PROJECT JENKINS: TURNING MONKEY NEURAL DATA INTO ROBOTIC ARM MOVEMENT, AND BACK

Andrii Zahorodnii MIT zaho@csail.mit.edu Dima Yanovsky MIT yanovsky@mit.edu

March 19, 2025



Figure 1: Project Jenkins. Leader arm velocities are computed via forward kinematics, then a transformer generates synthetic neural data. An MLP trained on real monkey neural data decodes it back into velocity space, commanding the follower arm's movement. (Monkey diagram adapted from [13]; robotic arm images from [1]).

ABSTRACT

Project Jenkins explores how neural activity in the brain can be decoded into robotic movement and, conversely, how movement patterns can be used to generate synthetic neural data. Using real neural data recorded from motor and premotor cortex areas of a macaque monkey named Jenkins, we develop models for decoding (converting brain signals into robotic arm movements) and encoding (simulating brain activity corresponding to a given movement).

For the interface between the brain simulation and the physical world, we utilized Koch v1.1 leader and follower robotic arms. We developed an interactive web console that allows users to generate synthetic brain data from joystick movements in real time. Our results are a step towards braincontrolled robotics, prosthetics, and enhancing normal motor function. By accurately modeling brain activity, we take a step toward flexible brain-computer interfaces that generalize beyond predefined movements.

To support the research community, we provide open source tools for both synthetic data generation and neural decoding, fostering reproducibility and accelerating progress. The project is available at https://www.808robots.com/projects/jenkins.



Figure 2: Our approach in action. The experimenter moves the leader robotic arm, its velocity is recorded and transformed to synthetic neural data using the encoder model. Then, a decoder that was only trained on real neural data is decoding the neural data back into movement velocities, which are passed through inverse kinematics to the follower robotic arm. For the full video, please see the project's webpage.

1 Introduction

Synthetic neural data generation and neuroprosthetic devices are active areas of research, sparked by advances in neuroscience and robotics [22, 4, 2, 15]. These fields have significant implications for brain-computer interfaces, rehabilitation, and simulation of brain dynamics for downstream tasks or gaining new understanding of the underlying neural mechanisms.

In this project, which we call "Project Jenkins," we explore such decoding and encoding of neural data from a macaque monkey named Jenkins. We used a publicly available dataset [5] containing neural firing patterns from Jenkins' motor and premotor cortical areas during a center-outreach task.

Generating synthetic neural activity enables researchers to test and refine decoding models without requiring continuous access to live neural recordings [12, 16], while neuroprosthetic advancements [18, 20, 21, 9, 7, 3, 8, 17] rely on robust encoding techniques to translate brain signals into precise motor commands.

Our aim was two-fold (Figure 1, 2):

- Decoding: Translate neural spiking data into predicted velocities for a robotic arm.
- Encoding: Generate synthetic neural activity corresponding to an intended robotic movement.

With this paper, we publish the developed open-source tools for both synthetic neural data generation and neural decoding, enabling researchers to replicate our methods and build upon them. Our full codebase and additional resources, including demonstration videos, can be found on the project's website: https://www.808robots.com/projects/jenkins.

2 Methods and Results

2.1 Data Acquisition

The neural data used in our project came from the primary motor cortex (M1) and caudal portion of the dorsal premotor cortex (PMd) areas of the brain of a rhesus macaque monkey, Jenkins. The dataset was published by Mark Churchland and others in 2021 [5]. In short, Jenkins was trained on multiple reaching tasks, where the goal of the task is to press on dots that randomly light up on the screen in front of it (Figure 3). Every time the monkey completes a trial successfully, he is rewarded with fruit juice. In our project, we used data from a center-outreach task, where the monkey always starts out with its hand in the middle of the screen, and the dots light up in one of 8 positions $(0^{\circ}, 45^{\circ}, 90^{\circ}, 135^{\circ}, 180^{\circ}, 225^{\circ}, 270^{\circ}, \text{ or } 310^{\circ})$.



Figure 3: A typical monkey reaching task (adapted from [13]). Jenkins starts with his hand at center, reaching out to one of 8 radial targets.



Figure 4: The decoding procedure.

