Proceedings Track

Neuromodulator-inspired task switching using a fixed weight-space transformation

*

Abstract

In biological systems, chemical neuromodulators provide a means to broadly reconfigure the circuit activity of large populations of neurons. This enables switching between distinct behavioral states with control of just a few non-local variables. Applying such control mechanisms to deep neural networks could enhance their flexibility. This approach examines how learning of switchable behavior can occur when the task transformation itself is constrained and networks are free to evolve around the transformation. We explore the implementation of chemical neuromodulator inspired control by demonstrating a network that can switch between AND, OR and XOR gates by applying a single randomly initialized affine transformation directly to its weights. We train this switchable network by co-optimizing the AND/OR/XOR neural network and its randomly transformed counterpart on different tasks. Geometrically, this can be described as finding two points in the functional solution spaces of the two networks that are linked by the transformation vector. We move towards extending this approach to more complex datasets by showing task-switching between two sets of 5 MNIST classes and comparing performance to task switching implemented as extra inputs.

Keywords: Neuromodulator, switchable, neuroAI

1. Introduction

1.1. Background

Neural circuits in biological systems can be flexibly reconfigured. Drawing inspiration from neuromodulatory behavior may lead to new strategies for flexible control of modern deep neural networks.

In biological circuits, chemical neuromodulators serve as important elements of neural control systems. They modify both the intrinsic properties of neurons and their synaptic efficiency, enabling the same system to provide different outputs for similar inputs in a context-dependent manner (Harris-Warrick and Marder, 1991). This is evident even in the simplest model circuits, such as the proctolin modulated phase patterns of the pyloric rhythm of the crustacean stomatogastric ganglion (Hooper and Marder, 1987). At higher levels of functional complexity, neuromodulators induce switching between multiple behavioral states such as roaming and dwelling (Flavell, 2013).

Neuromodulation allows for coordinated changes in activity through control of a single variable. In comparison to local rewiring through synaptic transmission of specific signals, this mechanism may allow for rapid re-configuration and large changes in functionality. Given the large, high-dimensional parameter spaces that neural circuits occupy (Kirst, Modes and Magnasco, 2017), it is impressive that this coarse tuning knob enables precise

---

*

switching between different, finely tuned functionality. Random traverses in parameter space are likely to result in aberrant behavior, as is frequently observed upon disruption of chemical neuromodulation (Marder, 2012). To meet the requirement that biological circuits either remain stable upon neuromodulation, or are switched to a desired target functionality, it is likely that neuromodulatory systems walk through a narrow set of constrained possible states (Fig 1A).

In biological systems, the physiological effects of the neuromodulator are to some degree predetermined. While neural circuits have many degrees of freedom that can be controlled either through local synaptic plasticity, or through genetic encoding of neural structures that evolve at the timescales of natural selection, the toolbox of available neuromodulators are much more limited. Their cellular-level physiological effects on biophysical subunits such as membrane channels appear to be conserved between individuals, and to an extent between species. This may suggest that neuromodulators are relatively static. Circuits may be optimized around the neuromodulator to provide stability or switchability in response to neuromodulation. Exploring the mechanisms behind how such circuits can be constructed could yield novel approaches towards artificial intelligence and insight into principles underlying computation in biological systems.

1.2. Contribution

Optimizing the circuit around the neuromodulator could allow for the creation of switchable circuits. We would like to explore this process using simple feed-forward neural networks as a computational abstraction. We draw inspiration from chemical neuromodulators to propose a simple single variable control mechanism for switchable, co-optimized neural networks. We perform interleaved training of an initial neural network, f and a transformed network f'. The weights of f' are given by $W' = MW$, transformed from the initial weights by a random, fixed transformation function M (Fig 1B). Biases are given by $b' = b + v$. During training, f and f' are co-optimized with respect to different objective functions (Fig 1C-E). During inference, applying M allows for switching between networks (Fig 1F).

Previous work has explored the role of neuromodulators in providing a form of uniform gain control to neural circuits. In these instances, all synaptic strengths are considered to be uniformly increased or decreased. This is computationally modeled as a scalar modulation of weights across the entire network, or through application of a low-rank modulation matrix (Costacurta et al., 2024, Tsuda et al., 2021). However, neurons express different receptor densities to neuromodulators (Salvan et al., 2023, Froudist-Walsh et al., 2023, Sun et al., 1996). This shifts the individual sensitivity to a particular neuromodulation in a heterogenous, per-neuron manner. The use of a square matrix provides a way to account for heterogenous gain factors across different neurons.

In contrast to recent meta-learning based approaches, which switch between different functions using learned task transformation networks (Lampinen and McClelland, 2020, Vecoven et al., 2020, Ha et al., 2016, Chalvidal et al. 2022), here we do not learn the task transformation but instead fix it from the outset. Naively, we would expect such an approach to yield poorer performance than learned task mappings. However, avoiding a learned task mapping perhaps allows for a closer analogy to biological systems. Evolution of new neuromodulators likely occurs on timescales longer than that of encoded circuits.

