, with the left and right hemisphere sensors being treated completely independently up until this point. For a target signal under the current stimulation setup, the animal is learning to converge on the simultaneous activation of the desired signal at both locations in the motor cortex. If we are interested in conditioning correlated release of potential, we do not want to provide positive feedback on a release of potential in only one of two regions. Therefore it is essential that the final model consider both channels simultaneously. The difficulty of this task is increased by the fact that the gain of the photosensor and offset between two channels base fluorescence reading is non-constant [17]. In order to effectively compare signals, they must

first be normalized and concatenated, such that they may be passed to a neural network or similarity algorithm.

For each channel in each subset, every value was subtracted from the mean value in the channel for that collection of 64 points, centering the distribution around zero. The normalized channel is ascertained after dividing each value by the global standard deviation of data acquired during that session to account for the non-constant sensor gain between trials. Mathematically, for some sequence $\mathbf{s}$ as a part of continuous stream m,

$$\mathbf{s}_{\text{norm}} = \frac{\mathbf{s} - \text{avg}(\mathbf{s})}{\sigma_m}.$$

Finally, the two channels require some mechanism of concatenation such that they can be passed into a single comparative model. For artificial neural models, the exact mechanism is not particularly important, as the method will be learned quickly as long as they are sufficiently distinct. For the purposes of this study, I simply offset the left channel by a value of four, such that the first 64 elements of the autoencoder input are centered around zero, and the remaining 64 elements are centered around four. This number was chosen so as to be small enough to not deal with bias terms in the network but large enough to not conflate the two channels.

3.3 Model Design

The computational model seeks to compare two signals: an observed signal and a target signal. The observed signal reflects the most recent selection of activity recorded from the sensors. The target signal reflects the signal we want our observed neuron group to produce. Electrostimulation is able to reinforce "desired" behavior [3], but that does not help to answer what is defined as a desirable signal. Much research in neurofeedback has been interested in inducing beta-band or gamma-band activity [22] [23], for which the Fourier Transform alone provides a good picture of desirable activity. Other methods, focusing on neural potentiation, have used peak-detection algorithms or binary thresholds [24] [25]. These techniques suffer in terms of generalizability and continuity. For this study, I propose the use of an autoencoder as a means of continuously comparing two signals for the purposes of neurofeedback. A model for this task requires a method of describing with a continuous variable exactly how close a signal is to a target signal. In theory, this would allow incremental improvements in quality, e.g., if one wished to reinforce a strong double peak. Instead of always rewarding that activity when it happens correctly, which happens to be considerably rare and unpredictable, the idea

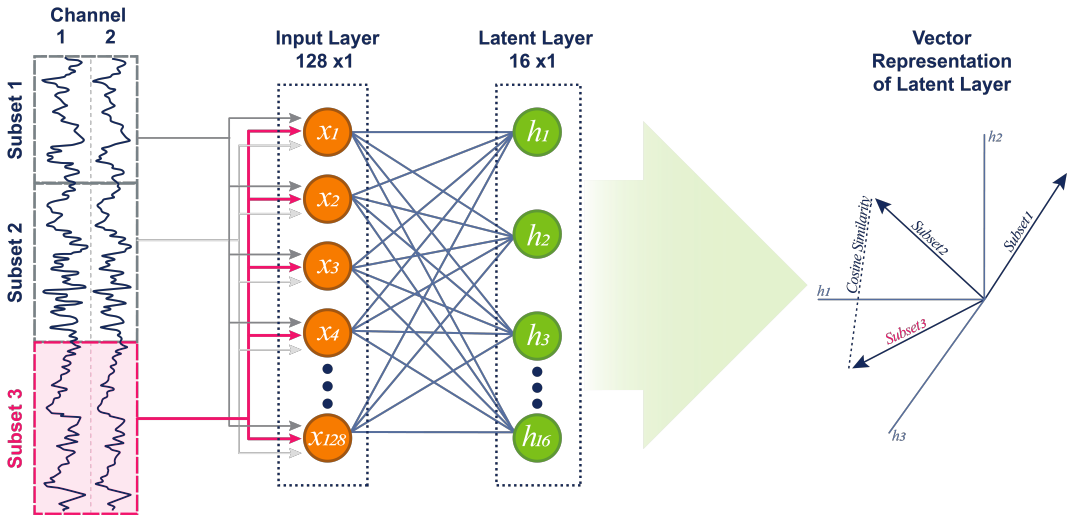

Figure 2. Visualization of headless autoencoder. Raw latent vectors considered in embeddings space.