During these trials, neural spiking activity was measured with a multielectrode array capturing signals from 192 distinct neurons. The final dataset contains thousands of trials across more than a dozen hours of recordings, accompanied by continuous measurements of Jenkins' hand position and velocity. In total, the spiking data was binned into discrete time intervals of 20 ms.

All in all, there are more than a dozen hours of brain recordings together with the tracking of the monkey arm as it pressed dots on the screen thousands of times.

2.2 Decoding: From Neural Data to Robot Movement

Decoding is where we convert monkey brain data into the movements of the robotic arm. Whenever Jenkins wants to move his arm, neurons in his brain activate, calculating the trajectory of his planned movement, and this activation signal is sent through his spinal cord into the muscles in the monkey's arm. This sequence of events in time means that the movement of the arm depends on the recent history of neural data, and our model should take this fact into account.

Feature Construction. To predict the monkey's hand movements from brain recordings, we split time into bins of 20 ms. For each bin, we recorded the total number of spikes from each of the 192 neurons in that time bin (Figure 4). This gives a vector $\mathbf{x}_t \in \mathbb{R}^{192}$ every 20 ms. We also recorded the average x and y velocity of Jenkins' hand (from motion-capture data) during that same bin, forming a velocity vector $\mathbf{v}_t = (v_x, v_y)$.

Neural Network Decoding Model. We employ a simple MLP (multilayer perceptron) with:

- Two hidden layers (sizes 256 and 128).
- ReLU nonlinearities.



Figure 5: Generating synthetic neural data.

• Input window of 50 bins (i.e., 1 second of neural history).

Thus, at each time step t, we feed $\mathbf{x}_{t-49}, \ldots, \mathbf{x}_t$ into the MLP to predict the current velocity \mathbf{v}_t . Despite its simplicity, this architecture performed effectively, achieving $R^2 \approx 0.9$ on a held-out test set.

Driving the Robotic Arm. We implemented a *follower* robotic arm using the Koch v1.1 design [1]. The arm has six servo motors arranged in a chain, and its end-effector position is determined by the servo angles. We decode velocity from the neural data, integrate it to obtain Cartesian coordinates (X, Y), using the exponentially moving average (EMA), and then apply inverse kinematics (via the ikpy library, [14]) to calculate servo angles for each motor. We defined a kinematic chain by specifying the servo motors' properties and the connecting link geometries. This chain configuration enables both forward kinematics (calculating end effector position from servo rotations) and inverse kinematics (determining required servo angles from desired coordinates). The EMA step is important to avoid the accumulation of errors over time. This loop yields continuous control of the robot in real time:

Neural Data $\implies \hat{\mathbf{v}} \implies (X, Y) \implies \text{IK} \implies \text{Servo Angles.}$

2.3 Encoding: From Movement to Synthetic Neural Data

While the decoding problem was relatively straightforward, the encoding stage proved to be considerably more challenging. The encoding model is designed to *generate* neural spiking patterns given a sequence of arm velocities (or positions).

Closed-Loop Simulation Challenge. The key difficulty lies in closed-loop simulation: to continuously produce simulated neural data, the model needs to take its own past outputs as new inputs, to produce future predictions based on what it has already said before. So if a model is a little bit off every time, then the errors from its output get fed back into the model, producing even larger errors in the next output, and so on.

Small errors can accumulate over time, destabilizing the generated signals. After several time bins, these errors would often blow up to large spiking rates or collapse to near-zero activity. Success required fiddling with multiple different ways of formatting the inputs, finding the right architecture (transformer or LSTM) and training hyperparameters, as well as training for a long time (400 epochs).

Transformer-based Encoder. Since the present movement of the arm depends on past neural data, it follows that the present neural data depends on the future (planned) movement of the arm. That is, to generate neural data, the model needs to know what the future arm movement will be. Accordingly, we process the data so that for every time bin, we input all of the past brain activity to the model, as well as the future arm movement velocities (Figure 5). Specifically, we provided:

• **Past neural activity**: Binned spike counts from previous time steps.

• Future arm movement: A "look-ahead" window of 40 bins (800 ms) of velocities.

