

Neuromodulators enable overlapping synaptic memory regimes and nonlinear transition dynamics in recurrent neural networks

Ben Tsuda ^{1,2,3,*}, Stefan C. Pate ¹, Kay M. Tye ⁴, Hava T. Siegelmann ⁵, Terrence J. Sejnowski ^{1,6,7,*}

¹ Computational Neurobiology Laboratory, The Salk Institute for Biological Studies, La Jolla, CA 92037, USA

² Neurosciences Graduate Program, University of California San Diego, La Jolla, CA 92093, USA

³ Medical Scientist Training Program, University of California San Diego, La Jolla, CA 92093, USA

⁴ Systems Neuroscience Laboratory, The Salk Institute for Biological Studies, La Jolla, CA 92037, USA

⁵ Biologically Inspired Neural & Dynamical Systems Laboratory, School of Computer Science, University of Massachusetts Amherst, Amherst, MA, 01003

⁶ Institute for Neural Computation, University of California San Diego, La Jolla, CA 92093, USA

⁷ Division of Biological Sciences, University of California San Diego, La Jolla, CA 92093, USA

* Correspondence: btsuda@salk.edu (B.T.), terry@salk.edu (T.J.S.)

Summary

Mood, arousal, and other internal neural states can drastically alter behavior, even in identical external circumstances — the proverbial glass half full or empty. Neuromodulators are critical in controlling these internal neural states, and aberrations in neuromodulatory processes are linked to various neuropsychiatric disorders. To study how neuromodulators influence neural behavior, we modeled neuromodulation as a multiplicative factor acting on synaptic transmission between neurons in a recurrent neural network. We found this simple mechanism could vastly increase the computational capability and flexibility of a neural network by enabling overlapping storage of synaptic memories able to drive diverse, even diametrically opposed, behaviors. We analyzed how local or cell-type specific neuromodulation changes network activity to support such behaviors and reproduced experimental findings of *Drosophila* starvation behavior. We revealed that circuits have idiosyncratic, non-linear dose-response properties that can be different for chemical versus electrical modulation. Our findings help explain how neuromodulation “unlocks” specific behaviors with important implications for neuropsychiatric therapeutics.

Introduction

Humans and animals exhibit a diversity of behaviors depending on neurological state. A bear in the wild spots an interesting mushroom (Figure 1A). Whether the bear devours or skeptically retreats depends on his level of hunger. A well-fed bear will forgo the risk of a poison mushroom, but a starved bear will not think twice about a fortunate mushroom meal. Most behaviors operate under a similar system, where emotion, arousal, and other internal neural states strongly influence decisions and behavior. Several recent works have sought to understand mechanisms underlying such neural states [Inagaki et al., 2012, Burgos-Robles et al., 2017, Patriarchi et al., 2018], universally revealing an integral role of neuromodulatory molecules. Yet how neuromodulators directly alter local and global network dynamics to select behaviors remains unclear.

Neuromodulators, which include serotonin, dopamine, acetylcholine, norepinephrine, melatonin, adenosine, and many neuropeptides, are linked to a variety of neurological states including emotional states, arousal states, cognitive states, and diverse behavioral states (e.g. hunger or motor protocols) [McCormick et al., 2020, Bacqué-Cazenave et al., 2020]. They are also linked to nearly

all psychiatric illnesses, with most treatment options designed to target neuromodulatory pathways [Avery and Krichmar, 2017, Katz and Edwards, 1999]. Neuromodulators act on brain circuitry by altering network properties that influence the manner of synaptic transmission whether by altering local synaptic strengths, neural excitability, or plasticity. Most neuromodulators induce their influence indirectly by binding to GPCRs on target cells and triggering a downstream signaling cascade, while a minority act directly on ion channels [Nadim and Bucher, 2012].

Pioneering work on neuromodulation showed that circuit dynamics can change with particular synapse-specific neuromodulation. Studies on the lobster pyloric network [Marder, 2012, Marder and Eisen, 1984, Johnson et al., 1995] and other systems [Katz et al., 1994, Jing et al., 2009] revealed that fine-tuned neuromodulation of individual neurons could lead to new patterns of activity underlying shifts in global network output and behavior. While specific tuning of neuromodulatory effects at each synapse could theoretically lead to entirely new network configurations with little resemblance to the non-neuromodulated network (save only that neurons retain connectivity patterns), neuromodulators acting on a particular receptor type often have a similar directional effect on synaptic properties. We considered whether the simpler case of neuromodulators acting nearly uniformly as synaptic weight amplifiers or dampeners within a local region of a neural circuit, could support different behaviors.

Different aspects of how neuromodulators may influence neural computations have been modeled including reduced models of neuromodulatory influence on Bayesian probability estimations [Yu and Dayan, 2005], neural network models of activation function modulation [Stroud et al., 2018, Vecoven et al., 2020], neuromodulatory network masking of effector networks [Beaulieu et al., 2020], and neuromodulator-specific biologically-guided circuit dynamics [Hasselmo et al., 1995]. Each of these models has provided insight into components of neuromodulatory function, yet how neuromodulation’s central role in modulating synaptic weights is utilized to control circuit dynamics and behavior remains unknown.

We adopted a complementary modeling approach, approximating neuromodulators as amplifiers/dampeners of effective synaptic weight, a well-established mechanism by which neuromodulators exert influence on neural circuits [Marder, 2012, Nadim and Bucher, 2012, Katz et al., 1994, Johnson et al., 1995]. In the first part of this paper, we explore the effect of this mechanism in a neural network, demonstrating its flexibility and robustness in supporting multiple behaviors. We show how neuromodulation establishes overlapping synaptic memories within a neural network that enable new behaviors even in response to identical stimuli and detail the conditions that support this phenomenon.

In the second part, we describe the mechanisms by which uniform scaling of synaptic weights shifts network behavior. We find our model reproduces and provides a mechanistic explanation for experimental findings on neuromodulation-dependent starvation behavior in *Drosophila*. We discover that individual neural circuits exhibit idiosyncratic, nonlinear transition dynamics between neuromodulated dynamical states. We investigate the characteristics of this “circuit-based sensitivity” in response to both chemical and electrical modulation and compare their relationship. These analyses provide a parsimonious explanation of how neuromodulators unlock state-dependent behaviors and suggest a powerful, previously unreported mechanism of therapeutic susceptibility and resistance with important implications for development of future therapeutics.

Results

Neuromodulation supports multiple weight regimes within the same synaptic connections. A predominant effect of many neuromodulators is alteration of synaptic transmission. Though neuromodulation in the brain is multi-faceted and nuanced depending on receptor and target cell types [Bacqué-Cazenave et al., 2020], we sought to understand how neuromodulatory effects on synaptic weights can alter circuit behavior, an important functional role of neuromodulators [Marder, 2012, Nadim and Bucher, 2012]. To do this, we adopted a simplified approximation of the aggregate effect on synaptic transmission as a uniform amplification or dampening shift in synaptic weights of all neuromodulation-targeted neurons. We first assessed if this simple uniform weight scaling could support divergent behaviors within a single neural network.

To assess if such a neuromodulatory effect could toggle network behaviors, we developed the modified Go-NoGo task (Figure 1B). For the classic Go-NoGo task, a network is given either a positive pulse or a null stimulus (+ or \emptyset in Figure 1B, respectively) and must generate a +1 or 0 output, respectively. In the modified Go-NoGo task, a network is trained to perform the classic Go-NoGo task (Behavior 1) in the absence of the neuromodulator’s effect, and the opposite behavior (Behavior 2: + \rightarrow 0, $\emptyset \rightarrow$ -1) given the same stimuli in the presence of the neuromodulator (see Methods). The modified Go-NoGo task presents a good assessment of the neuromodulated network’s ability to support different behaviors since a normal network presented with identical stimuli cannot support opposing behaviors.

We presented the modified Go-NoGo task to a neuromodulated recurrent neural network (RNN) in which all recurrent neural weights were dampened uniformly by a factor of 1/2 (Figure 1C). Surprisingly, we found the RNN was able to support the opposing behaviors, and independently trained models invariably replicated this finding, learning the task within a few thousand trials of training (Figure 1D). Neuromodulation thus was able to effectively separate synaptic memory regimes within a single network and access them through uniform scaling of weights to “unlock” specific behaviors (Figure S1). We found this result held over a wide range of neuromodulatory factors (Figure S2).

One may imagine that scaling of weights by a uniform factor may translate into a simple scaling of network activity or activities of individual neurons in the network. We found that global network activity (Figure 1E) and individual neuron activities (Figure 1F–G) did not follow such a simple relationship and were unpredictable. Individual neurons showed wide variation in altered activity patterns both in mean firing rate (Figure 1F) and in time-course activity pattern (Figure 1G) in a neuromodulation and stimulus dependent manner. Individual neuron activity profiles showed complex transformations under neuromodulation (Figure 1G) reminiscent of altered individual neural activity patterns reported in lobster stomatogastric ganglion and other lower organisms [Marder, 2012, Johnson et al., 1995, Johnson et al., 2011, Katz et al., 1994].

Neuromodulation of a network subpopulation can toggle global neural network states. In the brain, neuromodulators are released in specific regions – some tightly localized, others broadcast widely across brain regions – to influence local and global neural output. We asked if neuromodulation applied to specific subpopulations in our model could effectively toggle global network behavior. First, we assessed if neuromodulation of varying neural subpopulation sizes embedded within a RNN could support the modified Go-NoGo task (Figure 2A). We found that RNNs with neuromodulated subpopulations across the range of sizes tested from 100% to 10% of

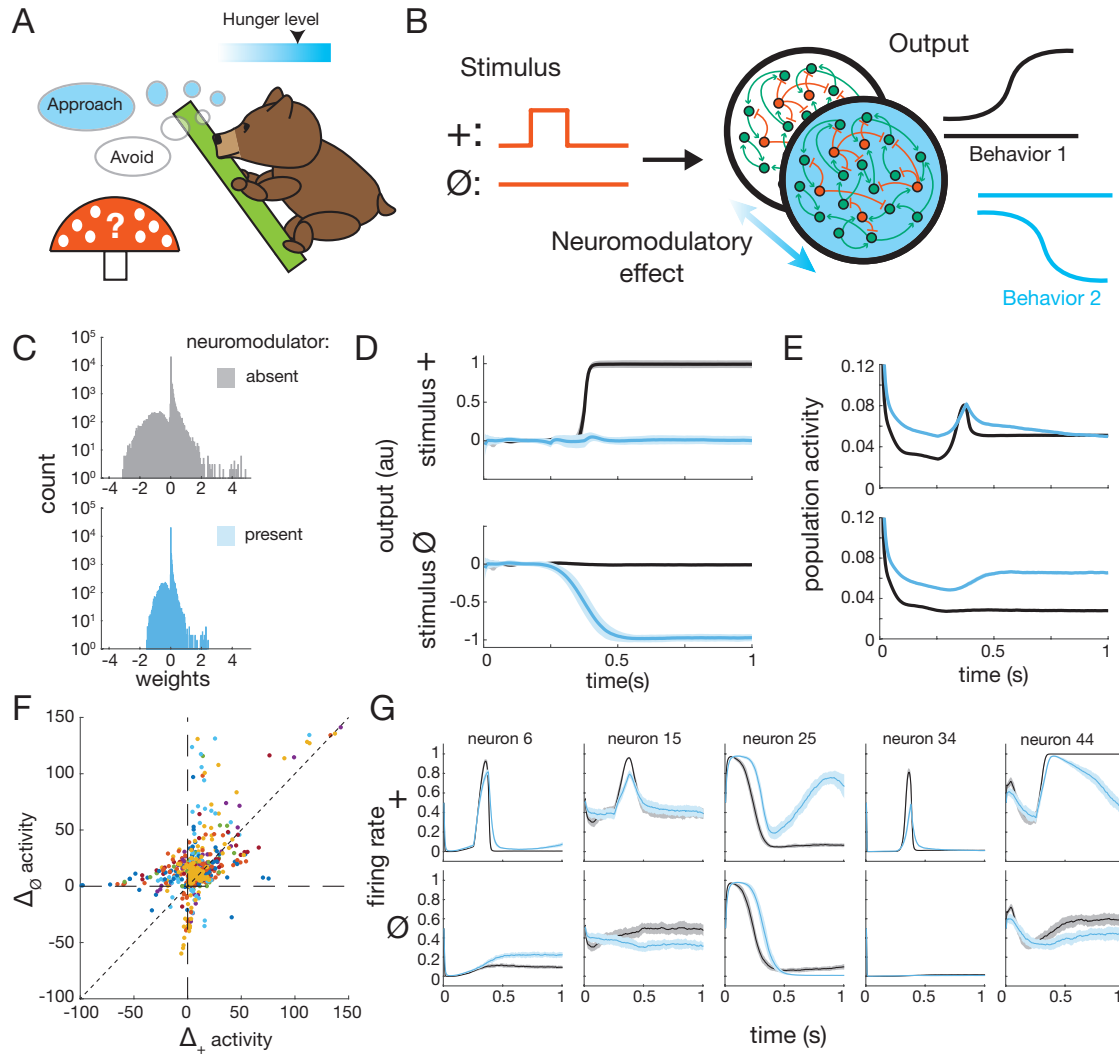


Figure 1. Neuromodulation weight scaling separates overlapping synaptic memory regimes. (A), Hunger level dictates a bear's behavior when encountering a mystery mushroom. (B) Modified Go-NoGo task. Given a stimulus (either + or \emptyset), in absence of neuromodulatory effect the recurrent neural network (RNN) should produce outputs from the Behavior 1 repertoire and in presence of neuromodulator, from Behavior 2. (C) Neuromodulatory effect uniformly scales all recurrent weights in the network, here with a factor 0.5. (D) Mean output (with shaded standard deviation) to + and \emptyset stimuli of 10 independently trained RNNs on the modified Go-NoGo with global neuromodulation factor 0.5. (E) Mean whole network population activity for example RNN over 100 trials after training. Activity under neuromodulation (blue) is not simple transform of activity without neuromodulation (black). (F) Difference of mean activity with and without neuromodulation (Δ activity) on + vs \emptyset stimulus trials for individual neurons. Each color represents an independently trained RNN (10 colors total). Points representing simple scaling of neural activity under neuromodulation would lie on dotted diagonal line. (G) Five example neurons activity patterns from an RNN show complex nonlinear transformations under neuromodulation.

neurons could consistently support the opposing behaviors (Figure 2B–C).

