



HAL
open science

Gain modulation of actions selection without synaptic relearning

Elif Köksal-Ersöz, Pascal Chossat, Frédéric Lavigne

► **To cite this version:**

Elif Köksal-Ersöz, Pascal Chossat, Frédéric Lavigne. Gain modulation of actions selection without synaptic relearning. 2024. hal-04418804v2

HAL Id: hal-04418804

<https://hal.science/hal-04418804v2>

Preprint submitted on 17 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Gain modulation of probabilistic selection without synaptic relearning

Elif Köksal-Ersöz^{1, 2}, Pascal Chossat^{3, 4}, Frédéric Lavigne^{3, 5}

¹Inria, Villeurbanne, France; ²Cophy Team, Lyon Neuroscience Research Center, Bron, France; ³MathNeuro Project-Team, Inria, Montpellier, France, ⁴Laboratoire Jean-Alexandre Dieudonné , Université Côte d'Azur, Nice, France; ⁵BCL, Université Côte d'Azur, CNRS, Nice, France.

Keywords: neuronal gain, short-term depression, local inhibition, latching dynamics, slow-fast dynamics

Abstract

Adaptation of behavior requires the brain to change goals in a changing environment. Synaptic learning has demonstrated its effectiveness in changing the probability of selecting actions based on their outcome. In the extreme case, it is vital not to repeat an action to a given goal that led to harmful punishment. The present model proposes a simple neural mechanism of gain modulation that makes possible immediate changes in the probability of selecting a goal after punishment of variable intensity. Results show

how gain modulation determine the type of elementary navigation process within the state space of a network of neuronal populations of excitatory neurons regulated by inhibition. Immediately after punishment, the system can avoid the punished populations by going back or by jumping to unpunished populations. This does not require particular credit assignment at the ‘choice’ population but only gain modulation of neurons active at the time of punishment. Gain modulation does not require statistical relearning that may lead to further errors, but can encode memories of past experiences without modification of synaptic efficacies. Therefore, gain modulation can complements synaptic plasticity.

1 Introduction

Adaptation of behavior requires the brain to select actions that bring benefits and to avoid those that bring costs. Selection of the most probably rewarded action requires updating the relation between actions and their rewards and punishments from past experience (Rescorla, 1971; Sutton and Barto, 2018; Behrens et al., 2007). In a stable environment, the probability of the outcomes is best estimated by experiences going back a long way to ensure exploitation of the rewarded actions. If the probabilistic structure of feedback changes, actions become prone to errors due to uncertainty on expected feedback. Exploration of different actions and relearning must then update the action-outcome probability (Cohen et al., 2007; Domenech et al., 2020) through statistical learning of trials and errors (Lazartigues et al., 2021; Rey et al., 2022; Lazartigues et al., 2023). This raises the question of the degree of recency of the experiences and of

the intensity of the outcomes to be taken into account.

Animal studies report that ancient and recent rewards are both memorized (Corrado et al., 2005; Fusi et al., 2007; Bernacchia et al., 2011). Changes in the rate of rewards and punishments can be adjusted by changes in a single learning rate to update synaptic efficacies (Behrens et al., 2007; Fusi et al., 2007; Nassar et al., 2010, 2012). Computational models have investigated a learning rate at synapses connecting a context to different actions that depends on the magnitude of the error signal to optimize the weighting of old and recent experiences over multiple timescales (Iigaya, 2016; Iigaya et al., 2019).

When changes in the environment are transient, different structures of action-feedback are associated to transiently alternating contexts. In this case, it is beneficial to adapt behaviors to the transient state without forgetting the previous state and hence without the need to relearn it. Synaptic relearning leads to forgetting of the previous probabilistic structure of action-outcome relations. The previously learned and forgotten environment has then to be learned again through repetition of trials and errors, even though at a fast rate. However, even more dramatically, in the extreme case of severe and dangerous punishment, the action that led to it should not be repeated. Some errors must not be made twice. This does not give time (or opportunity) for statistical relearning. Then is it possible to change actions without synaptic relearning? And, in the extreme case, without the need for any further learning trial?

Various cortical functions are reported to rely on the modulation of the input-output gain at the level of neurons, defined as the slope of the neurons transfer function (Salinas and Thier, 2000; Carandini and Heeger, 2012; Eldar et al., 2013; Ferguson and Cardin,

2020; Debanne et al., 2019). At the level of network behavior, gain modulation changes large scale network dynamics Shine et al. (2021) and the correlation of neuronal output activity (Doiron et al., 2016), giving the network the computational ability to change which neurons are activated by a context even though values of synaptic efficacy are fixed (Lavigne and Darmon, 2008). Further, a recent computational model has reported that a network can switch activation between sequences of neurons coding for items depending on the modulation of the gain of these neurons (Köksal Ersöz et al., 2022). Gain modulation was shown as efficient in changing the probability of activating an item or another in the network state space for a fixed synaptic matrix. Here we investigate the conditions under which neuronal gain alone enables switching from one action to another without the need for synaptic relearning, and possibly with immediate effect in case of punishment. This raises the question of the assignment of punishment to the sequence of neurons that led to the punished behavior (*e.g.*, Li and Huang (2020); Asaad et al. (2017); He et al. (2015); Friedrich et al. (2011); Liu et al. (2021)).

