This electronic thesis or dissertation has been downloaded from Explore Bristol Research, http://research-information.bristol.ac.uk

Author:
Zhu, Heng Wei

Title:
Cholinergic-mediated adaptive learning in cortical microcircuits

General rights
Access to the thesis is subject to the Creative Commons Attribution - NonCommercial-No Derivatives 4.0 International Public License. A copy of this may be found at https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode This license sets out your rights and the restrictions that apply to your access to the thesis so it is important you read this before proceeding.

Take down policy
Some pages of this thesis may have been removed for copyright restrictions prior to having it been deposited in Explore Bristol Research. However, if you have discovered material within the thesis that you consider to be unlawful e.g. breaches of copyright (either yours or that of a third party) or any other law, including but not limited to those relating to patent, trademark, confidentiality, data protection, obscenity, defamation, libel, then please contact collections-metadata@bristol.ac.uk and include the following information in your message:

•Your contact details
•Bibliographic details for the item, including a URL
•An outline nature of the complaint

Your claim will be investigated and, where appropriate, the item in question will be removed from public view as soon as possible.

Cholinergic-mediated adaptive learning in cortical microcircuits

By

Heng Wei Zhu

Department of Physiology, Pharmacology and Neuroscience
University of Bristol

A dissertation submitted to the University of Bristol in accordance with the requirements of the degree of Doctor of Philosophy in the Faculty of Life Sciences.

April 2023

Word count: 51656

ABSTRACT

With frequent exposure to novelty, it is imperative that the brain is capable of rapidly adjusting the appropriate synaptic connections to adapt its behavioural responses to new situations. Acetylcholine is a neuromodulator released during novelty and is important for synaptic plasticity and learning via its heterogeneous modulatory effects on different cell types. Here, we introduce a framework of cholinergic-mediated credit assignment (ChoCA) that models the cholinergic system (CS) based on adaptive methods used in machine learning. The CS serves as an adaptive module, integrating error signals to modulate synaptic updates derived from a backpropagation-like credit assignment mechanism in cortical microcircuits. The cholinergic signals are modelled at varying degrees of diffuseness to reflect the heterogeneity of the information transmitted. On standard pattern recognition tasks, our results show that diffuse cholinergic modulation is sufficient for effective adaptive learning. Moreover, the model postulates that cholinergic modulation of somatostatin (SST) interneurons is critical for credit assignment; specifically, it predicts that acetylcholine depresses the SST-to-pyramidal cell (PC) synapse. We also develop a novel plausible backprop-like model called Bursting Cortico-Cortical Networks which provides a mechanism for communicating error signals as bursting activity to the CS. Preliminary findings from *ex vivo* slice electrophysiology experiments reveal that while the excitability of SST/OLM interneurons increases with bath application of a cholinergic agonist, the inhibitory transmission of these interneurons onto CA1 PCs is reduced in line with the model's prediction. A dual optogenetics approach using the Chrimson and ChRger opsins was also developed to examine the effects of endogenous acetylcholine on this inhibitory transmission. Furthermore, optogenetic stimulation of these interneurons during a long-term potentiation (LTP) induction protocol prevents LTP indicating their importance in the regulation of synaptic plasticity. Overall, our work demonstrates the efficacy of coupling neuromodulatory signals with efficient credit assignment for adaptive learning in the brain.

ACKNOWLEDGEMENTS

It would not have been possible to complete this PhD if it wasn't for all the amazing people that have supported and guided me along the way. Thank you to the Neural Dynamics directors for giving me this opportunity and the Wellcome Trust / University of Bristol for financially supporting my PhD. I would like to give a huge thanks to my supervisors, Jack Mellor and Rui Ponte Costa, for their optimism, patience and guidance on the research that I wanted to pursue.

Thanks to all of the past and present members of the Mellor and Costa labs for the laughs, chats and the helpful discussions throughout the years. I would like to thank Paul Banks, Travis Bacon and Matt Udakis for teaching me how to patch those pesky cells, and in particular Paul Banks and Matt Udakis for being my ChatGPT for any Ephys questions. Thanks to Sarah Stuart for patiently teaching me how to do surgeries and Paul Anastasiades for his enthusiastic input on my project. I would also like to thank Michael Ashby for being a beacon of knowledge for any questions I had about optics and Jon Witton for the useful discussions on 2p imaging. I would like to especially thank Will Greedy for teaching me how to code, working alongside me during those graveyard shifts and generally being an absolutely great friend.

Thanks to my friends, particularly the ones that I have lived with during this PhD: Joe W, Dome, James, Adam, Joe U and Josh, for making sure that I was alive and helping to balance that very difficult concept of a work-life balance.

Thanks to my family for all of the encouragement, for making the time to visit me and coming in clutch with those care packages full of delicious food to keep me going. And finally to Laura, you have been the best partner I could have ever asked for. Even though I know how stressful your own PhD journey has been, you have only shown me your utmost love and support, and thank you just isn't enough to express how grateful I am.

COVID-19 STATEMENT

Research activities, particularly for PhD students, were significantly impacted by the COVID-19 pandemic. Due to a lack of consistent guidelines from authorities, sectors had to quickly make their own decisions in response to the crisis. In an effort to prioritize national research priorities, the Health Research Authority suspended all non-COVID research activities early on in the pandemic.

The University of Bristol closed its doors on 18th March 2020. The Animal Support Unit's (ASU) faced difficulties in light of high level University of Bristol decisions on closure and lockdown, prioritising compliance rather than consideration of management issues such as research and animal welfare. The ASU had to implement extreme measures, including ceasing all animal experiments and reducing or terminating most animal breeding colonies. The ASU was allowed to reopen at the end of July 2020 for researchers on a priority list, but the resumption of activity was slower than anticipated due to refurbishment works that continued until 12th October 2020. These refurbishment works highlighted existing issues with the ASU that further slowed down the process, negatively affecting the proper restart of research activities.

The lockdown and subsequent ASU closure resulted in the curtailment of several research activities, including:

- Breeding of the colonies: the ChAT-FlpO mouse line was delivered on the 27th October 2020 and due to strict breeding restrictions imposed by the ASU, the number of animals in the required mouse lines was kept to a minimum until the unit fully reopened.
- Development of a novel dual optogenetic approach to address the hypotheses of my project
- Data collection

The project was significantly impacted by the lockdown, as it had only started four months prior with limited preliminary data collected beforehand. The project relied on an exploratory approach that involved breeding two different transgenic mouse lines, one of which had not been characterised in the CNS before. This made it difficult to adapt the project once research activities were allowed to resume. The colonies required for the experiments and breeding had to be restored to appropriate numbers and ages, delaying experimental work until late May. This delay had a knock-on effect on the project, particularly as it was heavily reliant on the validation of the novel mouse line. Furthermore, the extended break from performing ex vivo electrophysiology resulted in additional time needed to reestablish smooth and efficient experiments upon returning to the lab. As a result, fewer electrophysiology experiments were performed as originally planned, and the focus of the thesis shifted more towards modeling. Overall, the experimental side of the project was greatly affected.

AUTHOR'S DECLARATION

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

SIGNED: .... DATE: ....

TABLE OF CONTENTS

Page

List of Tables **xiv**

List of Figures **xv**

1 Introduction **1**

1.1 Credit assignment 1
  1.1.1 Error backpropagation algorithm 2
  1.1.2 The credit assignment problem 3
  1.1.3 Biologically plausible backprop algorithms 4
1.2 The hippocampus 6
  1.2.1 Architecture and circuits 6
  1.2.2 Synaptic transmission and plasticity 9
  1.2.3 Dendritic integration 12
  1.2.4 Spatial navigation 13
1.3 Acetylcholine 16
  1.3.1 Acetylcholine in the hippocampus 16
    1.3.1.1 Septohippocampal pathway 16
    1.3.1.2 Acetylcholine receptors 16
  1.3.2 Learning and memory 18
  1.3.3 Cholinergic modulation of synaptic plasticity 19
  1.3.4 Theories of acetylcholine and cognitive processing 20
1.4 Aims and hypotheses 21

2 Cholinergic-mediated adaptive learning in a cortical microcircuit model **24**

2.1 Statement of contributions 24
2.2 Introduction 24
2.3 Adaptive credit assignment in cortical microcircuits 26
  2.3.1 Error-encoding Dendritic Network 26
  2.3.2 Cholinergic adaptive modulation methods 28
    2.3.2.1 Adagrad-ChoCA 28

2.3.2.2 RMSprop-ChoCA 29
2.3.2.3 Adadelta-ChoCA 30
2.3.2.4 Adam-ChoCA 30
2.3.2.5 Yi-ChoCA 30
2.3.2.6 Eve-ChoCA 31
2.3.3 Related work 31
2.4 Experimental details 32
2.4.1 Parameter requirements for each optimiser 32
2.4.2 Hyperparameter Search 32
2.4.3 Learning tasks 33
2.4.3.1 MNIST and F-MNIST 34
2.4.3.2 Spiral 34
2.4.3.3 Catastrophic forgetting tasks 34
2.5 Results 36
2.5.1 Synapse-specific neuromodulation is not required for efficient adaptive learning 36
2.5.2 The effect of removing cholinergic modulation on interneuron synaptic plasticity and learning 37
2.5.3 Adaptive learning increases sparsity and selectivity 38
2.5.4 The effect of adaptive learning on catastrophic forgetting 39
2.5.5 Neuromodulation of synaptic plasticity directly using cost information improves learning 40
2.5.5.1 Spirals 43
2.5.5.2 MNIST 43
2.5.5.3 F-MNIST 44
2.5.6 The effect of cost-based adaptive modulation of interneurons on learning . 45
2.5.7 The effect of cost-based adaptive learning on catastrophic forgetting 46
2.6 Discussion 47
2.6.1 Limitations 50
2.6.2 Closing remarks 51

3 Single-phase deep learning in cortico-cortical networks **52**
3.1 Statement of contributions 52
3.2 Introduction 52
3.3 Bursting Cortico-Cortical Networks 54
3.3.1 Burst Ensemble Multiplexing 54
3.3.2 Rate-based BurstCCN 55
3.3.3 Continuous-time BurstCCN 57
3.3.4 Spiking BurstCCN 57

3.3.5 Dalean BurstCCN . . . 58
3.3.6 Related work . . . 60
3.4 Experimental details . . . 61
3.4.1 Spiking XOR task . . . 61
3.4.2 Continuous-time input-output task . . . 61
3.4.3 CartPole-v0 reinforcement learning benchmark . . . 61
3.4.4 Rate-based model experiments . . . 62
3.4.4.1 Hyperparameter search . . . 62
3.4.4.2 Feedback plasticity on **Q** weights . . . 62
3.4.4.3 MNIST . . . 62
3.4.4.4 CIFAR-10 . . . 63
3.4.5 MNIST with Dalean model . . . 64
3.5 Results . . . 65
3.5.1 BurstCCN can learn with a single learning phase . . . 65
3.5.2 BurstCCN can learn with dynamic input-output . . . 66
3.5.3 Feedback plasticity rule facilitates alignment to backprop updates . . . 67
3.5.3.1 BurstCCN with weak feedback approximates backpropagation algorithm . . . 67
3.5.3.2 Learning **Q** feedback connections better approximates backprop-derived gradients . . . 68
3.5.3.3 Feedback plasticity on Y-weights . . . 69
3.5.4 BurstCCN learns image classification tasks with multiple hidden layers . 70
3.5.4.1 MNIST . . . 70
3.5.5 Symmetric MNIST . . . 71
3.5.5.1 CIFAR-10 . . . 72
3.5.6 Deep reinforcement learning with BurstCCN . . . 73
3.5.7 Dalean BurstCCN learns MNIST . . . 75
3.6 Discussion . . . 75
3.6.1 Limitations . . . 77
3.6.2 Closing remarks . . . 78

4 Cholinergic modulation of OLM interneurons and their regulation of LTP 79
4.1 Introduction . . . 79
4.2 Materials and methods . . . 80
4.2.1 Materials . . . 80
4.2.2 Animals . . . 80
4.2.3 Hippocampal slice preparation . . . 81
4.2.4 Electrophysiology . . . 81
4.2.4.1 Carbachol dose-response . . . 82

4.2.4.2 Somatostatin interneuron excitability 82
4.2.4.3 Associative LTP 83
4.2.5 Immunohistochemistry 83
4.2.6 Microscopy 84
4.2.7 Imaging analysis 84
4.2.8 Statistical analysis 85
4.3 Results 85
4.3.1 Characterisation of Chrna2-cre mouse line 85
4.3.2 Effect of cholinergic modulation on the inhibitory synaptic transmission from OLM and SST interneurons 88
4.3.3 Effect of cholinergic modulation on the intrinsic excitability of stratum oriens SST interneurons 90
4.3.4 OLM interneurons regulate aLTP in CA1 91
4.3.5 Effects of spiking, bursting and amount of depolarisation on aLTP 94
4.4 Discussion 96
4.4.1 Limitations 98
4.4.2 Future experiments 99
4.4.3 Closing remarks 99

5 Independent optical excitation of ChAT- and Chrna2-expressing neural populations 100
5.1 Introduction 100
5.2 Materials and methods 102
5.2.1 Materials 102
5.2.2 Animals 102
5.2.3 Surgical procedures 102
5.2.3.1 AAV vectors 103
5.2.4 Hippocampal slice preparation 103
5.2.5 Electrophysiology 103
5.2.5.1 Independent optical excitation of OLM interneurons 104
5.2.5.2 Endogenous acetylcholine release in ChAT-cre x Ai32 105
5.2.5.3 Effect of endogenous acetylcholine on dendritic inhibition 105
5.2.6 Two-photon imaging characterisation of ACh3.0 sensor 106
5.2.6.1 Carbachol dose-response 107
5.2.6.2 Detecting photostimulated release of endogenous acetylcholine . 107
5.2.7 Retrograde tracing of septo-hippocampal innervations 107
5.2.8 Characterisation of Chrna2-Cre x ChAT-FlpO mouse line 108
5.2.9 Characterisation of AAV2.9-ChAT-Cre-P2A-EGFP 108
5.2.10 Immunohistochemistry 108