Some neuromodulators affect neurons in a cell-type specific manner, for example selectively influencing activity of excitatory or inhibitory neurons with corresponding receptors [Wester and

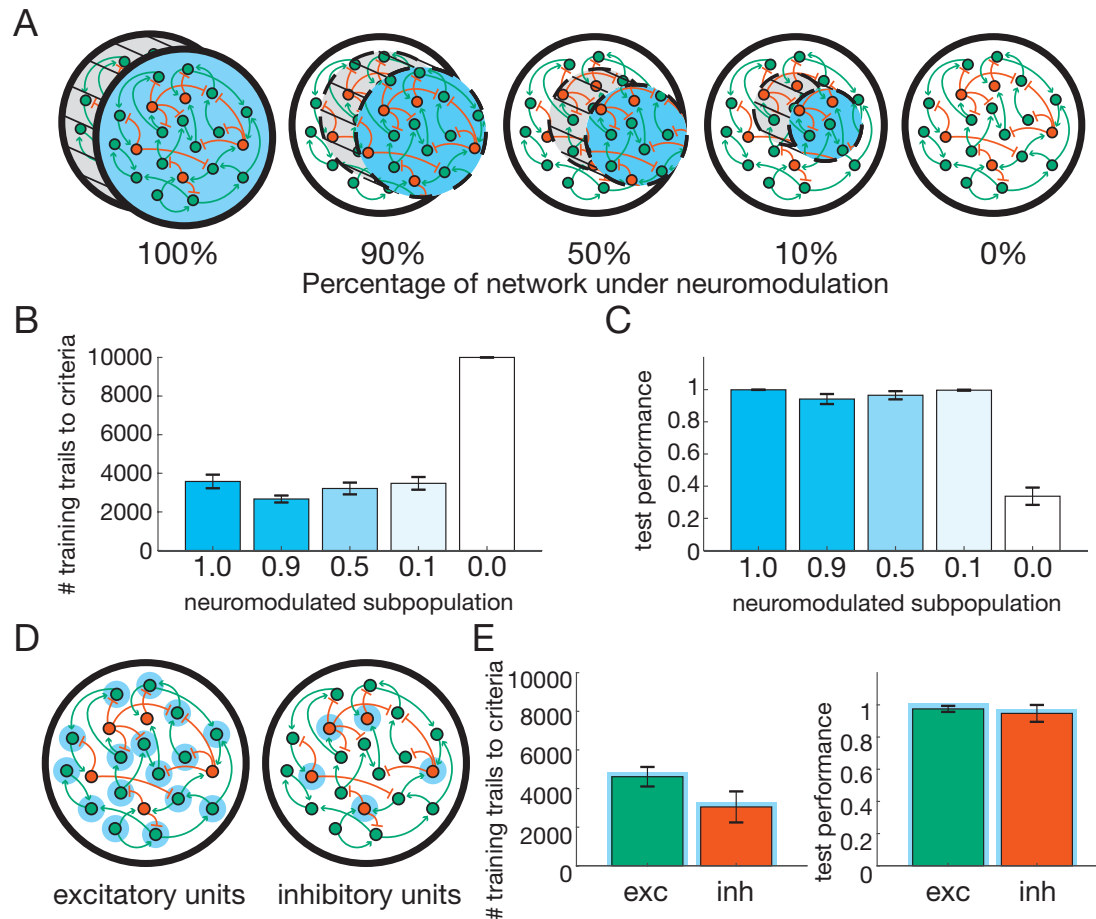


Figure 2. Targeted neuromodulation flexibly supports opposing behaviors. (A) A range of different sized neural subpopulations embedded within a RNN were neuromodulated. 100% was positive control demonstrated in Fig. 1; 0% was negative control. (B)–(C), RNNs with embedded neuromodulated subpopulations across the size spectrum could support the opposing behaviors of the modified Go-NoGo. (B) Number of training trials to reach stop criteria (see Methods). (C) Test performance (1 is 100% correct; see Methods). (D) Neuromodulation of exclusively excitatory or inhibitory neurons (blue annuli). (E) Excitatory or inhibitory neuromodulation supported learning of modified Go-NoGo.

McBain, 2014]. To assess if our model recapitulated this feature, we applied the neuromodulatory effect on excitatory or inhibitory neurons exclusively (Figure 2D). In both cases the RNN was able to support the opposing behaviors (Figure 2E).

A collection of neural states can be unlocked with multiple neuromodulators acting on multiple neural subpopulations. To assess the flexibility of this neuromodulatory mechanism, we asked whether multiple unique behaviors could be learned and unlocked from a single network through targeted neuromodulation. Using the modified Go-NoGo task extended to a third possible behavior ($+ \rightarrow -1$, $\emptyset \rightarrow +1$; see Methods), we found that both neuromodulation of distinct subpopulations by the same factor (e.g. factor of 2.5) or distinct neuromodulation levels applied to the whole network (e.g. factors of 0.5, 1, and 1.5) were able to support multiple behaviors (Figure 3A–B and 3C–D, respectively).

We then extended the modified Go-NoGo task to the maximum number of distinct behaviors given 3

possible outputs $(-1, 0, +1)$ for each of two stimuli $(+, \emptyset)$ ($P^R(3, 2) = 9$ behaviors; “9-behavior modified Go-NoGo”) and assessed if neuromodulation could support the 9 distinct behaviors (Figure 3E; see Methods). We found that neuromodulation of 9 non-overlapping subpopulations within a RNN with a single factor could support all 9 behaviors (Figure 3F). Furthermore, neuromodulation of overlapping subpopulations to varying amounts could also support the 9-behavior paradigm (Figure S3). We found that modulation of the entire RNN at different neuromodulation levels could consistently support 3 behaviors — 5/5 correct/total attempted — but failed to learn higher numbers of behaviors (4-behavior 0/5; 5-behavior 0/5; 9-behavior 0/5). This may relate to the magnitude of separation between neuromodulation levels, and also suggest a limit to whole RNN neuromodulation. In fact, though there is evidence in the brain of different levels of neuromodulation shifting to qualitatively distinct neural states, none report more than 3 levels of neuromodulation [Hasselmo et al., 1995]. Our finding supports this observation, suggesting neuromodulation of distinct subpopulations is a more robust and flexible mechanism than neuromodulation of one set of neurons at different neuromodulation strengths.

These results demonstrate neuromodulation’s flexibility and robustness shifting neural states to unlock behaviors from a single network. We next sought to understand how simple scaling of weights could lead to new complex activity dynamics supporting distinct behavior.

Neuromodulation alters individual neural output and global network cluster activity by shifting relative synaptic weight contributions. To gain insight into how neuromodulation alters neural dynamics to drive different behaviors, we first considered dynamics of single neurons in the RNN. As described previously, individual neurons could shift to qualitatively different activity patterns under neuromodulation (Figure 4A). Given that neural activity in a RNN is a complex function of all the activities of connected neurons tracing all the way back in time, we focused on the first time point in a trial. For any trial, at $t=0$ all activities are randomly initialized from a normal distribution (see Methods). As a result, at $t=1$, a neuron reacts only to the weighted inputs of its incoming connections, uncontaminated by propagating recurrent activity dynamics from past timepoints. We found that under neuromodulation, incoming weighted activity to a neuron could shift slightly which could translate into significant shifts in the neuron’s activity if occurring along the linear regime of its activation function (Figure 4B).

The shift in incoming weighted activity results from incoming excitatory and inhibitory tone. When the whole RNN was neuromodulated, although the ratio of excitatory to inhibitory weights (E/I ratio) for a neuron did not change (Figure 4C, top), the net difference between excitatory and inhibitory tone (E-I difference) was amplified (Figure 4C, bottom). Altering E-I difference in turn altered the net input current to a neuron causing an activity shift. We found this same effect occurs globally through the network, scaling the global E-I difference while maintaining a constant global E/I weight ratio (Figure 4D). The shifts in synaptic currents allow new dynamics that then recurrently feedback into the network to drive unique activity profiles and end behaviors.

To understand how individual neurons functionally related to different RNN behaviors, we evaluated each neuron’s selectivity to a given neuromodulation context and stimulus (see Methods). In a whole network-neuromodulated RNN trained on the modified Go-NoGo task, we found neurons across the range of possible selectivities (Figure 4E, top); some strongly selective for a particular stimulus-context (Figure 4E, bottom right) and others non-selective (Figure 4E, bottom left). Importantly, the neuromodulatory effect is applied to every neuron, so selectivity must emerge through recurrent dynamics within the neuromodulatory regime.

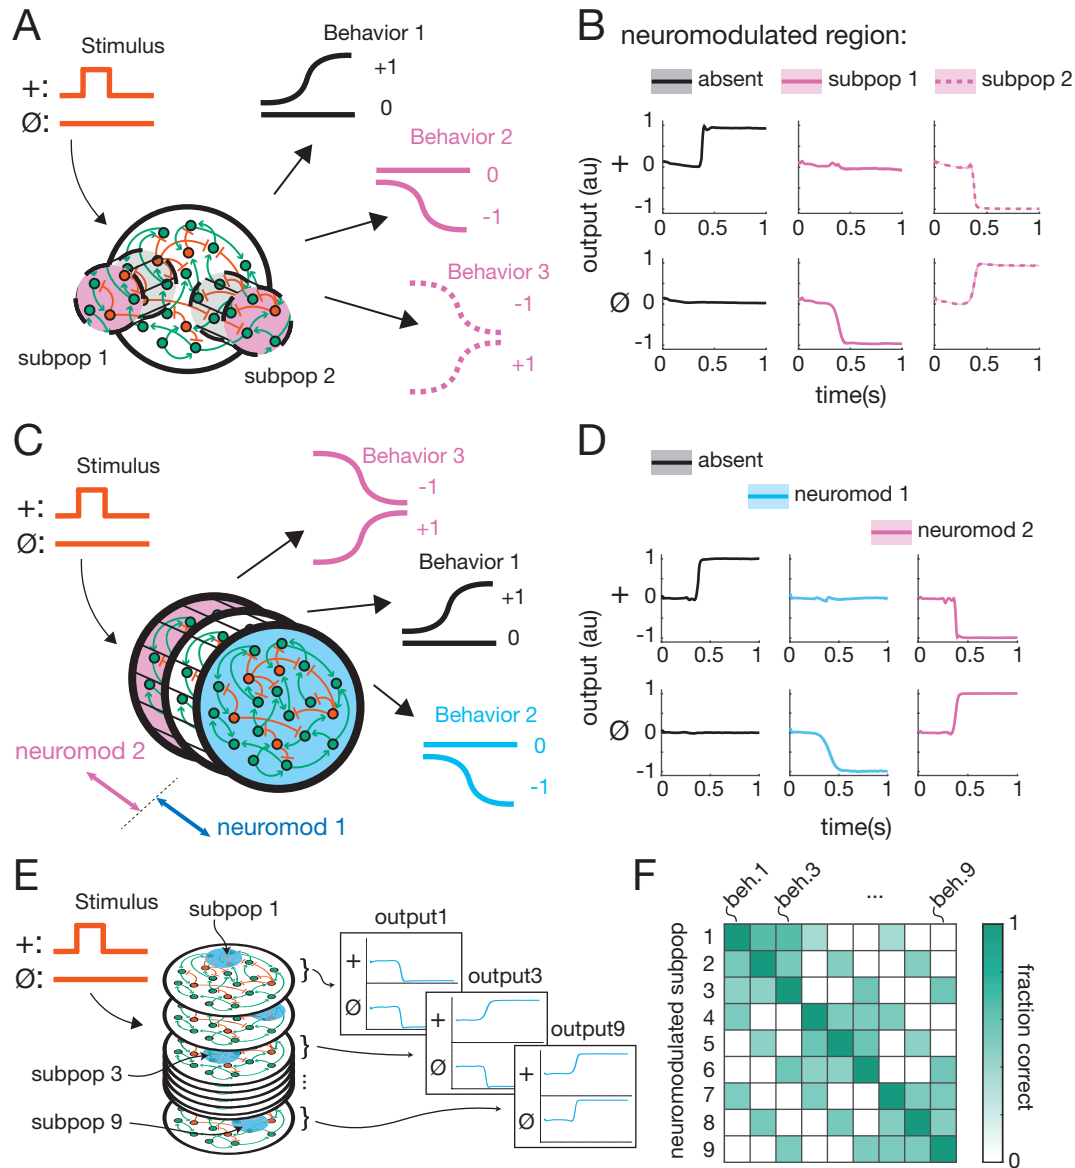


Figure 3. Multi-neuromodulator RNNs support multiple behaviors. (A) 3-behavior modified Go-NoGo. Neuromodulation (factor (f_{nm})=2.5) of either subpopulations (“subpop”) within a RNN (subpop1, subpop2) unlock unique behaviors (Behavior 2 (+ stimulus \rightarrow 0 output, $\emptyset \rightarrow$ -1), Behavior 3 (+ \rightarrow -1, $\emptyset \rightarrow$ +1), respectively). (B) Multi-subpop targeted RNN from (A) successfully learns task. Top row: RNN output to + stimulus when different subpops are neuromodulated; bottom row: same for \emptyset stimulus. (C) Same task where different levels of neuromodulation (neuromod 1 (f_{nm} =0.5), neuromod 2 (f_{nm} =1.5)) are applied to whole RNN. (D) Multi-neuromodulator RNN successfully learns task. (E) 9-behavior modified Go-NoGo task with unique neuromodulated subpops and example corresponding outputs. (F) RNN successfully learned task with 9 targeted subpops (each 10% of RNN, non-overlapping; f_{nm} =2.5). Application of neuromodulator to any subpop unlocked a specific behavior set (beh.) from the 9-behavior repertoire (fraction of trials correct is \approx 1 on diagonal; see Methods). Off-diagonal fraction correct (e.g. row 1, column 3) due to partial output overlap between behavioral sets (e.g. \emptyset stimulus output for task 1 and 3 as shown in part e).

We further asked how neurons across the selectivity spectrum may coordinate to drive different

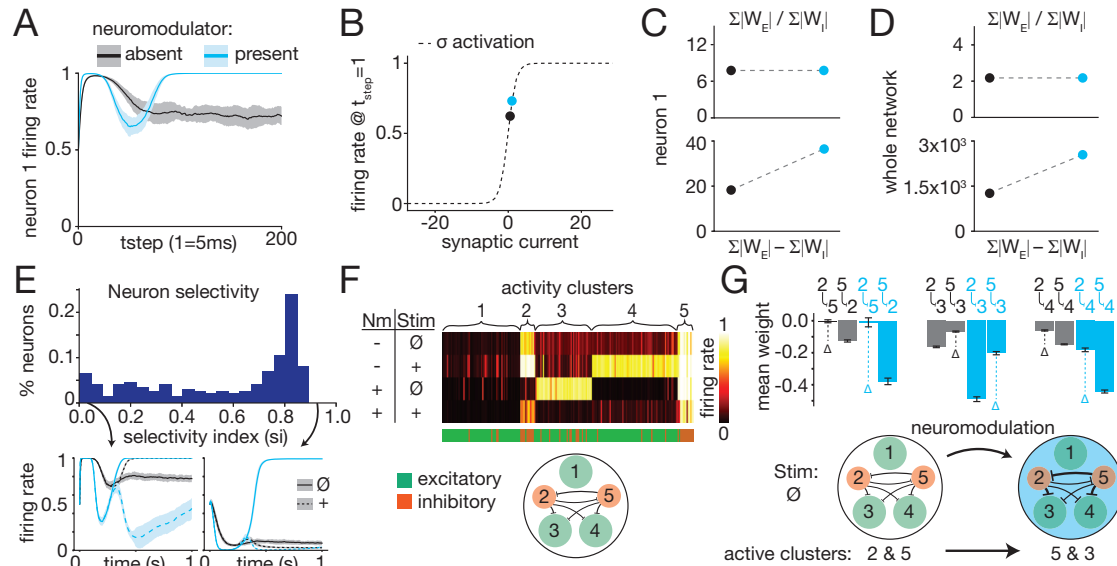


Figure 4. E-I difference scaling drives differential activity patterns. (A) Differential response of an individual neuron (neuron 1) in a whole-network neuromodulated RNN to same stimulus depending on neuromodulation presence. (B) At the start of a trial, neuromodulation causes neuron 1 to receive different synaptic current input, shifting its firing rate. (C) The different synaptic input under neuromodulation occurs due to amplification of the net difference in incoming excitatory and inhibitory weights; E/I balance is unchanged. (D) E-I difference across the whole network is also amplified; E/I remains unaltered. (E) Though all neurons in the RNN are influenced by the same neuromodulation, some exhibited activity selective for particular stimulus-neuromodulation combinations (high si; example neuron bottom right, si=0.89); others were less selective (low si; example neuron bottom left, si=0.01). (F) Neurons formed 5 clusters—2 predominantly inhibitory, 3 predominantly excitatory—whose activity coded for different stimulus-neuromodulation combinations. Each column of the heatmap is the mean firing rate of an individual neuron across conditions with excitatory/inhibitory identity labeled below. (Nm: neuromodulation present or absent; Stim: stimulus presented) (G) Amplification of relative weight differences (Δ s of mean weight) between inhibitory clusters drives cluster activity switches under neuromodulation: increased inhibition of clusters 2 and 4 by cluster 5 and disinhibition of cluster 3. In all panels light blue represents modulator present and grey/white modulator absent. Error bars are SEM.

output behaviors. To do this we ran k-means clustering using neurons' average firing rates across trials and found neurons formed five clusters, whose activity combinations coded unique stimulus-context pairs. Interestingly, two clusters were predominantly composed of inhibitory neurons and the other three predominantly excitatory neurons (Figure 4F).