While the receptor density of cells for different neuromodulators can vary throughout the lifetime of an individual, to first order, the effect of neuromodulation could be considered as fixed. Forcing this constraint, that the neuromodulator should not be learnt, may yield new insights. We explore this approach through switching between AND/OR/XOR gates as a model task.

2. Task Structure Mapping

We firstly map the task structure of the AND/OR/XOR/NOR gates by training the gates over 100 separate runs by stochastic gradient descent, with each training run consisting of 1000 epochs. We then flatten their parameters, consisting of the weights and biases of the linear layers of the neural network. We perform principal components analysis (PCA) on the parameters of the final epoch of each run, and plot each neural network as a point on the 3-D space of the first three principal components (Fig 2). We repeat this for neural networks with different degrees of complexity, including a network with two linear layers (Fig 2A) and four linear layers (Fig 2B). ReLU's are used as non-linearities apart from the final layer, where a sigmoid is used.

Networks with just two linear layers are distributed broadly in parameter space (Fig 2A). The properties of a mapping between two gates, such as AND/OR are characterized by considering the distribution of the cosine and Euclidean distances between trained AND networks and OR networks. The distributions of both gates are relatively broad. As the as the complexity of the network increases to four linear layers, the distributions cluster more sharply (Fig 2B), suggesting there are two or more ways to achieve these gates. The distribution of distances becomes more clearly defined. The XOR gate still occupies a distinct region of parameter space, which reflects the unique set of parameters needed for the non-linear separation required to achieve XOR gating. Additionally, the volume in parameter space the clusters seem to occupy is smaller as well, implying that the configurations which lead to this particular gate are more limited.

we will explore whether given some fixed transformation, we can evolve networks to achieve desired tasks on either side of the transformation.

Given that there seems to be only a limited angular and Euclidean distribution to switch between the AND/OR networks, a random transformation with some rotation and scaling would not be expected to land in this distribution. However, it is likely that there exists a much broader space of networks which functionally achieve the desired functionality but aren't local minima for the loss function of a particular gate.

3. Co-optimizing switchable neural networks

We follow the procedure outlined in Algorithm 1 and Figure 1B.

We summarize our approach as follows. We first generate a random affine transformation consisting of a modulation matrix M and a modulation vector of biases v. We train the neural network using stochastic gradient descent for some period of epochs on Function A, in this case an AND gate. We then flatten all the weights and biases into two separate vectors. We multiply the modulation matrix with the weight vector, and add the modulation vector to the bias vector and then update the model with these new modulated weight/bias

Table 1: Algorithm 1: Co-optimised switchable networks

```
**Data:** AND, OR, XOR gate inputs and outputs
**Initialize** feed-forward neural network $f$ with $b$ linear layers, $b - 1$ ReLU non-linear layers, 1 sigmoid non-linearity as the final layer.
**Generate** a random matrix, $T_w$ and a random vector $T_b$
**Repeat** for $n$ cycles
    Train $N_n$ to achieve AND outputs given AND inputs, for $m$ epochs using SGD
    Transform $N_n$ to $N_n$ by setting $W' = MW$, $b' = b + v$
    Train $f'$ to achieve OR outputs given OR inputs if $x_i > x_{i+1}$ for $m$ epochs using SGD
    De-transform $f'$ to $f$ by setting $W = M^{-1}W'$, $b = b' - v$
```

parameters. The model is then trained on the OR network for some period of epochs before being de-modulating by multiplying the weights by the inverse of the modulation matrix and subtracting the modulation vector from the bias. We use a network with two linear layers, as this provides the minimum network expressivity needed to reliably achieve XOR gates (and AND/OR) while minimizing overall complexity.

The model is trained for 10 epoch's on AND before being modulated and switching objective functions to OR. This rapidly changes the trajectory of the system in parameter space, resulting in jagged trajectories in parameter space become jagged, as shown in a representative example (Fig 3A). This is visible also in the loss (Fig 3B). The loss is calculated relative to its target objective at any epoch, and therefore quickly rises upon objective switching (Fig 3C), before a slow reduction upon re-training on the new objective. We find that for this AND/OR switching task the system can be co-optimized. The peak post-switch loss falls in both domains with training. The network is considered to succeed at the task if it achieves 100% accuracy on the outputs. After 500 switching cycles, or 2500 total epochs, the network alternates between achieving 100% accuracy on the AND and OR gates with each switch (Fig 3D). This strong performance is maintained for subsequent cycles.

The switched system requires more cycles to converge than simple unmodulated training of an AND gate (Fig 3E) as it cannot follow the simplest gradient descent pathway for a single gate. We plot clusters corresponding to the end-points of unmodulated training for AND/OR/XOR/NOR gates on the trajectory plots for both the modulated and unmodulated training. In both of these instances, its final location is near, but not within the clusters of parameters corresponding to previously trained networks.