2 Methods

The model has been directly inspired by (Köksal Ersöz et al., 2022), where the retrieving of multiple sequences in a collection of P learned states ξ^1, \dots, ξ^P has been investigated using the framework introduced in (Aguilar et al., 2017; Köksal Ersöz et al., 2020). In this model each learned state is a dynamically stable ‘pattern’ made of two active units coding for populations of neurons in the neural network, the others units being inactive. Moreover, these patterns can be destabilized under the effect of short

term synaptic depression (STD), hence allowing for dynamics of activation of patterns in the network state space.

In (Aguilar et al., 2017; Köksal Ersöz et al., 2020) any two consecutive learned patterns share one unit, so that the patterns form a chain of overlapping states $\xi^1 - \xi^2 - \dots - \xi^P$, where P is the number of patterns ξ^m . It has been shown that, under the effect of noise, the overlapping condition allows to produce a sequential stochastic dynamics, one state ‘jumping’ to the next with high probability. This behavior is called latching dynamics.

In (Köksal Ersöz et al., 2022) we analyzed the case when, as the system starting from ξ^1 reaches a given state ξ^m in the chain (the branching ‘node’), there is a choice among several continuing branches. In the simplest case, which is considered in the present work, the chain splits at the node into two chains $\xi^m, \xi^{m+1} \dots, \xi^q$ (branch 1, Br-1) and $\xi^m, \xi^{q+1}, \dots, \xi^P$ (branch 2, Br-2).

As shown in (Köksal Ersöz et al., 2022), as long as the connectivity matrix is symmetric, the probabilities that the dynamics starting from ξ^1 continues on either branch 1 or branch 2 are equal, but whenever the weights of connections from ξ^m to ξ^{m+1} are greater than those from ξ^m to ξ^{q+1} , the probability to continue on branch 1 is larger than the probability to continue on branch 2. The branching network which we numerically investigated is pictured in Figure 1A. Its parameters are $N = 10$, $P = 9$, $m = 4$ and $q = 7$.

The equations for the units are derived from (Amari, 1972), in which we have replaced the membrane potential u_i of each unit i by the activity $x_i = S(u_i) = 1/(1 + e^{-\gamma u_i})$. The variables are now the activities which take values in the interval $[0, 1]$ after

$S^{-1}(x_i)$ has been replaced by its polynomial expansion (we chose the simplest, linear approximation). The inhibition effects within the network are modeled by a term proportional to the averaged activity (Lerner et al., 2012).

The equation for unit i reads

$$\dot{x}_i = x_i(1 - x_i) \left(-\frac{4}{\gamma}x_i + \sum_{j=1}^N J_{i,j}x_j - \lambda \sum_{j=1}^N x_j - \lambda\nu_i x_i \right) + \eta \quad (1)$$

where γ is the gain and λ is the inhibitory coefficient. The coefficient ν_i accounts for the possibility of short-range inhibitory loops between excitatory neurons and inhibitory interneurons. This allows for selective self-inhibition of excitatory neurons, which depends on the number of afferent connections to excited units. In our case we take $\nu_i = 0$ for $i \neq m+1$ and $\nu_{m+1} = 1$ (unit $m+1$ receives inputs from units $m, m+1$ and $q+1$). Finally η is a noise term. We refer the reader to (Köksal Ersöz et al., 2022) for details. Note that $x_i = 0$ or 1 are always solutions of (1). Hence, any state such that $x_i = 0$ or $1, i = 1 \dots, N$ is a steady-state of the system. The learned states are stable solutions of this type.

The STD is expressed as follows. Let $J_{ij}(t)$ be the strength of the connection from units j to i at time t . We write $J_{ij}(t) = J_{ij}^{max} s_j(t)$, where J_{ij}^{max} is the connectivity matrix resulting from the learning process and $s_j(t)$ follow the STD law given in (Tsodyks and Markram, 1997), which is equivalent to

$$\tau_r \dot{s}_i = 1 - s_i - \rho s_i x_i. \quad (2)$$

The weights J_{ij}^{max} are computed from the simple Hebbian rule $J_{ij}^{max} = \sum_k \xi_i^k \xi_j^k$.

The parameter values of (1)-(2) are optimized from Köksal Ersöz et al. (2022) as $\lambda = 0.6, I = 0, \nu_4 = 1, \rho = 300, \tau_r = 1.2, \gamma = 10$ without punishment, $\gamma =$

{9, 5, 3.3, 2.5} with punishment (from weak (10%) to strong (75%) rates).