The same principle of amplification of relative E-I difference in inter-cluster synaptic weights helped explain cluster activity shifts associated with behavioral shifts under neuromodulation. For example, in the absence of neuromodulator the null stimulus led to balanced mutual inhibition between clusters 2 and 5 in the RNN and a resulting output of NoGo. When neuromodulator was applied, the slightly stronger inhibition from cluster 5 to cluster 2 became amplified, silencing cluster 2 (and 4) and disinhibiting cluster 3, driving output of AntiGo (Figure 4G, S4). To further examine the functional contribution of each cluster we performed progressive lesion studies which supported the stimulus-context dependence of each cluster, while demonstrating more complex underlying inter-cluster dynamics (Figure S4).

Although selectivity and clustering analysis provided some intuition for this particular RNN, independently trained RNNs on the same task exhibited unique clustering patterns, and different network configurations and task paradigms exhibited widely varying selectivity and clustering patterns (Figure S5). This motivated consideration of networks' population dynamics to understand more generalized features of the neuromodulatory mechanism.

Distinct global network activity manifolds with non-linear transition dynamics emerge from neuromodulation. Individual neurons in a network carry information that is often reflective of larger scale distributed representations or computations. Consideration of the entire network's dynamics can provide a more complete understanding of neural behavior, particularly when single neuron activities are difficult to interpret [Cunningham and Yu, 2014, McCormick et al., 2020, Mante et al., 2013]. To understand how neuromodulation led to network behavior shifts, we analyzed the coordinated activity of all neurons in the RNN in the absence and presence of neuromodulator. Neural activity trajectories for the same stimulus followed non-overlapping, stereotyped paths in activity space (Figure 5A). Through amplification of synaptic weights, neuromodulation effectively resets all the pins in the pinball machine, altering activity flow patterns through the RNN. To fully define the neuromodulation-dependent activity subspace, we generated random stimuli series drawn from a uniform distribution between 0 and 1 at each time point and fed this into the RNN with and without neuromodulator present. We found that the presence of neuromodulation defined a manifold non-overlapping with that of the absent neuromodulation in activity space (Figure 5B). The neuromodulation activity manifold provides access to specific behaviors while the non-neuromodulation manifold provides access to different behaviors.

Both activity manifolds described above derive from a common underlying neural network. This suggests that there must be a transition between the manifolds accessible through different amounts of neuromodulation. To characterize this manifold transition, we first trained a RNN on the modified Go-NoGo task with a global amplifying neuromodulation factor (f_{nm}) of 9. We then applied intermediate levels of neuromodulation to the RNN and observed its activity dynamics and output behavior. We found that intermediate amounts of neuromodulator ($f_{nm} \in (1:1:9)$) led to a smooth transition from trajectories on the non-neuromodulated manifold to the fully neuromodulated manifold (Figure 5C). Furthermore, output behaviors for intermediate neuromodulatory levels corresponded to intermediate outputs (Figure 5D). To characterize the transition in output behavior, we measured the output of the RNN at the midpoint of each trial for each level of neuromodulation. We found that increasing neuromodulator levels led to non-linear (exponential or sigmoidal) progression from non-neuromodulator behavior (Go: +1) to neuromodulatory behavior (NoGo: 0) (Figure 5E).

We then asked whether neuromodulatory transition dynamics were tightly constrained, defining a conserved property of neuromodulation, or might be more variable depending on individual network characteristics. To test this, we independently trained 29 RNNs and characterized their transition dynamics. All RNNs exhibited similar non-linear transition dynamics best fit by an exponential or sigmoid function (Figure S6). Surprisingly, both networks' sensitivities to neuromodulator and rates of transition varied drastically. To quantify this variability, we defined a "half maximal effective concentration" (EC50) as the amount of neuromodulator required to generate a half maximal output (see Methods). We found that the EC50 of individual networks trained with a full neuromodulation factor of 9 ranged from 2.1 to 6.5 (3.1x range; Figure 5F, left). Similarly, the rate of transition (steepness of the transition dynamics sigmoid) varied widely from 0.9 to 26.3 (Figure 5F, right). This result reveals a previously unknown phenomenon that may underlie drug sensitivity in

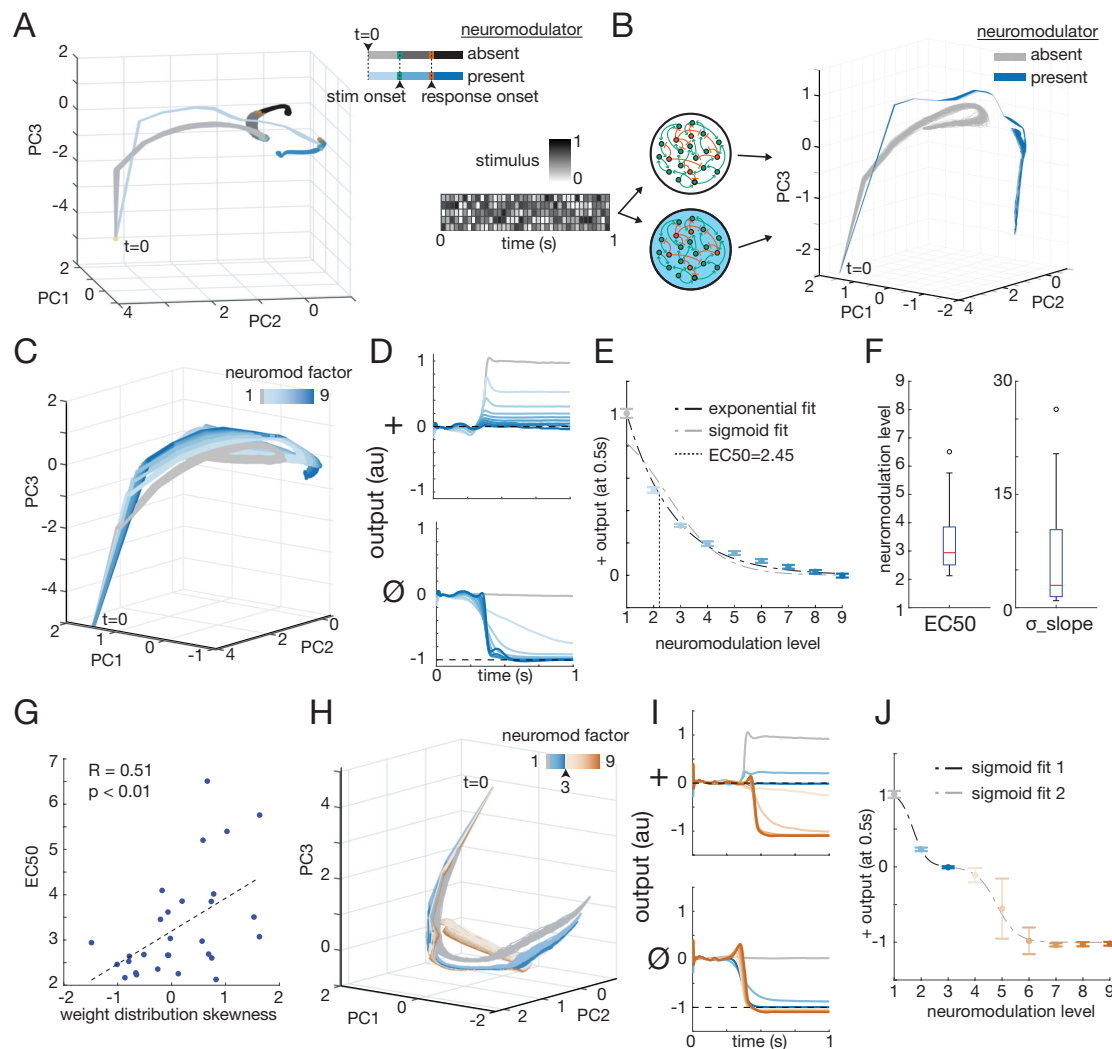


Figure 5. Neuromodulation creates separate activity manifolds with idiosyncratic nonlinear transition dynamics. (A) Global network activity dynamics after PCA-based dimensionality reduction for positive stimulus. Activity trajectories are non-overlapping. (B) Shotgun stimulus mapping of activity space inhabited by network in absence and presence of neuromodulation defines separable activity manifolds on which individual trajectories occur. (C) For RNN trained with neuromodulatory factor of 9, intermediate levels of neuromodulation lead to partial transitions toward full neuromodulation activity manifold. (D) Partial neuromodulation leads to intermediate output behaviors. (E) Transition from none to full neuromodulation behavior is non-linear (best-fit exponential & sigmoid shown) and defines network EC50. (F) 29 RNNs independently trained exhibit large variability in transition dynamics with EC50 ranging from 2.1 to 6.5 and sigmoid slope (σ_{slope}) from 0.9 to 26.3. (G) Network EC50 is positively correlated with skewness of global weight distribution. (H) For + stimulus, excessive neuromodulation pushes network activity into different space. (I) Over-neuromodulation can drive inappropriate behavior (top) or no change (bottom). (J) Transition dynamics for overmodulation is best fit by double sigmoid; first fits normal neuromodulation transition (blue) and second fits aberrant neuromodulation transition (orange). For all PCA, top 3 PCs accounted for 80–92% of activity variance.

the brain—“circuit-based sensitivity”—and may be critical to understanding and optimizing neuropsychiatric therapeutics.

Given that neuromodulation exerts its effect through amplification of weight differences, we posited that the global weight distribution of a network may relate to the idiosyncratic network sensitivities. The correlation between network EC50 and weight distribution skewness was significant ($p < 0.01$) and exhibited a positive trend ($R = 0.51$; Figure 5G), suggesting networks with more positively skewed weights (longer tail of strong excitatory weights) were less sensitive to neuromodulator. Future studies dissecting the underlying synaptic weight mechanism driving neuromodulator sensitivity will be vital to understanding and utilizing this feature clinically.

Excessive neuromodulation can also occur either pathologically or pharmacologically. To assess the effect of excessive neuromodulation in our model, we trained a RNN on the modified Go-NoGo task with a global amplifying factor of 3. We then applied neuromodulator factors ranging from 1 (no neuromodulation) to 3 (full normal neuromodulation) and beyond to 9 (3x over-neuromodulation). As before, we observed a smooth transition between non-neuromodulated and neuromodulated manifolds (Figure 5H). Interestingly, excessive neuromodulation could rapidly push neural dynamics off both manifolds into an adjacent activity space. This activity shift translated into inappropriate behavior (Figure 5I, top). The transition dynamics were defined by a double sigmoid; one sigmoid described the transition between non-neuromodulated and fully-neuromodulated and another sigmoid described the transition between fully-neuromodulated and excessively-neuromodulated conditions (Figure 5J). This activity is suggestive of attractor dynamics where excessive neuromodulation can push network activity into learned behaviors inappropriate in the given situation. This pattern mimics patients with neuromodulatory pathologies or those administered excessive doses of neuromodulatory drugs who shift into inappropriate behaviors (e.g. hypersexuality with excessive dopamine-replacement therapy, and in the other direction, tardive dyskinesia for antipsychotic dopaminergic antagonists) [Yamamoto et al., 2014, Quinn et al., 1983, Higley and Picciotto, 2014, Marder et al., 2004].

Neuromodulation-based activity manifolds explain starvation progression of dopamine-mediated *Drosophila* sugar sensitivity. We next sought to evaluate our model’s ability to recapitulate biological *in vivo* neuromodulation, and help explain experimental observations. We turned to results from Inagaki et al. 2012 which experimentally showed *Drosophila* starvation-related behavioral shifts were under the control of dopamine neuromodulation. Specifically, they reported that *Drosophila* sugar sensitivity, measured by proboscis extension reflex (PER) probability, increased with duration of starvation (fed, 1 day, 2 day starved) and the same relationship could be observed in fed flies when administered L-dopa in their diet (0, 3, 5 mg/ml) (Figure 6A).

We trained a RNN with a neuromodulated subpopulation (20% subpopulation) to reproduce the fed sugar sensitivity curve in the absence of neuromodulator (neuromodulator factor (f_{nm})=1), and the 2 day starved sensitivity curve in the presence of neuromodulator (f_{nm} =5). On each trial, the RNN output a probability of PER after a 500ms presentation of a sugar stimulus ranging from 6.25 to 800 mM. We independently trained 10 RNNs on this task and all reliably learned the sensitivity functions. We then tested the RNNs’ behaviors at an intermediate neuromodulator level (f_{nm} =3). Despite never being explicitly trained for any behavior at this intermediate neuromodulator level, we found the RNNs produced a shifted sensitivity curve very similar to that exhibited by 1 day starved flies and the flies fed an intermediate L-dopa concentration of 3 mg/ml (Figure 6B).