is to reward more frequent, but incorrect activity, and gradually increase the threshold for correctness based on the continuous numerical evaluation for similarity. This evaluation may be derived with an artificial neural model table to perform analysis that is (a) efficient enough to give the impression of continuity in the evaluation stream and (b) able to define a "loss" mechanism to show how the current target should descend towards the optimization signal. I hypothesized that an autoencoder would be an effective solution to this task. Autoencoder models take a higher-dimensional input stream such as an image or sequence, and, instead of performing inference on it, shrinks it to a narrow bottleneck layer, in which each weight represents some underlying property of the input data structure [26]. This layer is called the "latent layer" and represents the most compact representation the input that the model produces. The model then trains the weights of the encoder by attempting to reconstruct this bottleneck into the original input data with no external aid. Interpreted as a vector, latent vectors form the basis of the *latent space*, a inferential embedding space for the model containing the representation for all possible signals [26].

The artificial neural architectures that were tested take the averaged 128-point GCaMP signal and pass the signal through an autoencoder model, which concludes with a 16-node output layer that, in vector form, represents a 16-D vector in the latent space. In practice, the target signal, which by design may be completely arbitrary, can be encoded into a vector. The model evaluates quality of the observed signal at any given time by vectorizing it with the autoencoder and computing the cosine similarity of the target and observed vectors.

This similarity is a continuous variable that is able to robustly compare two neural signals

with an understanding of the types of signals that occur endogenously in the motor cortex. The continuity of the comparison facilitates new methods of feedback more nuanced than previous methods. In a particularly simple case, a threshold of acceptable similarity may be defined as a criterion for a reward, reinforcing a desirable signal. This can be re-applied with a low latency using a sliding window over the most recent datapoints. More advanced methods that utilize this continuity can be easily imagined: a particularly interesting technique may be a continuously applied reward or anti-reward to act as a titrated "hotter-colder mechanism for more effective reinforcement.

3.4 Implementation

Three autoencoder models were implemented in PyTorch based on this task: a 1-Layer LSTM [27], 3-Layer LSTM, and 2-Headed Transformer [28]. All models took a sequence of size 128 as an input and embedded the sequences in a 16-dimensional space. The cosine similarity of the pure signals (as opposed to the vectors) was used as a control. Signal cosine similarity is currently among the most widely used techniques for neuroscientific signal analysis [29].

3.5 Software Design

To use the artificial neural model for feedback applications, a novel software package titled *Rasa* was developed using the Rust programming language as a high-performance, multi-threaded, interactive solution to integrate analytical tools with memory-safe and accurate data-processing. The previously described Arduino-photosensor circuit is taken as an input, allowing data to be transferred to the computational software with a USB output bus at very high baud rate of 115,200. It was necessary to develop software that can log received data with an extremely low latency in order to ensure the robustness of a machine learning model and allow stimulation to be delivered with confidence closer to the observation of the neural signal. Furthermore, concerns about animal safety in response to neurofeedback dissuaded weakly structured software design and language use. These mandates demanded the use of a systems language with rigorous static analysis and high runtime performance. The novel software was designed with thread-safe software infrastructure to separate the visualization, input, and analytical components with synchronous channels of two-sided deques from which the most recent data points read from the Arduino can be served to the desired consumer.