3 Results and Discussion

The model presented here allows mathematical analysis and simulation of the effect of gain modulation on the selection of goals towards which to direct actions. We investigate a system of $N = 10$ units (x_i , for $i \in \{1, \dots, 10\}$) encoding 9 patterns (states) (Fig 1A), which we write (A,B, ..., I) for convenience (Fig 1B). The units are placed in a 3-node graph where units 1-3 are along branch 0, units 5-7 along branch 1, units 8-10 along branch 2 and unit x_4 is the branching node connected to x_3 , x_5 and x_8 (Fig 1A). Such network architecture could correspond to sequences of internal representations of successive goals to go. The branching synaptic architecture correspond to a classical Y-maze (Ghafarimoghadam et al., 2022) for which the network codes successive goals to the final goal at the end of branches, which can be rewarded or punished. The activation of two units side by side (*e.g.*, x_2 and x_3) code for a pattern corresponding to a network state (*e.g.*, B). A change in pattern activated would correspond to the activation of a new goal where to orient actions in the Y-maze. For the sake of clarity, We focus here on the elementary building block of network behavior at a single branching node (Köksal Ersöz et al., 2022). But the results presented here can be generalized to more complex networks involving 4-ways or more branchings and *branches in branches*).

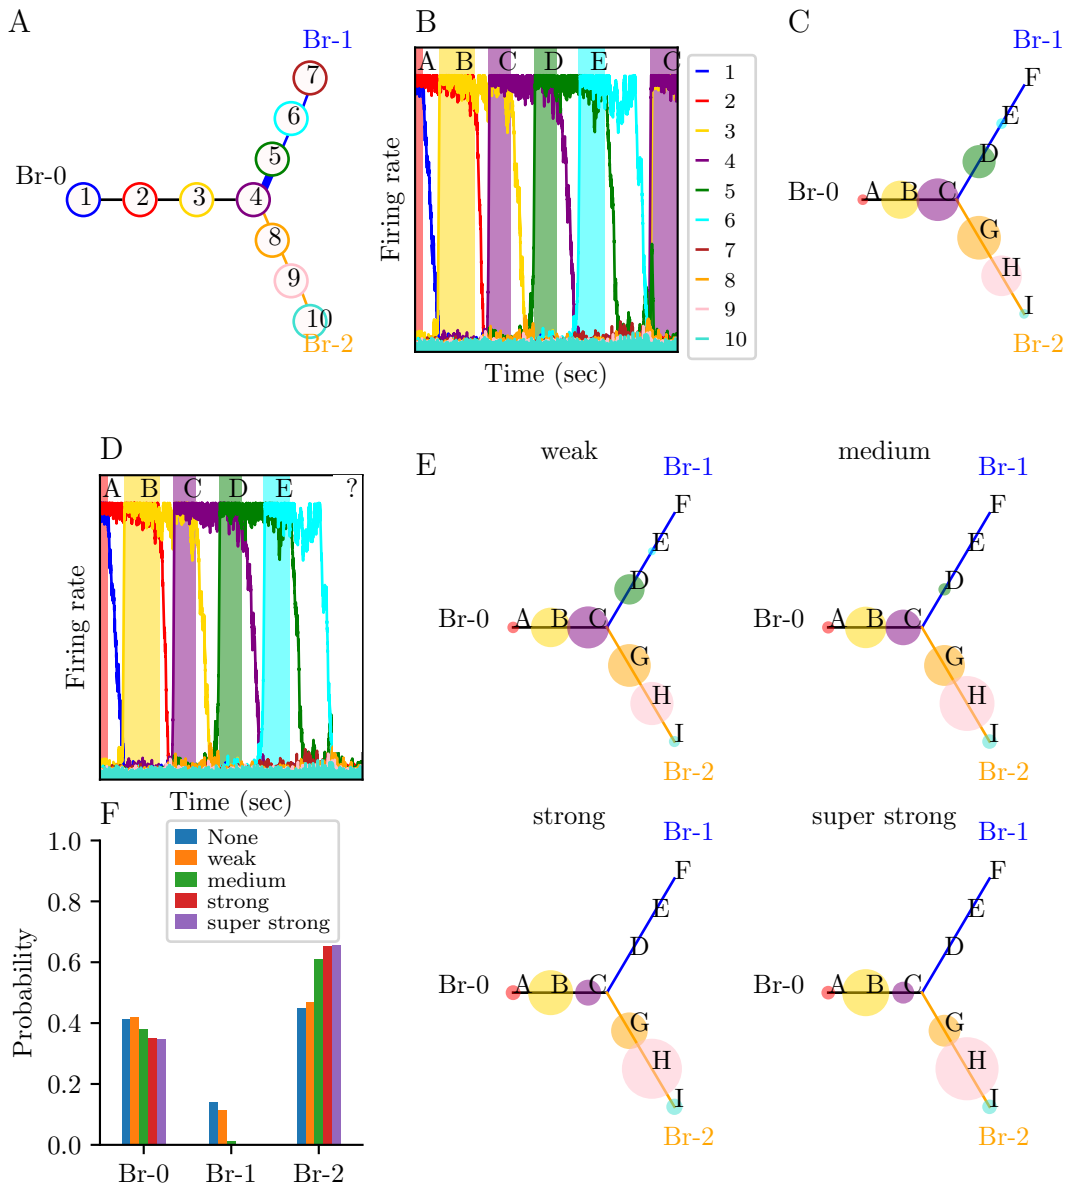


Figure 1: Branching behavior of a $N = 10$ units network at Trial T . (A) Network architecture of 10 units represented by numbered circles. Unit 4 is a branching node between three branches 0, 1 and 2. The synaptic efficacy between units 4 and 5 is 10% stronger than between units 4 and 8. (B) Network behavior is described from the activation of units 1 and 2 in the initial branch (Br-0). In this example, activation propagates from units 1 to 4 in a sequence, and (caption continues next page)