The RNNs’ behavior reliably mimicked the intermediate behaviors of flies *in vivo* because intermediate neuromodulation caused a shift in the RNN’s activity manifold between “fed” and “2 day starved/5 mg/ml L-dopa” manifolds. For example, for 100mM sugar, without neuromodula-

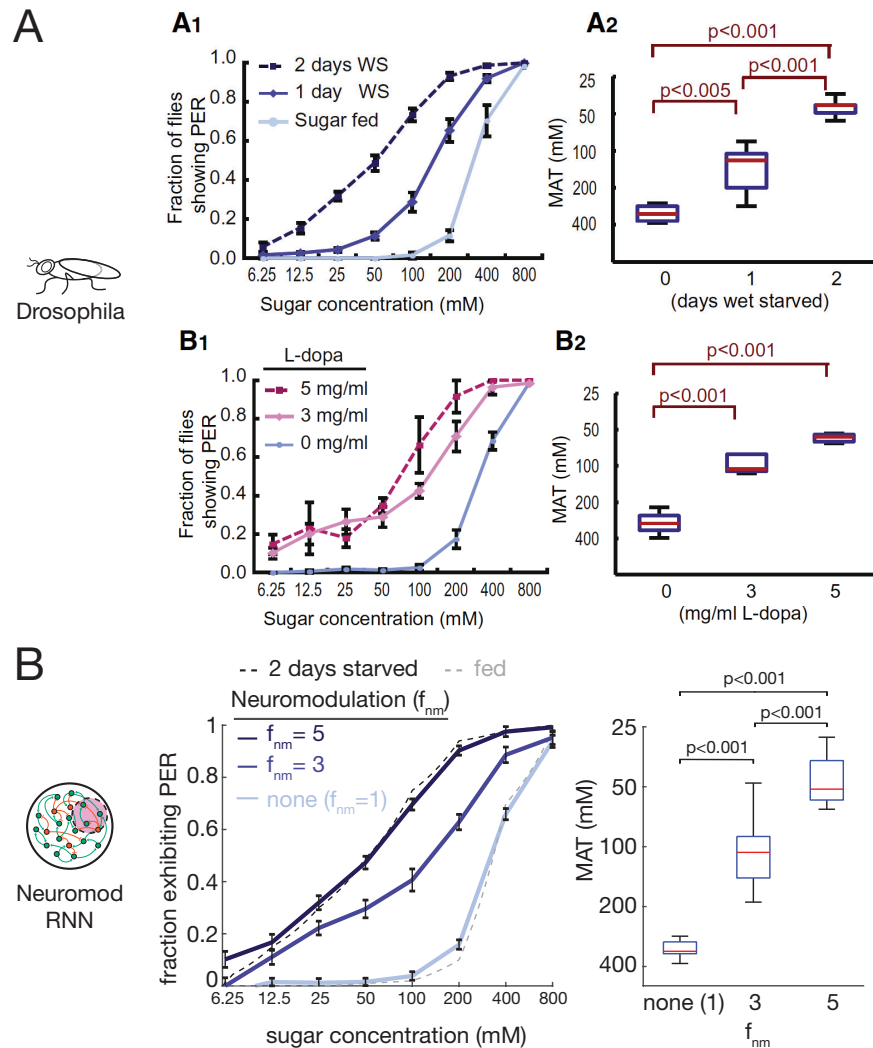


Figure 6. Neuromodulated RNN reproduces *Drosophila* sugar sensitivity behavior. (A) *Drosophila* sugar sensitivity behaviors from Inagaki et al. 2012 measured as PER behavior vs sugar concentration. Reprinted from Cell 148, Inagaki et al., Visualizing Neuromodulation In Vivo: TANGO-Mapping of Dopamine Signaling Reveals Appetite Control of Sugar Sensing, 583-595, 2012, with permission from Elsevier. (B) Neuromodulated RNNs trained on extremes of *Drosophila* sugar sensitivity (no neuromodulation factor ($f_{nm}=1$) for fed and $f_{nm}=5$ for 2 days starved) exhibit similar intermediate ($f_{nm}=3$; untrained) and extreme (no neuromodulation ($f_{nm}=1$) and $f_{nm}=5$; trained) behaviors ($n=10$; error bars are SEM; same statistical test as in Inagaki et al. 2012 for boxplots, see Methods).

tion RNNs produced non-overlapping activity traces to those under full neuromodulation ($f_{nm}=5$). Intermediate neuromodulation ($f_{nm}=3$) produced activity traces between absent and full neuromodulation traces (Figure 7A). The same observation held when output behavior was constant: the activity traces under intermediate neuromodulation producing PER output at 50% probability (mean acceptance threshold (MAT) from Inagaki et al. 2012) lay exactly between the traces for absent and full neuromodulation at MAT (Figure 7B). Viewed across the range of sugar concentrations presented, intermediate neuromodulation produced activity that defined a curtain in activity space that slotted neatly between absent and full neuromodulation activity curtains. Each neuromodulation level activity curtain led to a curved line attractor, representing end states that

generated the appropriate output behavior (Figure 7C).

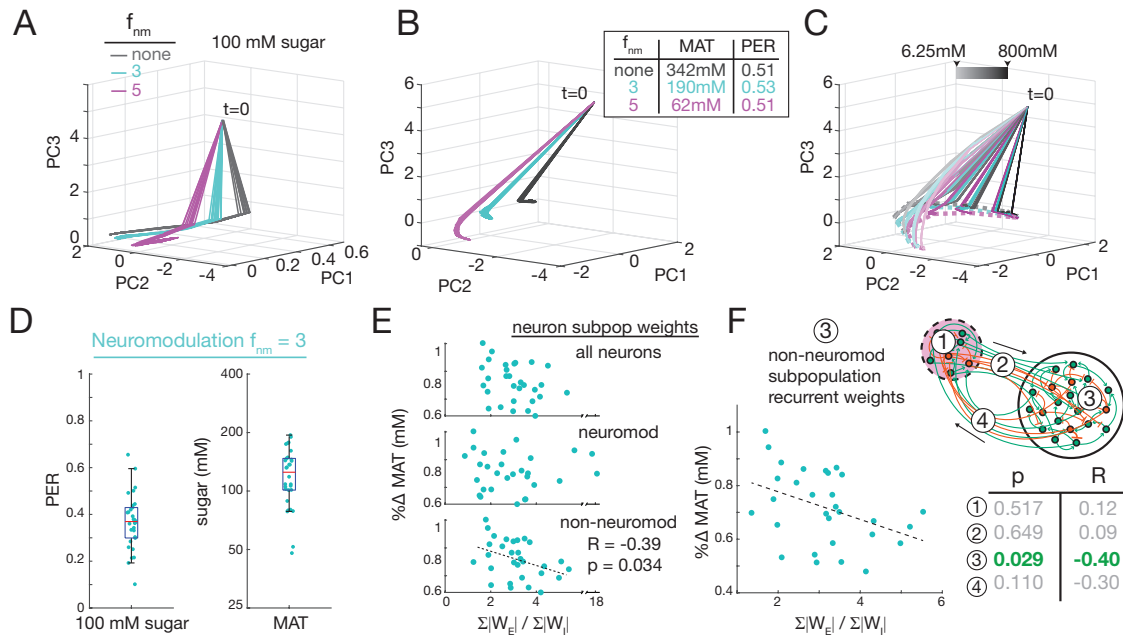


Figure 7. RNN models of *Drosophila* exhibit emergent transition behaviors and high variability of neuromodulator sensitivity. (A) For 100mM sugar, a RNN with intermediate neuromodulation ($f_{nm}=3$; cyan) generates activity between the neuromodulation extremes (none; grey and $f_{nm}=5$; magenta). (B) Activity trajectory at MAT sugar concentration for intermediate neuromodulation of a RNN lies between those for neuromodulation extremes. (C) Across sugar concentrations, intermediate neuromodulation trajectories (cyan) lay between neuromodulation extremes (grey and magenta), forming a 3-layer activity curtain ending on curved line attractors (dotted lines) defined by trajectory endpoints across the sugar spectrum. (D) Independently trained RNNs ($n=30$) exhibited high variability of PERs at 100mM sugar (0.10 to 0.66) and MATs (48 to 194mM). (E) Normalized MAT change (% Δ MAT) vs E/I ratio for whole RNN and subpopulations. % Δ MAT had significant negative correlation to E/I ratio of the non-neuromodulated neurons ($p<0.05$, $R=-0.39$). (F) % Δ MAT E/I ratio correlation was driven by recurrent weights within the non-neuromodulated subpopulation ③, rather than recurrent weights within the neuromodulated subpopulation ①, weights from neuromodulated to non-neuromodulated subpopulation ②, or weights from non-neuromodulated to neuromodulated subpopulation ④ as shown by Pearson correlation p-values (p) and coefficients (R) (for ①,②,④ correlation significance scores shown, scatter plots not shown).

We next asked whether our fly RNN models also exhibit behavioral variability derived from the “circuit-based sensitivity” we observed previously. To investigate this possibility, we independently trained 30 RNNs with 20% subpopulation under neuromodulation on the fed and 2-day starved sugar sensitivity behaviors ($f_{nm}=1$ and $f_{nm}=5$, respectively). We then applied intermediate neuromodulation ($f_{nm}=3$) to each model and assessed the RNNs’ PER probabilities at 100 mM sugar and MATs. We found profoundly idiosyncratic circuit-based sensitivities, with 100 mM sugar PER probability ranging from 0.10 to 0.66 (6.6x range) and MAT from 48 to 194 mM (4x range) (Figure 7D).

We asked if a RNN’s circuit sensitivity was related to the underlying excitatory-inhibitory weight distribution. To do this we compared network normalized MAT shifts (% Δ MAT; see Methods) to E/I weight balance across the whole network, within the neuromodulated subpopulation, and

within the non-neuromodulated subpopulation. While no relationship between $\% \Delta \text{MAT}$ and whole network or neuromodulated subpopulation weights was observed, we found a weak relationship to E/I balance in the non-neuromodulated subpopulation ($p=0.034$, $R=-0.39$; Figure 7E). To identify connectivity driving this effect, we parsed weights according to their source and destination within the network and found that the recurrent weights of the non-neuromodulated subpopulation were most strongly related to the network’s circuit sensitivity, with stronger inhibitory tone coinciding with increased sensitivity ($p=0.029$, $R=-0.40$; Figure 7F). Interestingly, this relationship is directionally consistent with that observed in the fully neuromodulated networks where positive weight distribution skewness (analogous to larger E/I) correlated with less sensitive networks (higher EC50, lower $\% \Delta \text{MAT}$). Further investigation detailing the driving factors underlying circuit sensitivity will be crucial. Importantly, our model makes two clear predictions: i) individual flies with similar fed (and 0 mg/ml L-dopa) and 2 day starved (and 5 mg/ml L-dopa) sugar sensitivity behaviors exhibit highly variable 1 day starved (and 3 mg/ml L-dopa) sugar sensitivity profiles (i.e. large 1 day (3 mg/ml L-dopa) MAT variability) and ii) this phenomena is due to idiosyncratic circuit-based sensitivity that is directly related to excitatory-inhibitory weight distributions in the relevant circuits. The first prediction is immediately testable, the second may require further investigation and technological innovation.

Electrical modulation shifts neural dynamics between manifolds through independent mechanism to neuromodulation. In addition to neuromodulation’s effect on synaptic transmission, other endogenous and exogenous influences can alter circuit dynamics through mechanisms that may be similar or distinct. For comparison, we considered the effect of electrical modulation, which influences circuit dynamics by altering current inputs into target neurons. Experimentally and clinically, electrical modulation is increasingly used in the form of optogenetics [Deisseroth, 2015], deep-brain stimulation (DBS) [Krauss et al., 2021] (Figure 8A), and transcranial magnetic stimulation [Risio et al., 2020, Gouveia et al., 2019]; the latter two providing much needed alternatives when traditional pharmacologic treatments fail. We used our model to ask whether electrical modulation alters network activity in an analogous manner to chemical neuromodulation or operates through an independent mechanism.

To test this we trained a RNN with a neuromodulated subpopulation (50% of neurons) on the previously described modified Go-NoGo task (output behavior +1 was called “approach”; 0 was “none”; -1 was “avoid”) (Figure 8A). After successful training, we assessed the RNN’s behavior when either no electrical stimulation, excitatory stimulation (as constant fixed input current of +1), or inhibitory stimulation (constant, fixed input current of -1) was given to all neurons in either the neuromodulation-targeted or an equal size random subpopulation (Figure 8A). While electrical stimulation to a random subpopulation did not affect the RNN’s performance, stimulation delivered to the neuromodulated subpopulation could significantly affect behavior in the absence and presence of neuromodulator (Figure 8B).

To determine if electrical stimulation induced random performance impairments or instead represented shifts to alternative learned behaviors, we characterized RNN output. Interestingly, electrical stimulation of neuromodulated neurons shifted behaviors directionally toward the opposing neuromodulation condition behavioral set. For example, inhibitory stimulation in the absence of neuromodulator decreased “approach” and increased “none” behavior to the positive input stimulus while decreasing “none” and increasing “avoid” behavior to the null stimulus. These shifted behaviors mirror the RNN’s normal behavior under neuromodulation (Figure 8C, top, rightmost bar graph compared to Figure 8D, top, left most bar graph). Likewise, the opposite electrical

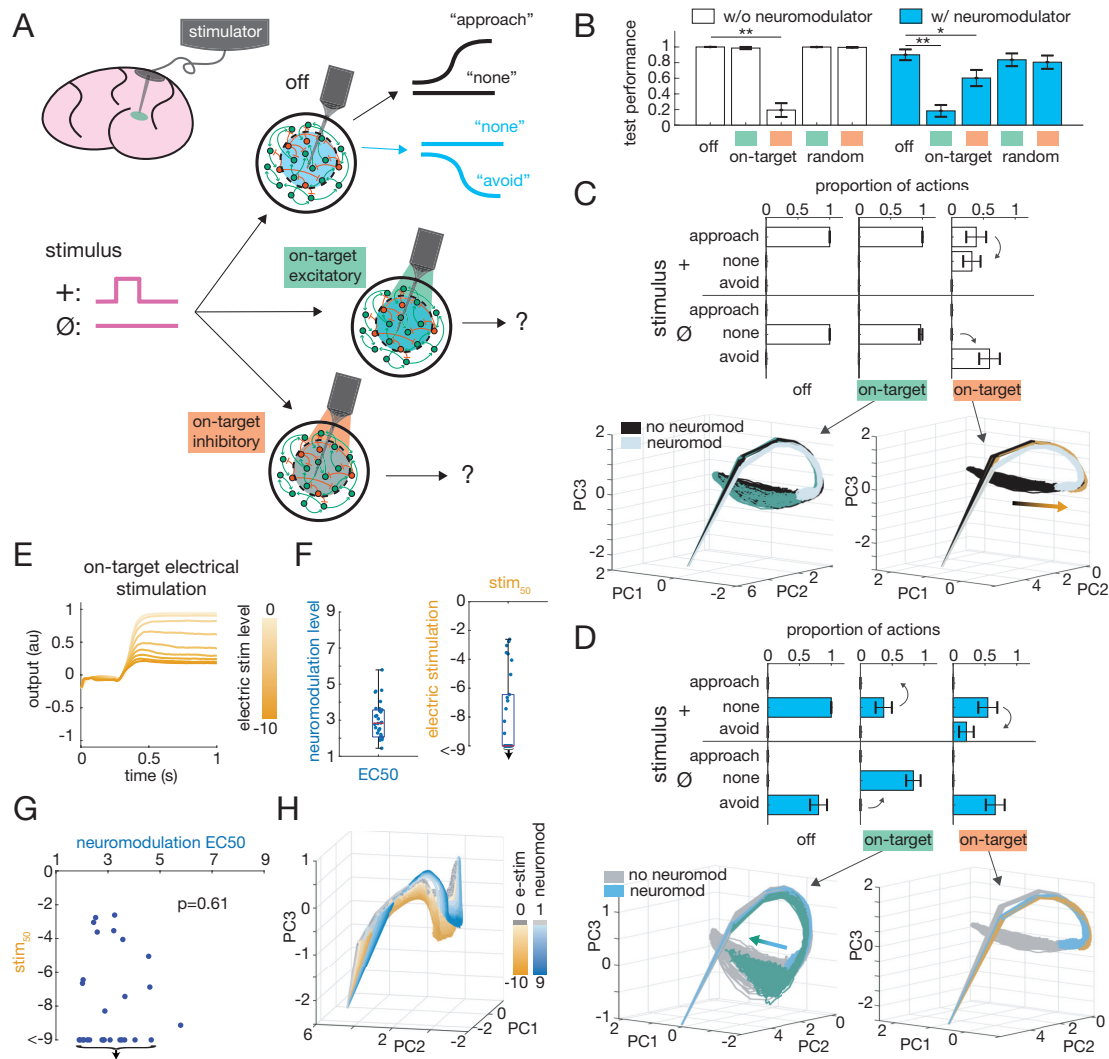


Figure 8. Targeted electrical modulation shifts network dynamics through independent mechanism. (A) Schematic of DBS and analogous electrical stimulation (e-stim) of a neuromodulation RNN. (B) Test performance was significantly impaired without neuromodulator effect for inhibitory on-target e-stim ($p=3.91\text{e-}11$) and with neuromodulator for excitatory and inhibitory on-target e-stim ($p=2.20\text{e-}08$ and $p=2.18\text{e-}02$, respectively). *: $p<0.05$, **: $p<0.01$. (C) Behavioral and activity shifts without neuromodulation. Top: behavioral shifts with on-target e-stim. Bottom: activity dynamics under e-stim. Activity trajectories without e-stim are in black (neuromodulated trajectories in light blue for reference); e-stim can shift trajectories off black toward neuromodulated activity space (orange arrow). (D) Same as (C) with neuromodulation. Top: behavioral shifts. Bottom: Activity trajectories without e-stim are in blue (non-neuromodulated trajectories in grey for reference); e-stim can shift trajectories off blue toward non-neuromodulated activity space (green arrow). (E) Output to + stimulus with increasing e-stim. (F) For 30 RNNs: Left: neuromodulation EC50. Right: electrical stim₅₀. Down arrows in (F),(G) indicate RNNs that did not achieve stim₅₀ at maximum stimulation. (G) No significant correlation between networks' EC50 and stim₅₀. (H) Modulation by neuromodulation and electrical modulation push activity trajectories toward different manifolds enabling independent transition dynamics. All error bars are SEM.

stimulation (excitatory) induced behavioral shifts in the neuromodulated condition toward non-neuromodulated behaviors (Figure 8D, top).

