Decoding EEG and LFP Signals using Deep Learning: Heading TrueNorth

Ewan Nurse^{*}
Department of Electrical &
Electronic Engineering
University of Melbourne
Parkville, Australia
enurse@student.unimelb.edu.au

Benjamin S. Mashford^{*}
IBM Research - Australia
Australia, VIC 3053
bmashfor@au1.ibm.com

Antonio Jimeno Yepes
IBM Research - Australia
Australia, VIC 3053
antonio.jimeno@au1.ibm.com

Isabell Kiral-Kornek
IBM Research - Australia
Australia, VIC 3053
isabeki@au1.ibm.com

Stefan Harrer[†]
IBM Research - Australia
Australia, VIC 3053
Department of Electrical &
Electronic Engineering
University of Melbourne
sharrer@au1.ibm.com

Dean R. Freestone^{†‡}
Department of Medicine
St. Vincent's Hospital
Melbourne, The University of
Melbourne
Fitzroy, Australia
deanrf@unimelb.edu.au

ABSTRACT

Deep learning technology is uniquely suited to analyse neurophysiological signals such as the electroencephalogram (EEG) and local field potentials (LFP) and promises to outperform traditional machine-learning based classification and feature extraction algorithms. Furthermore, novel cognitive computing platforms such as IBM's recently introduced neuromorphic TrueNorth chip allow for deploying deep learning techniques in an ultra-low power environment with a minimum device footprint. Merging deep learning and TrueNorth technologies for real-time analysis of brain-activity data at the point of sensing will create the next generation of wearables at the intersection of neurobionics and artificial intelligence.

CCS Concepts

•Applied computing → Life and medical sciences;
•Hardware → Neural systems; •Computing methodologies → Artificial intelligence;

^{*}These authors contributed equally to the presented work

[†]Corresponding authors

[‡]Additional authors are listed in Section 7

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CF'16, May 16-19, 2016, Como, Italy

© 2016 ACM. ISBN 978-1-4503-4128-8/16/05...\$15.00

DOI: <http://dx.doi.org/10.1145/2903150.2903159>

Keywords

Neuromorphic computing, brain computer interface, BCI, electroencephalography, EEG, convolutional neural network, CNN, deep learning

1. INTRODUCTION

Imagine a device implanted within the body that is small, efficient and robust, that can decode a person's cognitive state or intentions from their neural activity. Such decoding would facilitate groundbreaking technologies, enabling new ways for people to connect and interact with devices and the physical world. For example, such a tool will enable a completely natural interaction with exoskeletons, returning mobility to the physically disabled. Alternatively, this device will enable thought control of a computer without needing to move a muscle.

We have come a long way to achieve the goal of an implantable neural decoder. Neuroscientists are able to decode movement intentions with great precision from recordings of action potentials in the motor cortex [13, 12, 30, 10]. However, over time, and even day to day, the signals from the microelectrodes that record action potentials (neuron activity) change significantly, making decoding challenging. The instability of microelectrode systems has restricted the development of commercially available neural decoding devices. Furthermore, the computational demands of decoding devices have posed a major limitation in creating fully implantable systems.

A promising alternative to using spiking activity as a signal to decode one's intention is the local field potential (LFP) or intracranial electroencephalography (EEG). This alternative has the advantage of a more stable neural interface. However, decoding brain states from EEG/LFP signals is non-trivial. Complications arise due to the ambiguous mapping from individual neurons to the LFP. The EEG and the LFP are abstract signals, and despite efforts in the fields of biophysics and neuroscience, a causal understanding be-

tween the recorded data and the underlying cognitive states is not well established [26]. Ambiguities in mapping between the microscopic (single neuron) and mesoscopic (LFP/EEG) scales are due to the immense dimensionality of the problem. Solving the inverse problem (mapping from LFP to single neuron activity) based on biophysics requires an unrealistic amount of data. Furthermore, the complexity of the brain, and the emergent dynamics that arise due to this complexity, create the additional challenge of decoding inter-scale interactions [4, 5, 14].

The inability to model a causal link between the LFP/EEG signals and cognitive state has driven a pragmatic approach toward exploring signal features, such as spectral content of evoked potentials, in behavioral neuroscience. Even though this empirical approach of using spectral and other features has been somewhat successful in classifying cognitive states, there is still no principled model for decoding EEG/LFP signals. We believe deep learning methods can fill this void by bypassing a detailed bottom up description of the biophysics of EEG. In this paper we advocate for a top down data-driven approach to learning representations of cognitive states that are present in EEG/LFP signals.