5.2.11 Microscopy 110
5.2.12 Imaging analysis 110
5.2.13 Statistical analysis 110
5.3 Results 111
5.3.1 Characterisation of ChRger and Chrimson expression in OLM interneurons111
5.3.2 Independent optical excitation of OLM interneurons 113
5.3.3 Photostimulating endogenous release of ACh in ChAT-Cre x Ai32 mice . . 117
5.3.4 Characterisation of ACh3.0 sensor 118
5.3.5 Characterisation of viral channelrhodopsin expression in Chrna2-Cre x FlpO mice 126
5.3.6 Characterisation of ChAT-cre virus 129
5.3.7 Examining septo-hippocampal projections of Chrna2-expressing cells . . . 129
5.3.8 The effect of endogenous acetylcholine on OLM inhibitory transmission . . 131
5.4 Discussion 135
5.4.1 Limitations 136
5.4.2 Future experiments 138
5.4.3 Closing remarks 138

6 General discussion **139**

Bibliography **143**

LIST OF TABLES

TABLE **Page**

2.1 Number of additional parameters required by each adaptive method. 32
2.2 Learning rates obtained using hyperparameter search for the EDN model trained on MNIST and F-MNIST. 33
2.3 Learning rates obtained using Bayesian hyperparameter optimisation for the standard ANN models trained on the Spiral dataset. 33
2.4 Hyperparameters obtained using Bayesian hyperparameter optimisation for the 3 hidden layer EDN models trained on F-MNIST. 34

3.1 CartPole-v0 hyperparameters 62
3.2 Q-weight learning hyperparameters 63
3.3 MNIST hyperparameters 63
3.4 CIFAR-10 architecture 64
3.5 CIFAR-10 hyperparameters (random fixed **Y** feedback weight regime) 64
3.6 CIFAR-10 hyperparameters (**W**-**Y** symmetric weight regime) 64
3.7 Dalean BurstCCN hyperparameters for MNIST 65

5.1 The voltage outputs from the CED acquisition board and resultant irradiances for each LED through a 10x objective lens. 105
5.2 The voltage outputs from the CED acquisition board and resultant irradiances for a 470 nm LED through a 5x objective lens. 105

LIST OF FIGURES

FIGURE **Page**

1.1 How the backpropagation algorithm works. 3
1.2 The role of distal apical dendrites of pyramidal neurons in credit assignment. 5
1.3 Hippocampal anatomy and circuitry. 7
1.4 Circuit diagram illustrating some of the interneuron networks in the hippocampus. . 9
1.5 Types of dendritic regenerative events. 13
1.6 Place cells and remapping. 15
1.7 Illustrative overview of cholinergic pathways in the rodent brain. 17

2.1 Model of adaptive credit assignment in cortical microcircuits via cholinergic modulation (ChoCA). 26
2.2 The Spiral dataset. 35
2.3 Efficient adaptive learning without synaptic-specific neuromodulation. 36
2.4 Impact of removing cholinergic modulation of interneuron synapses on adaptive learning. 37
2.5 Adaptive learning yields a better use of network representation capacity by increasing sparsity and selectivity. 38
2.6 The impact of adaptive learning on catastrophic forgetting for the MNIST- and F-MNIST-split-five tasks. 39
2.7 Performance of adaptive methods on split MNIST. 41
2.8 Performance of adaptive methods on split F-MNIST. 42
2.9 Performance of cost-based adaptive methods on the Spirals dataset in standard deep neural networks. 43
2.10 Performance of cost-based adaptive methods on the MNIST dataset in the EDN model. 43
2.11 Performance of cost-based adaptive methods on the F-MNIST dataset in the EDN model. 44
2.12 Performance of cost-based adaptive methods on the F-MNIST dataset in a 3 hidden layer EDN model. 45
2.13 Impact of removing cost-based cholinergic modulation of interneuron synaptic plasticity on adaptive learning. 46
2.14 The impact of cost-based adaptive learning on catastrophic forgetting for the MNIST- and F-MNIST-split-five tasks. 47

2.15 Performance of cost-based adaptive methods on split MNIST and split F-MNIST. . . . 48

3.1 Bursting cortico-cortical networks (BurstCCN) for credit assignment through bursting activity. 54
3.2 Spiking BurstCCN does not require multi-phase learning to solve the XOR classification task. 65
3.3 BurstCCN can learn a dynamic non-linear regression task. 67
3.4 Feedback learning rule enables a close alignment with backprop and feedback alignment. 69
3.5 Feedback plasticity rule on **Y** weights facilitates updates to align with backprop and feedback alignment. 70
3.6 BurstCCN learns to classify handwritten digits (MNIST) with deep networks. 70
3.7 BurstCCN with **W**-**Y** symmetric weights. 72
3.8 BurstCCN with convolutional layers learns to solve natural image classification task (CIFAR-10). 72
3.9 BurstCCN learns a policy to perform well on the CartPole-v0 reinforcement learning benchmark. 74
3.10 BurstCCN with Dalean constraints learns to classify handwritten digits (MNIST). . . 76

4.1 OLM interneurons reside in the stratum oriens and innervate the stratum lacunosum moleculare of the hippocampus. 87
4.2 CCh causes a dose-dependent decrease in OLM and SST interneuron inhibition. . . . 89
4.3 CCh increases the excitability of SST interneurons. 91
4.4 OLM interneuron activity prevents theta-burst-induced aLTP in the SC pathway. . . 93
4.5 The number of spikes, bursts and amount of depolarisation during induction does not correlate with the expression of plasticity. 95

5.1 Chrimson and ChRger selectively express in stratum oriens OLM interneurons. . . . 112
5.2 The channel kinetic properties of ChR2, ChRger and Chrimson. 114
5.3 Characterisation of opsin violet (395 nm), blue (470 nm) and amber (590 nm) light sensitivities for two-colour excitation in acute hippocampal slices. 116
5.4 Photostimulation of cholinergic innervations onto OLM interneurons. 117
5.5 Fluorescence response of ACh3.0-expressing cells in acute hippocampal brain slices to the bath application of 100 μM CCh. 119
5.6 Fluorescence response of ACh3.0-expressing cells in acute hippocampal brain slices to different concentrations of bath applied CCh. 121
5.7 Detecting optogenetically evoked endogenous ACh release with ACh3.0 sensor in ChAT-Cre x Ai32 acute hippocampal brain slices. 123
5.8 Fluorescence detection of 1 Hz blue light (470 nm) photostimulation in wildtype C57BL/6J acute hippocampal brain slices. 126

5.9 AAV1-CAG-FLEXFRT-ChR2(H134R)-mCherry transfection of ChAT-expressing cells in the MS of Chrna2-Cre x FlpO mice. 127
5.10 FlpO-dependent ChR2 and GFP transfection of ChAT-expressing cells in the dorsal striatum of Chrna2-Cre x FlpO mice. 128
5.11 AAV2.9-ChAT-Cre-P2A-EGFP transfection of ChAT-expressing cells in the MS of C57BL/6J mice. 130
5.12 Chrna2-expressing neurons in MS do not project to dorsal hippocampus. 131
5.13 The effect of optogenetically evoked ACh release on OLM interneuron inhibition. . . . 133
5.14 Long-range inhibitory septo-hippocampal projections could potentially inhibit OLM interneurons. 134

CHAPTER 1

INTRODUCTION

1.1 Credit assignment

In neuroscience, credit assignment refers to the process of assigning credit to various neurons or synapses for a given behavioural outcome, and is a crucial component of learning in hierarchical neural networks. This process provides important information for synaptic plasticity mechanisms to appropriately modify the system to reach its desired output. Understanding how credit assignment occurs in the brain is critical for understanding how we learn from experience and adapt our behaviour in response to changing environments. Plasticity rules have traditionally been based on two-factor Hebbian updates which rely on the correlations between pre- and postsynaptic activity. However, this approach is inadequate for credit assignment since it overlooks the downstream effects of synapses on the network output (Richards and Lillicrap, 2019; Lillicrap et al., 2020). Therefore, coincident pre- and postsynaptic activity could provide an eligibility trace to determine which synapses are suitable for change but a third factor that represents downstream information in the form of neuromodulation or feedback must be involved to determine whether these synaptic changes occur (Gerstner et al., 2018). Credit assignment solely with a global neuromodulatory error signal is not sufficient for learning difficult tasks in deep networks where the effect of each synapse on the output is highly complex (Werfel et al., 2004; Hiratani et al., 2022). It is more efficient if a feedback error signal is present to directly inform the neuron about the impact of its synapses on the downstream network. The error backpropagation algorithm provides this feedback error signal to each neuron indicating how altering their synapses will affect the overall error of the network (Lillicrap et al., 2020; Rumelhart et al., 1986). This section will explore the error backpropagation algorithm and how it could potentially be implemented in the brain to address the credit assignment problem.

1.1.1 Error backpropagation algorithm

The error backpropagation algorithm (backprop) is a widely used method for training artificial neural networks (ANNs). An ANN is a computational model that is inspired by the structure and function of the biological neural networks in the human brain. It consists of a large number of interconnected processing elements, called neurons, typically organised into multiple layers with all-to-all connectivity between neurons of one layer and the next. These connections are referred to as weights and can be thought of as synaptic connections with varying synaptic strengths. When an input is provided to the network, the activities of the neurons are propagated forward through the weights to the final layer where the output of the network is produced. Backprop is an efficient way to adjust the weights of the connections between neurons in a network, allowing the network to learn and improve its performance over time. Backprop uses the chain rule from calculus to compute the gradient of a loss function with respect to the weights (Lillicrap et al., 2020; Rumelhart et al., 1986). The loss function, also known as the cost function or objective function, is a mathematical function that measures the difference between the predicted output of a model and the actual output or target value. Gradient descent is an optimisation approach that uses this gradient to directly update the weights in a way that reduces the error of the network's predictions (Fig. 1.1).

Backprop was first introduced in the 1970s by several researchers independently, including Paul Werbos, David Rumelhart, and Geoffrey Hinton (Rumelhart et al., 1986). However, it was not widely adopted until the 1980s when the parallel development of faster computers and larger datasets made it more practical for training larger and more complex neural networks. Today, backprop is the standard method for training many types of neural networks, including feedforward networks, recurrent networks, and convolutional networks (Goodfellow et al., 2016).

Backprop has several variations and extensions, depending on the type of neural network and the specific optimisation method used. For example, it can be extended for training recurrent neural networks on sequential data by using a method called "backpropagation through time" (Werbos, 1988). Additionally, there are variations of the optimisation method, such as stochastic gradient descent and adaptive learning rate methods, which control the learning rate of the weight updates (Duchi et al., 2011; Tieleman et al., 2012; Zeiler, 2012; Kingma and Ba, 2014; Shazeer and Stern, 2018; Anil et al., 2019; Yi et al., 2021; Hayashi et al., 2016). The learning rate determines the magnitude of the weight updates so controlling this hyperparameter is important for the convergence of the network to good local minima (i.e. good solutions). If the learning rate is too high, it may overshoot and oscillate around these minima, failing to converge. On the other hand, if the learning rate is too low, the optimisation may take a long time to converge or may get stuck in bad local minima.

Overall, the backpropagation algorithm is a fundamental tool for training neural networks and has been used to achieve state-of-the-art performance, comparable to humans, on a wide range of tasks, from image recognition to natural language processing.

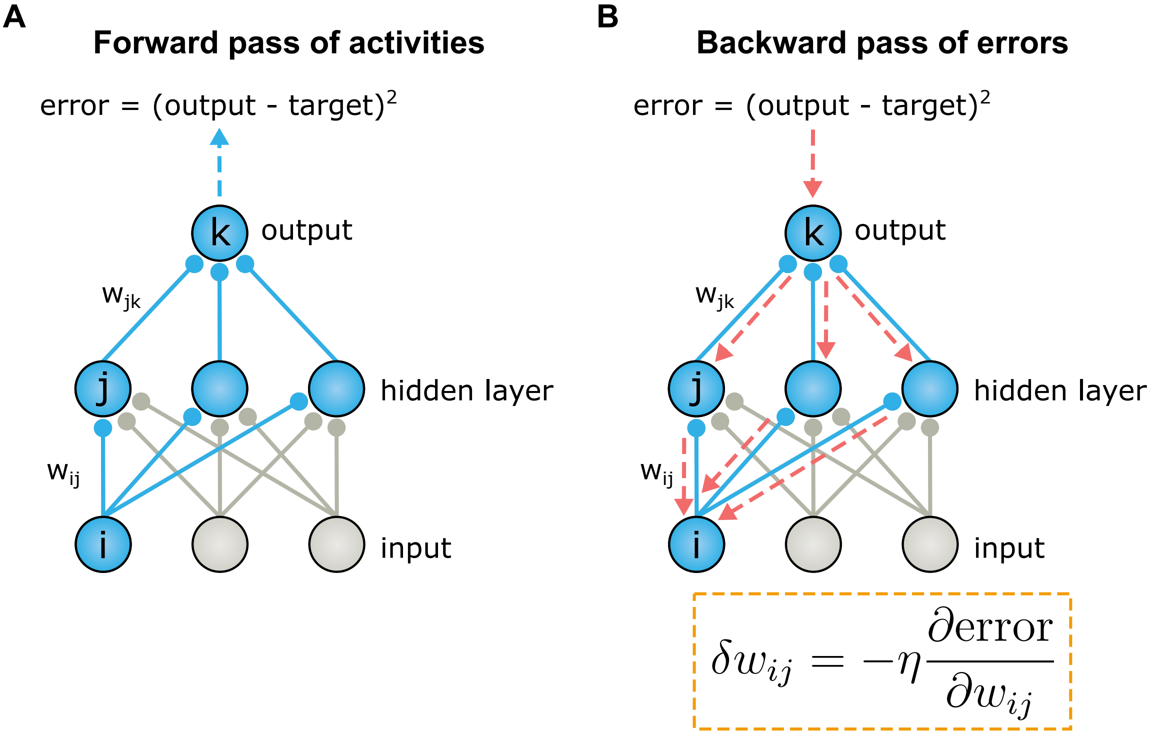
A
Forward pass of activities
error = (output - target)^2
k
output
w_{jk}
j
hidden layer
w_{ij}
i
input
B
Backward pass of errors
error = (output - target)^2
k
output
w_{jk}
j
hidden layer
w_{ij}
i
input
$\delta w_{ij} = -\eta \frac{\partial \mathrm{error}}{\partial w_{ij}}$

FIGURE 1.1: **How the backpropagation algorithm works.** (**A**) Suppose we have a feedforward neural network with a single hidden layer and we want to train it to predict the output of a function given some input. In the forward phase, an input is provided to the network and in this case it activates the neuron i. This activity is propagated forward through the weights of neuron i with the next layer neurons e.g. W_{ij}. This activity continues to be propagated through each layer of the network until it reaches the output layer where the final output is produced. (**B**) In the backward phase, the error between the predicted output and the actual output or target is calculated, and the weights in the network are adjusted to reduce this error. This is done by propagating the error backward through the network, layer by layer, and updating the weights using the calculated error. The update rule (equation in the dashed box) for each weight is based on the chain rule of calculus, which states that the derivative of a composite function is equal to the product of the derivatives of its individual components. The gradient of the error with respect to each weight $\frac{\delta \mathrm{error}}{\delta w_{ij}}$ is calculated using this rule, and the weights are adjusted in the opposite direction of the gradient to minimise the error. The learning rate, η, is used to scale the gradient to control the size of the weight updates to ensure convergence to the global minimum of the loss function.