To understand how electrical stimulation was affecting networks' behaviors, we analyzed network activity trajectories. We observed that when electrical stimulation did not change behavior, network activity overlapped with the unstimulated RNN's activity space (Figure 8C, bottom left, green traces) whereas when stimulation shifted behavior, network activity was pushed toward the activity manifold defined by the opposing neuromodulation-condition (Figure 8C, bottom right, orange traces) (Figure 8C–D, bottom). Targeted electrical modulation thus provides an alternative means to predictably influence circuit dynamics toward neuromodulation conditions. This result supports the observation that DBS clinically and optogenetics experimentally can manipulate network behavior to achieve similar end results as pharmacological manipulation.

This result from our model presented an opportunity to study a question critical to neuropsychiatric treatment: are mechanisms of chemical and electrical modulation for a given network correlated; or in clinically practical terms: if a patient fails multiple drug treatments, are they less likely to benefit from DBS?

We first trained 30 RNNs and tested the effects of graded electrical stimulation in a condition that had shown electrical sensitivity (on-target inhibitory stimulation without neuromodulator). Increasing the level of electrical input led to a graded transition in behavior very similar to the transition observed with graded administration of neuromodulator (Figure 8E). As before, RNNs exhibited idiosyncratic circuit-based sensitivity to neuromodulator (Figure 8F, left). To characterize electrical stimulation network sensitivity, we measured the amount of electrical stimulation that led to 50% of the fully-neuromodulated condition (stim_{50}) analogous to EC50 for neuromodulator levels. Interestingly, RNNs also exhibited idiosyncratic circuit-based sensitivity to electrical stimulation, with stim_{50} ranging from -2.6 input current to not achieving stim_{50} at the maximum stimulation we administered of -9 input current ($>3.5\times$ range) (Figure 8F, right).

We then asked if there was a relationship between a network's chemical and electrical circuit-based sensitivities. There was no significant correlation between EC50 and stim_{50} (Figure 8G). Since some RNNs did not reach stim_{50} output at the maximum electrical stimulation given, we also measured each RNN's output at maximum stimulation (input=-9) and similarly found no significant correlation to EC50 (Figure S7). The lack of correlation suggests that networks insensitive to chemical modulation may still be highly sensitive to electrical modulation and vice versa. Consistent with this, we found that electrical modulation progressively shifted population dynamics along a manifold transition distinct to neuromodulation, suggesting different possible rates of transition (Figure 8H). This may help explain why some patients who fail multiple drug treatments sometimes respond dramatically to DBS. By utilizing an independent mechanism to drug treatment, DBS exploits a parallel circuit-sensitivity to achieve therapeutic efficacy.

Discussion

Neuromodulation is pervasive in the brain and drives unique neural function in both health and disease. Using an RNN model, we showed how neuromodulation's effect on synaptic weights can separate distinct synaptic memory regimes within a single RNN, each tuned to unlock unique outputs and behaviors. We observed emergent neural activity transformations reminiscent of prior single-neuron neuromodulation work [Marder, 2012, Johnson et al., 1995, Johnson et al., 2011, Katz et al., 1994] and revealed how robust and flexible this neuromodulatory system is across network

configurations. We showed how neuromodulation affects the network from individual neurons to circuit clusters to global network activity dynamics. Through this analysis we discovered a feature of networks previously unreported to our knowledge: “circuit-based sensitivity,” which helps explain the clinical observation of high variability in drug and other therapeutic response [Furukawa et al., 2019] (alongside more standard explanations like enzyme variant-dependent drug metabolism and clearance rates). Application of our model to *in vivo* experiments with starved *Drosophila* recreated experimental findings and provided both explanation for the neural mechanism and concrete predictions for future experimentation. Finally, we used our model to investigate electrical modulation of circuits and relate chemical and electrical modulation properties. Together these results elucidate a fundamental mechanism of neuromodulation in the brain, provide concrete predictions, and define a previously unknown feature of circuits that may be critical to therapeutic success.

Neuromodulators in the brain also affect neural excitability, plasticity, and indirectly downstream circuit activity [Newman et al., 2012, Nadim and Bucher, 2012, Marder, 2012, Weele et al., 2018]. Other models have focused on other aspects of neuromodulation [Newman et al., 2012, Yu and Dayan, 2005, Stroud et al., 2018, Beaulieu et al., 2020]. Our model of synaptic weight scaling shares some similarities to previous models, particularly of neuromodulatory effects on intrinsic excitability [Stroud et al., 2018]. Both formulations lead to increased network flexibility and versatility, yet they operate through different mechanisms (see Supplemental Appendix A). Extending our model to incorporate other functions of biological neuromodulation (e.g. intrinsic excitability) and understanding their relationships will help bridge the divide to the full complexity of observations in the brain and to other models of neuromodulation that are highly tailored to specific functional profiles in specific brain areas of specific neuromodulators [Newman et al., 2012, Yu and Dayan, 2005].

Our formulation of the neuromodulatory effect on synaptic weights is a simplification. Elaborating our model to support differential weight modulation (e.g. via neuromodulator receptor subtypes on specific cell-types) [Harris-Warrick and Johnson, 2010, Bacqué-Cazenave et al., 2020], neuromodulator multiplexing [Wester and McBain, 2014, Li et al., 2018], and metamodulation [Nadim and Bucher, 2012, Ribeiro and Sebastiao, 2010, Katz and Edwards, 1999] will likely lead to even more sophisticated network behavior. Our model also uses arbitrary neuromodulation levels, whereas the brain likely uses specific levels for optimal functionality. Future investigation into which levels are optimal and methods of learning these will be important.

The “circuit-based sensitivity” we observed is related to the well-known variability of independently trained neural networks [Mehrer et al., 2020]. Our neuromodulation model allows study for the first time of the unique transition dynamics that emerge from neuromodulated networks and motivates several avenues of future research. Identification of the circuit parameters that control the transition dynamics will be critical for understanding and use in therapeutic optimization. The idiosyncratic nature of circuit-based sensitivity aligns exactly with current efforts in precision medicine calling for the need to consider each patient as an idiosyncratic individual — here we provide hard evidence for this claim and its particular importance in neuropsychiatric treatment [Provenza et al., 2019]. Fully understanding the relationship between chemical and electrical modulation and sensitivity is also crucial. Although our simplified model suggests how the modes of modulation influence dynamics (see Supplemental Appendix B), further analytical and experimental investigation into their relationship as network dynamics evolve over time is necessary.

In our analysis of electrical-like stimulation of RNNs we used a constant stimulation form (constant,

fixed stimulation of neuromodulated subpopulation) that we observed could shift network behavior. To fully characterize how networks respond to electrical stimulation, further exploration with different target regions and varied wave forms is warranted. Such experimentation in clinical use has led to major advancements in efficacy of DBS [Krauss et al., 2021] and TMS [Risio et al., 2020]. Future work utilizing wave forms and target areas mimicking and extending from clinically-used parameters will help fully characterize the effects of electrical stimulation on neural networks.

Finally, our work provides interesting directions for machine learning (ML). By separating synaptic memory regimes in a single network, we demonstrate how a network can have drastically increased capacity, supporting a library of unique behaviors for overlapping external situations. Furthermore, each behavior can be rapidly accessed through targeted application of the relevant neuromodulatory factor. High capacity, compact networks with high-speed access to different output modes presents a promising component for future ML algorithm development. The separation of memory regimes also effectively splits a single RNN into multiple computers, possibly achieving super-Turing capability [Cabessa and Siegelmann, 2014]. Future theoretical assessment of neuromodulated RNNs' capacity will establish if this simple mechanism is sufficient to exceed the Turing limit.

Acknowledgements

We thank Y. Chen and R. Kim for helpful discussions and feedback on the manuscript. We thank J. Aldana for support with computing resources. This work was supported by the Kavli Institute for Brain and Mind at University of California San Diego, Office of Naval Research Grant N00014-16-1-2829, Defense Advanced Research Projects Agency Grant HR0011-18-2-002, National Institute of Mental Health Grant R37-MH102441, and National Center for Complementary and Integrative Health Pioneer Award DP1-AT009925.

Author contributions

B.T., K.M.T., H.T.S., and T.J.S. formulated the ideas. B.T. and S.C.P. performed the simulations and analyses. B.T., S.C.P., K.M.T., H.T.S., and T.J.S. wrote the manuscript.

Declaration of interests

The authors declare no competing interests.

References

- [Avery and Krichmar, 2017] Avery, M. C. and Krichmar, J. L. (2017). Neuromodulatory Systems and Their Interactions: A Review of Models, Theories, and Experiments. *Frontiers in Neural Circuits* 11, 108.
- [Bacqué-Cazenave et al., 2020] Bacqué-Cazenave, J., Bharatiya, R., Barrière, G., Delbecque, J.-P., Bouguiyoud, N., Giovanni, G. D., Cattaert, D. and Deurwaerdère, P. D. (2020). Serotonin in Animal Cognition and Behavior. *International Journal of Molecular Sciences* 21, 1649.

- [Beaulieu et al., 2020] Beaulieu, S., Frati, L., Miconi, T., Lehman, J., Stanley, K. O., Clune, J. and Cheney, N. (2020). Learning to Continually Learn. In 24th European Conference on Artificial Intelligence (ECAI), (Giacomo, G. D., Catalá, A., Dilkina, B., Milano, M., Barro, S., Bugarín, A. and Lang, J., eds), vol. 325, pp. 992–1001, IOS Press, Santiago de Compostela, Spain.
- [Burgos-Robles et al., 2017] Burgos-Robles, A., Kimchi, E. Y., Izadmehr, E. M., Porzenheim, M. J., Ramos-Guasp, W. A., Nieh, E. H., Felix-Ortiz, A. C., Namburi, P., Leppla, C. A., Presbrey, K. N., Anandalingam, K. K., Pagan-Rivera, P. A., Anahtar, M., Beyeler, A. and Tye, K. M. (2017). Amygdala inputs to prefrontal cortex guide behavior amid conflicting cues of reward and punishment. *Nature Neuroscience* 20, 824–835.
- [Cabessa and Siegelmann, 2014] Cabessa, J. and Siegelmann, H. T. (2014). The Super-Turing Computational Power of Plastic Recurrent Neural Networks. *International Journal of Neural Systems* 24, 1450029.
- [Cunningham and Yu, 2014] Cunningham, J. P. and Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience* 17, 1500–1508.
- [Deisseroth, 2015] Deisseroth, K. (2015). Optogenetics: 10 years of microbial opsins in neuroscience. *Nature Neuroscience* 18, 1213–1225.
- [Furukawa et al., 2019] Furukawa, T. A., Cipriani, A., Cowen, P. J., Leucht, S., Egger, M. and Salanti, G. (2019). Optimal dose of selective serotonin reuptake inhibitors, venlafaxine, and mirtazapine in major depression: a systematic review and dose-response meta-analysis. *Lancet Psychiatry* 6, 601–609.
- [Gouveia et al., 2019] Gouveia, F. V., Gidyk, D. C., Giacobbe, P., Ng, E., Meng, Y., Davidson, B., Abrahao, A., Lipsman, N. and Hamani, C. (2019). Neuromodulation Strategies in Post-Traumatic Stress Disorder: From Preclinical Models to Clinical Applications. *Brain Sciences* 9, 45.
- [Harris-Warrick and Johnson, 2010] Harris-Warrick, R. M. and Johnson, B. R. (2010). Checks and balances in neuromodulation. *Frontiers in Behavioral Neuroscience* 4, 47.
- [Hasselmo et al., 1995] Hasselmo, M. E., Schnell, E. and Barkai, E. (1995). Dynamics of Learning and Recall at Excitatory Recurrent Synapses and Cholinergic Modulation in Rat Hippocampal Region CA3. *Journal of Neuroscience* 15, 5249–5262.
- [Higley and Picciotto, 2014] Higley, M. J. and Picciotto, M. R. (2014). Neuromodulation by acetylcholine: examples from schizophrenia and depression. *Current Opinion in Neurobiology* 29, 88–95.
- [Inagaki et al., 2012] Inagaki, H. K., de Leon, S. B.-T., Wong, A. M., Jagadish, S., Ishimoto, H., Barnea, G., Kitamoto, T., Axel, R. and Anderson, D. J. (2012). Visualizing Neuromodulation In Vivo: TANGO-Mapping of Dopamine Signaling Reveals Appetite Control of Sugar Sensing. *Cell* 148, 583–595.
- [Jing et al., 2009] Jing, J., Gillette, R. and Weiss, K. R. (2009). Evolving Concepts of Arousal: Insights from Simple Model Systems. *Reviews in the Neurosciences* 20, 405–427.
- [Johnson et al., 2011] Johnson, B. R., Brown, J. M., Kvarta, M. D., Lu, J. Y. J., Schneider, L. R., Nadim, F. and Harris-Warrick, R. M. (2011). Differential Modulation of Synaptic Strength

- and Timing Regulate Synaptic Efficacy in a Motor Network. *Journal of Neurophysiology* 105, 293–304.
- [Johnson et al., 1995] Johnson, B. R., Peck, J. H. and Harris-Warrick, R. M. (1995). Distributed Amine Modulation of Graded Chemical Transmission in the Pyloric Network of the Lobster Stomatogastric Ganglion. *Journal of Neurophysiology* 74, 437–452.
- [Katz and Edwards, 1999] Katz, P. S. and Edwards, D. H. (1999). Metamodulation: the control and modulation of neuromodulation. In *Beyond Neurotransmission: Neuromodulation and its Importance for Information Processing*, (Katz, P. S., ed.), chapter 10, pp. 349–381. Oxford University Press Oxford, UK.
- [Katz et al., 1994] Katz, P. S., Getting, P. A. and Frost, W. N. (1994). Dynamic neuromodulation of synaptic strength intrinsic to a central pattern generator circuit. *Nature* 367, 729–731.
- [Kim et al., 2019] Kim, R., Li, Y. and Sejnowski, T. J. (2019). Simple framework for constructing functional spiking recurrent neural networks. *PNAS* 116, 22811–22820.
- [Kim and Sejnowski, 2020] Kim, R. and Sejnowski, T. J. (2020). Strong inhibitory signaling underlies stable temporal dynamics and working memory in spiking neural networks. *Nature Neuroscience* 24.
- [Krauss et al., 2021] Krauss, J. K., Lipsman, N., Aziz, T., Boutet, A., Brown, P., Chang, J. W., Davidson, B., Grill, W. M., Hariz, M. I., Horn, A., Schulder, M., Mammis, A., Tass, P. A., Volkmann, J. and Lozano, A. M. (2021). Technology of deep brain stimulation: current status and future directions. *Nature Reviews Neurology* 17, 75–87.
- [Laje and Buonomano, 2013] Laje, R. and Buonomano, D. V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nature Neuroscience* 16, 925–933.
- [Li et al., 2018] Li, X., Bucher, D. and Nadim, F. (2018). Distinct Co-Modulation Rules of Synapses and Voltage-Gated Currents Coordinate Interactions of Multiple Neuromodulators. *Journal of Neuroscience* 38, 8549–8562.
- [Mante et al., 2013] Mante, V., Sussillo, D., Shenoy, K. V. and Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84.
- [Marder, 2012] Marder, E. (2012). Neuromodulation of Neuronal Circuits: Back to the Future. *Neuron* 76, 1–11.
- [Marder and Eisen, 1984] Marder, E. and Eisen, J. S. (1984). Electrically Coupled Pacemaker Neurons Respond Differently to Same Physiological Inputs and Neurotransmitters. *Journal of Neurophysiology* 51, 1362–1374.
- [Marder et al., 2004] Marder, S. R., Essock, S. M., Miller, A. L., Buchanan, R. W., Casey, D. E., Davis, J. M., Kane, J. M., Lieberman, J. A., Schooler, N. R., Covell, N., Stroup, S., Weissman, E. M., Wirshing, D. A., Hall, C. S., Pogach, L., Pi-Sunyer, X., J. Thomas Bigger, J., Friedman, A., Kleinberg, D., Yevich, S. J., Davis, B. and Shon, S. (2004). Physical Health Monitoring of Patients With Schizophrenia. *The American Journal of Psychiatry* 161, 1334–1349.
- [McCormick et al., 2020] McCormick, D. A., Nestvogel, D. B. and He, B. J. (2020). Neuromodulation of Brain State and Behavior. *Annual Review of Neuroscience* 43, 391–415.