Techniques from machine learning are commonly used to construct classifiers of (spectral) features in EEG/LFP analysis. We believe that deep learning will further advance the field by removing the need to explicitly define and hand code features, thus providing a complete decoding framework (Figure 1). The standard analysis pipeline for LFP/EEG signals is as follows: data acquisition, pre-processing, feature extraction, classification, and prediction of signal class [33]. Examples of typical features include power in specified frequency bands, signal entropy and variance. The majority of research has focused on finding optimal features for a given application, thus this step has become a bottleneck in EEG/LFP brain machine interface development. Feature extraction has been thought to be a necessary step in dimension reduction prior to classification. However, with the advent of more sophisticated machine learning algorithms, such as convolutional neural networks (CNNs), the traditional hand coded feature extraction bottleneck can be bypassed [19].

Figure 1: (a) Traditional decoding method. A signal is acquired from the brain, from which features are extracted. These features are then used to assign an output class for the given signal. **(b). New deep learning method,** which avoids explicit feature extraction and classification, instead using a convolutional neural network to directly map the input signal to the output.

There is a massive opportunity to improve EEG/LFP decoding using deep learning. Deep learning has brought upon breakthroughs in problems such as face, object, and speech recognition [28, 18, 11]. A good illustration of breakthroughs is handwritten digit recognition, which improved vastly through deep learning methods. The EEG/LFP decoding problem has common aspects to digit recognition, including segmentation, inter-subject variation, and intra-subject variation. Just as a model should be able to segment background noise and image content, EEG systems must be able to disregard background signal fluctuations. Inter-subject variation in anatomy and health can significantly affect the EEG signal between individuals, just as individuals can write digits in widely different ways. Intra-subject variation in the execution of a task can produce markedly different responses from the brain, just as an individual can write the same sentence in varying ways.

The EEG/LFP decoding problem can be cast into similar problems faced in computer vision that avoid the need for hand-coded features by learning representations of features directly from data. Importantly, the quality of the features that are learnt improve with larger datasets. Learning from data also provides a systematic approach for finding features, as opposed to a trial and error hand-coded approach. Interestingly, the learnt rules can be studied to understand what signal properties are important for decoding, which can advance our knowledge about the processes that take place in the brain. Moving away from constrained models that require specialist, a-priori understanding of the data and towards generalized deep learning methods will provide robust, subject-specific models that expand our understanding and improve performance of neural decoding.

There is a growing wave of interest in applying machine learning approaches to EEG/LFP analysis [35]. Multi-layer perceptrons (an early type of supervised artificial neural network) have been used to classify a hand-squeeze task at high accuracy, with analysis demonstrating that features were created at the anticipated electrode locations and signal frequencies [22]. Deep learning has been applied successfully to a finger-flexion task, with data recorded from EEG electrodes on the surface of the brain [32]. Deep belief networks have similarly out-performed traditional features in anomaly detection from EEG [34].

Implementation of deep learning algorithms on hardware has typically involved physically large, high-power GPUs. This has made implementation of these powerful models for real-time EEG analysis difficult, as medical devices need to be low power, generate minimal heat and be deployable in autonomous, always-on mode. The recent appearance of a new generation of low-power systems based on traditional GPU and hardware technology [29, 3, 21, 23, 25] and neuromorphic platforms spearheaded by IBM's TrueNorth chip [20, 24] is key to overcoming this hurdle by providing the cognitive analytical power of the most advanced deep learning algorithms on an ultra-low-power system the size of a postage stamp. Throughout the rest of this paper, we will demonstrate how a combination of a EEG/LFP data, deep learning methods and TrueNorth have the potential to make our hypothesized neural decoding device a reality in the near future.

Figure 2: TrueNorth board deployed to decode neural signals during movement to control robotic arm.

2. DEEP LEARNING & TRUENORTH

Despite decades of combined industrial and academic research and development, the human brain is still the most powerful computer in existence. In an effort to build a processor with brain-like functionality, while at the same time operating more efficiently than any other brain-inspired platform, IBM has recently introduced the TrueNorth chip. TrueNorth is a novel neuromorphic computing platform consisting of 1 million neurons that operate with ultra low power consumption [20].