The model is trained to predict how many times each neuron will spike in the current 20 ms time bin. We formulate this problem as a 9-way classification: either a neuron will spike 0 times (be quiet), or 1, 2, ..., 7, or 8+ times.

The model was trained to output a 9-category distribution for each neuron, representing the number of spikes: $\{0, 1, \ldots, 7, 8+\}$ We train our encoding model to output what it thinks are the probabilities of each neuron's number of spikes being in any of our defined categories. This discrete classification formulation, akin to next-token prediction in language models [19, 11], helped stabilize training.

Training Procedure We trained for approximately 400 epochs, carefully tuning hyperparameters (learning rate, batch size, dropout) to avoid divergence. We found best results by setting learning rate to 0.0005 and no weight decay ($\lambda_{wd} = 0$). The training took under 4 hours on an off-the-shelf GPU with 12GB of RAM. We experimented with an LSTM-based approach [10] but eventually found a transformer architecture to be more robust.

2.4 Closing the Loop: Robot Movement to Neural Data and Back to Movement

To accurately record the robotic arm's movement, we assembled a Koch v1.1 leader arm, sampling its (x, y, z) coordinates at a frequency of 50 Hz. Since our study focuses exclusively on two-dimensional data, we discarded the z-coordinate. We then differentiated the (x, y) positions to obtain the velocities (V_x, V_y) , which serve as input for our encoding model.

After the model generates neural data corresponding to the movement inputs, we apply the decoding procedure to transform this synthetic neural data back into velocity components (V_x, V_y) . To reconstruct reliable spatial coordinates (x, y) from these velocities, we employ an exponential decay filter $(\lambda_{decay} = 0.95)$, which mitigates the compounding of integration errors over time.

Finally, the filtered positional data is passed to the Koch follower arm, where inverse kinematics algorithms compute the necessary servo rotations, enabling the follower arm to accurately replicate the original movements. The resulting system can be seen in Figure 2 and on the project website.

2.5 Interactive Interface

To enable users without direct robotic hardware experience to interactively generate neural data from movement, we developed a web application that allows joystick manipulation. This app records the velocities from joystick input and processes them through our transformer model to produce synthetic neural data, visualizing the output in real-time directly in the browser. Notably, the transformer model runs entirely within the user's browser rather than on a remote server. To achieve this, we converted our PyTorch model into a browser-compatible .onnx format using ONNX Runtime [6], ensuring efficient local execution.

3 Conclusion and Future Directions

Project Jenkins demonstrates the feasibility of translating real neural activity into robotic arm movement and generating synthetic neural data to accompany or predict such movement. While our decoding model was successful and robust, the encoding model required more intricate architectures (transformers) and careful training to produce stable spike patterns over time.

In practice, these methods can be extended to broader applications, such as human BCIs, prosthetics, and motor rehabilitation. Although our data was limited to eight primary reaching directions, preliminary experiments suggest generalization to more complex trajectories (e.g., drawing circles). Future work will focus on:

- Testing on extended movement repertoires and more complex tasks.
- Improving the robustness of the encoding model in closed-loop simulation.
- Exploring real-time human-interface prototypes that adapt these neural decoders.

For more details, code, and demonstration videos, please visit the project webpage: https://www.808robots.com/projects/jenkins.

References

[1] Koch v1.1 robotic arm images. Retrieved from https://github.com/jess-moss/koch-v1-1.