- [Mehrer et al., 2020] Mehrer, J., Spoerer, C. J., Kriegeskorte, N. and Kietzmann, T. C. (2020). Individual differences among deep neural network models. *Nature Communications* *11*, 5725.
- [Nadim and Bucher, 2012] Nadim, F. and Bucher, D. (2012). Neuromodulation of neurons and synapses. *Current Opinion in Neurobiology* *29*, 48–56.
- [Newman et al., 2012] Newman, E. L., Gupta, K., Climer, J. R., Monaghan, C. K. and Hasselmo, M. E. (2012). Cholinergic modulation of cognitive processing: insights drawn from computational models. *Frontiers in Behavioral Neuroscience* *6*, 24.
- [Patriarchi et al., 2018] Patriarchi, T., Cho, J. R., Merten, K., Howe, M. W., Marley, A., Xiong, W.-H., Folk, R. W., Broussard, G. J., Liang, R., Jang, M. J., Zhong, H., Dombeck, D., von Zastrow, M., Nimmerjahn, A., Gradinaru, V., Williams, J. T. and Tian, L. (2018). Ultrafast neuronal imaging of dopamine dynamics with designed genetically encoded sensors. *Science* *360*, eaat4422.
- [Provenza et al., 2019] Provenza, N. R., Matteson, E. R., Allawala, A. B., Barrios-Anderson, A., Sheth, S. A., Viswanathan, A., McIngvale, E., Storch, E. A., Frank, M. J., McLaughlin, N. C. R., Cohn, J. F., Goodman, W. K. and Borton, D. A. (2019). The Case for Adaptive Neuromodulation to Treat Severe Intractable Mental Disorders. *Frontiers in Neuroscience* *13*, 152.
- [Quinn et al., 1983] Quinn, N. P., Lang, A. E., Marsden, C. D., Parkes, J. D. and Toone, B. (1983). Dopa Dose-Dependent Sexual Deviation. *British Journal of Psychiatry* *142*, 296–298.
- [Ribeiro and Sebastiao, 2010] Ribeiro, J. A. and Sebastiao, A. M. (2010). Modulation and meta-modulation of synapses by adenosine. *Acta Physiologica* *199*, 161–169.
- [Risio et al., 2020] Risio, L. D., Borgi, M., Pettorruso, M., Miuli, A., Ottomana, A. M., Sociali, A., Martinotti, G., Nicolo, G., Macri, S., di Giannantonio, M. and Zoratto, F. (2020). Recovering from depression with repetitive transcranial magnetic stimulation (rTMS): a systematic review and meta-analysis of preclinical studies. *Translational Psychiatry* *10*, 393.
- [Stroud et al., 2018] Stroud, J. P., Porter, M. A., Hennequin, G. and Vogels, T. P. (2018). Motor primitives in space and time via targeted gain modulation in cortical networks. *Nature Neuroscience* *21*, 1774–1783.
- [Trujillo et al., 2020] Trujillo, C. A., Adams, J. W., Negraes, P. D., Carromeu, C., Tejwani, L., Acab, A., Tsuda, B., Thomas, C. A., Sodhi, N., Fichter, K. M., Romero, S., Zanella, F., Sejnowski, T. J., Ulrich, H. and Muotri, A. R. (2020). Pharmacological reversal of synaptic and network pathology in human MECP2-KO neurons and cortical organoids. *EMBO Molecular Medicine* *13*, e12523.
- [Vecoven et al., 2020] Vecoven, N., Ernst, D., Wehenkel, A. and Drion, G. (2020). Introducing neuromodulation in deep neural networks to learn adaptive behaviors. *Plos One* *15*, e0227922.
- [Weele et al., 2018] Weele, C. M. V., Siciliano, C. A., Matthews, G. A., Namburi, P., Izadmehr, E. M., Espinel, I. C., Nieh, E. H., Schut, E. H. S., Padilla-Coreano, N., Burgos-Robles, A., Chang, C.-J., Kimchi, E. Y., Beyeler, A., Wichmann, R., Wildes, C. P. and Tye, K. M. (2018). Dopamine enhances signal-to-noise ratio in cortical-brainstem encoding of aversive stimuli. *Nature* *563*, 397–401.

- [Wester and McBain, 2014] Wester, J. C. and McBain, C. J. (2014). Behavioral state-dependent modulation of distinct interneuron subtypes and consequences for circuit function. *Current Opinion in Neurobiology* 29, 118–125.
- [Yamamoto et al., 2014] Yamamoto, K., Shinba, T. and Yoshii, M. (2014). Psychiatric symptoms of noradrenergic dysfunction: A pathophysiological view. *Psychiatry and Clinical Neurosciences* 68, 1–20.
- [Yu and Dayan, 2005] Yu, A. J. and Dayan, P. (2005). Uncertainty, Neuromodulation, and Attention. *Neuron* 46, 681–692.

Methods

Modified Go-NoGo tasks. The classic Go-NoGo task has two possible stimuli (positive pulse referred to as positive stimulus or +; no pulse, also referred to as null stimulus or \emptyset). The agent is trained to give a positive output (+1, “Go”) for the positive stimulus and zero output (0, “NoGo”) for the null stimulus. In the modified Go-NoGo task we added a second possible behavioral set: NoGo for the positive stimulus and negative output (-1, “AntiGo”) for the null stimulus. The network was trained on the classic Go-NoGo behavior in the absence of neuromodulator and on the new NoGo-AntiGo behavior in the presence of neuromodulator. The 3 behavior and 9 behavior variants of the modified Go-NoGo task followed a similar paradigm with additional added behaviors. In the 3 behavior version, a third behavior of positive stimulus \rightarrow AntiGo, null stimulus \rightarrow Go was added. In the 9 behavior version all possible stimulus \rightarrow output response pairs were added. As such, Behavior 1 was + \rightarrow AntiGo, $\emptyset \rightarrow$ AntiGo; Behavior 2 was + \rightarrow NoGo, $\emptyset \rightarrow$ AntiGo; Behavior 3 was + \rightarrow Go, $\emptyset \rightarrow$ AntiGo; ...; Behavior 9 was + \rightarrow Go, $\emptyset \rightarrow$ Go. In our simulations each trial lasted 200 timesteps of 5ms each for a total represented duration of 1 second.

Neuromodulatory neural network model. For our simulations we used a continuous rate recurrent neural network (RNN) model with biologically plausible parameters similar to RNNs in prior works [Kim and Sejnowski, 2020, Trujillo et al., 2020]. Consistent with biological neural networks, we implemented Dale’s Law such that each neuron was either excitatory or inhibitory. For all our simulations we used a RNN with $N = 200$ neuron units, 80% excitatory and 20% inhibitory. In the RNN each neuron can be connected to any other neuron with probability p_{con} (p_{con} was initialized at 0.8 in our simulations), and each neuron receives weighted input from connected neurons to produce a firing rate governed by the neural dynamical equation

$$\tau \frac{d\mathbf{x}}{dt} = -\mathbf{x} + W\mathbf{r} + W_{in}u + N(0, 0.1) \quad (1)$$

where $\tau \in R^{1 \times N}$ is the synaptic decay time constant for the N neurons in the network, $\mathbf{x} \in R^{1 \times N}$ is the synaptic current variable for the N neurons, $W \in R^{N \times N}$ is the matrix of synaptic weights between all N neurons, $\mathbf{r} \in R^{1 \times N}$ is the output firing rates of the N neurons in the network, $W_{in} \in R^{1 \times N}$ are weights associated with external input u , and $N(0, 0.1)$ is added noise drawn from a normal distribution with mean 0 and variance 0.1. The output firing rate for the neurons is given by an elementwise nonlinear transfer function transformation of the synaptic current variable. In our network we used the standard logistic sigmoid function as implemented by prior models [Kim et al., 2019, Kim and Sejnowski, 2020]:

$$\mathbf{r} = \frac{1}{1 + e^{-\mathbf{x}}} \quad (2)$$

The synaptic connectivity matrix W was randomly initialized from a normal distribution with zero mean and standard deviation $g/\sqrt{N \cdot p_{con}}$, where g is the gain. We set $g = 1.5$ as previous studies have shown that networks operating in a high gain regime ($g \geq 1.5$) support rich dynamics analogous to those of biological networks [Kim and Sejnowski, 2020, Kim et al., 2019, Laje and Buonomano, 2013]. The synaptic decay time constants were randomly initialized to a value in the biologically plausible range of 20–100 ms. As in Kim et al. 2019, we used the first-order Euler approximation method to discretize equation (1) for the simulations; for neuron i :

$$x_{i,t} = (1 - \frac{\Delta t}{\tau})x_{i,t-1} + \frac{\Delta t}{\tau}(\sum_j W_{ji}r_{ji,t-1} + W_{ui}u_{t-1}) + N(0, 0.1) \quad (3)$$

Output was generated by taking all recurrent network neurons' activities and passing them through a weighted output unit

$$o_{network} = W_{out}\mathbf{r} + b_{out}$$

where $W_{out} \in R^{N \times 1}$ are the neural output weights and b_{out} is the output unit's bias term. RNNs were trained by backpropagation through time using AdamOptimizer with a least square error objective function.

To apply an amplifying or dampening neuromodulatory effect, target neurons' weights were scaled by the neuromodulatory factor. For whole network neuromodulation this effect was applied to all neurons in the RNN; for subpopulation neuromodulation the effect was applied only to the selected subpopulation of neurons.

For the modified Go-NoGo task, RNNs were trained until one of two possible stop criteria was met: 1) average trial least square error over the last 50 trials was under a threshold of 1, or 2) 10,000 training trials was reached. Performance on the task was then assessed by evaluating the percentage of test trials that matched the following performance criteria: for Go trials, output was required to reach 1.0 ± 0.2 by timestep 120 (full trial was 200 timesteps); for NoGo trials, output was required to be 0.0 ± 0.2 and for AntiGo trials -1.0 ± 0.2 at timestep 120.

Neuromodulation of multiple subpopulations and multiple levels. For neuromodulation of non-overlapping subpopulations, same-sized groups of neurons were chosen randomly without any overlap and neuromodulator applied to each for a given behavior. For overlapping subpopulations, groups of neurons were chosen randomly allowing overlap (Figure S3).

Neuromodulation at different levels ("multi-factor networks") was done by applying different neuromodulation factors (f_{nm}). For the 9-behavior modified Go-NoGo this was done using factors $\in [1 : 1 : 9]$, i.e., for Behavior 1 no factor was applied ($f_{nm} = 1$), for Behavior 2 $f_{nm} = 2$, for Behavior 3 $f_{nm} = 3$, et cetera.

To test networks across the range of subpopulation sizes with single or multiple neuromodulator factors on n-behavior modified Go-NoGo tasks, stop criteria were adjusted to account for the increased behaviors: 1) average trial least square error over last $n \times 25$ trials was under threshold of 1, or 2) 15,000 training trials was reached. Performance on the tasks was assessed as before. For overlapping subpopulations, overlap was quantified in two ways. For each network, the number of neurons neuromodulated in 2 or more subpopulations was measured (Figure S3D). Overlap was also quantified by measuring the average number of neuromodulated subpopulations a neuron in

the network was a member of (Figure S3E).

Functional clustering and selectivity index. In order to examine whether trained models contained functionally specialized neurons, we grouped neuron activities by combination of sub-task (which maps one-to-one with neuromodulatory state) and stimulus given (“stimulus-subtask combinations”). We averaged activity of each neuron over time and trials within each group. This resulted in a matrix of time-trial averaged neuron activities with a number of rows equal to the stimulus-subtask combinations, and number of columns equal to the number of neurons. Using k-means, we clustered neurons with similar activity levels across stimulus-subtask combinations. We computed a silhouette score to find the optimal number of clusters, which for the RNN in Figure 4, was 6. The silhouette score computed for 5 and 6 clusters differed only by 0.7% and the additional cluster was very small and similar to an existing cluster, so we conducted further analysis with 5 clusters for simplicity.

To measure the selectivity of individual neurons for particular stimulus-subtask combinations, we calculated a “selectivity index” (si) for each neuron j :

$$si_j = \frac{\bar{r}_j^{max} - \bar{r}_j^{second_max}}{\bar{r}_j^{max}}$$

where \bar{r}_j is the average firing rate of neuron j over the trial duration, \bar{r}_j^{max} indicates the maximum \bar{r}_j across all the stimulus-subtask combinations and $\bar{r}_j^{second_max}$ indicates the second highest \bar{r}_j over all the stimulus-subtask combinations. The selectivity index thus captures a normalized approximation of how uniquely active a neuron was for a given stimulus-subtask combination.

Network population dynamics. To represent whole network population activity dynamics we sought a low dimensional representation of whole population activity. We used principal component analysis (PCA) since the leading components capture the largest projections of activity variability, which we hypothesized would effectively separate our neuromodulatory conditions if large differences occurred [Cunningham and Yu, 2014]. We found this was the case. We found qualitatively similar results using multidimensional scaling which finds projections designed to best preserve distances in high-dimensional activity space. For our figures we display the first 3 PCs, as these captured a large amount of the activity variance (80–92% explained across the analyses) and effectively represent the activity dynamic differences in the analyses.