1.1.2 The credit assignment problem

With frequent exposure to novelty, it is imperative that the brain is capable of rapidly adapting in order to execute appropriate responses to new situations. The brain learns by modifying the strength of the synaptic connections between neurons. However, it remains unclear how the brain assigns accountability, or credit, to synapses to achieve learning and this is referred to as the "credit assignment problem" Minsky (1961). This is a critical problem in both machine learning (ML) and neuroscience. In ML, credit assignment is essential for optimising the performance of

ANNs, while in neuroscience, it is crucial for understanding how the brain learns from experience and adapts to changing environments.

In the context of supervised learning, where the desired output for a given input is known, backprop is a typical solution to the credit assignment problem in neural networks. Gradients computed using this algorithm can be seen as a way of assigning credit or blame to each weight for the error in the output. However, in the context of reinforcement learning, where agents learn from trial and error in an environment, credit assignment is more complex. The agent's actions affect not only the immediate reward but also the future rewards, and the relationship between the agent's actions and the rewards can be highly non-linear and non-trivial. As this encompasses a temporal component, traditional backprop alone, which acts hierarchically in ANNs, is not sufficient to solve the credit assignment problem in reinforcement learning. One approach is to integrate backprop with other methods such as temporal difference learning which uses the difference between predicted and actual rewards to update the value of actions (Sutton, 1988). This can take the form of actor-critic architectures, where an actor network generates actions and a critic network assigns credit to those actions (Konda and Tsitsiklis, 2003).

In neuroscience, the use of backprop has also been studied as a possible mechanism for credit assignment in the brain. Recent research has shown that backprop-trained models produce neuronal responses which are consistent with the activity observed from neurons in the posterior parietal, primary motor and visual cortices (Lillicrap et al., 2020). For example, a study by Güçlü and van Gerven (2015) demonstrated that the internal representations reconstructed from each layer of an ANN trained on over a million natural images with backprop were representative of the increasingly complex features found along the visual ventral stream. Despite outperforming alternative frameworks in accurately reproducing cortical activity patterns and mounting evidence that deep neural networks trained with backprop can help explain neural data, the extent to which the brain uses backprop for credit assignment is still a topic of ongoing research as its direct translation into physiology harbours numerous biological implausibilities (Crick, 1989; Lillicrap et al., 2020; Cadieu et al., 2014).

1.1.3 Biologically plausible backprop algorithms

Recent developments have begun to capture the properties of backprop with more biologically plausible implementations by introducing concepts from neuroscience including dendritic computation, interneurons and bursting (Guerguiev et al., 2017; Sacramento et al., 2018; Payeur et al., 2020). Downstream feedback is a key feature in backprop for credit assignment. The properties of distal apical dendrites of pyramidal neurons have implicated their involvement in credit computations. To begin with, they are a major recipient of feedback projections from higher cortical areas such as the entorhinal cortex (EC) back to the CA1 region of the hippocampus (Fig. 1.2A). In addition, it is well established that they are important for the generation of dendritic spikes, such as plateau potentials, and burst activity; thus, they are critical for updating of synaptic connection weights within the network (Richards et al., 2019; Sjöström and Häusser, 2006; Bittner et al., 2015; Williams and Fletcher, 2019; Pike et al., 1999; Larkum et al., 2009). In pyramidal cells, the coincidence of somatic spiking and synaptic input to the apical dendrite or sufficient depolarisation of the apical dendrite alone can generate dendritic spikes including plateau potentials (Fig. 1.2B) (Larkum and Zhu, 2002; Gambino et al., 2014). These regenerative dendritic events induce burst firing in the pyramidal neuron and studies have demonstrated that apical driven plateau potentials control the formation of CA1 place cells and determine whether plasticity in basal dendrites occurs (Bittner et al., 2015). Given these properties of the apical dendrite, it has been proposed that they integrate higher-order feedback error signals that drive dendritic spikes and bursting to regulate synaptic plasticity. This has resulted in the development of computational models that use the properties of apical dendrites to implement credit assignment for synaptic plasticity (Guerguiev et al., 2017; Sacramento et al., 2018; Payeur et al., 2020, 2021). However, these models still contain biological implausibilities and omit some important factors for plasticity such as neuromodulation. Overall, feedback prediction errors propagated via dendrites could govern the directionality of plasticity with neuromodulation potentially acting as a gating mechanism for the process.

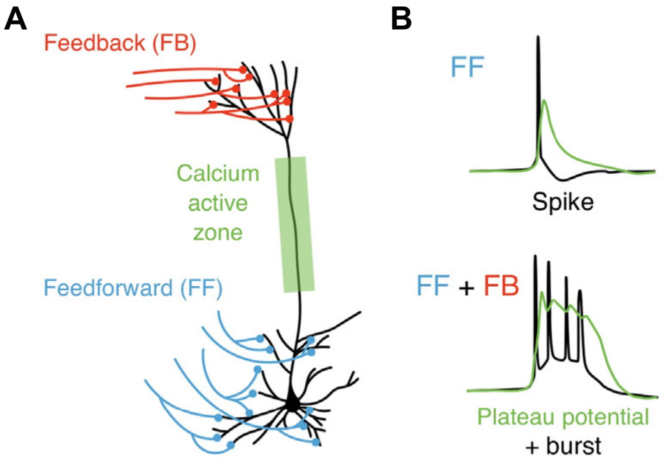

FIGURE 1.2: **The role of distal apical dendrites of pyramidal neurons in credit assignment.** (**A**) Schematic of a representative layer 5 neocortical pyramidal neuron. The basal and distal apical dendrites are thought to be electrotonically separated compartments and generally receive different inputs. Feedfoward (FF) inputs predominantly innervate the basal compartment whereas feedback (FB) inputs mostly target the distal apical compartment. There is an area on the apical dendrite that possesses an abundance of calcium channels (green) which enables the propagation of regenerative dendritic events such as plateau potentials. (**B**; top) FF inputs alone trigger regular spiking activity which can lead to backpropagating action potentials arriving into the apical dendrite (green traces) but this is not enough to elicit a plateau potential. (bottom) However, the coincident activity from FF and FB inputs is strong enough to activate the voltage-dependent calcium channels in the calcium active zone to generate a plateau potential. This plateau potential results in burst firing at the soma and synaptic plasticity occurs. Adapted from Richards and Lillicrap (2019).

1.2 The hippocampus

The hippocampus, named after its resemblance to a seahorse, is a structure located deep in the medial temporal lobes of the brain. Within the hippocampal formation, there exist distinct subregions encompassing the hippocampus and the entorhinal cortex. Each of these subregions plays a role in representing specific types of information, including spatial information (Brickman et al., 2011; Seok and Cheong, 2020). As a result, the interconnected networks within the hippocampus possess the essential elements required for hierarchical information organisation similar to that of ANNs used in machine learning (McKenzie et al., 2014; Collin et al., 2017). The hippocampus has been one of the most thoroughly studied brain regions, particularly in relation to memory function. The famous case study of patient H.M., who suffered from an inability to form new declarative memories following the removal of his hippocampus and neighbouring temporal lobe regions, paved the way for extensive research concerning the hippocampus's involvement in learning and memory. As a result of this research, long-term potentiation (LTP) was discovered in the hippocampus and has become the leading model for the cellular basis of memory (Bliss and Lømo, 1973; Bliss and Collingridge, 1993; Malenka and Bear, 2004). In addition, the discovery of place cells, head direction cells, and grid cells in the rodent hippocampal formation provided a foundation for the idea that the hippocampus is crucial in memory formation by creating a spatiotemporal framework that binds together the various sensory, emotional, and cognitive components of an experience (Moser et al., 2015; O'Keefe, 1976; O'Keefe et al., 1998; Kim and Cho, 2020; Sweatt, 2004). This framework enables the storage and retrieval of this experience as a memory.

This section will review the architecture of the hippocampus, the function of the hippocampus in learning and memory, synaptic plasticity in the hippocampus and the importance of dendritic processing for hippocampal synaptic plasticity.

1.2.1 Architecture and circuits

The general architecture of the hippocampal anatomy remains remains fairly consistent across species (Clark and Squire, 2013). In rodents, the hippocampus is situated just below the neocortex and has a shape similar to that of a cashew (Fig. 1.3). It is primarily composed of three major regions: the dentate gyrus (DG), the Cornu Ammonis (CA)3 region, and the CA1 region. The hippocampal trisynaptic loop is the main pathway of information flow through the hippocampus: axons from layer II of the entorhinal cortex (EC) transmits sensory information to the DG via the perforant pathway (PP); the mossy fibre (MF) projections from DG granule cells relay this information to the CA3 which propagates this forward to the CA1 via the Schaffer collateral (SC) projections; CA1 pyramidal neurons then send projections back to layers IV and V of the EC via the subiculum (Fig. 1.3). CA3 pyramidal neurons also have recurrent connectivity to other ipsilateral and contralateral CA3 pyramidal neurons and these projections are known as the association and commissural pathways, respectively. Additionally, CA1 and subiculum receive direction projections from layer III of the EC which is known as the temporammonic (TA) pathway. These different projection pathways are found in particular layers or strata of the hippocampus and each layer is composed of cells with distinct morphologies and functions. The hippocampus has a laminar organisation consisting of the stratum oriens (SO), stratum pyramidale (SP), stratum radiatum (SR), stratum lacunosum-moleculare (SLM), stratum lucidum (only present in CA3) and stratum granulosum (only present in the DG). The soma of hippocampal pyramidal cells are found in the SR: the basal and apical dendrites of these cells extend into the SO and SLM, respectively. The SC pathway resides in the SR and predominantly innervates the proximal apical dendrite of CA1 pyramidal cells, whereas the TA pathway projects through the SLM and innervates the distal apical dendrite of CA1 pyramidal cells.

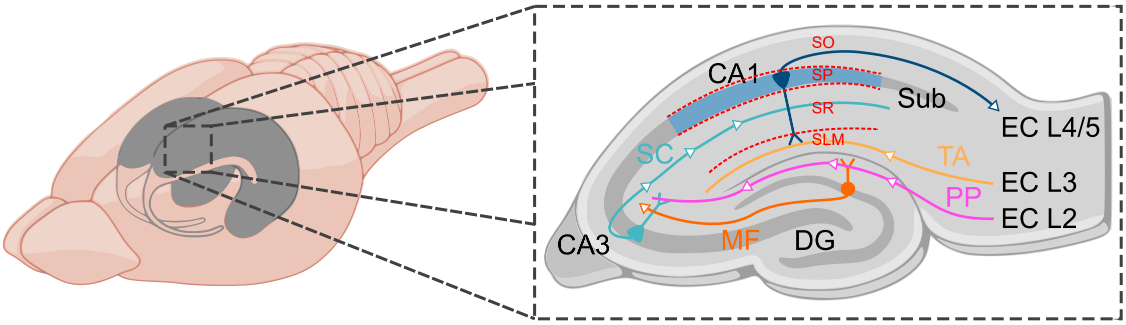

FIGURE 1.3: **Hippocampal anatomy and circuitry.** Schematic of a rodent brain displaying the location of hippocampi (grey) directly beneath the neocortex. Inset represents a transverse section of the hippocampus and highlights the classic hippocampal trisynaptic pathway with arrows to indicate the direction of information flow. DG, dentate gyrus; MF, mossy fibre pathway; SC, Schaffer collateral pathway; TA, temporammonic pathway; PP, perforant pathway; EC, entorhinal cortex; L, layer; Sub, subiculum; SO, stratum oriens; SP, stratum pyramidale; SR, stratum radiatum; SLM, stratum moleculare-lacunosum.

Although only accounting for approximately 10-15% of the total neurons in the hippocampus, GABAergic inhibitory interneurons are a diverse neuronal population that plays an extremely important role in regulating cellular and circuit function (Pelkey et al., 2017). In CA1, some of the most thoroughly researched interneuron types include the parvalbumin-expressing (PV), cholecystokinin-expressing (CCK) and somatostatin-expressing (SST) cells. These cell types are anatomically and morphologically different from each other and their axons target distinct postsynaptic domains allowing for a diverse range of modulatory effects. PV and CCK basket cells are typically located around the SP and preferentially target the perisomatic region of CA1 pyramidal cells. PV-expressing cells demonstrate fast-spiking behaviour without frequency adaptation, while CCK basket cells exhibit notable spike frequency adaptation. In addition, unlike PV expression, CCK expression is also found in several types of dendrite-targeting inhibitory interneurons (Fig. 1.4). The state of PV and CCK feedforward inhibition has been demonstrated to be important in the regulation of information flow into the CA1 circuit (Udakis et al., 2020; Palacios-Filardo et al., 2021; Pouille and Scanziani, 2001; Pouille et al., 2013). Another interneuron population that comprises of dendrite-targeting interneurons are the SST cells. The SST molecular marker is detected in a variety of interneurons but the Oriens Lacunosum-Moleculare (OLM) cells are the most extensively studied. The cell bodies and dendrites of these cells reside in the SO and they have dense axonal arborisation in the SLM. OLM cells are primarily excited by local pyramidal cells and provide inhibitory feedback to these cells (Honoré et al., 2021). Additionally, they receive input from cholinergic afferents that originate from the septum and diagonal band of Broca (Leão et al., 2012; Lovett-Barron et al., 2014). Local inhibitory interneurons, particularly those expressing vasoactive intestinal peptide (VIP), also inhibit OLM cells (Tyan et al., 2014). OLM cells predominantly target the distal apical dendrite of CA1 pyramidal cells but they also target other interneuron types including basket cells, bistratified cells and cells located in the SR that innervate the proximal apical dendrite. This allows the OLM cells to differentially modulate the excitatory afferents from the CA3 and EC onto CA1 pyramidal cells by disinhibiting and inhibiting excitatory inputs from the SC and TA pathways, respectively. It has been demonstrated that the capacity of this interneuron subtype to control the excitatory inputs into CA1 is crucial for synaptic plasticity in CA1 and their activity modulates encoding of object and fear-related memories (Leão et al., 2012; Siwani et al., 2018; Lovett-Barron et al., 2014). These dendrite-targeting interneurons have also been incorporated into a plausible model of backprop-like credit assignment as a critical component for the generation of error signals that drive synaptic plasticity (Sacramento et al., 2018). Furthermore, incorporating the properties of interneurons into such models has addressed key biological implausibilities associated with directly translating backprop into biological systems. These include challenges related to how the derivative of activation functions are computed, two-phase learning, the need for a distinct error network and non-local learning rules (Greedy et al., 2022; Sacramento et al., 2018; Payeur et al., 2021; Crick, 1989).