We study the effect of the switching period. A smaller switching period of 3 epochs results in a smoother trajectory, (Fig 3F) as the system can't evolve too far from the pre-switch objective before being switched back to that objective. A longer switching period of 30 cycles leads to a much more jagged trajectory (Fig 3G), as the network can evolve and deviate significantly from the pre-switch objective function after modulation. Notably, in this representative example, it reaches 100% performance in fewer cycles than at low switch periods. This is despite the peak loss, stabilizing at 0.2 appearing much higher than the low-period cases, which stabilize at 0.15. The trajectory deviates significantly from the mapped clusters of unmodulated gates, while maintaining good performance. This implies that for this problem, there are many medium-to-low loss regions of parameter space which achieve functionally acceptable switchable performance that are not local loss minima for individual functions trained for a single task. Switchable networks may traverse surfaces on the loss landscapes corresponding to the two individual objective functions that are not local/global minima. The trajectory seems to follow a plane of the parameter space, which may be perpendicular to two parallel but separated regions of parameter space corresponding to AND/OR. The possibility of switchability given some fixed neuromodulator likely depends on the distribution of distances between these non-minimal geometric surfaces. This would be expected to relax the previously suggested quite challenging requirement, that the neuromodulator should fall within the distribution of distances between local/global minima. Mapping such a system would therefore benefit from sampling of a broader parameter space, perhaps using recently reported Bayesian methods such as BAX (Chitturi et al., 2024).

This network structure and approach also allows for switching to more complex network functions such as a switchable AND/XOR gate (Fig 4). In a representative example, the training trajectory appears noisier, and finally concentrates near the mapped XOR region. This may be because there are fewer regions in parameter space that satisfy XOR, so it cannot traverse large planes like AND/OR (Figure 3G). In contrast to training for AND/OR switching, training a switchable AND/XOR gate has a 30 % success rate (Fig 4C). This implies a greater sensitivity to initial conditions. We note that this is consistent with reported results from training of an unmodulated XOR with a simple two-layer network structure. Here training is tricky and often fails even without modulation (McClelland and Rumelhart 1988). Using more complex network structures which could have larger appropriate regions of parameter space could yield improved performance.

4. Performance on MNIST

We carried out task-switching between two sets of 5 MNIST classes, with an extra input encoding the task as a comparison baseline. We repeated these experiments as a function of number of parameters in the network by changing hidden layer size while keeping the same CNN architecture. We found that the peak performance of the neuromodulated task switching case was lower than the extra input baseline. As network size became much smaller (below 10'000 parameters, or hidden size ¡ 10), neuromodulated task switching performance degraded only mildly, reaching while the extra input case dropped more rapidly. This may suggest that extra inputs are excellent when resources are plentiful, but are more brittle when capacity becomes too low to separate tasks. Task switching on the other hand is robust but constrained. In contexts where information is processed in a physical systems (e.g. the brain) with constraints on physical space and number of connections, resource constraints may become important. Further exploration of neuromodulator inspired switching may be of benefit for low computational resource environments such as physical computing, edge computing or neuromorphic computing.

5. Discussion, Limitations and Outlook

Neural network models of biological systems and cognition have lead to many insights into the nature of information processing in biological systems (Yamins, Wang et al., 2021). Switchable networks that recapitulate features of neuromodulation may provide insight into the nature of information processing in biological systems. For example, neuromodulation enables state dependent switching of attractive/aversive responses to food based on hunger, based on gain control of opposing pathways (Vogt et al., 2021). It would be interesting to observe whether training of neuromodulator inspired networks from a naïve state (Wang et al. 2021) could lead to creation of similar opposing pathways. Similarly, constraints that circuits be stable and switchable under different types of neuromodulation may be responsible for the small number of observed stomatogastric ganglion circuits given the large number of circuits that generate plausible rhythms (Prinz and Marder, 2004). Neuromodulator-constrained biophysical circuit or neural network models could allow these hypotheses to be evaluated.

Biological systems are often switched by multiple neuromodulators to achieve multiple functions. We posit that the diverse toolkit of neuromodulators found in biology provides a range of stereotyped affine (or other) transformation directions which allow networks to move towards different parts of parameter space. In this work, we only show switching between a set of two gates at a time. As the number and complexity of tasks increases, there is a risk that random fixed transformations may never connect multiple separated target surfaces in parameter space. This challenge would be ameliorated by access to multiple networks to traverse key directions in some lower-dimensional parameter space. This presents neuromodulators as an interesting circuit design tool and is consistent with the way neuromodulators such as serotonin often have opposing actions in different brain regions. Neuromodulators may simply provide access to an appropriate direction in parameter space for a given transformation. A small library of orthogonal neuromodulators could span different directions, allowing for parameter-free optimization to choose an appropriate neuromodulator for some task switching process. The precision required by the transformation would be relaxed by the implication that switchable networks may only need to find acceptably good loss surfaces, rather than local or global loss minima with tighter clustering. Future work may look to develop networks that are constrained to switch between multiple different objective functions by different neuromodulators.

The modern discipline of deep learning has descended from developments in cognitive neuroscience (Mcclelland and Rumelhart, 1994). Neuromodulatory systems have been broadly influential in many important architectures (Mei, Muller and Ramaswamy, 2022). For example, the reward prediction error, signalled by the dopamine system, has been an inspiration for deep reinforcement learning (Schultz, Dayan and Montague, 1997). Further exploration of neuromodulator-like control like systems that enable transitions between states or dynamic reconfiguration of network properties could enable alternate neural network control strategies. Principles governing such reconfigurable circuits could find use in future deep learning ”meta-networks” to allow for flexible and adaptive processing.