EC50. The EC50 of a network was defined as the level of neuromodulation that led to half the output of full neuromodulation. For the results reported, we used EC50 calculated for the positive stimulus. For this stimulus in the modified Go-NoGo task, a non-neuromodulated network outputs +1 and a fully neuromodulated network outputs 0 (measurements for output level were taken at 0.5 s through the trial); the EC50 for the network in this case is the amount of neuromodulator required to output 0.5. The EC50 was calculated by fitting a sigmoid curve to the progression of output (from +1 to 0 in this case) with increasing neuromodulation level (Figure 5E)

$$output = 1 - \frac{1}{1 + e^{a \cdot f_{nm} + b}}$$

where f_{nm} is the neuromodulation level. EC50 neuromodulation level was calculated by finding the intersection of the sigmoid and the half-maximal output; for half-maximal output of 0.5, EC50 = $-b/a$. Sigmoid curves were fit using a least squares fit.

***Drosophila* sugar sensitivity task.** We implemented a computational version of Inagaki et al. 2012 to train our network models. During training, models were presented with a constant sugar concentration (external input proportional to sugar concentration) for 100 timesteps (equivalent to 500 ms) and trained to output a probability of PER. For fed and 2-day starved training we used a piece-wise linear approximation estimated from Inagaki et al. 2012. To compare boxplots of MAT, one-way ANOVA followed by *t*-test with Bonferroni correction was used as in Inagaki et al. 2012.

MAT. Analogous to the analysis done for flies in Inagaki et al. 2012, for each RNN a sigmoid was fit

$$PER = \frac{1}{1 + e^{-a \cdot \log_2 \frac{x_{sugar}}{MAT}}}$$

where *a* is the slope of the sigmoid. When *PER* = 0.5 then $x_{sugar} = MAT$. Sigmoid curves were fit using a least squares fit.

For intermediate neuromodulatory level ($f_{nm}=3$) MAT variability analysis, a normalized change in MAT (%ΔMAT) was calculated:

$$\% \Delta MAT = \frac{MAT_{f_{nm}=3} - MAT_{f_{nm}=1}}{MAT_{f_{nm}=5} - MAT_{f_{nm}=1}}$$

where $MAT_{f_{nm}=x}$ is the RNN’s MAT with neuromodulation factor *x*. %ΔMAT gives a network normalized metric for how much the intermediate neuromodulation ($f_{nm}=3$) moved the fly from no neuromodulation ($f_{nm}=1$) to full neuromodulation ($f_{nm}=5$) sensitivity.

Electrical modulation. We administered electrical modulation as an external current applied for the duration of the trial. “On-target” modulation was applied to the neuromodulated neuron population and “random” modulation was applied to a randomly selected group of neurons of equal size; these could include both neuromodulated or non-neuromodulated neurons. All neurons (both excitatory and inhibitory) within the selected subpopulation were given identical external current stimulation. For fixed electrical modulation simulations, a current of magnitude 1 was applied (+1 for excitatory stimulation; -1 for inhibitory).

For graded electrical modulation, networks that did not achieve stim₅₀ at maximum stimulation (-9 units) were assigned a surrogate stim₅₀ value of -10 to calculate correlation to EC50. To account for possible missed correlation due to this substitution, correlation of EC50 to output at maximum stimulation (-9) was also calculated (Figure S7).

Data and code availability

The code and RNN models in this work will be made available at <https://github.com/tsudacode/neuromodRNN>

Supplemental Information

Supplementary Figures

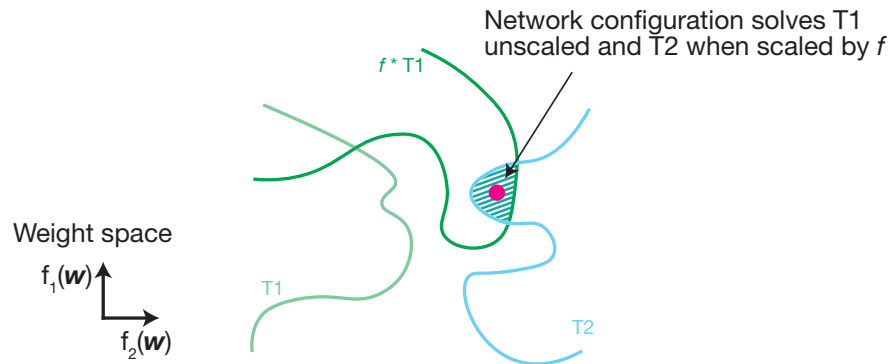


Figure S1. Relationship between neuromodulated weight configurations. For the contradictory end behaviors in response to shared stimuli, as in the modified Go-NoGo task, a single network without neuromodulation cannot simultaneously learn both tasks as depicted in this schematic by the lack of overlap between weight space that solves task 1 ($T1$) and task 2 ($T2$). By scaling weights in the network by a factor f , neuromodulation allows overlap between the $f \cdot T1$ and $T2$ spaces. The network solves task 1 when weights are unscaled ($T1$), and task 2 when weights are scaled ($f \cdot T1 \cap T2$).

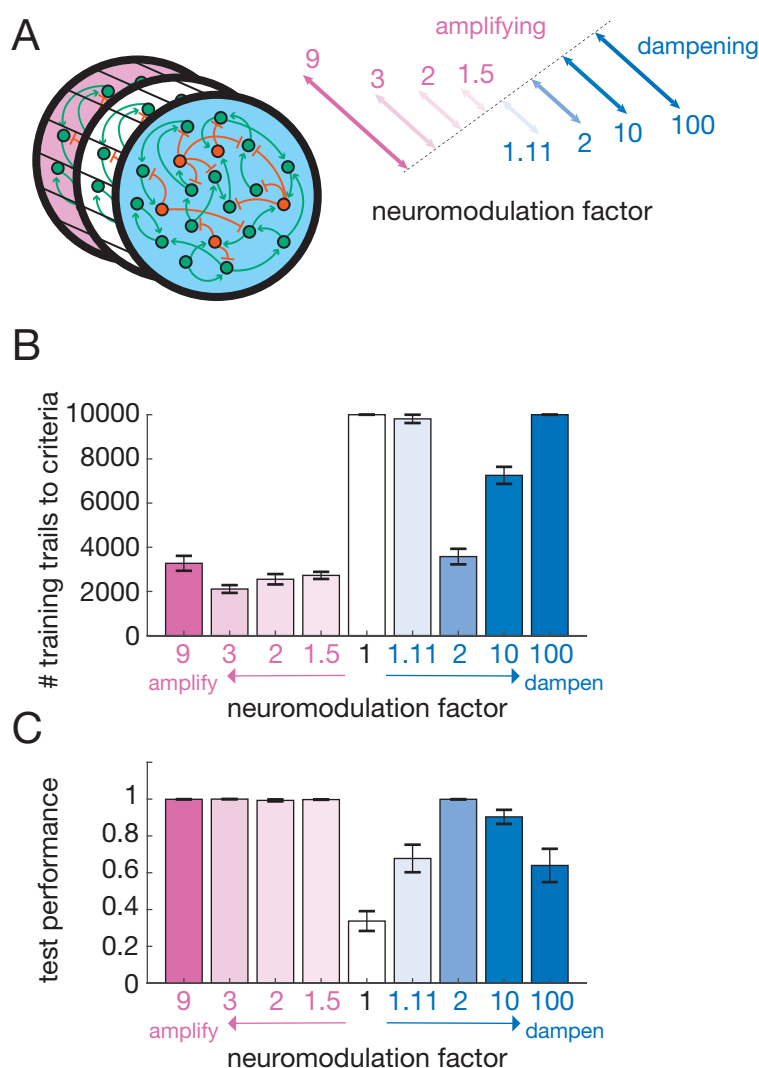


Figure S2. Neuromodulation weight scaling mechanism works over a range of factors. (A) A range of neuromodulation factors were tested on the modified Go-NoGo task. Factors were applied to weights initialized as described in Methods. (B) All amplifying factors tested supported task learning in a similar number of training trials. Dampening factors that were either too small (e.g. 1.11) or too large (e.g. 100) led to longer training. (C) All amplifying factors tested perfect task performance. Dampening factors that were either too small or too large led to impaired performance, though better than without neuromodulation (factor=1). Extreme strong dampening effectively silences all transmission between neurons, impairing information flow in the network. Too little scaling (e.g. 1.11 factor dampening) did not create enough separation to distinguish the behaviors.

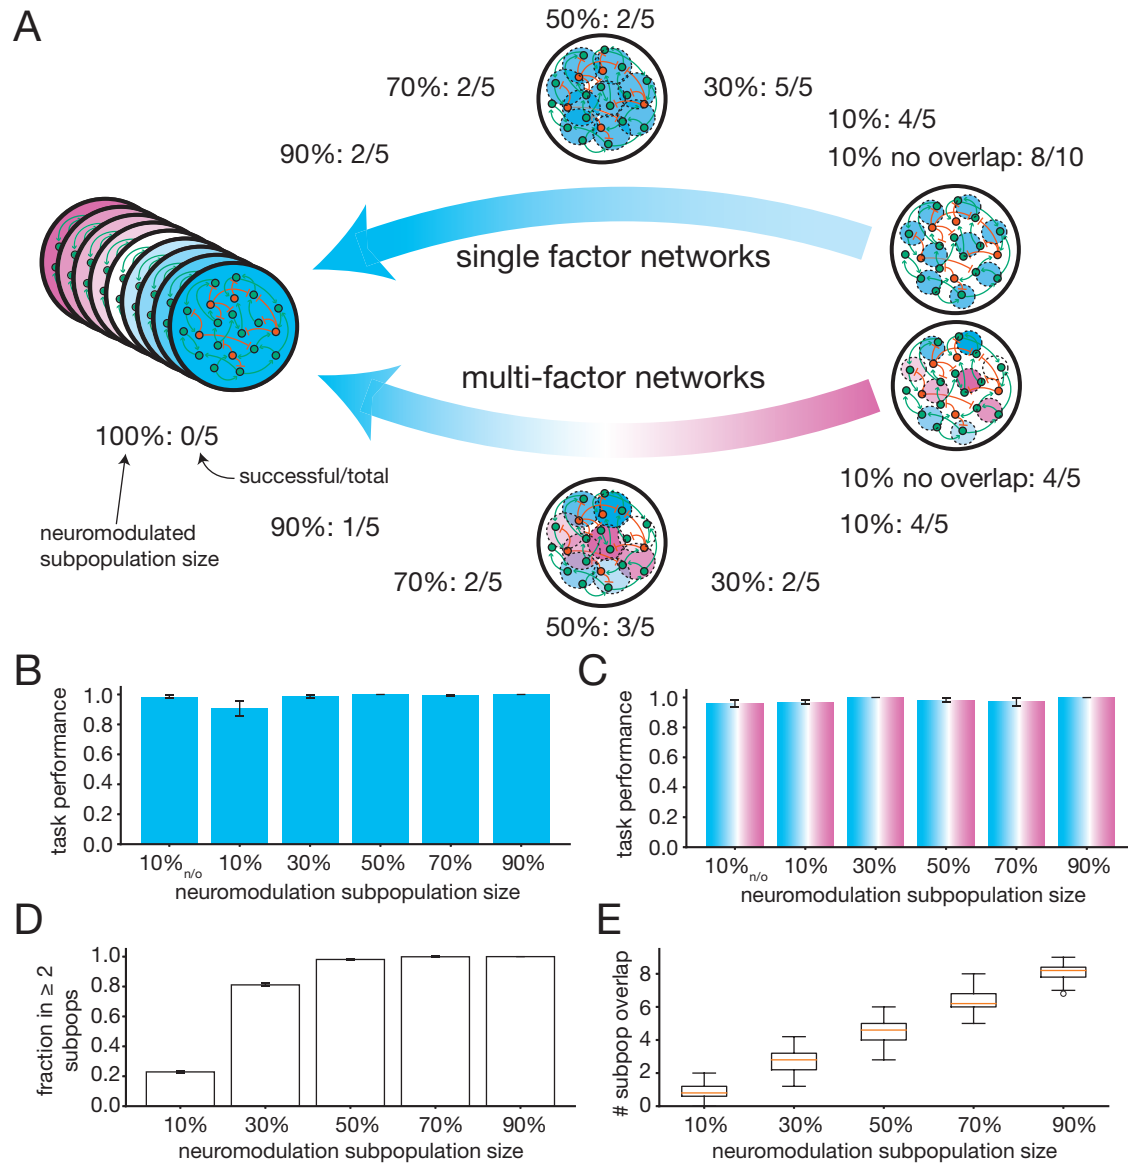


Figure S3. Spectrum of neuromodulation subpopulation size and factor. (A) Different conformations of neuromodulation of a single network support multi-behavior task (9-behavior modified Go-NoGo). Neuromodulation with a single factor (e.g. $f_{nm} = 2.5$) of non-overlapping and overlapping subpopulations across the spectrum of sizes could learn the full 9-behavior task, with larger overlapping subpopulations less consistently learning the full task ("correct/total" refers to independent networks that achieved successful training loss criteria over total attempted). Neuromodulation of the full network with different factors ($f_{nm} \in [1:1:9]$) consistently supported the 3 behavior task (5/5 correct/total), but not >3 behavior tasks (4-behavior 0/5; 5-behavior 0/5; 9-behavior 0/5 correct/total). Subpopulation neuromodulation with different factors ($f_{nm} \in [1:1:9]$) could also learn the 9-behavior task with overlapping and non-overlapping subpopulations, similarly exhibiting less consistent learning with larger overlapping subpopulations. (B) Single factor networks test performance on 9-behavior task across networks that achieved training loss criteria. (C) Same as (B) but for multi-factor networks. (D) Fraction of neurons neuromodulated for ≥ 2 conditions across range of neuromodulated subpopulation sizes. (E) Mean number of conditions a neuron was neuromodulated for across range of subpopulation sizes. For (D)–(E), 5 replicates per condition. All error bars are SEM.