The hippocampus also receives a variety of neuromodulatory inputs from different brain regions, which play important roles in learning and memory. Some of the major neuromodulators released in the hippocampus include: acetylcholine released primarily from cholinergic projections of the medial septum, dopamine released from dopaminergic projections of the ventral tegmental area, serotonin released from serotonergic projections of the raphe nuclei and noradrenaline released from noradrenergic projections of the locus coeruleus (Palacios-Filardo and Mellor, 2019; Ruivo et al., 2017; Sara, 2009; Rosen et al., 2015; Ohmura et al., 2010). These neuromodulators act on various receptors located on different types of neurons in the hippocampus, including pyramidal neurons, interneurons, and glial cells, and can modulate synaptic transmission and plasticity. One major focus of this thesis is to investigate the effects of acetylcholine on OLM interneurons and synaptic plasticity, with a particular emphasis on using a model that has similarities to the hippocampal microcircuit architecture (Sacramento et al., 2018).

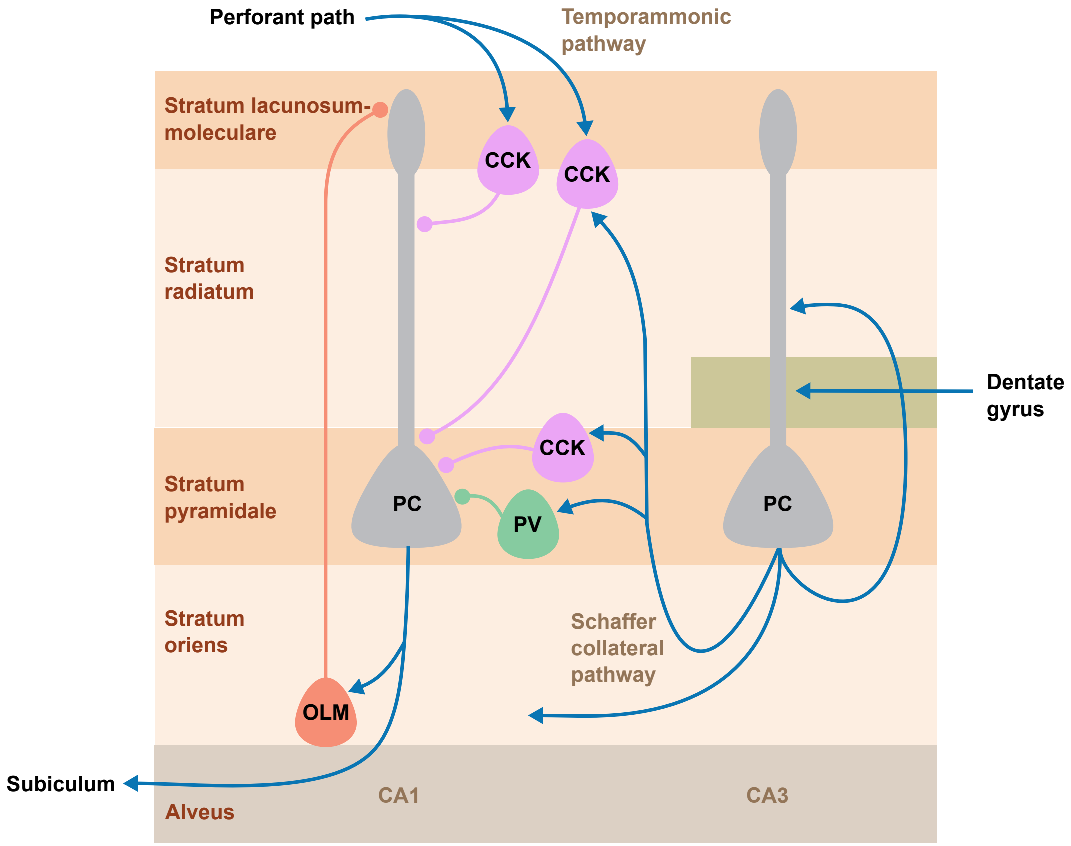
Perforant path
Temporammonic pathway
Stratum lacunosum-moleculare
CCK
CCK
Stratum radiatum
Dentate gyrus
CCK
Stratum pyramidale
PC
PV
PC
Stratum oriens
Schaffer collateral pathway
OLM
Subiculum
CA1
CA3
Alveus

FIGURE 1.4: **Circuit diagram illustrating some of the interneuron networks in the hippocampus.** Fast-spiking parvalbumin (PV; green) and regular-spiking cholecystokinin (CCK; purple) interneurons project perisomatically onto CA1 pyramidal cells (PCs; grey) and mediate feedforward inhibition. Dendrite-targeting CCK interneurons located at the border of the stratum lacunosum-moleculare (SLM) and stratum radiatum are innervated by the perforant path fibres of the entorhinal cortex and project onto the dendrite of CA1 PCs. Oriens Lacunosum-Moleculare (OLM; orange) interneurons receive inputs from CA1 PCs providing feedback inhibition onto their distal dendrites in the SLM.

1.2.2 Synaptic transmission and plasticity

Synaptic transmission is the process through which neurons communicate with one another by releasing and receiving chemical and electrical signals at specialised structures known as synapses. Synapses are found between the axons of one neuron and the dendrites or cell bodies of another neuron. The majority of synapses are chemical, and they interact using chemical messengers known as neurotransmitters. The release of neurotransmitters from the presynaptic neuron acts to excite or inhibit the postsynaptic neuron by binding onto distinct receptors. Excitatory and inhibitory synaptic transmission is mediated by glutamate and γ-aminobutyric acid (GABA) release, respectively. Both glutamate and GABA bind onto ionotropic and metabotropic receptors on the postsynaptic membrane. Activation of ionotropic glutamate receptors, such as AMPA, NMDA and kainate receptors, results in the influx of positively charged ions, such as sodium and calcium, into the postsynaptic cell. This influx of ions depolarises the postsynaptic membrane and generates an excitatory postsynaptic potential (EPSP), which can trigger an action potential if it reaches the threshold for firing. On the other hand, activation of GABA receptors, $GABA_A$ and $GABA_B$, leads to the influx of negatively charged ions, such as chloride, into the postsynaptic cell. This influx of ions hyperpolarises the postsynaptic membrane and generates an inhibitory postsynaptic potential (IPSP), reducing the probability of the postsynaptic neuron firing an action potential. Metabotropic glutamate receptors are G protein-coupled receptors that modulate neuronal excitability through intracellular signalling pathways. Therefore, the postsynaptic effect is dependent on the specific signalling pathway that is activated.

The strength and efficacy of synapses can be modified through a process called synaptic plasticity, which is thought to underlie many forms of learning and memory. Distinct patterns of neural activity induce different forms of synaptic plasticity, which are governed by diverse molecular mechanisms (Citri and Malenka, 2008; Basu and Siegelbaum, 2015; Bliss et al., 2014; Nicoll, 2017). The discovery of LTP provided the first piece of evidence to support the idea that synapses are plastic in the mammalian brain (Bliss and Lømo, 1973). LTP is the most extensively researched form of long-term synaptic plasticity and remains one of the most appealing cellular models for learning and memory. It has been observed *in vitro* and *in vivo* at synapses across the brain (Nicoll, 2017). LTP exists in multiple forms but perhaps the most well-known is the homosynaptic Hebbian LTP which was first discovered by Bliss and Lømo (1973) who demonstrated, in anaesthetised rabbits, that a brief 100 Hz tetanic stimulation of the PP inputs to DG can cause an long-lasting increase in synaptic transmission at the PP-DG synapses. Typically, LTP induced in hippocampal synapses is via a postsynaptic mechanism and dependent on the NMDA receptor (NMDAR-LTP). The exception is LTP at mossy fibres synapses in the CA3 as this is believed to be NMDAR-independent and expressed presynaptically. The events that lead to NMDAR-LTP are as follows: repeated rapid glutamate release, as seen during a high-frequency tetanus, from the presynaptic terminal binds onto postsynaptic AMPA and NMDA receptors; initially, the influx of Na^+ and K^+ ions occurs through the AMPA receptor channel, while the NMDA receptor channel is blocked by Mg^{2+} ions; the influx of cations depolarises the postsynaptic neuron causing the release of the voltage-dependent Mg^{2+} block, allowing the entry of Ca^{2+} through the NMDA receptor; Ca^{2+} influx triggers a cascade of signalling events that result in the activation of various kinases and phosphatases, including calcium/calmodulin-dependent protein kinase II (CaMKII) and protein phosphatase 1 (PP1), which upregulate AMPA receptor trafficking; the insertion of new AMPA receptors or the increase in the conductance of existing receptors leads to an increase in the amplitude of postsynaptic currents (Nicoll, 2017). Synaptic plasticity is bidirectional: low frequency stimulation (LFS) can induce long-term depression (LTD) or depotentiate LTP via the activation of NMDARs and group I metabotropic glutamate receptors which trigger downstream signalling pathways that reduce the probability of presynaptic glutamate release and cause the endocytosis of AMPA receptors (Collingridge et al., 2010).

As *in vivo* recordings became more common, it became clear that the conventional high frequency tetanic induction protocol did not accurately reflect the normal physiological firing range of hippocampal pyramidal neurons (Bliss and Collingridge, 1993; Rose and Dunwiddie, 1986). Single unit recordings of CA1 pyramidal cells in awake, freely moving rates indicate that the firing rate of these cells range from 5-20 Hz (Waters and Smith, 2002). In addition, these cells exhibit spontaneous complex burst spiking activity which consists of high frequency trains of typically four to five action potentials particularly during behavioural states important for learning and memory (Grienberger et al., 2014). These observations led to the development of more physiologically-relevant induction protocols including the theta-burst and primed-burst stimulation protocols (Larson et al., 1986; Rose and Dunwiddie, 1986). The theta-burst stimulation paradigm (TBS) encapsulates the bursting activity seen in the hippocampus and the rhythmic hippocampal theta oscillations from the synchronous firing of hippocampal neurons (Larson and Munkácsy, 2015). *In vitro*, this stimulation pattern of burst spiking at a repeated theta frequency produces robust LTP, as effective as tetanic stimulation, in CA1 pyramidal cells.

Another more biologically plausible induction protocol is the Spike Timing Dependent Plasticity (STDP) protocol. This is an asymmetric form of Hebbian learning that is induced by precise temporal correlations between the spikes of pre- and postsynaptic neurons. LTP is induced when a presynaptic spike occurs just before a postsynaptic spike within a specific time frame, while the opposite sequence leads to LTD (Bi and Poo, 2001; Shouval et al., 2010). Similar to STDP, associative LTP (aLTP) requires the precise timing of coordinated inputs such as those from the SC and TA pathways to the CA1 region (Jarsky et al., 2005; Dudman et al., 2007; Takahashi and Magee, 2009). This coordinated activity induces large, regenerative dendritic events known as plateau potentials that cause bursting in the soma and ultimately LTP (Bittner et al., 2015). aLTP can be heterosynaptic as demonstrated by Dudman et al. (2007): simultaneous stimulation of the SC and TA pathways only produced LTP at the SC synapses.

LTP continues to be an attractive phenomenon for explaining the cellular processes underlying learning and memory. Pharmacological studies infusing APV, an NMDA receptor antagonist, into the hippocampus demonstrates that disrupting the mechanism of LTP results in impaired learning (Morris et al., 1986). The impairment on learning and memory is further highlighted in a study by Giese et al. (1998) which shows that knockin mice containing the point mutant T286A of CaMKII prevents LTP as the downstream CaMKII signalling pathways involved in LTP are disrupted. Furthermore, LTP has been correlated with fear conditioning and inhibitory avoidance learning, and saturating LTP with repeated tetanic stimulation interferes with memory (Moser et al., 1998; Whitlock et al., 2006; Rogan et al., 1997). Martinez et al. (2002) have also demonstrated *in vivo* in anaesthetised rats that aLTP can occur in hippocampal CA3. More specifically, they found that subthreshold LTP stimulation to either the commissural/associational CA3 projections or perforant pathway projections could trigger aLTP at their respective CA3 synapses if this stimulation was delivered in conjunction with high-intensity activation of the other pathway. Together, these data show that LTP can occur *in vivo* and suggest that it shares some of the mechanisms that have been identified in *ex vivo* experiments. This supports the use of induction protocols used in *ex vivo* experiments as reliable models for studying LTP *in vivo*.