Understanding the information content of such switchable systems compared to non-switchable systems may enable approaches to more compact neural networks. Given that a small affine transformation is not difficult to store, a network which switches from AND/OR essentially compresses an AND network and an OR network into a single network. This may be of benefit in situations where complexity is expensive. For instance, in biology, there is a physical, energetic cost to every computation and maintenance of every synapse. For very common, fundamental responses such as sleep-wake transitions (Flavell et al., 2013), changes in peristaltic rhythms (Hooper and Marder, 1987), aversive/attractive responses towards food (Vogt et al., 2021) and mating, encoding large-scale transitions in a small number of global variables through neuromodulation may have advantages in efficiency and robustness. Such compactness could be valuable when implementing deep-learning inspired methods in constrained computing platforms. For instance, many emerging applications, such as control of microrobots, operate under significant constraints on both energy usage and physical wiring complexity, while unconventional physical neuromorphic computing approaches (Lee et al., 2024) have limited degrees of freedom for reconfiguration. Encoding complex computational transitions into simple physical transitions could enable task-adaptive networks in such contexts. Efficiency and robustness should be be carefully quantified in these different contexts to understand the specific advantages of globally switched, neuromodulator-like control in comparison to task switching via other approaches such as extra-inputs or transfer learning.

Finally, biological systems commonly have many recurrent connections. Recent work has shown that neuromodulator inspired synaptic scalar or low-rank gain control allows for gating of network responses and generation of diverse behaviors. However, neuromodulators in practice act differentially on different biological neurons based on their expressed receptor density. A matrix type modulation, similar to the affine transformation performed in this study, may better reflect biological neuromodulation. Applying the methods of this study to more biologically relevant neural network models such as recurrent neural networks (Tsuda et al., 2021, Costacurta et al., 2024, Miconi et al. 2018) may better capture the complex behavioral changes achieved by chemical neuromodulation in biological systems.

References

Harris-Warrick, R.M. and Marder, E. (1991). Modulation of Neural Networks for Behavior. Annual Review of Neuroscience, 14(1), pp.39–57.

Flavell, Steven W., Pokala, N., Macosko, Evan Z., Albrecht, Dirk R., Larsch, J. and Bargmann, Cornelia I. (2013). Serotonin and the Neuropeptide PDF Initiate and Extend Opposing Behavioral States in C. elegans. Cell, 154(5), pp.1023–1035.

Kirst, C., Modes, C.D. and Magnasco, M.O. (2017). Shifting attention to dynamics: Self-reconfiguration of neural networks. Current Opinion in Systems Biology, 3, pp.132–140.

Hooper, S.L. and Marder, E. (1987). Modulation of the lobster pyloric rhythm by the peptide proctolin. Journal of Neuroscience, 7(7), pp.2097–2112.

Marder, E. (2012). Neuromodulation of Neuronal Circuits: Back to the Future. Neuron, 76(1), pp.1–11.

McClelland, James L., and David E. Rumelhart. 1988. Explorations in parallel distributed processing: a handbook of models, programs, and exercises. Cambridge, Massachusetts: MIT Press

Tsuda, B., Pate, S.C., Tye, Kay M., Siegelmann, H.T., Sejnowski, T.J. (2021). Neuromodulators generate multiple context-relevant behaviors in a recurrent neural network by shifting activity flows in hyperchannels bioRxiv : pages 2021–05, https://doi.org/10.1101/2021.05.31.446462

Costacurta, J. C., Bhandarkar, S., Zoltowski, D. M., Linderman, S. W. (2024). Structured flexibility in recurrent neural networks via neuromodulation. bioRxiv : the preprint server for biology, pages 2024.07. https://doi.org/10.1101/2024.07.26.605315

Salvan, P., Fonseca, M., Winkler, A.M. et al. Serotonin regulation of behavior via large-scale neuromodulation of serotonin receptor networks. Nat Neurosci 26, 53–63 (2023).

Froudist-Walsh, S., Xu, T., Niu, M. et al. Gradients of neurotransmitter receptor expression in the macaque cortex. Nat Neurosci 26, 1281–1294 (2023)..

Sun ZY, et al. Differential distribution of functional receptors for neuromodulators evoking short-term heterosynaptic plasticity in Aplysia sensory neurons. J Neurosci. 1996 Dec 1;16(23):7540-9.

Lampinen, A.K. and McClelland, J.L. (2020). Transforming task representations to perform novel tasks. Proceedings of the National Academy of Sciences, 117(52), pp.32970–32981.

David Ha, Andrew Dai, and Quoc V Le. (2016). Hypernetworks. arXiv preprint arXiv:1609.09106,

Chalvidal, M., Serre, T., VanRullen, R. (2022). Meta-reinforcement learning with self-modifying networks. Advances in Neural Information Processing Systems, 35, 7838-7851.