Figure 1: continues to units 5 and 6 in branch 1 (Br-1), thanks to the stronger efficacy between units 4 and 5. Line colors correspond to the activation of units and colored areas correspond to patterns coded by these units. (C) The network embeds 9 patterns coded by two connected units which follow each other in a sequence (letters A to I). Unit 4 is part of patterns C, D, and G, themselves parts of branches 0, 1 and 2, respectively. Circles size is proportional to the probability of activation of the patterns. In this same example as in (B), the sequential activation propagates from pattern A to pattern E in Br-1, again thanks to the stronger efficacy between units 4 and 5 of pattern D. Pattern D is the last possible pattern to activate in Br-1. After reaching it, the network randomly jumps to another pattern (here pattern C). (D) If punishment is given at the activation of pattern E, the network behavior after this pattern depends on the punishment rate. In this example, the network stops without activating any pattern. (E) Probability of activation of the patterns in the 3 branches immediately after punishment during the punished trial T , as a function of the level of punishment. Circles size is proportional to the probability of activation of the patterns immediately after punishment. (F) Probability of activation of any pattern in each of the 3 branches for all levels of punishment (1000 simulations for each level). When the punishment is weak (10% decrease in neuronal gain), the system can still activate patterns D or E but in only in 11% of the trials. For a medium punishment (50% decrease in gain), this ratio decreases to 1%. Whereas for a strong punishment (66% decrease in gain), the system does not activate patterns along Br-1 anymore. Instead, the system activates patterns along branch 2 (Br-2) with a 40% increase from weak to strong punishment.

The synaptic coupling coefficient between units 4 and 5 (branch 1) was 10% stronger than between units 3 and 4 (branch 0), and units 4 and 8 (branch 2). We study if and how gain modulation can change the sequence of neurons activated during a punished trial and on the immediately following trial. For simulations, the system was initialized at pattern A and punishment was applied to units 5 and 6 (coding for pattern E) when they became activated during the trial T . Punishment was assigned to neurons active at the time of feedback (Asaad et al., 2017), assumed to depend on punishment signaling (noradrenaline, serotonin and/or dopamine; (Tanaka et al., 2009; Cohen et al., 2012; Oleson et al., 2012; Cohen et al., 2015; Michely et al., 2022) that decreases neuronal gain (Henze et al., 2000; Bandyopadhyay and Hablitz, 2007; Noudoost and Moore, 2011; Seillier et al., 2017).

Results show that punishment-dependent gain modulation changes the probability of activating the punished and unpunished patterns. Further, the effects of punishment observed immediately after punishment *during* the punished trial T (Fig 1E, 1F) is maintained at the following trial $T + 1$ (Fig 2A, 2B).

Three main network behaviors are observed depending on the level of punishment:

- In the absence of punishment, the gain is the same in all units in the two branches.

In that case, the stronger synaptic connection between the branching units 4 and unit 5 drives the network behavior. It induces more frequent activation of branch 1 by activating patterns D then E after the initial sequence A-B-C (Figure 1B).

The model reproduces an exploitation strategy that increases the probability of reward (Cohen et al., 2007; Domenech et al., 2020).

- After punishment of medium intensity, the gain is decreased in the units 5 and 6