- [2] Puya Afshar and Yoky Matsuoka. Neural-based control of a robotic hand: Evidence for distinct muscle strategies. volume 2004, pages 4633 – 4638 Vol.5, 01 2004.
- [3] Manfredo Atzori, Matteo Cognolato, and Henning Müller. Deep learning with convolutional neural networks applied to electromyography data: A resource for the classification of movements for prosthetic hands. *Frontiers in Neurorobotics*, 10, 2016.
- [4] M Burrow, J Dugger, D Humphrey, DJ Reed, and LR Hochberg. Cortical control of a robot using a time-delay neural network. In *Proceedings of International Conference on Rehabilitation Robotics ICORR. Bath, UK*, pages 83–86, 1997.
- [5] Mark Churchland, John P. Cunningham, Matthew T. Kaufman, Justin D. Foster, Paul Nuyujukian, Stephen I. Ryu, and Krishna V. Shenoy. Neural population dynamics during reaching, 2024. Data set.
- [6] ONNX Runtime developers. Onnx runtime. https://onnxruntime.ai/, 2021. Version: x.y.z.
- [7] Aaron Fleming, Wentao Liu, and He (Helen) Huang. Neural prosthesis control restores near-normative neuromechanics in standing postural control. *Science Robotics*, 8(83):eadf5758, 2023.
- [8] Vikash Gilja, Paul Nuyujukian, Cindy A. Chestek, John P. Cunningham, Byron M. Yu, Joline M. Fan, Mark M. Churchland, Matthew T. Kaufman, Jonathan C. Kao, Stephen I. Ryu, and Krishna V. Shenoy. A high-performance neural prosthesis enabled by control algorithm design. *Nature Neuroscience*, 15(12):1752–1757, 2012.
- [9] Anirban Gupta, Nikolaos Vardalakis, and Fabian B. Wagner. Neuroprosthetics: from sensorimotor to cognitive disorders. *Communications Biology*, 6:14, 2023.
- [10] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. Neural Computation, 9(8):1735–1780, 1997.
- [11] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Aditya Ramesh, and Dario Amodei. Scaling laws for neural language models. arXiv preprint arXiv:2001.08361, 2020.
- [12] Jaivardhan Kapoor, Auguste Schulz, Julius Vetter, Felix Pei, Richard Gao, and Jakob H. Macke. Latent diffusion for neural spiking data, 2024.
- [13] M. T. Kaufman, M. M. Churchland, S. I. Ryu, and K. V. Shenoy. Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience*, 17:440–448, 2014.
- [14] Pierre Manceron. IKPy.
- [15] Dailin Marrero, John Kern, and Claudio Urrea. A novel robotic controller using neural engineering frameworkbased spiking neural networks. *Sensors*, 24(2):491, 2024.
- [16] Ryota Nakajima, Arata Shirakami, Hayato Tsumura, Kouki Matsuda, Eita Nakamura, and Masanori Shimono. Deep neural generation of neuronal spikes. *bioRxiv*, 2023.
- [17] Max Ortiz-Catalan, Jan Zbinden, Jason Millenaar, Daniele D'Accolti, Marco Controzzi, Francesco Clemente, Leonardo Cappello, Eric J. Earley, Enzo Mastinu, Justyna Kolankowska, Maria Munoz-Novoa, Stewe Jönsson, Christian Cipriani, Paolo Sassu, and Rickard Brånemark. A highly integrated bionic hand with neural control and feedback for use in daily life. *Science Robotics*, 8(83):eadf7360, 2023.
- [18] Chethan Pandarinath, Paul Nuyujukian, Chris H. Blabe, Brian L. Sorice, Jason Saab, Francis R. Willett, Leigh R. Hochberg, Krishna V. Shenoy, and Jaimie M. Henderson. High performance communication by people with paralysis using an intracortical brain-computer interface. *eLife*, 6:e18554, 2017.
- [19] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. *OpenAI*, 2018.
- [20] Francis R. Willett, David T. Avansino, Leigh R. Hochberg, Jaimie M. Henderson, and Krishna V. Shenoy. High-performance brain-to-text communication via handwriting. *Nature*, 593:249–254, 2021.
- [21] Francis R. Willett, Emily M. Kunz, Chaofei Fan, David T. Avansino, Garret H. Wilson, Chandramouli Chandrasekaran, Leigh R. Hochberg, Krishna V. Shenoy, and Jaimie M. Henderson. A high-performance speech neuroprosthesis. *Nature*, 620:1031–1036, 2023.
- [22] Ruohan Zhang, Sharon Lee, Minjune Hwang, Ayano Hiranaka, Chen Wang, Wensi Ai, Jin Jie Ryan Tan, Shreya Gupta, Yilun Hao, Gabrael Levine, Ruohan Gao, Anthony Norcia, Li Fei-Fei, and Jiajun Wu. Noir: Neural signal operated intelligent robots for everyday activities, 2023.