4 Results

4.1 Model Evaluation

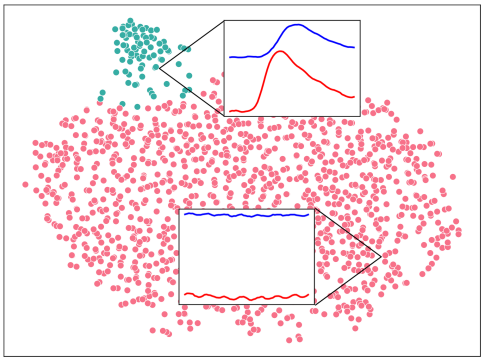

Figure 3. 1,325 Latent space vectors mapped to PCA embeddings, further mapped to t-SNE embeddings, clustered with K-Means to visualize continuous network of signals and strong peak outliers. Two regions selected from flat and peaked areas of distribution. Blue line: right cortex channel. Red line: left cortex channel. Vertical offset for clarity.

In order to evaluate the quality of each architecture, an evaluation set of target signals was created. This serves to validate that the model is indeed capable of identifying the target signal successfully and with a high statistical power. This test does not take advantage of the continuity aspect, but serves as a baseline needed to start using the method experimentally. To generate the feature set, a medium-sized dataset of latent vectors were clustered provisionally with the SciPy implementation of the Principal Component Analysis (PCA) algorithm to form an initial 2D representation. The PCA points were then passed through a t-distributed Stochastic Neighbor Embedding (t-SNE) algorithm and finally clustered with K-Means, as shown in Figure 3. Correlated spikes between the left and right motor cortex are one of the most recognizable features that occur endogenously and were selected as the point of comparison for the models. To be explicit: this clustering did not influence the creation of the models in any way. It only served as a semi-automated method for generating the feature set. Finally, the feature set was inspected manually.

Each architecture was evaluated based on the percent correct:

Model	Parameters	Latency (ms)	Correct (%)
Cos. Similarity	0	3.72	32.5
1-Layer LSTM	12961	2.72	92.4
3-Layer LSTM	29857	4.08	90.7
Transformer	277297	4.77	98.7

The latency for each system was approximately equal. Latency was averaged between 5000 samples per model, performed on a Google Colab V100 GPU instance. The latency of all models was considered within tolerance for use for real-time neurofeedback. Despite the large parameter size and comparatively high latency, a transformer model was chosen as the ideal design due to its high accuracy, which demonstrates that the latent vectors form a more meaningful representation of the data using this architecture.

4.2 Software Evaluation

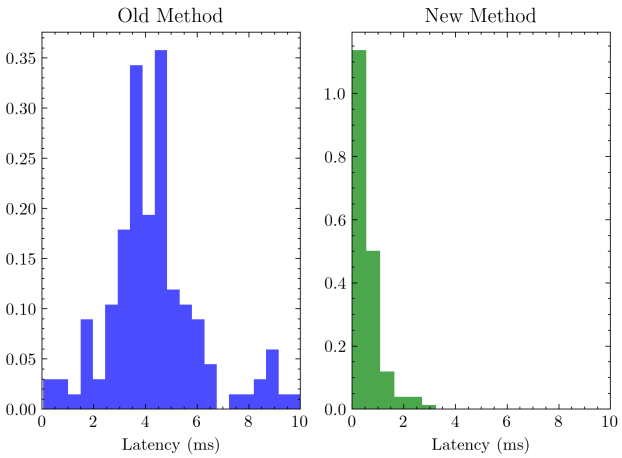

Figure 4. Latency after 5 minutes by software. ($n = 100, p = 2.6 \cdot 10^{-34}$).

To determine the effectiveness of the novel software, both the new and old software were tested after five minutes of idling, eliminating cold-start bias and exemplifying the long-term performance benefits of optimized data structures provided by the new implementation. In order to evaluate the latency of the software, an artificial signal would be supplied to the datastream buffer of the respective software every 100 milliseconds by a separate independent timer thread. Following this, the active thread would exhaustively search for the signal among the 128 most recent entries, recording the difference in time between the most recent signal and the beginning of the search. Bias was eliminated by controlling for the search time in Rust ($n = 5; \mu = 0.133$ ms) and Python ($n = 5; \mu = 0.805$ ms).