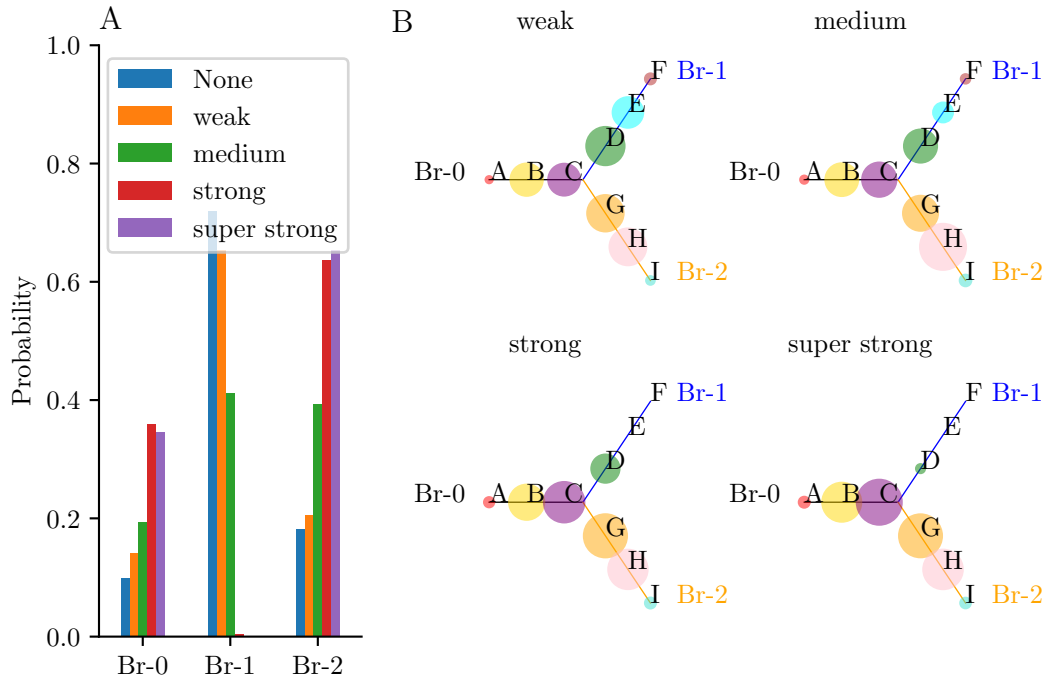


Figure 2: Branching behavior of a $N = 10$ units network (see Figure 1A-C) at Trial $T + 1$. (A) Probability of activation of patterns in the 3 branches during the punished trial T , for all levels of punishment (1000 simulations for each level). In the absence of any punishment ('None'), the system takes branch 2 (Br-2) in 71% of the trials. For a weak punishment, the system activates branch 1 (Br-1) in 65% of trials and Br-1 in 20% of trials. Moderate punishment equalizes the probability between the two branches (41% for Br-1 and 39% for Br-2). For strong punishment, the networks activates the Br-2 only (0% for Br-1 vs 64% for Br-2). (B) Probability of activation of the patterns in the 3 branches at trial $T + 1$ as a function of the level of punishment. Circles size is proportional to the probability of activation of the patterns. The effect of punishment at trial T is maintained in the trial $T + 1$. Strong punishment at trial prevents from activating the punished branch. Activation either goes back to Br-0 or jumps to Br-2.

(coding for pattern E) active at time of punishment. This decreases the probability of recalling patterns D or E in the punished branch 1. This occurs immediately during the punished trial T (Fig 1E) and in the following trial $T + 1$ (Fig 2B). The lower gain in the punished branch 1 makes neurons populations less responsive to input activity coming from the initial branch 0 (patterns A-B-C). Given the stronger synaptic efficacy between the punished branch 1 and the branching node 4, the system still activates the punished branch 1 but with lower probability (Fig 2A). The balanced probabilities to select the two branches corresponds to an exploration strategy to search for the most rewarded or less punished goals (Cohen et al., 2007; Domenech et al., 2020).

- After strong punishment, the gain is strongly decreased in the punished units 5 and 6. After the first regular sequence A-B-C (Fig 2A), the activation of patterns along branch 1 is stopped immediately at trial T (Fig 1F) and is avoided at trial $T + 1$ (Fig 2B), although this branch 1 is the most strongly (synaptically) associated to the branching node. In that case, the network can go back or switch directly to branch 2. The model reproduces an avoidance strategy that prevents from further strong punishment.

In the model presented in Figures 1 and 2, punishment decreased the gain of units (5-6) directly connected to the branching unit (4) and impacted the choice behavior at unit 4 without exploration of the branch 1. However, sequences of activation of units preceding a feedback can be longer, leading to punishment of distant units not directly connected to the branching unit 4 (*e.g.*, 6-7). In that case the system does not ‘know’ which branch was punished when arriving at the branching unit (because unit 5 connected to unit 4

was not punished). This scenario was tested in a network with longer branches (Fig 3).

When punishment arrived at patterns distant from the branching node (Fig 3A), gain was decreased on units 6 and 7 and remained unchanged in all other units of the punished branch that were not active at the time of punishment (here unit 5). Results show that such distant assignment of punishment is sufficient to switch branch at trials T and $T + 1$ (Fig 3B). The system can still activate the unpunished unit 5 at the beginning of the punished branch, but switches branch before arriving at the punished units (6-7). Then the network can begin in the punished branch but switches to and fully activates the other (unpunished) branch. This navigation process makes possible to switch behavior after punishment. It is not exclusive of synaptic eligibility traces He et al. (2015) but does not need assignment of feedback to the whole sequence of units that were activated before the punishment. This way the punishment assignment obeyed a simple mechanism combining punishment with only units active at the time of feedback. Such simple mechanism of assignment is made efficient even in long sequences of patterns, thanks to the network's ability to navigate forwards, backwards or jump branches in its state space, depending on the gain of the neuron population.

Interestingly, different navigation processes can be generated by different types of sequential activation in the phase space of the network. The punished state can be avoided by an activity that can either

- go back to the starting branch 0 coding for the context,
- stay at the pattern preceding the punished pattern in branch 1,
- jump directly to the unpunished branch 2.