Due to its novel non-von-Neumann architecture TrueNorth is capable of performing a fundamentally new type of computing which replaces conventional processing steps with brain-inspired operations. This creates a uniquely suitable hardware environment for the inference step of deep learning algorithms. Unlike conventional processors, neuromorphic chips process spiking data in real-time following an event-driven massive parallel scheme. This eliminates the need for moving data back-and-forth between memory and processing units during computation and is the main reason for the unprecedented power efficiency of neuromorphic systems. In particular TrueNorth is capable of analyzing large amounts of data more efficiently than conventional state of the art platforms - yet, it only consumes an unprecedented 70 mW at full capacity. This is orders of magnitude lower not only than what is consumed by any conventional von-Neumann type platform but also by other existing neuromorphic systems.

2.1 Cognitive Computing on TrueNorth

TrueNorth is exceptionally well suited for implementation of complex neural network architectures to process a data stream and predict the category of a feature in it. Programming a TrueNorth chip thereby entails first specifying its configuration through programming, training and testing a neural network with a pre-existing dataset and then deploying the network onto the chip in form of so called corelets [2]. A typical workflow for the deployment of a new TrueNorth network is shown here (Figure 3). After the trained network is deployed, the chip is then capable of processing measured input in streaming at the point of sensing (Figure 2). There is no inherent limitation as to which type of data can be processed by TrueNorth. In particular TrueNorth is not only capable of analyzing LFP/EEG data using deep learning technology, but, by doing so, it will also allow for implementing a stable analytical system providing the reliability and repeatability required of med-

Figure 3: TrueNorth deployment pipeline.

ical devices. Building on these insights we have assessed the feasibility of analyzing EEG data by first developing a convolutional neural network in Caffe which we successfully trained and tested using EEG data. We used this result as a basis for building a TrueNorth-compatible neural network that uses local receptive field inner product layers. Detailed results are described in the following sections and strongly suggest that deep learning-enabled EEG data analysis using the TrueNorth chip will fundamentally change the way brain activity data can be utilized for neurobionics applications in the future.

3. BRAIN COMPUTER INTERFACING WITH TRUENORTH

3.1 Description of Dataset

EEG data was acquired from a participant who performed a self-paced hand squeeze task. The onset of movement were time-locked to the EEG by simultaneously collecting electromyography (EMG) using the same acquisition system. The surface EMG electrodes were attached to the forearms. Data was sampled at 1 kHz, bandpass filtered between 0.1-100 Hz and notch filtered at 50 Hz. The participant was seated and asked to alternate clenching their left and right

hands at their own pace with their eyes open, focusing their eyes on a point on the wall. Recordings were undertaken in four minute sessions, with rests between each session. Data was recorded for a total of approximately 30 minutes. Recorded EEG data was decimated to 250 Hz prior to training the network.

Movement detection was achieved by extracting the root mean squared (RMS) power from the EMG signal in order to label the EEG data into the various classes. RMS power was calculated in sliding windows of 500 ms. When the RMS value exceeded a manually set threshold, a movement was detected. This subject-specific threshold was determined prior to the trial start by visual inspection of the baseline RMS power compared to during hand-clenches. 546 left and 563 right hand squeezes were recorded, respectively.

Using the EMG labels, the EEG data was separated into epochs, 100 ms before the onset of clenching and ending 300 ms after the onset of clenching. A refractory period of 300 ms was applied after a clench was detected to prevent a single clench being detected twice. Time periods in between hand squeezes were also used as examples of no-movement signals; i.e., times when the participant was not clenching either hand for at least 500 ms. The same window length (400 ms) was used for the no-movement EEG epochs. Data collected is available at: <https://github.com/EwanNurse/A-Generalizable-BCI-using-Machine-Learning-for-Feature-Discovery>.

3.1.1 Convolutional Neural Networks

The use of deep convolutional neural networks (CNNs) for image classification tasks has increased rapidly in recent years due to the demonstrated high accuracy made possible by modern network designs and training methods. Convolutional neural networks employ a biologically-inspired design that differs from fully-connected NNs by utilizing spatially-local correlation, forming a network of local connectivity between neurons in adjacent layers. The use of shared weights in convolutional layers also reduces the system resources, including computer memory, required for the network. During the training phase, the network forms a set of ‘filters’ that respond to a particular input pattern.