The distribution of latencies observed from the experiment was evaluated with a two-sample T-Test under the hypothesis that the new method has a lower latency than the previous method. The test evaluated to a p-value significant past $\alpha = 0.0001$, indicating a strong confidence in the improved performance of the novel software.

4.3 *In Vivo* Evaluation

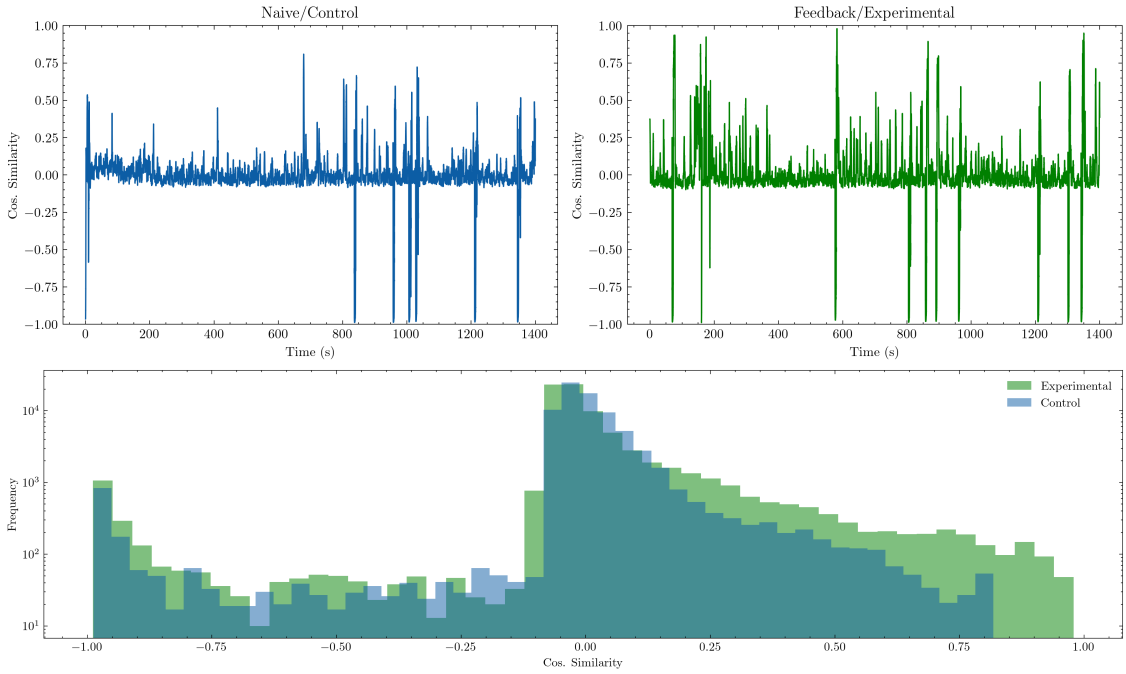

Figure 5. Top: Similarity values over time for control and experimental trial. Bottom: distribution of similarity scores over time. $p = 0.0021$

The model and software were integrated into the software package *Rasa* and run on the McGovern Institute 9.4T fMRI Imaging Setup for an EPI scan under neurofeedback. This study seeks to demonstrate that this method of neurofeedback will, at a minimum, act equally in place of the existing neurofeedback setup. As such, it would increase the generality of the methods with no decrease in learning rate. Feedback was administered with the use of a deep-brain electrode shown in Figure 1 as a response to desirable motor cortex activity measured through the *Rasa* model. For simplicity, we rewarded correlated peaks of electrical activity in the two neuronal ensembles considered by the two optical fibers. For the temporarily paralyzed rat, the neural signals and neurofeedback were monitored for two subsequent 20-minute sessions, the first of which generated a "naive" dataset under no stimulation and the second of which was performed with neurofeedback monitoring, with a reward delivered when the observed signal had a computed cosine similarity greater than 0.7

to the target signal. During the second session, the model and software achieved a correlated peak above the threshold 15 times and was stimulated 13 times (discrete peaks separated from one another by less than 16 seconds removed were noted, but not rewarded in order to prevent overstimulation and preserve the health of the animal). The choice of 0.7 as a threshold was arbitrary and a hyperparameter of the experiment. As such, its value is controllable inside of the *Rasa* software, allowing researchers to control the frequency and lenience with which reward to be administered.