1.2.3 Dendritic integration

Dendritic integration is a process in which synaptic inputs from multiple sources are integrated within the dendrites of a neuron. It is an essential part of neural processing and allows neurons to perform complex computations. Dendrites were initially believed to be passive cables that simply receive and transmit information to the soma (Yuste and Tank, 1996). However, work conducted by Wong et al. (1979), which involved intradendritic recordings in hippocampal pyramidal neurons, revealed the occurrence of regenerative events in dendrites following intracellular current injections and synaptic stimulation. Dendritic spikes, which allow for non-linear integration of inputs, enable neurons to enhance and tune their response to specific inputs (Burger et al., 2023; Augusto and Gambino, 2019; Larkum et al., 2022). They have the potential to alter the relationship between input and output in neurons, thereby enabling inputs to the distal dendrites to exert greater influence on the initiation of action potentials at the soma (Branco and Häusser, 2011). There are several types of dendritic spikes that are mediated by different types of voltage-gated ion channels, such as sodium or calcium, or are dependent on ligand-gated ion channels, particularly NMDARs (Fig. 1.5) (Larkum et al., 2022; Schiller et al., 2000; Antic et al., 2010; Larkum et al., 2009). Unlike calcium and sodium dendritic spikes, NMDAR-dependent spikes (as well as plateau potentials) cannot be triggered solely by membrane depolarization, but rather also require the presence of glutamate and glycine. Suprathreshold NMDAR-dependent spikes/plateaus in layer V pyramidal cells do not increase in amplitude with increasing stimulus intensity, but instead exhibit nearly linear growth in duration, which could potentially change the time window for integration. In addition, depolarisation along the same dendrite can reduce the glutamate threshold for NMDA spikes which enables interactions between distant synapses (Major et al., 2013). It has been shown in *ex vivo* hippocampal slice experiments that the coincident inputs of the SC and TA pathways onto the proximal and distal regions of CA1 pyramidal cells, respectively, facilitate the generation of dendritic spikes (Jarsky et al., 2005; Dudman et al., 2007; Takahashi and Magee, 2009). The temporal correlation of the two inputs determines the generation of these dendritic events which induce LTP. This form of aLTP also drives the formation and stabilisation of place fields (Bittner et al., 2015, 2017). Furthermore, a study conducted by Sheffield et al. (2017) using virtual reality and two-photon imaging in mice found that regenerative dendritic events occur during the formation of place fields. The prevalence of these events was higher across the dendritic arbour when the animal traversed novel environments compared to familiar environments. They also found that exposure to a novel environment is accompanied by a short time window of reduced dendritic inhibition in CA1, facilitating branch spiking and creating a window of plasticity for place field formation. This finding, along with a study by Kim et al. (1995), which used intradendritic recordings from neocortical pyramidal neurons to demonstrate that IPSPs can diminish and ultimately block dendritic sodium and calcium spikes elicited by current injections, provides evidence that inhibition plays a crucial role in regulating dendritic responses.

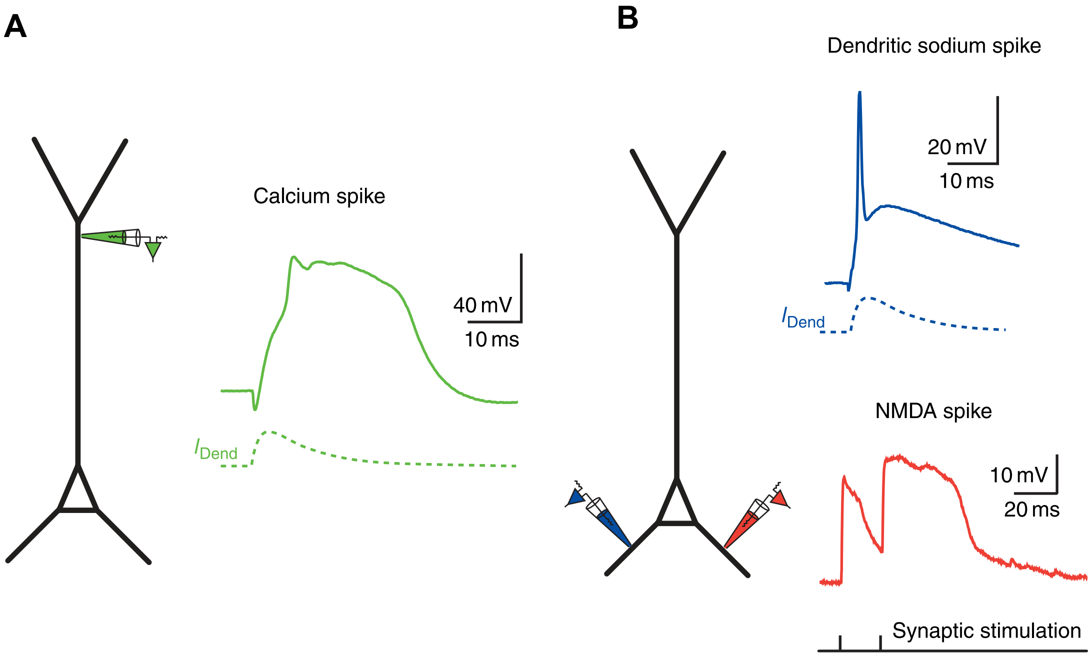

FIGURE 1.5: **Types of dendritic regenerative events.** Schematic of the different dendritic spikes observed in neocortical pyramidal neurons. Electrodes at the distal (green) and basal (red and blue) dendrites indicate the sites of membrane voltage recordings. (**A**) Local dendritic depolarisation of the apical dendrite can produce calcium spikes (green waveform) mediated by voltage-gated Ca^{2+} channels. (**B**) Dendritic sodium spikes (blue waveform) experience significant attenuation as they propagate towards the soma and can be observed in both apical and basal dendrites. Strong synaptic activation can lead to the generation of N-methyl-D-aspartate (NMDA) spikes (green waveform). Adapted from Waters et al. (2008).

1.2.4 Spatial navigation

The importance of the hippocampus in episodic memory is well-established but it serves a much more complex function beyond being a mere memory storage unit. The discovery of place cells, neurons that preferentially fire at specific locations in an environment known as "place fields", has brought great attention to the hippocampus' role in spatial memory (Fig. 1.6A-B) (O'Keefe and Dostrovsky, 1971; O'Keefe, 1976). Place cells do not exhibit topography, indicating that two adjacent cells may not have neighbouring place fields (O'Keefe et al., 1998). Within the place field, place cells fire spikes in bursts at a high frequency, while they remain mostly inactive outside of the place field (Bures et al., 1997). The place fields are allocentric in nature, which implies that they are defined with respect to the outside world (world-centred) instead of the body (egocentric) (O'Keefe and Burgess, 1996; Bures et al., 1997). By representing space based on the environment instead of the individual, place fields can function effectively as neural maps of the environment (Jeffery et al., 2004). Subsequent lesion studies have supported the hypothesis that the hippocampus implements cognitive maps (O'Keefe and Nadel, 1979). For example, hippocampal lesions in rats during the Morris water maze and T maze alternation tasks revealed that hippocampal damage results in deficits to spatial navigation and memory (Morris et al., 1982; Cohen et al., 1971). It has also been shown that place cells respond to non-spatial characteristics of the environment, such as odours, tactile inputs and timing, which suggests that the hippocampus encodes both spatial and non-spatial information to form these cognitive maps (Moser et al., 2015).

Variations in firing rates of spatially stable place cells have been hypothesised to represent non-spatial components of an experience. This phenomenon is known as "rate remapping," which involves no changes in the population of active cells or the location of their respective place fields, but instead significant changes in firing rates (Fig. 1.6C) (Leutgeb et al., 2005). These rate fluctuations can communicate features of experiences that fluctuate from one trial to the next or between specific segments of a trial. They can also reflect minor differences in sensory or motivational stimuli on a second-to-second basis within the same trial. The magnitude of rate changes may differ across conditions, possibly depending on the degree of dissimilarity between the conditions. If there are major changes between conditions, such as a complete novel environment, "global remapping" occurs. Unlike rate remapping, global remapping is an all-or-none phenomenon that results in changes to both the firing rate and place fields of place cells (Fig. 1.6D) (Colgin et al., 2008). However, it remains unclear what the underlying mechanisms are for initiating remapping. Experimental studies have demonstrated that the presence of plateau potentials is important for the formation and stabilisation of place cells (Bittner et al., 2015, 2017). Computational modeling has also suggested that the interaction between somatic and dendritic inhibition facilitates these processes (Pedrosa and Clopath, 2020). Furthermore, studies have shown that noradrenergic inputs from the locus coeruleus may induce place cell reorganisation, and that lesions to the cholinergic projections into the hippocampus and entorhinal cortex can significantly reduce the probability of remapping (Kaufman et al., 2020; Newman et al., 2012). Therefore, it is possible that neuromodulation could control remapping through their regulatory effects on local inhibition in the CA1 circuit.

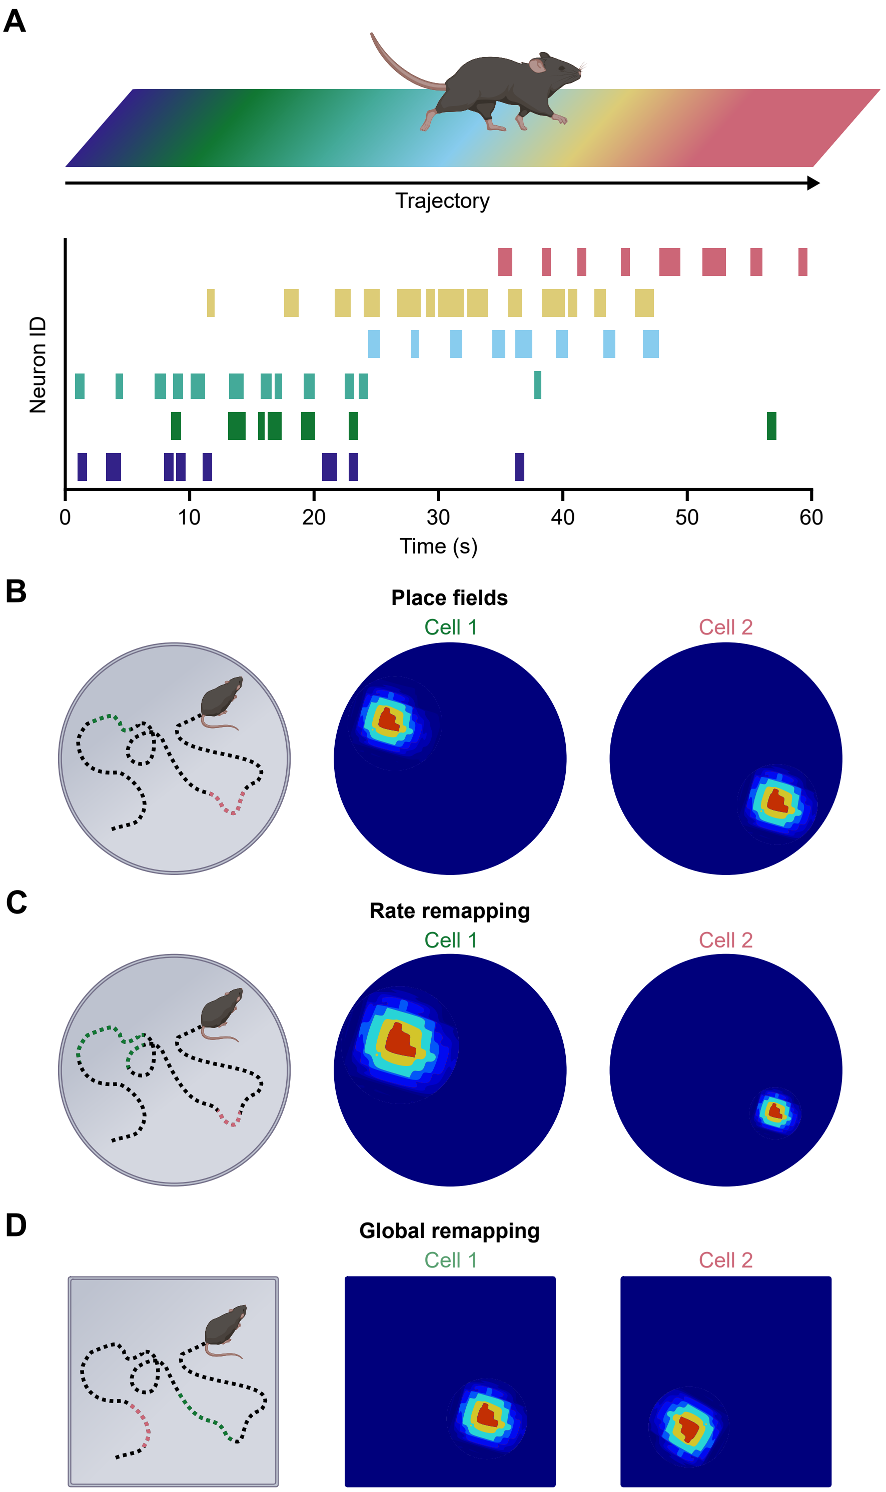
A
Trajectory
Neuron ID
0
10
20
30
40
50
60
Time (s)
B
Place fields
Cell 1
Cell 2
C
Rate remapping
Cell 1
Cell 2
D
Global remapping
Cell 1
Cell 2

FIGURE 1.6: **Place cells and remapping.** (**A**; top) As the mouse traverses along a linear track environment, (bottom) place cells are activated - each represented by a different colour associated with its position of firing on the track. (**B**; left) Trajectory of mouse movement in circular arena with green and pink dots representing spike firing of cells 1 and 2, respectively. Colour-coded rate maps showing place field of (middle) cell 1 and (right) cell 2 (dark blue and red represent 0 Hz and peak firing rates, respectively).

(**C**) Same as (**B**) but slight changes to the environment induces rate remapping of place fields: place fields remain unchanged but there are variations in the firing rates. (**D**) Mouse is placed in a novel environment, triggering global remapping which involves changes in both firing rates and place fields.

1.3 Acetylcholine

Neuromodulators are a class of signalling molecules in the central nervous system that can modulate the activity of neurons and neural circuits. Under particular behavioural states, such as novelty and fear, these neurotransmitters are released and can influence a variety of cellular processes, including synaptic transmission and plasticity, over extended periods of time beyond what is typical for classical neurotransmitters. Neuromodulation of neuronal activity is accomplished by the coordinated actions of several central modulators, including acetylcholine, serotonin, dopamine, norepinephrine, and various neuropeptides. The emphasis of this section will be on the function of acetylcholine in memory and learning, specifically within the hippocampus.