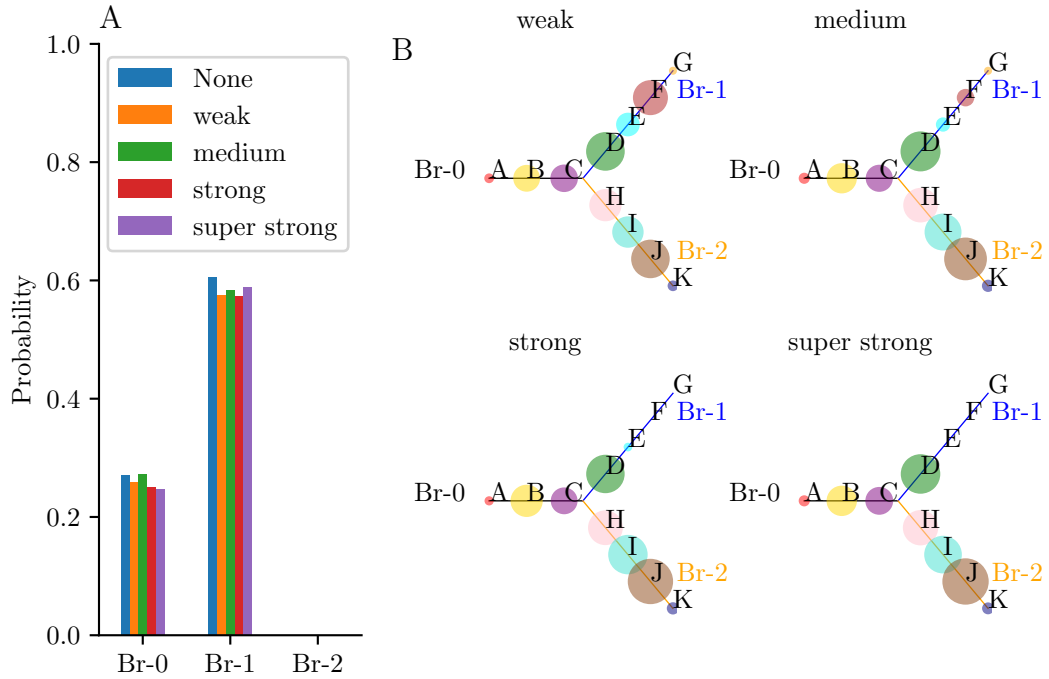


Figure 3: Branching behavior of a $N = 12$ units network at Trial $T + 1$. (A) Probability of activation of patterns in the 3 branches immediately after punishment *during* the punished trial T , for all levels of punishment (1000 simulations for each level). Punishment does not prevent from activating pattern D in branch 1 (Br-1) at trial T , due to the stronger synaptic connection with the branching node 4. (B) Probability of activation of the patterns in the 3 branches at trial $T + 1$ *following* the punished trial T , as a function of the level of punishment. Circles size is proportional to the probability of activation of the patterns. Increased punishment decreases the probability of activation of patterns E and F in Br-1 at trial $T + 1$. Strong punishment at trial T even prevents from activating the punished branch again at trial $T + 1$. The network activates pattern D at the beginning of Br-1, but then either goes back to branch 0 (Br-0) or jumps to branch 2 (Br-2).

The model exhibits elementary navigation processes within the network's state space (*building blocks* (Köksal Ersöz et al., 2022)) that depend on neuronal gain. Such processes allow the network to adapt its behavioral strategy to synaptically learned reward and to the level of punishment. Taken as a whole, results show that gain modulation switches the network behavior between exploitation and exploration behaviors:

- an approach exploitation behavior: a synaptically driven activation of the goal learned as rewarded (branch 1) rather than the goal learned as punished (branch 2) (2A, no punishment),
- an exploration behavior: a balanced selection between the goal synaptically learned as rewarded and that has been recently weakly punished (branch 1) and the goal synaptically learned as punished (branch 2) (Fig 2B and Fig 3B, weak punishment),
- an avoidance exploitation behavior: a gain driven blocking of the strongly punished branch 1 and switch to the learned as less punished branch 2 (Fig 2B and Fig 3B, strong punishment).

The avoidance strategy after strong punishment allows to not repeat harmful errors during trial and error learning. We hope that these can provide a framework for modeling and experimental approaches investigating the effects of punishment on gain modulation and goal selection without synaptic relearning.

Conclusion

From a learning point of view, the present results indicates that neuronal gain can embed knowledge on the relation between goals and outcomes. This way the value of neuronal gain contributes to store memories of past experiences (Zhang and Linden, 2003) without changes in the synaptic matrix. Such gain-based neural learning could complete synaptic learning in the alternation between exploitation, exploration and avoiding strategies. Synaptic learning allows knowledge to change rapidly and/or slowly depending on the volatility of the environment (Iigaya et al., 2019). Given that gain modulation alters neuronal excitability, it could provide an alternative means of storing knowledge at the microscopic neuronal level, in addition to synapses (Daoudal and Debanne, 2003; Abraham et al., 2019; Debanne et al., 2019). The present model shows that local gain decrease by punishment signaling (Henze et al., 2000; Bandyopadhyay and Hablitz, 2007; Noudoost and Moore, 2011; Seillier et al., 2017) is able to change network behavior at the macroscopic level without synaptic relearning. In the framework of the effects of gain on the macroscopic network behavior, intrinsic gain learning raises a number of points for further study: to what extent the effects of gain on network behavior differ from those of synaptic learning? under which conditions of intensity and frequency reward or punishment affect transient and/or long-term behavior? and what are the interactions between the variations in gain and synaptic efficacy when both change at the same time? Further, in the present model, patterns learned in the synaptic matrix are assumed to code for goals that orient actions. Model refinements could couple this network with a network whose patterns code for the sequences of actions between the different goals.