This training session provides strong evidence that the use of stimulation under the novel system provides improvement over no stimulation. However, due to the small training size and lack of controls, it did not show strong evidence of improvement over the previous method ($p = 0.20$ + No unbiased comparison).

5 Discussion

The method showed promise in our initial experiments as a method of defining neurofeedback qualification. Unfortunately, due to time constraints and the high demand for use of the 9.4T fMRI, this experiment was only able to be performed once. During the two-week *in vivo* portion of this study, there were no opportunities to use the machine during weekdays between the hours of 8AM-5PM, which were the only times available due to laws involving supervision of minors in bio-hazardous laboratories. This prevented the execution of the initial experimental design, which was to test the novel method's unique ability to encode more complex signals by attempting to reinforce the embedding of a non-correlated peak, a signal that does not generally occur endogenously within rats (constituted 1 of 13,709 sequences in unabridged dataset). This more complex task represents the fundamental advantage of the autoencoder model and loss mechanism. Despite this, the full experiment is scheduled now on the fMRI/Photosensor setup in early 2024, likely occurring over the course of 5-8 multi-hour training sessions to account for the increased signal complexity and additional robustness checks. Even though the central experiment was not executed, the *Rasa* software system and use of ML-based signal analysis for simple signals is beginning to underscore the wider neurofeedback procedure outside the scope of this experiment in isolation.

6 Conclusions

This paper outlines the developmental process of the software system *Rasa* and its constituent transformer model, which represents a novel technique for reinforcing cognition modulation based on closed-loop neural signals. We succeeded in demonstrating the performance effectiveness of the software, which has served to aid in more robust data analysis, and improving the ability to detect subtler, and shorter neuronal signals. Furthermore, we developed three unsupervised models that demonstrated the ability to identify arbitrary neural signals and reward them based on the literal positioning of the embedding vectors.

7 Proposed Experiment

Despite a paucity of *in vivo* data under the novel stimulation paradigm, the model and software have shown experimental promise. As such, I propose additional experiments that are entirely possible given the current capabilities of *Rasa*. This study provides software and methods towards robust neurofeedback; this begins with conditioning signals that never occur endogenously within the brain. Under previous training paradigms, such a signal would not be possible to train, as it would by definition never occur in the observed GCaMP signal. If we wish to train a completely artificial signal, imagine a moving threshold throughout the course of training. As the animal begins to learn from feedback, the threshold will be very low (i.e., any signal resembling the form of the target will be rewarded). As it continues to produce this, the standard for reward will grow. The use of an autoencoder for the signal analysis gives a very robust way to measure how the characteristic of a signal compares to a target, provided knowledge of the spectrum of signals the animal's brain can produce. Essentially, this experiment proposes a stochastic gradient descent to the signal within the brain through the continuous similarity metrics provided by *Rasa*.

8 Code Availability

The source code for *Rasa*, as described in this paper, is available on GitHub under an MIT license.

9 Acknowledgements

I would like to thank Kevin Chung, Dr. Alan Jasanoff, and the Jasanoff Lab for mentoring me on this project. I am grateful to the Center for Excellence in Education and the Massachusetts Institute of Technology for sponsoring and hosting the Research Science Institute. Finally, I would like to especially express my gratitude to the McGovern Institute for Brain Research for the use of its highly specialized equipment and the Department of Defense for its generous sponsorship.

References

[1] Jose del.R. Millan. On the need for on-line learning in brain-computer interfaces. In *2004 IEEE International Joint Conference on Neural Networks (IEEE Cat. No.04CH37541)*, volume 4, pages 2877–2882 vol.4, July 2004.

[2] Matthew K. Hogan, Gillian F. Hamilton, and Philip J. Horner. Neural stimulation and molecular mechanisms of plasticity and regeneration: A review. *Frontiers in Cellular Neuroscience*, 14, 2020.