1.3.1 Acetylcholine in the hippocampus

1.3.1.1 Septohippocampal pathway

The primary source of cholinergic input to the hippocampus is through the septohippocampal pathway (Fig. 1.7). This pathway encompasses the medial and lateral septal nuclei as well as the nucleus of the diagonal band of Broca (DBB). The medial septum is reciprocally connected to the hippocampus via the fimbria-fornix pathway and provides innervation to all regions of the hippocampus, including the stratum oriens of CA1 and CA3. Immunohistochemical analysis has revealed that septal cholinergic, GABAergic, and glutamatergic neurons project to the hippocampus. These projections target both pyramidal cells and GABAergic interneurons (Teles-Grilo Ruivo and Mellor, 2013). While the structure of the cholinergic projection from the septum to the hippocampus is well understood, there is limited knowledge regarding the functional output of these neurons. It remains unclear whether the septohippocampal cholinergic system should be considered as a cohesive unit that diffusely innervates the entire hippocampus or if cholinergic neurons act independently to target specific cell types in distinct hippocampal regions. Recent studies have provided evidence to suggest that cholinergic signalling is heterogeneous and dependent on behavioural states (Lohani et al., 2020; Záborszky et al., 2018).

1.3.1.2 Acetylcholine receptors

When ACh is released into the hippocampus, it binds to and activates both nicotinic and muscarinic acetylcholine receptors (nAChRs and mAChRs, respectively). These receptors are classified into subtypes that produce diverse effects and are differentially expressed in various cell types throughout the hippocampus (Dannenberg et al., 2017). nAChRs are ionotropic receptors that

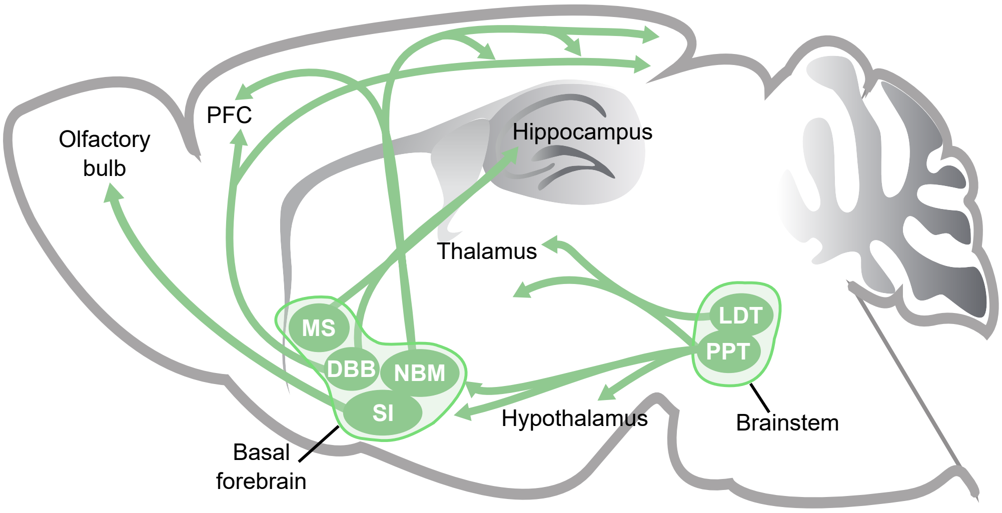

FIGURE 1.7: **Illustrative overview of cholinergic pathways in the rodent brain.** The basal forebrain and the brainstem are the two major sources of cholinergic projections. The basal forebrain cholinergic system consists of the medial septum (MS), vertical limbs of the diagonal band of Broca (DBB), the nucleus basalis of Meynert (NBM) and the substantia innominata (SI). The brainstem cholinergic system consists of the laterodorsal tegmental (LDT) and pedunculopontine tegmental (PPT) nuclei. These cholinergic sources innervate various cortical and subcortical brain regions as indicated by the arrows.

mediate the flow of cations. They consist of five protein subunits and are formed from the homomeric or heteromeric combinations of α2-10 and β2-4 subunits. The $\alpha 7$, $\alpha 4\beta 2$, and $\alpha 3\beta 4$ subtypes of nAChRs are mainly expressed in the hippocampus (Alkondon and Albuquerque, 2004). In CA1, pyramidal neurons and interneurons typically express $\alpha 7$ and $\alpha 4\beta 2$ nAChRs. Pharmacological studies using nAChR agonists have demonstrated that activation of nAChRs mediates cell excitability: the presynaptic probability of neurotransmitter release and postsynaptic fast excitatory synaptic transmission (Gray et al., 1996; Albuquerque et al., 2009; Dannenberg et al., 2017). These receptors are highly sensitive to desensitisation (Haam and Yakel, 2017). Upon ACh-induced depolarisation of inhibitory interneurons in the stratum radiatum, the desensitisation of these receptors can decrease inhibition, providing a mechanism for bidirectional control of neuronal excitability by indirectly exciting pyramidal neurons (Frazier et al., 1998). Furthermore, cholinergic inputs from the medial septum directly stimulate CA1 OLM interneurons via nAChRs. These interneurons target the distal apical dendrites of CA1 pyramidal neurons in the SLM. By activating these interneurons, cholinergic signaling could regulate the flow of information into CA1 by altering the relative weights of the SC and TA inputs (Leão et al., 2012).

mAChRs are metabotropic receptors with five distinct subtypes M_1-M_5. M_1, M_3 and M_5 receptors are coupled to G_q proteins. G_q protein activation triggers the activation of phospholipase C which raises inositol trisphosphate and diacylglycerol levels, causing calcium influx and increased neuronal excitability. In contrast, M_2 and M_4 receptors are coupled to $G_{i/o}$ proteins. Activation of $G_{i/o}$ proteins inhibits the activation of cyclic-adenosine monophosphate, leading to downstream signalling pathways that result in reduced cell activity (Eglen, 2006). The hippocampus predominantly expresses the M_1, M_2 and M_4 subtypes of muscarinic receptors (Dannenberg et al., 2017). However, the expression of the M_2 subtype is limited to interneurons and is absent in pyramidal cells (Levey et al., 1995).

M1 receptors have a vital role in the modulation of pyramidal cell excitability in response to transient ACh application in slices (Gulledge and Kawaguchi, 2007). Previous pharmacological studies using intracellular recordings in guinea pig hippocampal slices revealed that the application of ACh to the perisomatic region of CA1 pyramidal neurons resulted in a slow and prolonged depolarisation and an increase in spike frequency (Benardo and Prince, 1982). This depolarisation occurred due to the blockade of voltage-dependent K conductance. Other studies performed on rat hippocampal slices have also shown that pyramidal cells experience this slow depolarising effect but via cholinergic modulatory suppression of a leak potassium current (Cole and Nicoll, 1984). Local transient ACh application at the soma of CA1 pyramidal cells results in membrane hyperpolarisation and inhibition of action potential generation whereas tonic application of carbachol, an AChR agonist, depolarises these cells via the activation of M_1 and M_3 receptors (Dasari and Gulledge, 2011). Basket cells and OLM interneurons exhibit increased excitability upon activation of M1 and M3 receptors (Lawrence et al., 2006; Cea-del Rio et al., 2010). However, this increase in excitability does not necessarily result in increased inhibitory transmission (Pitler and Alger, 1992). A study by Szabó et al. (2010) demonstrated that activation of presynaptic M2 receptors can reduce the amplitude of IPSCs produced by PV interneurons onto CA3 pyramidal cells. Additionally, postsynaptic M1 and M_3 receptors can trigger the synthesis of endocannabinoids in pyramidal cells, which can reduce GABA release from CCK basket cells via type I cannabinoid receptor activation.

Studies have demonstrated that the cholinergic suppression of synaptic transmission in the SC pathway is a result of presynaptic M4 receptor activation (Dasari and Gulledge, 2011; Palacios-Filardo et al., 2021). Recently, it has also been shown that presynaptic M3 receptor activation reduces excitatory and feedforward inhibitory transmission in the temporammonic pathway (Palacios-Filardo et al., 2021). Together, these data suggests that cholinergic modulation is capable of gating information flow into CA1 via its heterogeneous effects on different cell types and pathways.

1.3.2 Learning and memory

It is well-established that ACh plays a critical role in learning and memory. Acetylcholinesterase inhibitors such as Donepezil, Rivastigmine and Galantamine are commonly used drugs to improve memory function in Alzheimer's disease patients by inhibiting the degradation of ACh (Haam and Yakel, 2017). Selective lesion studies in monkeys and rodents using Ig-saporin have demonstrated that the loss of cholinergic modulation leads to severe impairments in working and visual recognition memory (Turchi et al., 2005; McGaughy et al., 2005). Additionally, electrolytic lesions in the medial septum and selective cholinergic lesions in the hippocampus and medial septum result in impairments in spatial memory (Pang and Nocera, 1999; Winson, 1978; Cai et al., 2012). Lesions to septohippocampal projections in rats impairs novel place field formation and remapping, and causes place fields to become less stable (Shapiro et al., 1989; Ikonen et al., 2002).

During spatial learning and object exploration, ACh levels increase in the brain (Dannenberg et al., 2017). When scopolamine, an AChR antagonist, is infused into the hippocampus and medial septum, animal performance on spatial memory tasks is impaired (Blokland et al., 1992; Elvander et al., 2004). Intracerebroventricular injection of scopolamine also degrades hippocampal place cell firing (Brazhnik et al., 2003). The concentration of ACh is important as carbachol impairs memory when infused into the medial septum (Elvander et al., 2004). Additionally, the timing of ACh release is critical. For instance, during a reward-based navigation task in mice, optogenetic stimulation of cholinergic neurons impaired memory formation when activated at the goal location, but not during navigation (Jarzebowski et al., 2021). In summary, these findings indicate that the temporal dynamics of ACh release in the brain is crucial for the formation of place cells and spatial memory.

1.3.3 Cholinergic modulation of synaptic plasticity

Cholinergic receptor activation not only alters the intrinsic properties of neurons, but it also influences synaptic plasticity. Depending on the spatiotemporal activation of nAChR, LTP is either induced or inhibited: nAChR-mediated activation of presynaptic glutamatergic synapses induces LTP but nAChR-mediated interneuron activation can inhibit this (Ge and Dani, 2005; Ji et al., 2001). In addition, the activation of M_1 receptors has been shown to facilitate LTP induction via a number of mechanisms: M_1 receptor-mediated inhibition of small conductance calcium-activated potassium channels, voltage-activated Kv7 and Kv4.2 potassium channels resulting in the enhancement of NMDAR activity, bursting activity and dendritic excitability, respectively (Buchanan et al., 2010; Petrovic et al., 2012; Losonczy et al., 2008). This muscarinic receptor dependent control of LTP has also been shown *in vivo* by Ovsepian et al. (2004). They demonstrate that subthreshold LTP stimulation of CA1 basal dendritic synapses paired with medial septum stimulation is sufficient to induce LTP and this form of LTP induction is inhibited by muscarinic receptor antagonists. Furthermore, optogenetically-evoked release of ACh has been shown to independently increase dendritic excitability leading to the generation of plateau potentials and bursting activity which provides another mechanism for inducing LTP (Williams and Fletcher, 2019). The effect of ACh on synaptic plasticity is bidirectional as ACh has been shown to depress SC synapses and muscarinic activation receptor can switch LTD to LTP, suppress both LTP and LTD and potentiate LTP (Brzosko et al., 2017; Sugisaki et al., 2011). Taken together, this suggests that the spatiotemporal activation of cholinergic receptors determines the direction of synaptic plasticity.

1.3.4 Theories of acetylcholine and cognitive processing

While there is extensive knowledge about how acetylcholine physiologically affects cellular and network dynamics, there is yet a unified theory on how these properties contribute to cognitive function. Computational modelling has been used to integrate these seemingly disconnected areas of knowledge in an attempt to gain a better understanding of how certain aspects of cholinergic modulation may play a role in cognitive processing. Since the environment is constantly changing, animals must develop the ability to effectively acquire new information, assimilate feedback and adapt their behaviour in order to survive. Zannone et al. (2018) proposed that this behavioural flexibility could be achieved through the interplay of ACh and dopamine. As ACh has been shown to bias synaptic plasticity of SC synapses towards depression, this enables learning from negative outcomes. Specifically, when a negative outcome occurs, the synapses associated with the unrewarding choices are depressed, which suppresses those choices and encourages systematic exploration. This ACh-dependent LTD enables learning from negative outcomes by facilitating the suppression of previously learned sequences of actions that no longer yield rewards. On the other hand, dopamine is capable of converting this depression to potentiation which reinforces certain actions associated with reward (Brzosko et al., 2017). These observations were combined to form a plasticity rule which was tested on a simplified network model of a radial arm maze test. Cholinergic modulation facilitated flexible learning by suppressing unrewarding actions and promoting systematic exploration of the maze (Zannone et al., 2018). Prince et al. (2021) also demonstrated in a spiking neural network model of CA3 that ACh and noradrenaline positively shifts the E-I balance at CA3 synapses via the inhibition of feedforward inhibition in the mossy fibre pathway, allowing for LTP and ensemble formation at CA3-CA3 synapses. However, unlike noradrenaline, ACh had opposite effects on cellular excitability and basal synaptic strength which enhanced the robustness of ensemble formation and the amount of allowable overlap between ensembles. These results suggest that acetylcholine and noradrenaline play distinct roles in facilitating memory formation in the CA3 network.

Ideas in cognitive neuroscience and statistical inference offer an alternative perspective on how ACh may promote learning and behavioural adaptability to environmental changes. Here, ACh is thought to provide a measure of uncertainty; this can be framed as prediction errors, which ACh has recently been found to encode (Sturgill et al., 2020). Yu and Dayan (2005) proposes ACh signals expected uncertainty which refers to the amount of uncertainty that an organism anticipates in the environment. Noradrenaline complements ACh in this system by signalling unexpected uncertainty which refers to situations where an individual encounters unpredicted events or stimuli, which deviate from their prior expectations or beliefs about the environment. These two neuromodulatory systems allow the brain to update its internal model of the world. When there is a high level of expected uncertainty, the brain needs to allocate more attentional resources to monitoring the environment, in order to update its predictions and adjust behavior accordingly. Experimental observations support the role of ACh in selectively gating top-down versus bottom-up processing: ACh can promote the processing of external stimuli by potentiating the incoming afferent projections that convey bottom-up sensory input and inhibiting the intrinsic projections that represent top-down learned expectations (Newman et al., 2012).

Another view of the role of ACh in adaptive learning, described by Doya (2002) in the context of reinforcement learning, proposes that ACh controls the learning rate of weight updates. There is mounting evidence to suggest that ACh is important in regulating the strength of synapses or weights (see Section 1.3.3) and it has also been shown that the modulatory effects of ACh on memory update is dependent on the presence of prediction errors (Krawczyk et al., 2021).