In its current state, the present models shows that gain-based neuronal learning enables a modulation of knowledge activated in memory in a less synapse-dependent way though not altering knowledge previously stored in the synapses. In the case of punishment, gain-based learning could then give the system the necessary time for synaptic relearning without repeating harmful errors.

Acknowledgements

The authors gratefully thank Pascal Benquet and Gianluigi Mongillo for insightful discussions on a previous version of this article. This work was supported by the Hebbian ANR-project (ANR-23-CE28-0008).

References

- Abraham, W.C., Jones, O.D., & Glanzman, D.L. (2019). Is plasticity of synapses the mechanism of long-term memory storage? *NPJ science of learning*, 4(1), 9.
- Aguilar, C., Chossat, P., Krupa, M., & Lavigne, F. (2017). Latching dynamics in neural networks with synaptic depression. *PLOS ONE*, 12(8), e0183710.
- Amari, S. (1972) Characteristics of random nets of analog neuron-like elements. *IEEE Transactions on systems, man, and cybernetics*, 5, 643–57.
- Asaad, W.F., Lauro, P.M., Perge, J.A., & Eskandar, E.N. (2017). Prefrontal neurons encode a solution to the credit-assignment problem. *Journal of Neuroscience*, 37(29), 6995–7007.

- Bandyopadhyay, S., & Hablitz, J. (2007). Dopaminergic Modulation of Local Network Activity in. *J Neurophysiol*, 97, 4120–4128.
- Behrens, T.E., Woolrich, M.W., Walton, M.E., & Rushworth M.F. (2017). Learning the value of information in an uncertain world. *Nature neuroscience*, 10(9),1214–1221.
- Bernacchia, A., Seo, H., Lee, D., & Wang, X.J. (2011). A reservoir of time constants for memory traces in cortical neurons. *Nature neuroscience*, 14(3), 366–372.
- Carandini, M., & Heeger, D.J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1), 51–62.
- Cohen, J.Y. Amoroso, M.,W., & Uchida N. (2015) Serotonergic neurons signal reward and punishment on multiple timescales. *elife*, 4, e06346.
- Cohen, J.Y., Haesler, S., Vong, L., Lowell, B.B.,& Uchida, N. (2012). Neuron-type-specific signals for reward and punishment in the ventral tegmental area. *Nature*, 482(7383), 85–88.
- Cohen, J.D., McClure, S.M., & Yu, A.J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933–942.
- Corrado, G.S., Sugrue, L.P., Seung, H.S., & Newsome, W.T. (2005). Linear-nonlinear-Poisson models of primate choice dynamics. *Journal of the experimental analysis of behavior*, 84(3), 581–617.
- Daoudal, G., & Debanne, D. (2003). Long-term plasticity of intrinsic excitability: learning rules and mechanisms. *Learning & memory*, 10(6), 456–465.

- Debanne, D., Inglebert, Y., & Russier, M. (2019). Plasticity of intrinsic neuronal excitability. *Current opinion in neurobiology*, 5, 73–82.
- Doiron, B., Litwin-Kumar, A., Rosenbaum, R., Ocker, G.K., & Josić, K. (2016). The mechanics of state-dependent neural correlations. *Nature neuroscience*, 19(3), 383–393.
- Domenech, P., Rheims, S., & Koechlin, E. (2020). Neural mechanisms resolving exploitation-exploration dilemmas in the medial prefrontal cortex. (2020). *Science* 369(6507), eabb0184.
- Eldar, E., Cohen, J.D., & Niv, Y. (2013). The effects of neural gain on attention and learning. *Nature neuroscience*, 16(8), 1146–1153.
- Ferguson, K.A., & Cardin, J.A. (2020). Mechanisms underlying gain modulation in the cortex. *Nature Reviews Neuroscience*, 21(2), 80–92.
- Friedrich, J., Urbanczik, R., & Senn, W. (2011). Spatio-temporal credit assignment in neuronal population learning. *PLoS computational biology*, 7(6), e1002092.
- Fusi, S., Asaad, W.F., Miller, E.K., & Wang, X.J. (2007). A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales. *Neuron*, 54(2), 319–333.
- Ghafari Moghadam, M., Mashayekh, R., Gholami, M., Fereydani, P., Shelley-Tremblay, J., Kandezi, N., Sabouri, E., & Motaghinejad, M. (2022). A review of behavioral methods for the evaluation of cognitive performance in animal models: Current techniques and links to human cognition. *Physiology & behavior*, 244, 113652.