Although traditionally applied to image datasets, CNNs have shown good results when applied to data types containing spatio-temporal information, such as audio (human speech [11], music [7]) and video [17].

3.2 Methods

Experimental investigations on the classification of EEG signals using a CNN were performed with the open-source software Caffe [16]. An EEG dataset from a single subject (participant D) was selected for analysis. The dataset was split into 964 training set samples and tested on 145 validation set samples, with each sample being the EEG signal that was recorded when the subject squeezed either their left hand or right hand. The no-squeeze cases were not included in the dataset. Each signal consisted of a 2D matrix with dimensions 46×100 (Figure 5), with the vertical axis corresponding to the 46 measurement channels (corresponding to 46 scalp electrodes) and the horizontal axis corresponding to 100 datapoints across the time domain. These signals were first filtered to remove outlier spikes and then resized to a set of 46×46 pixel images for conversion into a Caffe-compatible dataset.

Figure 5: Six randomly selected data epochs, with each image representing one action by the subject (either left hand squeeze or right hand squeeze).

The parameter settings used in our CNN model for 46×46 pixel input images are explained below. The first convolutional layer uses 40 kernels of size 5×5 with a stride of 1 pixel, followed by a max pooling layer that uses a 2×2 kernel and stride of 2. The second convolutional layer takes as input the output of the first convolutional layer and filters it with 100 kernels of size 5×5 , followed by a max pooling layer that uses a 2×2 kernel and stride of 2. The fully connected (FC) layer has 500 neurons. We train the CNN with stochastic gradient descent at a learning rate of 0.01. The learning rate is reduced by a factor of 10 at every 10 epochs.

In order to visualize the convolutional filters, the two-dimensional convolution of each 5×5 filter matrix was taken with each 46×46 training example, and the magnitude of each element was taken. The mean was then calculated across each training example and time to determine the average spatial weight magnitude distribution of the 40 filters in the first layer. This was done separately for the two classification outputs. Hence, each of the outputs represents the average spatial response of each of the filters to the left and right training classes, respectively.

3.3 Results

3.3.1 Classification Results

During the period over 30 training epochs (Figure 6), the training loss rapidly decreases to below 0.1, indicating the CNN found a good fit to the training data, with a lower loss value on the validation set. The final classification accuracy on the validation set was 81%.

The dataset was then transferred to codebase that trains a model within the constraints of the TrueNorth architecture, as discussed elsewhere [20]. The TrueNorth hardware has a maximum 256 axons (inputs) per neuron and the connectivity of the network is required to operate within these limits. Recently, a new off-line training approach for creating a TrueNorth compatible network via back-propagation was developed [8] and the same method was employed here to re-train on the same dataset. The maximum classification accuracy we currently achieve via the TrueNorth-compatible

Figure 4: (a) Illustration of convolutional neural network architecture implemented in Caffe (b) Visualization of the 40 filter kernels in the first convolutional layer after training (c) Visualization of mean spatial intensity response (squared average weights) of the 40 filter kernels to training examples from the left hand-squeeze class and (d) right hand-squeeze class.

Figure 6: Plot of loss and classification accuracy of network during training in Caffe software.

network is 76%. On-going work is focused on further improving the classification performance and to demonstrate real-time operation with spiking input datastreams.

3.3.2 Network Analysis

CNNs are capable of making complex, high-level features from the EEG data. As shown in Figure 4 (b) after just one stage of convolution it is difficult to infer what type of information the kernels are extracting. This difficulty is compounded because unlike an image, EEG data is abstract and cannot be directly interpreted. Nevertheless it is possible to work backwards from the CNN kernels to infer what spatial information is relevant to signal classification. For each of the training examples, the temporal aspects of the first feature (convolutional) layer was averaged, leaving the spatial information. The spatial information was then averaged across the training examples to reveal important spatial features corresponding to the original 46 channels, showing which channels strongly activate the first layer of the CNN. The weights were computed separately for training examples of left and right hand squeezes (Figure 4 (c) and (d) respectively, show weight intensity, the squared weight values). It can be seen that the 40 filters correspond to distinct regions of the brain, predominately covering the motor cortex. Notably, left and right hand-squeezes show clear differences in spatial activation patterns. The information in these heat maps can be used to infer the areas of the brain that are useful in the classification task.