Thomas Miconi, Aditya Rawal, Jeff Clune, and Kenneth O Stanley.(2018) Backpropamine: training self-modifying neural networks with differentiable neuromodulated plasticity. ICLR.

Chitturi, S.R., Ramdas, A., Wu, Y. et al. Targeted materials discovery using Bayesian algorithm execution. (2024) npj Comput Mater 10, 156.

Prinz, A.A., Bucher, D. and Marder, E. (2004). Similar network activity from disparate circuit parameters. Nature Neuroscience, 7(12), pp.1345–1352.

Vogt, K., Zimmerman, D.M., Schlichting, M., Hernandez-Nunez, L., Qin, S., Malacon, K., Rosbash, M., Pehlevan, C., Cardona, A. and Samuel, A.D.T. (2021). Internal state configures olfactory behavior and early sensory processing in Drosophila larvae. Science Advances, 7(1).

Wang, P.Y., Sun, Y., Axel, R., Abbott, L.F. and Yang, G.R. (2021). Evolving the olfactory system with machine learning. Neuron, 109(23).

Mei, J., Muller, E. and Ramaswamy, S. (2022). Informing deep neural networks by multiscale principles of neuromodulatory systems. Trends in Neurosciences, 45(3), pp.237–250.

Schultz, W., Dayan, P. and Montague, P.R. (1997). A Neural Substrate of Prediction and Reward. Science, 275(5306), pp.1593–1599.

Vecoven, N., Ernst, D., Wehenkel, A. and Drion, G. (2020). Introducing neuromodulation in deep neural networks to learn adaptive behaviours. PLOS ONE, 15(1), p.e0227922.

Lee, O., Wei, T., Stenning, K.D. et al. Task-adaptive physical reservoir computing. (2024). Nat. Mater. 23, 79–87