[3] Hui Ye, Jenna Hendee, Joyce Ruan, Alena Zhirova, Jayden Ye, and Maria Dima. Neuron matters: Neuromodulation with electromagnetic stimulation must consider neurons as dynamic identities. *Journal of NeuroEngineering and Rehabilitation*, 19(1):116, Nov 2022.

[4] Ranganatha Sitaram, Tomas Ros, Luke Stoeckel, Sven Haller, Frank Scharnowski, Jarrod Lewis-Peacock, Nikolaus Weiskopf, Maria Laura Blefari, Mohit Rana, Ethan Oblak, Niels Birbaumer, and James Sulzer. Closed-loop brain training: The science of neurofeedback. *Nature Reviews Neuroscience*, 18(2):86–100, Feb 2017.

[5] Ghazaleh Soleimani, Michael A. Nitsche, Til Ole Bergmann, Farzad Towhidkhah, Ines R. Violante, Romy Lorenz, Rayus Kuplicki, Aki Tsuchiyagaito, Beni Mulyana, Ahmad Mayeli, Peyman Ghobadi-Azbari, Mohsen Mosayebi-Samani, Anna Zilverstand, Martin P. Paulus, Marom Bikson, and Hamed Ekhtiari. Closing the loop between brain and electrical stimulation: Towards precision neuromodulation treatments. *Translational Psychiatry*, 13(1):279, Aug 2023.

[6] Stefanie Enriquez-Geppert, René J. Huster, and Christoph S. Herrmann. EEG-neurofeedback as a tool to modulate cognition and behavior: A review tutorial. *Front Hum Neurosci*, 11:51, February 2017.

[7] Hengameh Marzbani, Hamid Reza Marateb, and Marjan Mansourian. Neurofeedback: A comprehensive review on system design, methodology and clinical applications. *Basic Clin Neurosci*, 7(2):143–158, April 2016.

[8] Tsai-Wen Chen, Trevor J. Wardill, Yi Sun, Stefan R. Pulver, Sabine L. Renninger, Amy Baohan, Eric R. Schreiter, Rex A. Kerr, Michael B. Orger, Vivek Jayaraman, Loren L. Looger, Karel Svoboda, and Douglas S. Kim. Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499(7458):295–300, July 2013.

[9] Golan Karvat, Artur Schneider, Mansour Alyahyay, Florian Steenbergen, Michael Tangermann, and Ilka Diester. Real-time detection of neural oscillation bursts allows behaviourally relevant neurofeedback. *Communications Biology*, 3(1), Feb 2020.

[10] Robert J. Schafer and Tirin Moore. Selective attention from voluntary control of neurons in prefrontal cortex. *Science*, 332(6037):1568–1571, 2011.

[11] Kelly B. Clancy, Aaron C. Koralek, Rui M. Costa, Daniel E. Feldman, and Jose M. Carmena. Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nature Neuroscience*, 17(6):807–809, 2014.

[12] Brendan Parsons and Jocelyn Faubert. Enhancing learning in a perceptual-cognitive training paradigm using EEG-neurofeedback. *Scientific Reports*, 11(1):4061, Feb 2021.

[13] Lydia Hellrung, Matthias Kirschner, James Sulzer, Ronald Sladky, Frank Scharnowski, Marcus Herdener, and Philippe N. Tobler. Analysis of individual differences in neurofeedback training illuminates successful self-regulation of the dopaminergic midbrain. *Communications Biology*, 5(1):845, Aug 2022.

[14] Robert W. Thatcher. Chapter 56 - Neurofeedback's first fifty years: My story. In James R. Evans, Mary Blair Dellinger, and Harold L. Russell, editors, *Neurofeedback*, pages 443–452. Academic Press, 2020.

[15] Helmuth Adelsberger, Olga Garaschuk, and Arthur Konnerth. Cortical calcium waves in resting newborn mice. *Nature Neuroscience*, 8(8):988–990, 2005.