3.4 Discussion

We have presented a new framework for decoding EEG/LFP signals using convolutional neural networks (CNNs).

These results suggest that, in future, this scheme could be used to create networks that operate in real time on TrueNorth as the analysis platform for a fully implantable neural prosthesis.

The EEG/LFP decoding problem was motivated by the need for robust communication channels from the brain that can be developed for brain machine interface technology. Currently, intracranial EEG/LFP electrodes form the most stable interface for neural recordings enabling chronic and stable recordings over years [27, 6], providing a compelling motivation for using this type of data. However, decoding EEG/LFP signals is challenging compared to other recording modalities, due to the abstract nature of the signal. We have shown how the CNN overcomes this challenge by learning representations of the important features of the data.

The classification accuracy obtained from TrueNorth in this study is 76%. This is not as high as previous results on this set of 86% [22]. It should be noted however that this prior study extensively explored the classifier hyperparameter space, which we are yet to undertake with this implementation.

The results presented in this paper demonstrate that the CNN can perform state of the art decoding. The power of the CNN was utilized by transforming the EEG data into a spatiotemporal image of the EEG, following an approach from an earlier study [22]. It is important to note that this data was not optimal for developing the network and further improvements in decoding accuracy can be realized. Figure 6 showed that the validation loss function did not improve with the training loss function. This indicates that more training data will further improve performance. Furthermore, scalp EEG data was used and the higher fidelity intracranial EEG signal, which will be used in an implantable device, is expected to improve performance.

A key point of difference between object recognition from images and EEG/LFP decoding is in the abstract nature of the EEG/LFP signal and the variability of neural dynamics. A single training example of recorded EEG shows no obvious signature that represents the execution of a movement (or any other output). The inability to visualize the brain's intention from the signal limits the use of hand-coded features and motivates a machine learning approach. CNNs are particularly effective at classifying highly variable data [1], as a result of the properties of locality and weight sharing that underlie the CNN architecture. Deep CNNs can implicitly learn important features of the data, such as the well-known work using large image databases [18], where the kernels learned by the first convolutional network layer were shown to contain a set of 'edge-detectors'. Similarly for our data, we see a set of kernels formed, but it is not apparent how they operate on the EEG signal. The visualization of these kernels in Figure 4 (b) shows no obvious spatial or temporal structure. The indistinct kernel structure highlights that there is no systematic way to hand-code the same feature maps that the CNN has discovered from the data. The inferred spatial structure of the feature maps shown in Figures 4 (c) and (d), demonstrates that the networks are extracting spatially relevant signals with dominant sidedness focused on the motor regions.

As mentioned above, we expect improvements in decod-

ing accuracy with an increase in the size of the training dataset. Improvements with larger data sets are another point of differentiation from the traditional hand-coded feature approach. The bottleneck with the CNN lies in the quality and size of the dataset, rather than on the choice of features. In an implantable device, more data is easier to acquire than better hand-coded features, adding weight to the argument for using a CNN. Conventional GPU/CPU platforms are not suitable for implementing CNNs on an implantable device due to power consumption requirements. Our results suggest that TrueNorth could be used to overcome this limitation by offering an ultra-low power environment for implementing deep neural networks.

The neural network that was used for decoding is able to run on one single TrueNorth chip with a power consumption of less than 70 mW. The implication of this extremely low level of power consumption is that the algorithm can be implemented in a fully implantable system without compromising classification performance. This contrasts with modern state of the art systems that had to make algorithmic compromises due to power limitations [6].

Previous TrueNorth applications have been able to undertake classification at a rate of 30 input frames per second [20]. This is comparable with state-of-the-art EEG/LFP classification rates [15, 31] that used much physically larger and power-intensive hardware. Hence, the main time delay is due to the length of data analysed in each input frame (400 ms for this study). Previous work has shown an input window of 300 ms can still provide natural control of a 3D cursor [31], hence a delay of 400 ms is not anticipated to greatly affect natural BCI control. Further work is required to understand how changing the duration of the input window affects classification accuracy.

This study provides a proof of principle for using TrueNorth as a decoding platform and further studies are required for full validation. Real time testing for brain machine interfaces is a much more difficult problem than offline decoding. Nevertheless, we have demonstrated the feasibility of decoding motor signals using TrueNorth, thus further work towards online validation is warranted.