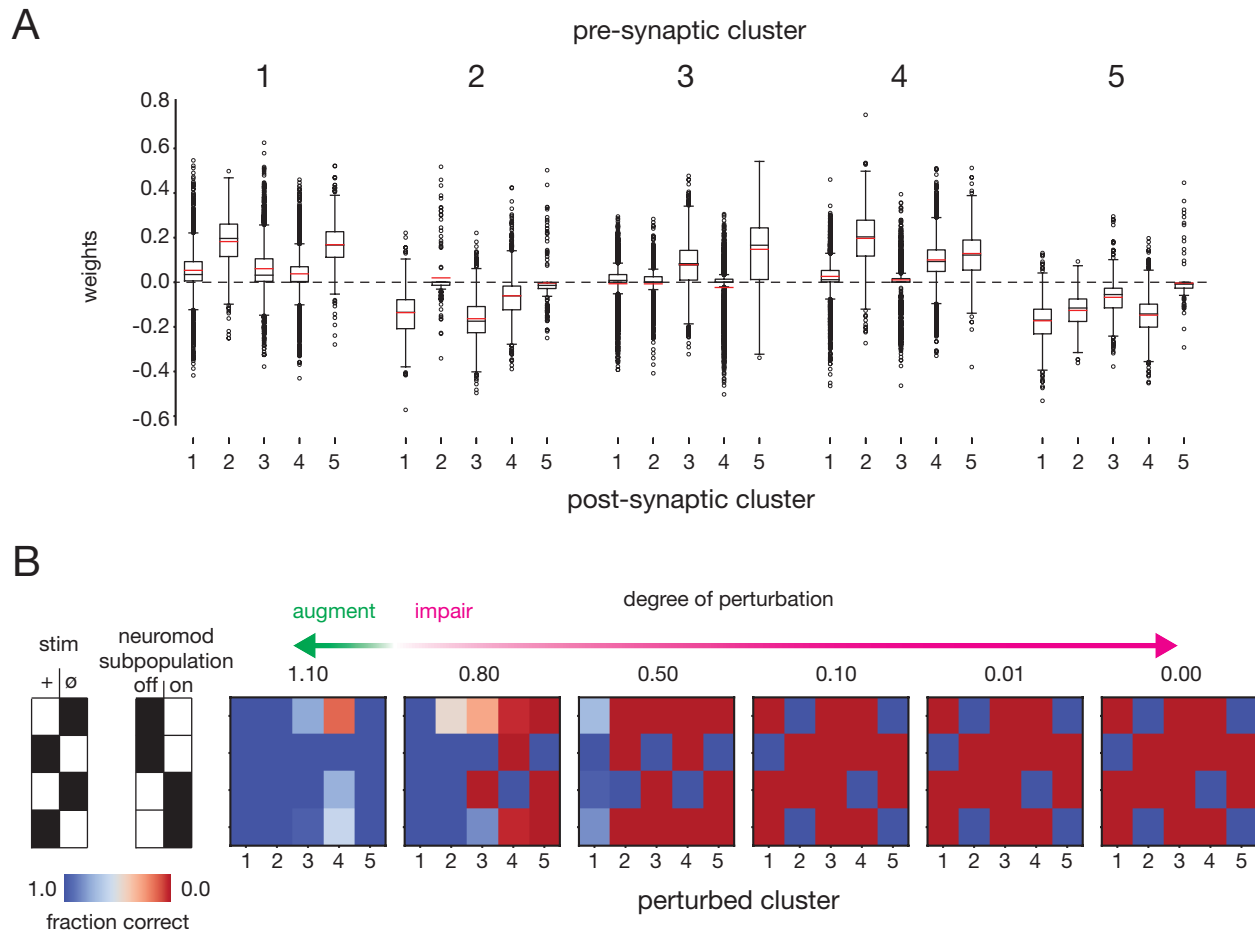


Figure S4. Inter-cluster weights and functional clustering for network in Fig. 4. (A) Weight distribution from each cluster (pre-synaptic) to other clusters (post-synaptic). Clusters 2 and 5 were predominantly composed of inhibitory neurons resulting in inhibitory tone on downstream clusters, while clusters 1, 3, and 4 were predominantly excitatory. (B) Augmentation (perturbation=1.1x) and impairment (i.e., synaptic lesions; perturbation $\in [0.80, 0.50, 0.10, 0.01, 0.00]x$) of specific clusters led to impairment of performance. For most clusters, increased degree of synaptic impairment led to decreased performance, with higher sensitivity for stimulus-context associated clusters (Fig. 4f; e.g. cluster 4 for + stim, neuromodulation off). In some stimulus-contexts, strong impairment or full lesion of clusters led to recovery of performance (perturbation ≤ 0.10 for clusters 2 and 5 with null stimulus without neuromodulation or positive stimulus with neuromodulation) suggestive of a counter-intuitive switch in underlying cluster dynamics.

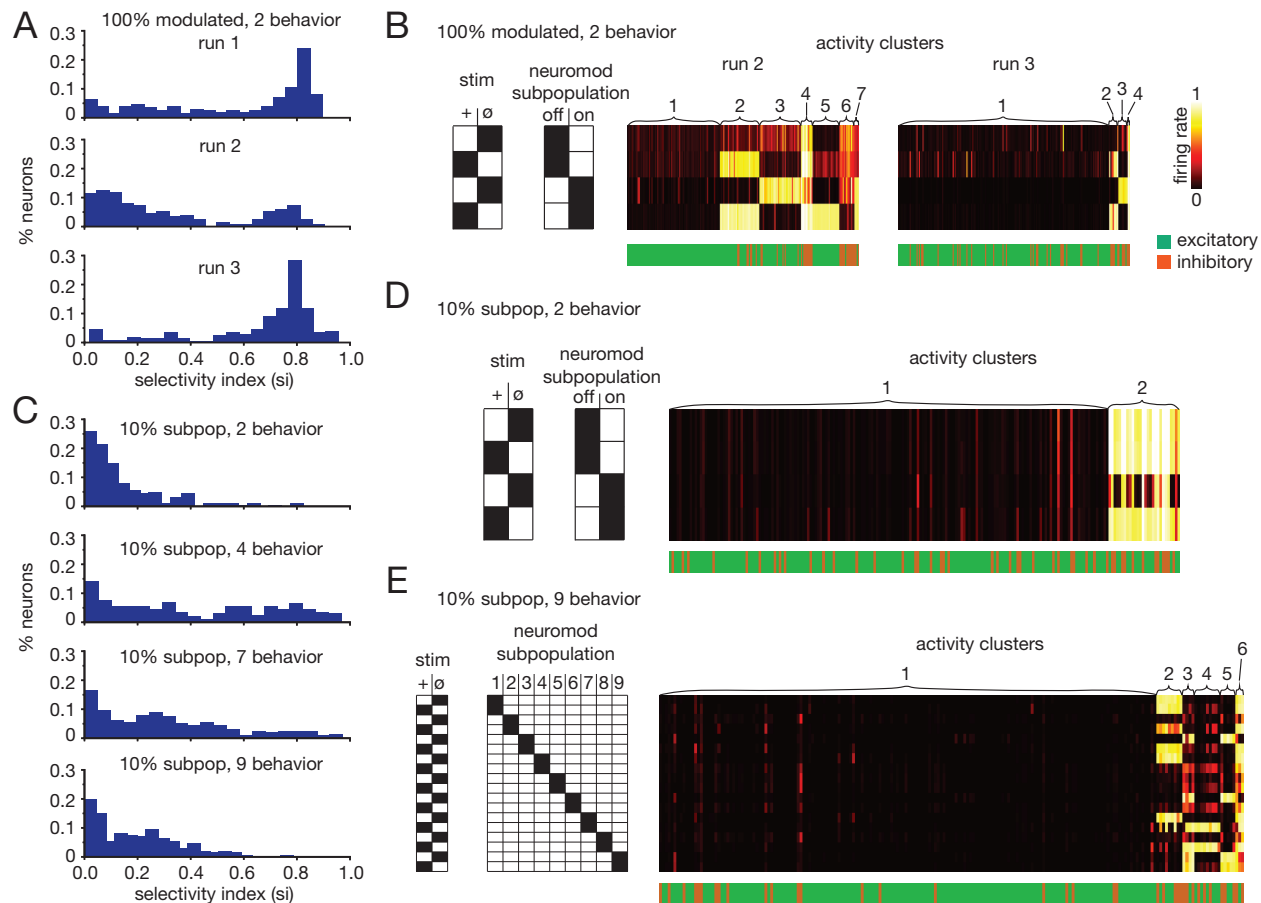


Figure S5. Independent networks have unique clustering patterns & different network configurations exhibit less selective neurons and complex, overlapping clustering profiles. (A) Neural selectivity profiles for 3 networks trained independently on the same condition (100% modulated, 2-behavior). All histograms exhibit high selectivity peak. (B) Each network from (A) exhibits a unique neural activity clustering pattern including number of clusters and excitatory-inhibitory composition of clusters (see Fig. 4f for run 1 cluster heatmap). (C) Selectivity index histograms for example networks with different configurations and tasks. Compared to a network in which all neurons were neuromodulated trained on the 2-behavior modified Go-NoGo, network configurations with smaller neuromodulated subpopulations and more behaviors exhibited less selective neurons and more non-selective neurons. (D) A network with 10% neuromodulated subpopulation trained on the 2-behavior modified Go-NoGo formed 2 activity clusters, showing high overlap across stimulus-context conditions. (E) The stimulus-context clustering profile for a network with 10% neuromodulated subpopulations (non-overlapping) trained on the 9-behavior modified Go-NoGo task. Cluster profiles are complex and highly overlapping. (D)–(E) cluster heatmaps and neuron type labels are same as (B).

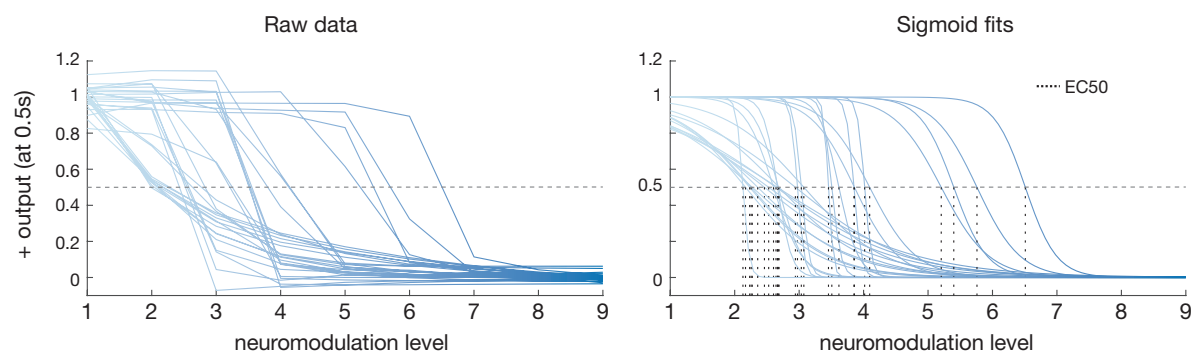


Figure S6. Variability in network neuromodulatory transitions. Left: Output to positive stimulus at time point 0.5s of 29 networks at varying levels of neuromodulation. Each network was independently trained with amplifying neuromodulation factor 9 on the modified Go-NoGo task and then tested at intermediate neuromodulation levels. Right: Sigmoid fits to raw data with EC50 of each curve indicated by dotted vertical line. EC50s ranged from 2.1 to 6.5. Slope of sigmoids (σ_{slope}) ranged from 0.9 to 26.3.

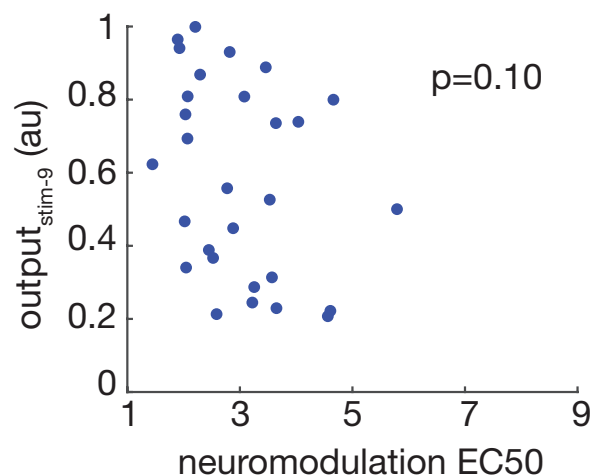


Figure S7. Neuromodulation EC50 versus maximum response to electrical modulation. Some networks did not reach an output level equivalent to half the response of maximal neuromodulation (EC50) even at the highest level of electrical modulation of -9 units. To assess all networks electrical modulation sensitivity in comparison to neuromodulation sensitivity, network output at maximum electrical stimulation is compared to EC50. Like EC50 vs stim₅₀, there is no significant correlation ($p=0.10$).

Supplemental Appendix

A. Relationship to isolated gain modulation

Gain modulation changes the slope of the unit activation function (the neuron’s “intrinsic excitability”) by changing a gain parameter g :

$$r = f(x; g)$$

where r is the firing rate and x is the synaptic current variable. For the sigmoid activation function this corresponds to:

$$r = \frac{1}{1 + e^{-gx}}$$

To see the effective change on the activation function of amplifying or dampening the weights we can compare the effect of gain modulation to weight modulation on the equation that governs neural dynamics:

$$\tau \dot{x}_i = -x_i + \sum_j W_{ji} r_{ji} + W_{ui} u + N(0, 0.1)$$

If we assume no input and no noise, we get a simplified equation describing neuron dynamics:

$$\tau \dot{x}_i = -x_i + \sum_j W_{ji} r_{ji}$$

Gain modulation g gives:

$$\tau \dot{x}_i = -x_i + \sum_j W_{ji} r_{ji}(g)$$

while weight modulation with a neuromodulation factor m gives

$$\tau \dot{x}_i = -x_i + \sum_j m \cdot W_{ji} r_{ji}(1)$$

To compare each type of modulation, we can consider the modified terms:

$$\begin{aligned} 1 \cdot r(g) &= \frac{1}{1 + e^{-gx}} && \text{(gain effect)} \\ m \cdot r(1) &= \frac{1}{\frac{1}{m} + \frac{1}{m} \cdot e^{-x}} && \text{(weight effect)} \\ &= \frac{1}{\frac{1}{m} + e^{-x - \ln m}} \end{aligned}$$

These effects are equivalent when

$$g = -\frac{\ln(\frac{1-m}{m} + e^{-x - \ln m})}{x} \quad (1)$$

So generally gain modulation, g , is only equivalent to a weight modulation by m if the gain term is precisely the time varying function of x defined by equation (1). Furthermore, for some values of m , there is no equivalent g for certain values of x . E.g. for $m = 2$, g is defined by equation (1) only for $x < 0$. For arbitrary values of x with fixed, constant g and m :

$$r_{ji}(g) \neq m \cdot r_{ji}(1)$$

except when $g = m = 1$, i.e., when there is no modulation. Thus, weight neuromodulation and gain modulation operate through different effects.

B. Chemical and electrical modulation

Chemical (neuromodulation) and electrical modulation in our network operate in directionally similar manner with qualitatively different effects. This can be seen by inspecting the equation governing each neuron i 's synaptic current variable and thereby the network's activity dynamics:

$$\tau \dot{x}_i = -x_i + \sum_j W_{ji} r_{ji} + W_{ui} u + N(0, 0.1)$$

For any given neuron we can break up the terms by whether an input neuron is in the neuromodulated subpopulation or not:

$$\tau \dot{x}_i = -x_i + \sum_k W_{ki} r_{ki} + \sum_q W_{qi} r_{qi} + W_{ui} u + N(0, 0.1)$$

where k is the index for non-neuromodulated neurons and q is the index for neuromodulated neurons.

Neuromodulation in our model acts by scaling the target neurons outgoing weights by a factor f :

$$\tau \dot{x}_i = -x_i + \sum_k W_{ki} r_{ki} + f \cdot \sum_q W_{qi} r_{qi} + W_{ui} u + N(0, 0.1)$$

Electrical stimulation acts by adding exogenous synaptic current to target neurons. For a given neuron i in the non-neuromodulated subpopulation, electrical stimulation of the neuromodulated subpopulation is felt through altered incoming firing rates:

$$\tau \dot{x}_i = -x_i + \sum_k W_{ki} r_{ki} + \sum_q W_{qi} r_{qi, Estim} + W_{ui} u + N(0, 0.1)$$

and a neuron i in the neuromodulated subpopulation is additionally affected through direct stimulation:

$$\tau \dot{x}_i = -x_i + \sum_k W_{ki} r_{ki} + \sum_q W_{qi} r_{qi, Estim} + W_{ui} u + u_{Estim} + N(0, 0.1)$$

Thus we can see that both chemical and electrical stimulation act through the same term in the equation that governs neural synaptic currents yet in different manners. Neuromodulation directly scales presynaptic weighted inputs from neuromodulated neurons, whereas electrical stimulation acts by altering the firing rate of presynaptic neuromodulated neurons with an additional direct influence on the synaptic current if the neuron of interest is in the electrically stimulation subpopulation.

The similarities of these forms of modulation (acting through the same terms of the synaptic current equation) indicates why they can have similar affects on network output in some circumstances. Nevertheless, the differences in how they affect the synaptic current equation are propagated through the recurrent connections of the network at each time step which drives the distinct dynamical changes seen under chemical versus electrical modulation.