[16] Henry Lutcke, Masanori Murayama, Thomas Hahn, David J. Margolis, Simone Astori, Stephanie Zum Alten Borgloh, et al. Optical recording of neuronal activity with a

genetically-encoded calcium indicator in anesthetized and freely moving mice. *Frontiers in Neural Circuits*, 4(9), 2010.

[17] Christina K. Kim, Samuel J. Yang, Nandini Pichamoorthy, Noah P. Young, Isaac Kauvar, Joshua H. Jennings, Talia N. Lerner, Andre Berndt, Soo Yeun Lee, Charu Ramakrishnan, et al. Simultaneous fast measurement of circuit dynamics at multiple sites across the mammalian brain. *Nature Methods*, 13(4):325–328, 2016.

[18] Yaroslav Sych, Maria Chernysheva, Lazar T. Sumanovski, and Fritjof Helmchen. High-density multi-fiber photometry for studying large-scale brain circuit dynamics. *Nature Methods*, 16(6):553–560, 2019.

[19] Zhifeng Liang, Yuncong Ma, Glenn D.R. Watson, and Nanyin Zhang. Simultaneous GCaMP6-based fiber photometry and fMRI in rats. *Journal of Neuroscience Methods*, 289:31–38, 2017.

[20] Stephen H. Scott. Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5(7):532–545, 2004.

[21] Rudolf Nieuwenhuys, Leo M.G. Geeraedts, and Jan G. Veening. The medial forebrain bundle of the rat. I. General introduction. *Journal of Comparative Neurology*, 206(1):49–81, 1982.

[22] Maciej J. Szul, Sotirios Papadopoulos, Sanaz Alavizadeh, Sébastien Daligaut, Denis Schwartz, Jérémie Mattout, and James J. Bonaiuto. Diverse beta burst waveform motifs characterize movement-related cortical dynamics. *Progress in Neurobiology*, 228:102490, 2023.

[23] Shenghong He, Emilie Syed, Flavie Torrecillos, Gerd Tinkhauser, Petra Fischer, Alek Pogosyan, Erlick Pereira, Keyoumars Ashkan, Harutomo Hasegawa, Peter Brown, and Huiling Tan. Beta Oscillation-Targeted neurofeedback training based on subthalamic LFPs in Parkinsonian patients. *Int IEEE EMBS Conf Neural Eng*, 2019:81–84, March 2019.

[24] Golan Karvat, Artur Schneider, Mansour Alyahyay, Florian Steenbergen, Michael Tangermann, and Ilka Diester. Real-time detection of neural oscillation bursts allows behaviourally relevant neurofeedback. *Communications Biology*, 3(1):72, Feb 2020.

[25] Beatrix Barth, Kerstin Mayer-Carius, Ute Strehl, Sarah N. Wyckoff, Florian B. Haeussinger, Andreas J. Fallgatter, and Ann-Christine Ehlis. A randomized-controlled

neurofeedback trial in adult attention-deficit/hyperactivity disorder. *Scientific Reports*, 11(1):16873, Aug 2021.

[26] Diederik P. Kingma, Max Welling, et al. An introduction to variational autoencoders. *Foundations and Trends in Machine Learning*, 12(4):307–392, 2019.

[27] Nur Izzati Ab Kader, Umi Kalsom Yusof, Mohd Nor Akmal Khalid, and Nik Rosmawati Nik Husain. A review of long short-term memory approach for time series analysis and forecasting. In Mohammed A. Al-Sharafi, Mostafa Al-Emran, Mohammed Naji Al-Kabi, and Khaled Shaalan, editors, *Proceedings of the 2nd International Conference on Emerging Technologies and Intelligent Systems*, pages 12–21, Cham, 2023. Springer International Publishing.

[28] Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun. "Transformers in time series: A survey," arxiv:2202.07125, 2023.

[29] Elizabeth P. Chou and Shen-Mou Hsu. Cosine similarity as a sample size-free measure to quantify phase clustering within a single neurophysiological signal. *Journal of Neuroscience Methods*, 295:111–120, 2018.