4. OUTLOOK

On-going work is focused on achieving high classification accuracy of EEG signals on the TrueNorth hardware and demonstrating real-time operation on streaming data. Ultimately, it is expected that trained classification models will be robust enough to handle data that is highly variable, such as data collected over different days where disparity in electrode positioning is inevitably introduced.

The improvement of deep learning classifiers is heavily reliant on the use of larger sets of data, to allow for the production of more generalizable and complex features. However, in the case of neural systems, significantly increased data collection is only possible using implanted recording systems, which require low-power, small processing hardware. The operating parameters of TrueNorth mean it is ideally suited to neurobionics applications, and we have shown that its classification performance is state of the art.

Success in motor decoding for brain-computer interfacing motivates use of this technique for other EEG applications, such as epileptic seizure prediction. Although recent results have demonstrated that seizure prediction is possible in implanted devices [6], the implementation of more com-

plex models could greatly improve prediction results [9].

5. CONCLUSION

We have programmed, trained and tested a convolutional neural network for analyzing labeled EEG data using the deep learning platform Caffe. EEG data was taken from patients squeezing their hands with the two data labels being left hand and right hand squeeze patterns. After achieving state of the art classification accuracy of 81% we re-built the neural network into a configuration that is designed to operate within the TrueNorth neuromorphic architecture. We then trained this constrained network using the same EEG dataset and achieved 76% classification accuracy, demonstrating that the neural network can be run on one single TrueNorth chip at a maximum peak power consumption of only 70 mW. Current work focuses on optimizing the TrueNorth compatible neural network towards exceeding state of the art testing performance and in-depth benchmarking of the chip during operation. TrueNorth is ideally suited to neurobionics applications due to its low physical profile and power consumption, and potential for real-time, high accuracy classification.

6. ACKNOWLEDGMENTS

The authors would like to thank the IBM SyNAPSE Team lead by Dr. Dharmendra Modha at the IBM Almaden Research Center as well as the Precision Medicine Research Team at IBM Research - Australia for fruitful discussions. We would also like to acknowledge the support and guidance of Prof. David Grayden, Prof. Mark Cook, Prof. Anthony Burkitt, Prof. Terry O'Brien and Prof. Patrick Kwan (all University of Melbourne), and Dr. Qiang Chen.

7. ADDITIONAL AUTHORS

Additional authors: Philippa Karoly, Department of Electrical & Electronic Engineering, The University of Melbourne, email: p.karoly@student.unimelb.edu.au; Jianbin Tang, IBM Research-Melbourne, Australia, VIC 3053, email: jb-tang@au1.ibm.com; Hariklia (Lili) Deligianni, IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, email: lili@us.ibm.com