- He, K., Huertas, M., Hong, S.Z., Tie, X., Hell, J.W., Shouval, H., & Kirkwood, A. (2015). Distinct eligibility traces for LTP and LTD in cortical synapses. *Neuron*, 88(3), 528–538.
- Henze, D.A., González-Burgos, G.R., Urban, N.N., Lewis, D.A., & Barrionuevo, G. (2000). Dopamine increases excitability of pyramidal neurons in primate prefrontal cortex. *Journal of neurophysiology*, 84(6), 2799–2809.
- Iigaya, K. (2016) Adaptive learning and decision-making under uncertainty by meta-plastic synapses guided by a surprise detection system. *elife*, 5, e18073.
- Iigaya, K., Ahmadian, Y., Sugrue, L.P., Corrado, G.S., Loewenstein, Y., Newsome, W.T., & Fusi, S. (2019). Deviation from the matching law reflects an optimal strategy involving learning over multiple timescales. *Nature communications*, 10(1), 1466.
- Köksal Ersöz, E., Aguilar, C., Chossat, P., Krupa, M., & Lavigne, F. (2020). Neuronal mechanisms for sequential activation of memory items: Dynamics and reliability. *PLOS ONE*, 15(4), e0231165.
- Köksal Ersöz, E., Chossat, P., Krupa, M., & Lavigne, F. (2022). Dynamic branching in a neural network model for probabilistic prediction of sequences. *Journal of Computational Neuroscience*, 50, 537–557.
- Lavigne, F., & Darmon, N. (2008). Dopaminergic neuromodulation of semantic priming in a cortical network model. *Neuropsychologia*, 46, 3074–87.
- Lazartigues, L., Mathy, F., & Lavigne, F. (2021). Statistical learning of unbalanced exclusive-or temporal sequences in humans. *PLoS ONE*, 16(2), e0246826.

- Lazartigues, L., Mathy, F., & Lavigne, F. (2023). Probability, dependency, and frequency are not all equally involved in statistical learning. *Experimental Psychology*, 69(5), 241-52.
- Lerner, I., Bentin, S., & Shriki, O. (2012). Spreading activation in an attractor network with latching dynamics: automatic semantic priming revisited. *Cognitive Science*, 36, 1339–82.
- Li C, Huang H. (2020). Learning credit assignment. *Physical Review Letters*, 125(17), 178301.
- Liu, Y.H., Smith, S., Mihalas, S., Shea-Brown, E., & Sümbül, U. (2021). Cell-type-specific neuromodulation guides synaptic credit assignment in a spiking neural network. *Proceedings of the National Academy of Sciences*, 118(51), e2111821118.
- Michely, J., Eldar, E., Erdman, A., Martin, I.M., & Dolan, R.J. Serotonin modulates asymmetric learning from reward and punishment in healthy human volunteers. *Communications Biology*. 2022; 5(1):812.
- Nassar, M.R., Rumsey, K.M., Wilson, R.C., Parikh, K., Heasly, B., & Gold, J.I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature neuroscience*, 15(7), 1040–1046.
- Nassar, M.R., Wilson, R.C., Heasly, B., & Gold, J.I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*, 30(37), 12366–12378.

- Noudoost, B., & Moore, T. (2011). Control of visual cortical signals by prefrontal dopamine. *Nature*, 474(7351), 372–375.
- Oleson, E.B., Gentry, R.N., Chioma, V.C., & Cheer, J.F. (2012). Subsecond dopamine release in the nucleus accumbens predicts conditioned punishment and its successful avoidance. *Journal of Neuroscience*, 32(42), 14804–14808.
- Rescorla, R.A.. (1971). Variation in the effectiveness of reinforcement and nonreinforcement following prior inhibitory conditioning. *Learning and motivation*, 2(2), 113–123.
- Rey, A., Fagot, J., Mathy, F., Lazartigues, L., Tosatto, L., Bonafos, G., Freyermuth, J., & Lavigne, F. (2022). Learning Higher-Order Transitional Probabilities in Nonhuman Primates. *Cognitive Science*, 46(4), e13121.
- Salinas, E., & Thier, P. (2000). Gain modulation: a major computational principle of the central nervous system. *Neuron*, 27(1), 15–21.
- Seillier, L., Lorenz, C., Kawaguchi, K., Ott, T., Nieder, A., Pourriahi, P., & Nienborg, H. (2017). Serotonin decreases the gain of visual responses in awake macaque V1. *Journal of Neuroscience*, 37(47), 11390–405.
- Shine, J.M., Müller, E.J., Munn, B., Cabral, J., Moran, R.J., & Breakspear, M. (2021). Computational models link cellular mechanisms of neuromodulation to large-scale neural dynamics. *Nature neuroscience*, 24(6), 765–776.
- Sutton, R.S., & Barto, A.G. Reinforcement learning: An introduction. MIT press; 2018.

- Tanaka, S.C., Shishida, K., Schweighofer, N., Okamoto, Y., Yamawaki, S., & Doya, K. (2009). Serotonin affects association of aversive outcomes to past actions. *Journal of Neuroscience*, 29(50), 15669–15674.
- Tsodyks, M.V., & Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, 94, 719–23.
- Zhang, W., & Linden, D.J. (2003) The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nature Reviews Neuroscience*, 4(11), 885–900.