Taken together, these findings indicate that the physiological properties of ACh offer diverse mechanisms to facilitate efficient learning and behavioural adaptability in response to changing environments. Furthermore, by integrating the concept of prediction errors into cholinergic modulation, researchers have gained new perspectives on the contributions of ACh to cognitive function and behaviour.

1.4 Aims and hypotheses

This literature review indicates that while the backprop algorithm is a good solution to the credit assignment problem, its direct translation into the brain faces biological implausibilities. Recent models have been developed to address this issue by incorporating known physiological properties of the brain. However, these models still contain some biological implausibilities and fail to account for important factors that affect plasticity such as neuromodulation. It has been shown that ACh can have varied effects on different types of cells, offering distinct mechanisms for synaptic plasticity. One particular cell type of interest are the OLM interneurons located in the hippocampus as their property of targeting distal dendrites suggests that they may have a significant role in regulating synaptic plasticity. Although *ex vivo* pharmacological experiments have shown that ACh increases the intrinsic excitability of these interneurons, the impact of ACh on the inhibitory transmission of OLM interneurons onto CA1 pyramidal cells is not yet fully understood. The overall objective of this thesis is to investigate whether the brain could employ a backprop-like approach for credit assignment and how ACh might interact with the cortical circuitry to modulate this process. The specific aims are outlined below:

1. To develop a model for cholinergic neuromodulation in a biologically plausible deep learning network that is consistent with experimental findings.
2. To develop a neural network model by integrating known physiological properties of cortical networks to approximate the backpropagation algorithm.
3. To test the effect of cholinergic modulation on dendrite-targeting OLM interneurons in the hippocampal CA1 circuit.

4. To assess whether OLM interneurons are capable of regulating the induction of associative LTP in the hippocampus.

Aim 1: To develop a model for cholinergic neuromodulation in a biologically plausible deep learning network that is consistent with experimental findings.

A long-standing debate of the cholinergic system is the dynamics of acetylcholine release. Early work suggested cholinergic modulation acted globally across the brain with a homogeneous effect (Descarries, 1998; Lohani et al., 2020). However, more recently there has been evidence to indicate the contrary (Lohani et al., 2020). It is not well understood what benefits either observation would provide for the functionality of the brain. Therefore, building a biologically plausible deep network with cholinergic modulation capable of performing tasks to a similar degree of accuracy as humans helps us to understand the consequences of altering the behaviour of the cholinergic system. In Chapter 2, we propose a model of the cholinergic system as an adaptive module that incorporates error information and uses it to modify synaptic updates obtained from a biologically plausible backprop-like model (Sacramento et al., 2018). This model also generated a prediction about how ACh would affect a certain synapse which could be tested experimentally in Chapter 4.

Aim 2: To develop a neural network model by integrating known physiological properties of cortical networks to approximate the backpropagation algorithm.

As the use of the backpropagation algorithm in artificial neural networks has proven to enable learning of non-trivial tasks to human-like levels of accuracy, there has been a resurgence of interest into whether the brain employs a similar method for credit assignment. Recent work has attempted to develop biologically plausible models of backprop but many implausibilities remain (Sacramento et al., 2018; Payeur et al., 2021, 2020). Therefore, constructing a model to further address the biological implausibilities while maintaining the high performance qualities demonstrated by backprop takes us one step closer to understanding whether backprop in the brain is a reasonable concept. This novel model is introduced in Chapter 3 and incorporates the concept of burst ensemble multiplexing which provides a clear mechanism for the framework of cholinergic modulation developed in Aim 1 (Naud and Sprekeler, 2018; Payeur et al., 2021).

Aim 3: To test the effect of cholinergic modulation on dendrite-targeting OLM interneurons in the hippocampal CA1 circuit.

The hippocampal circuitry is heavily involved in the encoding of memories, particularly episodic memories that involve the spatial context provided by place cells residing in CA1. Cholinergic modulation is thought to gate hippocampal synaptic plasticity via a number of proposed mechanisms including the regulation of excitatory and feedforward inhibitory activity from the local CA3 circuitry and external entorhinal cortex (Palacios-Filardo et al., 2021). Another population of cells thought to be relevant in regulating the activity in the CA1 circuit is the local feedback

OLM interneurons. This subpopulation of SST cells almost exclusively reside in the stratum oriens of CA1 and target the distal dendrite of CA1 pyramidal cells (PCs) (Leão et al., 2012). This indicates that they have the ability to also regulate information flow to the CA1 circuit and thus carries the potential to regulate the formation of new memories. Although there have been studies to investigate the effect of acetylcholine on the excitability of these interneurons, it remains unclear how this influences the final inhibitory output onto CA1 PCs (Pelkey et al., 2017; Zheng et al., 2011; Lawrence et al., 2006; Lawrence, 2008). The model developed in Aim 1 predicts the effect of ACh modulation on this OLM-PC synapse, which is then tested pharmacologically in Chapter 4 and also with a dual optogenetics approach in Chapter 5 using whole-cell patch clamp electrophysiology.

Aim 4: To assess whether OLM interneurons are capable of regulating the induction of associative LTP in the hippocampus.

It is known that the coincident activity of the Schaffer collateral and temporammonic inputs in the hippocampus is crucial for the generation of plateau potentials that result in bursting activity for the induction of LTP (Bittner et al., 2015). As previously mentioned, the distal dendrite-targeting property of OLM interneurons make them an obvious cell population that could play an important role in mediating these coincident events and therefore hippocampal synaptic plasticity. A study by Leão et al. (2012) performing extracellular field recordings demonstrated this, as the activation of these interneurons modulated synaptic plasticity in CA1. However, this technique lacked cellular resolution, thus, could not provide details of how it was affecting the properties of the CA1 PCs to achieve this modulatory effect. Additionally, in this study the LTP induction protocol was applied to either the Schaffer collateral or temporammonic pathways but never both so did not address the question of how the OLM interneurons regulated associative plasticity. Chapter 4 uses *ex vivo* whole-cell patch clamp electrophysiology to investigate the effects of OLM inhibitory transmission on the induction of associative LTP.

CHAPTER 2

CHOLINERGIC-MEDIATED ADAPTIVE LEARNING IN A CORTICAL MICROCIRCUIT MODEL

2.1 Statement of contributions

This work was done in collaboration with a PhD student, Will Greedy, from the Department of Computer Science at the University of Bristol. Data involving the Spiral dataset and the Eve and Yi optimisers were predominantly produced by me. Will and I contributed equally to the rest of this project.

2.2 Introduction

The brain can adapt remarkably well to new situations. Some of the properties that enable it to do so include robustness to noise, rapid learning and generalisation (Viswanathan et al., 2016; Cole et al., 2013; Asok et al., 2019). Advancements in machine learning have demonstrated that adaptive learning methods applied to artificial neural networks are capable of achieving these properties (Kingma and Ba, 2014; Liu et al., 2019; Keskar and Socher, 2017; Zeiler, 2012). These methods rely on fine control of individual synaptic learning rates for gradients often derived using the backprop algorithm.

In neuroscience, it is well documented that specific neuromodulators control learning (Rasmusson, 2000; Hasselmo, 2006; Lovett-Barron et al., 2014). In particular, there is extensive research to support a strong role of the cholinergic system in learning and memory (Rasmusson, 2000; Hasselmo, 2006; Ballinger et al., 2016). There is also growing evidence to suggest that the effect of the cholinergic system on learning is produced by heterogeneous modulation of both pyramidal neurons and interneurons within various neural circuits (Klausberger and Somogyi, 2008; Lovett-Barron et al., 2014; Schmid et al., 2016; Urban-Ciecko et al., 2018; Leão et al., 2012; Haam et al., 2018; Williams and Fletcher, 2019). In particular, a study by Lovett-Barron et al. (2014) demonstrates that the cholinergic modulation of somatostatin-positive (SST) interneurons in the hippocampus is important for contextual fear conditioning. These dendrite-targeting feedback interneurons function to have an important role in cortical microcircuits. They can regulate the generation of plateau potentials (Royer et al., 2012; Lovett-Barron et al., 2012, 2014) which are involved in the learning of novel spatial representations (Bittner et al., 2015).

Building upon these findings, we introduce a framework which explores the idea of the cholinergic system as a module that receives error information from and modulates backprop-like credit assignment across the brain. This is consistent with the presence of cholinergic innervations throughout the brain (Obermayer et al., 2017; Disney and Higley, 2020) and how the cholinergic signal could be representing an unsigned prediction error (Sturgill et al., 2020). We incorporate this framework within an implementation of the error-encoding dendritic microcircuit model proposed by Sacramento et al. (2018) which approximates the backprop algorithm using biologically plausible mechanisms. This model consists of local feedback SST interneurons which control the encoding of error signals in the apical dendrites of neighbouring pyramidal neurons and closely reflects the CA3-CA1-EC circuitry involved in spatial memory. Here, we propose a model of cholinergic-mediated credit assignment (ChoCA) in cortical microcircuits in which the cholinergic system computes modulation factors, akin to adaptive learning methods from machine learning, by accumulating local error information which can in principle be decoded through multiplexing of bursting activity (Payeur et al., 2020, 2021; Naud and Sprekeler, 2018).

Neuromodulators are widely believed to operate diffusely, therefore not having the level of specificity often required by machine learning adaptive learning rules (Obermayer et al., 2017; Záborszky et al., 2018; Disney and Higley, 2020). However, it remains unclear exactly how spatially specific the neuromodulation provided by the cholinergic system is (Obermayer et al., 2017; Záborszky et al., 2018; Disney and Higley, 2020). Therefore, we examined the impact of modulating credit assignment with adaptive learning at different levels (synaptic, neuron, layer and network). Additionally, motivated by the importance of modulating SST interneurons during learning (Lovett-Barron et al., 2014), we tested whether it was necessary to apply the same modulation to the interneurons within our model. We showed that modulation at the level of the synapse was not necessary for adaptive learning methods to perform well. Additionally, we demonstrated that it was important to modulate interneurons within the network as the lack of modulation resulted in a significant reduction in performance. Finally, we showed that these adaptive methods can yield increased sparsity and task-specific selectivity but they do not mitigate catastrophic forgetting.

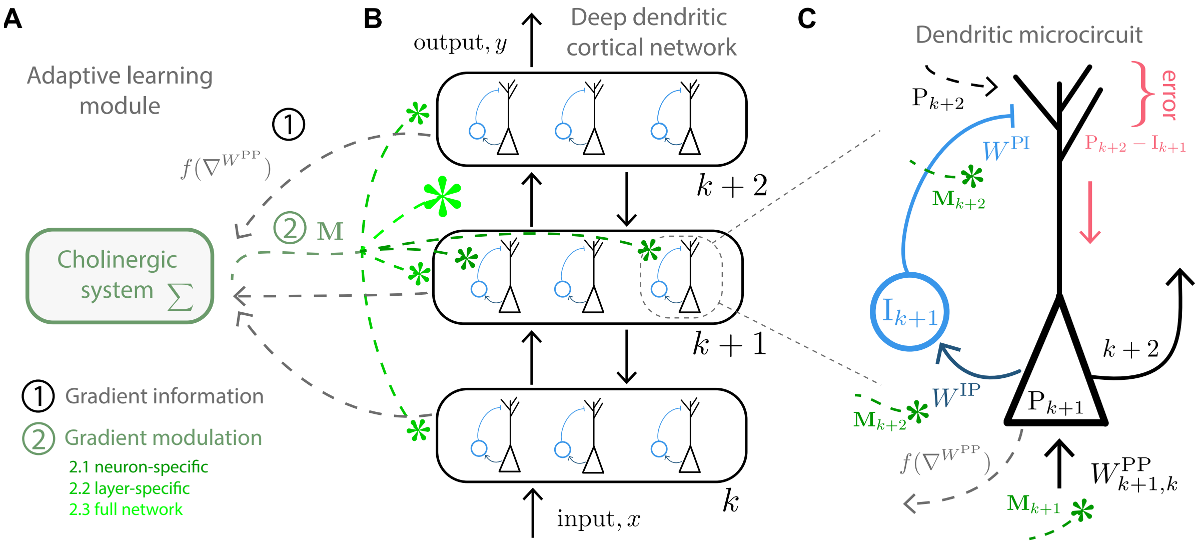
A
Adaptive learning module
①
$f(\nabla^{W^{\mathrm{PP}}})$
② M
Cholinergic system $\sum$
① Gradient information
② Gradient modulation
2.1 neuron-specific
2.2 layer-specific
2.3 full network
B
output, y
Deep dendritic cortical network
$k+2$
$k+1$
k
input, x
C
Dendritic microcircuit
P_{k+2}
error
W^{PI}
$\mathrm{P}_{k+2} - \mathrm{I}_{k+1}$
M_{k+2}
I_{k+1}
$k+2$
W^{IP}
M_{k+2}
P_{k+1}
$f(\nabla^{W^{\mathrm{PP}}})$
$W^{\mathrm{PP}}_{k+1,k}$
M_{k+1}

FIGURE 2.1: **Model of adaptive credit assignment in cortical microcircuits via cholinergic modulation (ChoCA).** (**A**) The cholinergic adaptive module receives weight gradient information $f(\nabla^{\mathbf{W}^{\mathrm{PP}}})$ across a multi-layered cortical network. These gradients are then integrated at one of three levels of specificity (neuron, layer or full network) and a signal $\mathbf{M}$ is generated to modulate credit assignment at the same level of specificity. (**B**) Simplified representation of an error-encoding dendritic cortical network with multiple layers receiving cholinergic neuromodulation (Sacramento et al., 2018). (**C**) Simplified schematic of an error-encoding dendritic microcircuit with cholinergic modulation. In this example, each neuron receives neuron-specific modulation (indicated by the green colour code; cf. (A)) from the cholinergic system (green terminals). This in turn modulates learning at all three connections ($\mathbf{W}^{\mathrm{PP}}_{k+1,k}$, $\mathbf{W}^{\mathrm{IP}}$ and $\mathbf{W}^{\mathrm{PI}}$) and receives a function of the gradients with respect to the incoming weights of P_{k+1}, $f(\nabla^{\mathbf{W}^{\mathrm{PP}}_{k+1,k}})$.

2.3 Adaptive credit assignment in cortical microcircuits

First, we introduce the Sacramento et al. (2018) model with cholinergic modulation, followed by a description of the different adaptive methods we considered.