8. REFERENCES

- [1] O. Abdel-Hamid, A.-r. Mohamed, H. Jiang, and G. Penn. Applying convolutional neural networks concepts to hybrid nn-hmm model for speech recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pages 4277–4280. IEEE, 2012.
- [2] A. Amir, P. Datta, W. P. Risk, A. S. Cassidy, J. Kusnitz, S. K. Esser, A. Andreopoulos, T. M. Wong, M. Flickner, R. Alvarez-Icaza, et al. Cognitive computing programming paradigm: a corelet language for composing networks of neurosynaptic cores. In *Neural Networks (IJCNN), The 2013 International Joint Conference on*, pages 1–10. IEEE, 2013.
- [3] Auviz systems. <http://auvizsystems.com/products/auvizdnn/>.
- [4] M. Breakspear and C. J. Stam. Dynamics of a neural system with a multiscale architecture. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1457):1051–1074, 2005.
- [5] D. R. Chialvo. Emergent complex neural dynamics. *Nature physics*, 6(10):744–750, 2010.
- [6] M. J. Cook, T. J. O'Brien, S. F. Berkovic, M. Murphy, A. Morokoff, G. Fabinyi, W. D'Souza, R. Yerra, J. Archer, L. Litewka, et al. Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: a first-in-man study. *The Lancet Neurology*, 12(6):563–571, 2013.
- [7] S. Dieleman, P. Brakel, and B. Schrauwen. Audio-based music classification with a pretrained convolutional network. In *ISMIR*, pages 669–674, 2011.
- [8] S. K. Esser, R. Appuswamy, P. Merolla, J. V. Arthur, and D. S. Modha. Backpropagation for energy-efficient neuromorphic computing. In *Advances in Neural Information Processing Systems*, pages 1117–1125, 2015.
- [9] D. R. Freestone, P. J. Karoly, A. D. Peterson, L. Kuhlmann, A. Lai, F. Goodarzi, and M. J. Cook. Seizure prediction: Science fiction or soon to become reality? *Current Neurology and Neuroscience Reports*, 15(11):1–9, 2015.
- [10] A. P. Georgopoulos, A. B. Schwartz, and R. E. Kettner. Neuronal population coding of movement direction. *Science*, 233(4771):1416–1419, 1986.
- [11] A. Graves, A.-r. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 6645–6649. IEEE, 2013.
- [12] L. R. Hochberg, D. Bacher, B. Jarosiewicz, N. Y. Masse, J. D. Simeral, J. Vogel, S. Haddadin, J. Liu, S. S. Cash, P. van der Smagt, et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375, 2012.
- [13] L. R. Hochberg, M. D. Serruya, G. M. Friehs, J. A. Mukand, M. Saleh, A. H. Caplan, A. Branner, D. Chen, R. D. Penn, and J. P. Donoghue. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442(7099):164–171, 2006.
- [14] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.
- [15] G. Hotson, D. P. McMullen, M. S. Fifer, M. S. Johannes, K. D. Katyal, M. P. Para, R. Armiger, W. S. Anderson, N. V. Thakor, B. A. Wester, et al. Individual finger control of a modular prosthetic limb using high-density electrocorticography in a human subject. *Journal of neural engineering*, 13(2):026017, 2016.
- [16] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. B. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM Multimedia*, volume 2, page 4, 2014.
- [17] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1725–1732. IEEE, 2014.

- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [19] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [20] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, et al. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197):668–673, 2014.
- [21] Movidius. <http://www.movidius.com/solutions/vision-processing-unit>.
- [22] E. S. Nurse, P. J. Karoly, D. B. Grayden, and D. R. Freestone. A generalizable brain-computer interface (bci) using machine learning for feature discovery. *PLoS ONE*, 10(6):e0131328, 06 2015.
- [23] Nvidia jetson tk1 whitepaper. http://developer.download.nvidia.com/embedded/jetson/TX1/docs/jetson_tx1_whitepaper.pdf.
- [24] J. Pavlus. The search for a new machine. *Scientific American*, 312(5):58–63, 2015.
- [25] Qualcomm zeroth. <https://www.qualcomm.com/invention/cognitive-technologies/zeroth>.
- [26] S. J. Schiff. *Neural control engineering: the emerging intersection between control theory and neuroscience*. MIT Press, 2012.
- [27] K. A. Sillay, P. Rutecki, K. Cicora, G. Worrell, J. Drazkowski, J. J. Shih, A. D. Sharan, M. J. Morrell, J. Williams, and B. Wingeier. Long-term measurement of impedance in chronically implanted depth and subdural electrodes during responsive neurostimulation in humans. *Brain stimulation*, 6(5):718–726, 2013.
- [28] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1701–1708. IEEE, 2014.
- [29] Teradeep. <http://teradeep.com/nmx.html>.
- [30] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz. Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198):1098–1101, 2008.
- [31] W. Wang, J. L. Collinger, A. D. Degenhart, E. C. Tyler-Kabara, A. B. Schwartz, D. W. Moran, D. J. Weber, B. Wodlinger, R. K. Vinjamuri, R. C. Ashmore, et al. An electrocorticographic brain interface in an individual with tetraplegia. *PloS one*, 8(2):e55344, 2013.
- [32] Z. Wang, S. Lyu, G. Schalk, and Q. Ji. Deep feature learning using target priors with applications in ecog signal decoding for bci. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 1785–1791. AAAI Press, 2013.
- [33] J. Wolpaw and E. W. Wolpaw. *Brain-computer interfaces: principles and practice*. Oxford University Press, 2012.
- [34] D. Wulsin, J. Gupta, R. Mani, J. Blanco, and B. Litt. Modeling electroencephalography waveforms with semi-supervised deep belief nets: fast classification and anomaly measurement. *Journal of neural engineering*, 8(3):036015, 2011.
- [35] M. Yang, S. A. Sheth, C. A. Schevon, G. M. M. II, and N. Mesgarani. Speech reconstruction from human auditory cortex with deep neural networks. In *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.