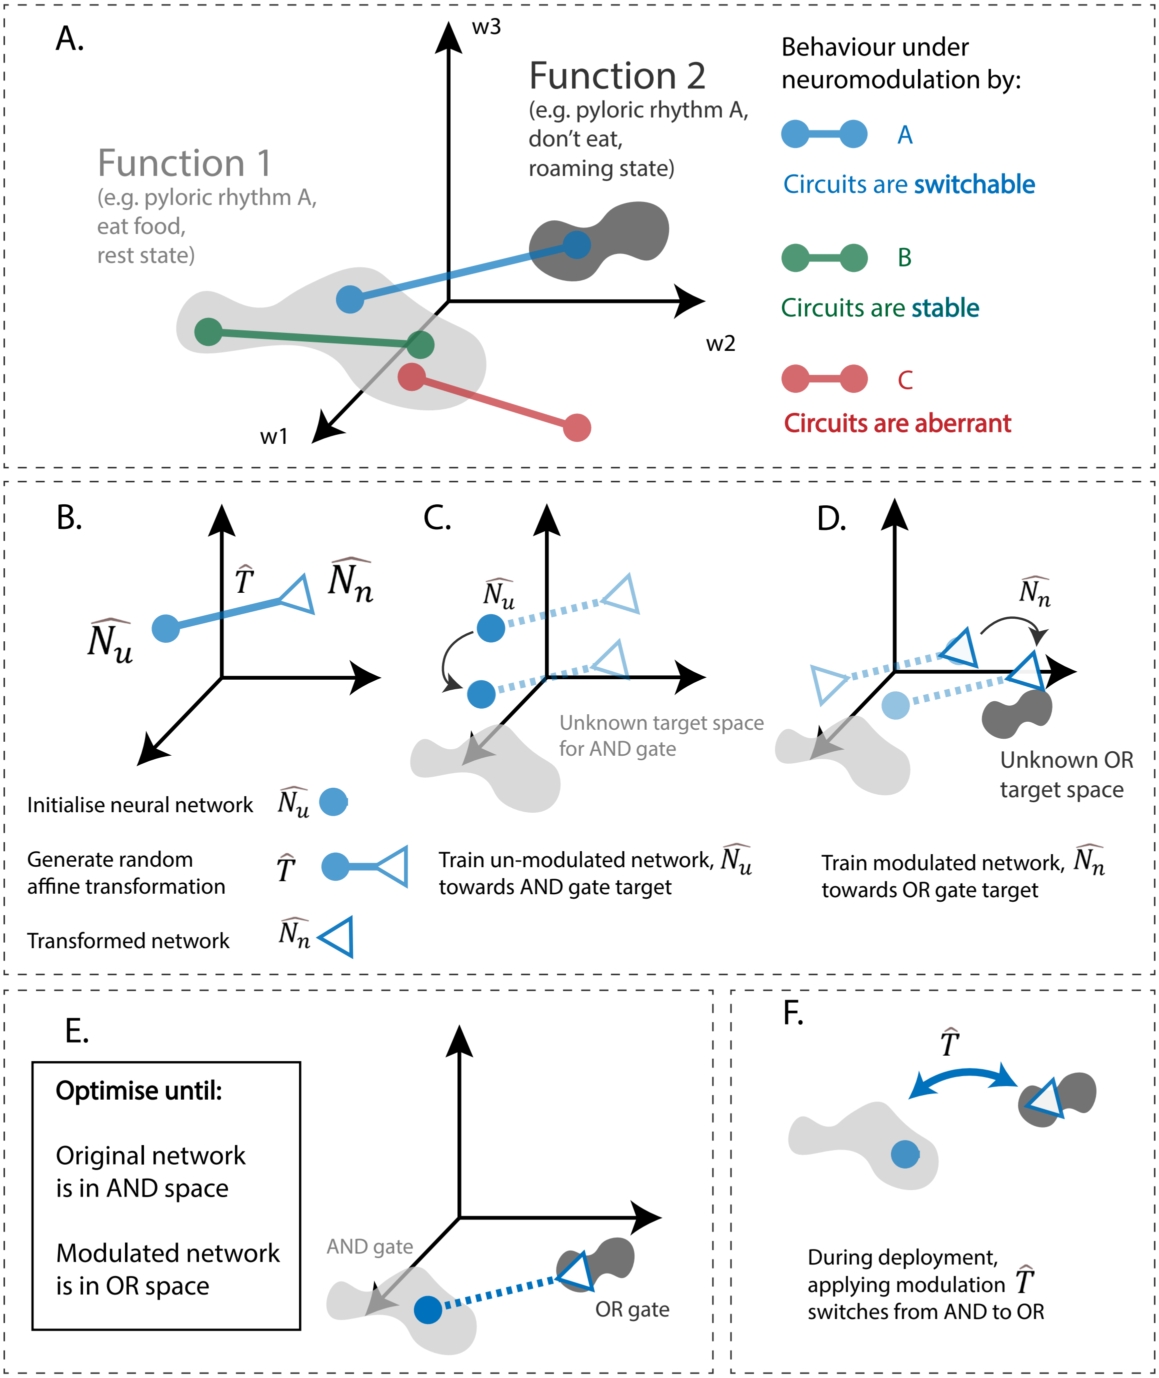

Figure 1: (A) Biological neuromodulators enable switching between different circuit states through constrained parameter space traversals. (B) Our approach uses a fixed transformation matrix M to switch between networks. (C-E) During training, networks f and f' are co-optimized for different objective functions. (F) During inference, applying M allows switching between network functions.

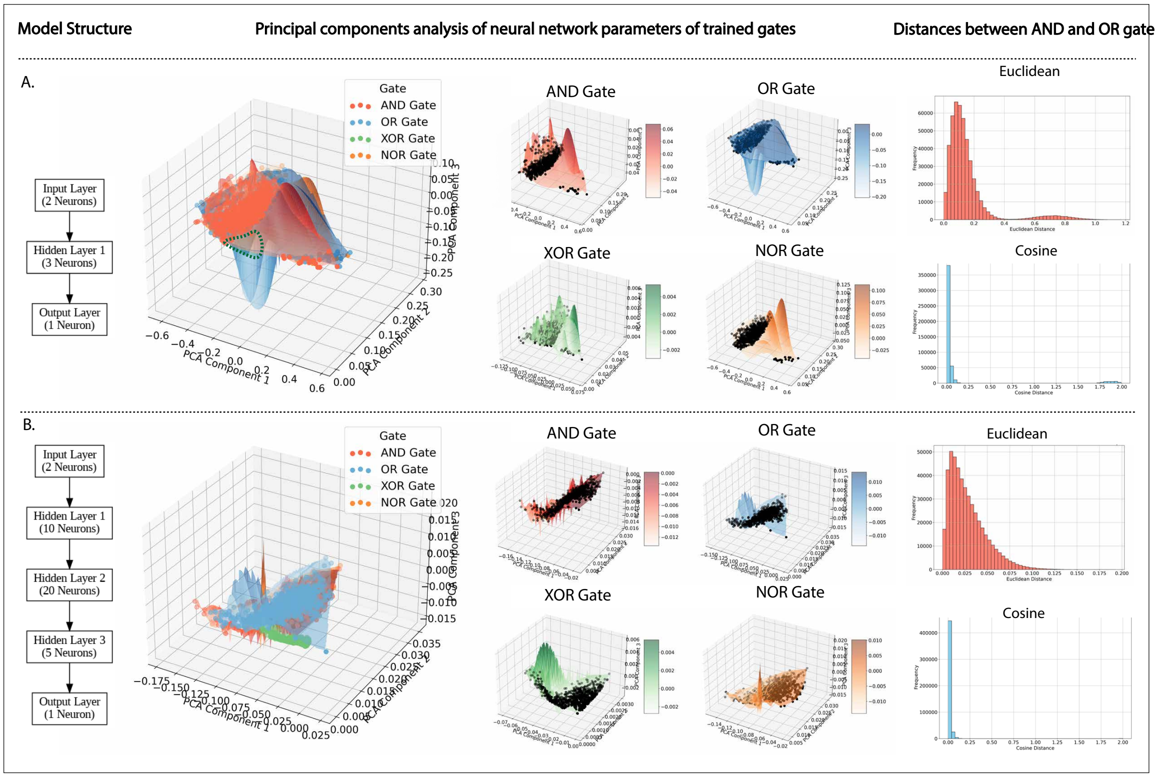

Figure 2: PCA visualization of neural network parameters after training for different logical gates. (A) Two-layer networks showing a broad distribution in in parameter space for AND and OR, but XOR occupying a distinct region. (C) Four-layer networks exhibit sharper, more concentrated clusters with well-defined distance distributions.

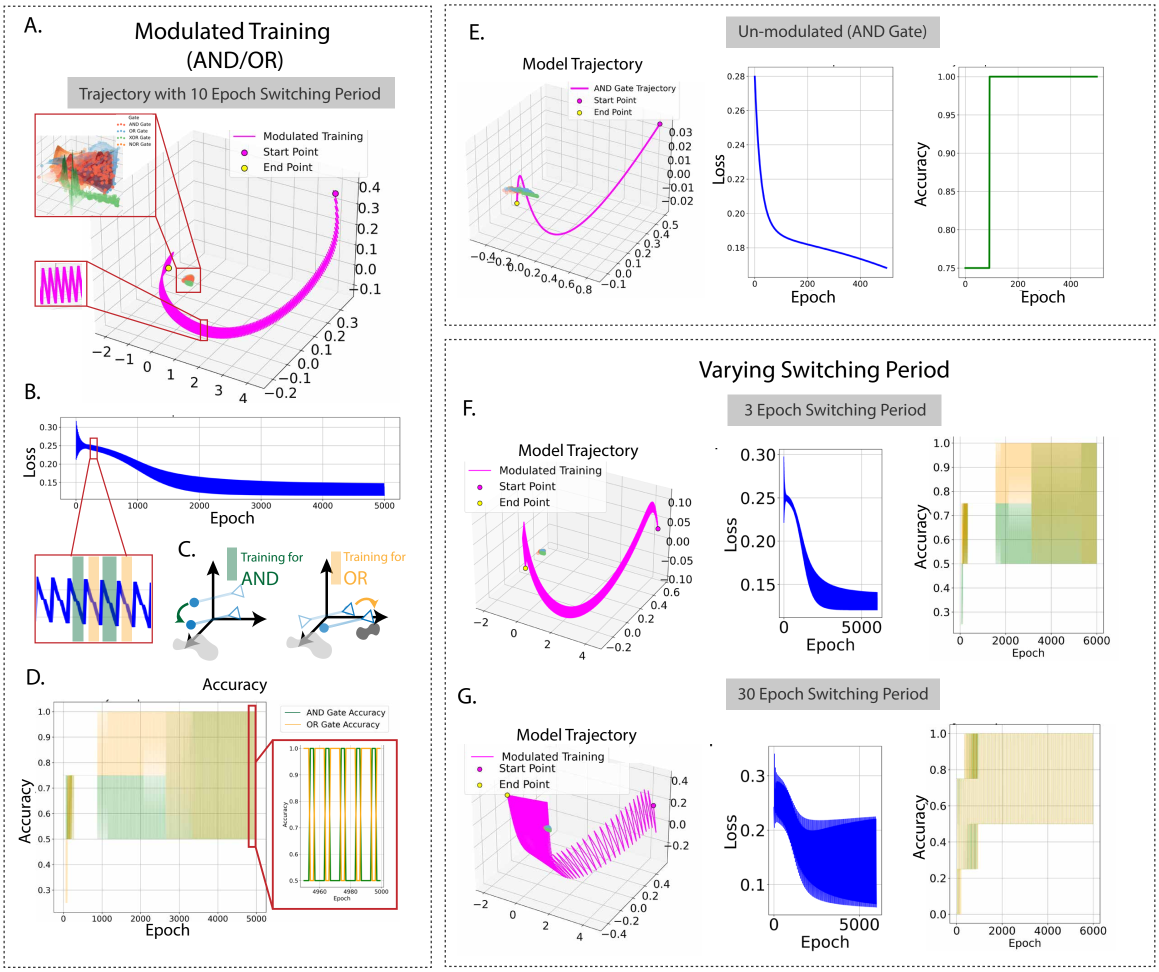

Figure 3: Training dynamics of switchable AND/OR networks. (A) Parameter space trajectory showing jagged path during switching. (B) Loss evolution during training. (C) Loss spikes upon objective switching followed by gradual reduction. (D) Network accuracy alternating between AND and OR gates. (E) Comparison with unmodulated training. (F) Effect of shorter switching period (3 epochs) resulting in smoother trajectory. (G) Effect of longer switching period (30 epochs) showing more pronounced deviations.

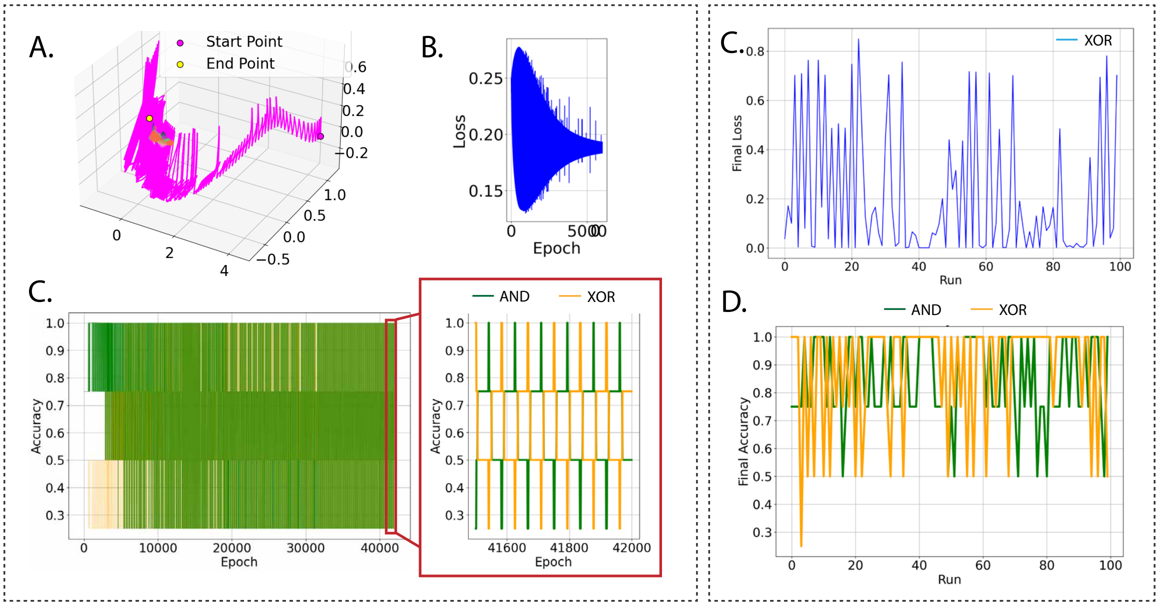

Figure 4: Switching between AND and XOR gates. (A) Training trajectory concentrates near the XOR region in parameter space. (B) Loss evolution during AND/XOR switching. (C) Success rate comparison showing approximately 30% success for AND/XOR switching compared to the higher success rate of AND/OR switching.

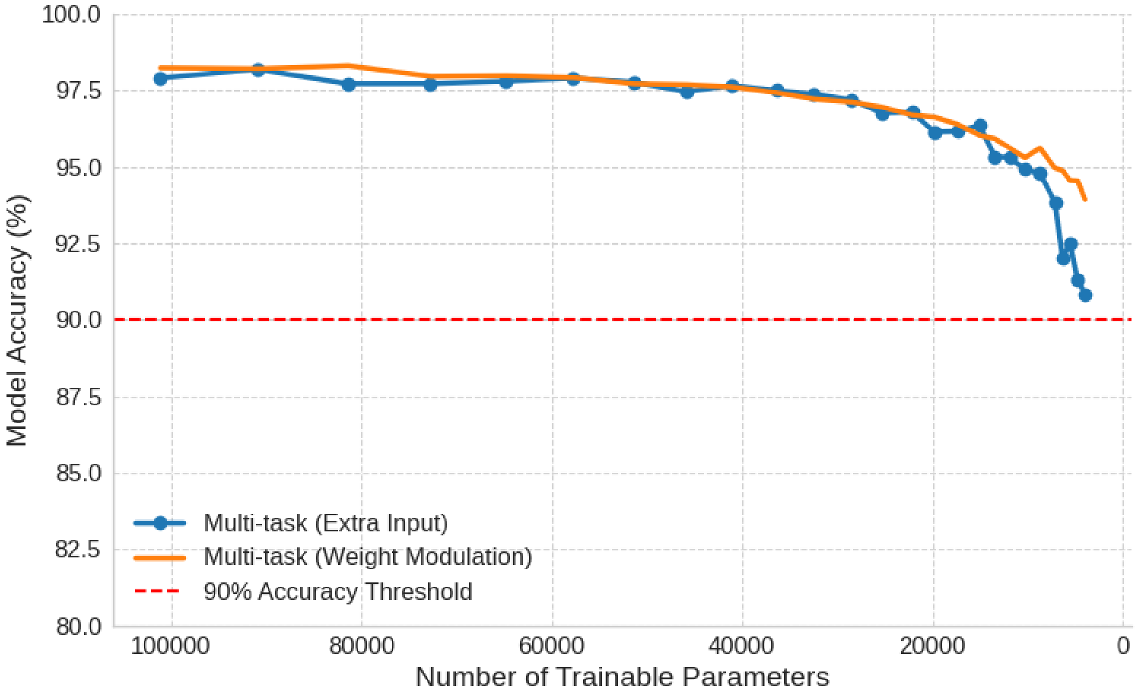

Figure 5: Model accuracy vs. parameter Count for two multi-task learning strategies. The plot compares the average classification accuracy on two distinct 5-class MNIST tasks for a model using an explicit task indicator (Extra Input) versus a model using neuromodulator-inspired weight modulation. Performance is shown as the number of network parameters is reduced.