2.3.1 Error-encoding Dendritic Network

Our model builds on Error-encoding Dendritic Networks (EDNs) (Sacramento et al., 2018). Here we focus on the elements of EDNs that are key to our model. This is a layered network model that consists of two cell types: pyramidal cells (P) and SST-interneurons (I) (Fig. 2.1). Each pyramidal cell is modelled with three simple dendritic compartments: somatic u, basal v_B and apical v_A, which have been observed to be key dendritic integration sites (Spruston, 2008; Larkum, 2013). Within layer k, the basal and apical compartments receive bottom-up ($\mathbf{W}^{\mathrm{PP}}_{k,k-1}$) and top-down connections ($\mathbf{W}^{\mathrm{PP}}_{k,k+1}$), respectively. Consistent with experimental observations (Urban-Ciecko and Barth, 2016; Williams and Holtmaat, 2019), SST-interneurons, I, target apical compartments v_A of P thereby enabling the generation of error signals within pyramidal cells by aiming to cancel top-down input onto these compartments (Fig. 2.1C). These interneurons have connections, $\mathbf{W}^{\mathrm{IP}}_{k,k}$, from neighbouring pyramidal cells that they project back to through $\mathbf{W}^{\mathrm{PI}}_{k,k}$.

We now outline how the learning rules of EDNs are modulated by the cholinergic system. The learning rules for the three connection-types mentioned above are given by

$$\nabla^{\mathbf{W}^{\mathrm{PP}}_{k,k-1}} = \left(\phi(\mathbf{u}^{\mathrm{P}}_k) - \phi(\hat{\mathbf{v}}^{\mathrm{P}}_{\mathrm{B},k})\right)\left(\mathbf{r}^{\mathrm{P}}_{k-1}\right)^T, \tag{2.1}$$

$$\Delta\mathbf{W}^{\mathrm{PP}}_{k,k-1} = \mathbf{M}_k \odot \left[\eta^{\mathrm{PP}}\nabla^{\mathbf{W}^{\mathrm{PP}}_{k+1}}\right], \tag{2.2}$$

$$\Delta\mathbf{W}^{\mathrm{IP}}_{k,k} = \mathbf{M}_{k+1} \odot \left[\eta^{\mathrm{IP}}\left(\phi(\mathbf{u}^{\mathrm{I}}_k) - \phi(\hat{\mathbf{v}}^{\mathrm{I}}_k)\right)\left(\mathbf{r}^{\mathrm{P}}_k\right)^T\right], \tag{2.3}$$

$$\Delta\mathbf{W}^{\mathrm{PI}}_{k,k} = \mathbf{M}_{k+1} \odot \left[\eta^{\mathrm{PI}}\left(\mathbf{v}_{\mathrm{rest}} - \mathbf{v}^{\mathrm{P}}_{\mathrm{A},k}\right)\left(\mathbf{r}^{\mathrm{I}}_k\right)^T\right], \tag{2.4}$$

where $(\cdot)^T$ denotes vector transpose, $\mathbf{r}_k = \phi(\mathbf{u}_k)$ are the layer k firing rates and $\mathbf{M}_k$ is the *adaptive modulation factor*. Modulation of the interneuron connections is done by $\mathbf{M}_{k+1}$ because the interneurons need to have the same dimensionality as layer $k+1$ (Sacramento et al., 2018). This factor is a matrix applied element-wise $\odot$ to layer k synaptic weight updates. The same modulation factor is applied to all learning rules to preserve the delicate learning dynamics required for successful approximate backprop in the model (Fig. 2.1C) (Sacramento et al., 2018). The different terms in Eqs. 2.1 and 2.4 represent the neuron's prediction error $\mathbf{v}^{\mathrm{P}}_{\mathrm{A},k} \approx \mathrm{error}_k$ ($\mathbf{v}_{\mathrm{rest}} = 0$). We can interpret the impact of M as modulating the error directly through the modulation of the apical dendrite excitability (Williams and Fletcher, 2019). In Eq. 2.3, $\mathbf{M}_k$ modulates the connection directly, consistent with strong modulation of PC-onto-SST synapses (Urban-Ciecko et al., 2018). To compute $\mathbf{M}_k$, a function of the gradient used for feedforward updates, $f(\nabla^{\mathbf{W}^{\mathrm{PP}}_{k,k-1}})$, is received by the cholinergic system (Fig. 2.1A; see next section for more details).

For computational efficiency, we used the simplified EDN configuration and dynamics described by Sacramento et al. (2018). In this case, the network is initialised in its optimal self-predicting configuration in which interneurons perfectly cancel top-down inputs. In addition, the model follows simplified dynamics with two simulation steps before synaptic learning rules are applied. For error backpropagation a weak supervised nudging signal ($\lambda = 0.1$) is given to the last layer. Finally, we also learn the neuronal biases which follow similar learning rules as Eqs. 2.2 and 2.3. We tested this network in a range of tasks using the MNIST and Fashion-MNIST (F-MNIST) datasets. As in Sacramento et al. (2018) target classes are one-hot encoded. During testing, the class is determined by selecting the most active output neuron.

Here we summarise the two-step update of the simplified EDN model (Sacramento et al., 2018). In the forward pass, after presenting an example image, all pyramidal cells and interneurons are updated using their feedforward activations such that $\mathbf{u}^{\mathrm{P}}_k = \mathbf{v}^{\mathrm{P}}_{\mathrm{B},k} = \mathbf{W}^{\mathrm{PP}}_{k,k-1}\phi(\mathbf{u}^{\mathrm{P}}_{k-1})$ and $\mathbf{u}^{\mathrm{I}}_k = \mathbf{v}^{\mathrm{I}}_k = \mathbf{W}^{\mathrm{IP}}_{k,k}\phi(\mathbf{u}^{\mathrm{P}}_k)$. In the second phase, output neuron potentials, $\mathbf{u}^{\mathrm{P}}_N$, are nudged towards a target output, $\mathbf{u}^{\mathrm{trgt}}_N = \phi^{-1}(\psi(\mathbf{r}^{\mathrm{trgt}}_N))$, where ϕ^{-1} is the inverse transfer function, $\mathbf{r}^{\mathrm{trgt}}_N$ is the one-hot encoded target vector and ψ is a softening function that transforms values of 0 and 1 to 0.2

and 0.8, respectively. The strength of this nudging is determined by the mixing factor, λ_N, in $\mathbf{u}_N^{\mathrm{P}} = (1-\lambda_N)\mathbf{v}_{\mathrm{B},N}^{\mathrm{P}} + \lambda_N \mathbf{u}_N^{\mathrm{trgt}}$. The updated somatic potentials are then used to update the previous layer interneurons with a weak teaching signal, $\mathbf{u}_{\mathbf{k}}^{\mathrm{I}} = (1-\lambda_{\mathrm{I}})\mathbf{v}_k^{\mathrm{I}} + \lambda_{\mathrm{I}}\mathbf{u}_{k+1}^{\mathrm{P}}$. These new interneuron potentials are used to update the pyramidal neurons' apical compartment potentials within the same layer following, $\mathbf{v}_{\mathrm{A},k}^{\mathrm{P}} = \mathbf{W}_{k,k+1}^{\mathrm{PP}}\phi(\mathbf{u}_{k+1}^{\mathrm{P}}) + \mathbf{W}_{k,k}^{\mathrm{PI}}\phi(\mathbf{u}_k^{\mathrm{I}})$. Finally, the somatic compartments of the pyramidal cells are updated to include the original feedforward input as well as the new error signal encoded within the apical compartments, $\mathbf{u}_k^{\mathrm{P}} = \mathbf{v}_{\mathrm{B},k}^{\mathrm{P}} + \lambda_k \mathbf{v}_{\mathrm{A},k}^{\mathrm{P}}$. This process is repeated for all previous layers which are revisited in reverse order. After all compartment potentials have been computed, the synaptic weight updates are performed using Eqs. 2.2-2.4 and the biases are updated using

$$\nabla^{\mathbf{b}_{k,k-1}^{\mathrm{PP}}} = \left(\phi(\mathbf{u}_k^{\mathrm{P}}) - \phi(\hat{\mathbf{v}}_{\mathrm{B},k}^{\mathrm{P}})\right), \tag{2.5}$$

$$\Delta\mathbf{b}_{k,k-1}^{\mathrm{PP}} = \mathbf{M}_k^{\mathbf{b}} \odot \left[\eta^{\mathrm{PP}} \nabla^{\mathbf{b}_{k+1}^{\mathrm{PP}}}\right], \tag{2.6}$$

$$\Delta\mathbf{b}_{k,k}^{\mathrm{IP}} = \mathbf{M}_{k+1}^{\mathbf{b}} \odot \left[\eta^{\mathrm{IP}} \left(\phi(\mathbf{u}_k^{\mathrm{I}}) - \phi(\hat{\mathbf{v}}_k^{\mathrm{I}})\right)\right], \tag{2.7}$$

where $\mathbf{M}_k^{\mathbf{b}}$ refers to the bias-specific modulation factor accumulated using $f(\nabla^{\mathbf{b}_{k,k-1}^{\mathrm{PP}}})$, similar to the weight updates described in Section 2.3.1. In the case of network level modulation, $\mathbf{M}_k^{\mathbf{b}} = \mathbf{M}_k$.

2.3.2 Cholinergic adaptive modulation methods

Below we describe the adaptive modulation methods that we propose to be abstract models of the function provided by the cholinergic system.

2.3.2.1 Adagrad-ChoCA

Adagrad was introduced by Duchi et al. (2011) to facilitate learning through modulation of learning rates as a function of the gradient magnitude. Here, together with the original modulation variable $G_t \in \mathbb{R}^{N_K \times N_J \times N_S}$ [1] and inspired by the cholinergic system, we consider less specific and more biologically plausible variants in which the gradients are collapsed across synaptic weights,

---

[1] Note that for simplicity this is a more general case in a fully connected network with the same number of neurons across layers.

$G^S_t \in \mathbb{R}^{N_K \times N_J}$, neurons, $G^{J,S}_t \in \mathbb{R}^{N_K}$, or layers, $G^{K,J,S}_t \in \mathbb{R}^{1\times 1}$, as follows

$$G_t = G_{t-1} + (\nabla^{\mathbf{W}^{\text{PP}}})^2 \qquad \text{weight-specific,} \tag{2.8}$$

$$G^S_t = G^S_{t-1} + \frac{1}{N_S}\sum_{s=0}^{N_S}(\nabla^{\mathbf{W}^{\text{PP}}_{k,j,s}})^2 \qquad \text{neuron-specific,} \tag{2.9}$$

$$G^{J,S}_t = G^{J,S}_{t-1} + \frac{1}{N_S N_J}\sum_{j=0}^{N_J}\sum_{s=0}^{N_S}(\nabla^{\mathbf{W}^{\text{PP}}_{k,j,s}})^2 \qquad \text{layer-specific,} \tag{2.10}$$

$$G^{K,J,S}_t = G^{K,J,S}_{t-1} + \frac{1}{N_S N_J N_K}\sum_{k=0}^{N_K}\sum_{j=0}^{N_J}\sum_{s=0}^{N_S}(\nabla^{\mathbf{W}^{\text{PP}}_{k,j,s}})^2 \qquad \text{full network,} \tag{2.11}$$

$$\mathbf{M}_t = \left(\sqrt{G^*_t + \epsilon}\right)^{-1}. \tag{2.12}$$

where N_S, N_J and N_K denote the number of parameters, neurons and layers, respectively. Here, G^*_t denotes any level of modulation $\{G_t, G^S_t, G^{J,S}_t, G^{K,J,S}_t\}$. $\mathbf{M}_t$ is then used to modulate learning rules in EDNs as $\Delta\mathbf{W} = \mathbf{M}_t \odot \eta \nabla^{\mathbf{W}}$ (see Eqs. 2.2-2.4). Note that we only store the necessary dimensions for G^*_t which are $N_k \times N_j \times N_s$, $N_k \times N_j$, N_k and 1, for synapse-specific, neuron-specific, layer-specific and the full network, respectively. These are then broadcasted in PyTorch so that element-wise $\odot$ multiplication can be performed in Eqs. 2.2-2.4. Finally, similar operations are performed for biases, and the respective G^*_t are kept in memory.

We argue that it is necessary for the simultaneous modulation of all three learning rules to be carried out by a separate system (Eqs. 2.2-2.4; Fig. 2.4 & 2.13). Biologically, the information (mean of the squared gradients) required for this type of modulation at the neuron-specific, layer-specific or network level can in principle be decoded from neural activity sent to the cholinergic system. Recently, Payeur et al. (2020) have proposed biologically plausible models capable of encoding gradient information through the multiplexing of event rates and burst probabilities, allowing for both feedforward activity and error signals to be decoded from neural activity.

2.3.2.2 RMSprop-ChoCA

One of the key problems with Adagrad is that G does not decay; thus, $\mathbf{M}$ only decreases over time and ultimately leads to no parameter updates. Tieleman et al. (2012) introduced RMSprop, in which the adaptive component is modelled as a running average, as follows

$$G_t = \rho G_{t-1} + (1-\rho)(\nabla^{\mathbf{W}^{\text{PP}}})^2, \tag{2.13}$$

$$\mathbf{M}_t = \left(\sqrt{G^*_t} + \epsilon\right)^{-1}. \tag{2.14}$$

Similarly to Adagrad, we consider more biologically plausible variants for RMSprop in which G is also collapsed across synapses, neurons, layers or the full network.

2.3.2.3 Adadelta-ChoCA

Similar to RMSprop, Adadelta (Zeiler, 2012) keeps a running average of the gradient magnitude:

$$G_t = \rho G_{t-1} + (1-\rho)(\nabla^{\mathbf{W}^{\mathrm{PP}}})^2, \tag{2.15}$$

$$\mathbf{M}_t = \frac{\sqrt{Z_t + \epsilon}}{\sqrt{G_t^* + \epsilon}}. \tag{2.16}$$

Adadelta also includes another factor that depends on the weight update directly, $Z_t = \rho Z_{t-1} + (1-\rho)(\mathbf{M}_t \odot \nabla^{\mathbf{W}^{\mathrm{PP}}})^2$, which is collapsed in the same way as G when we are considering more biologically plausible variants across the different levels of modulation.

2.3.2.4 Adam-ChoCA

A more recent and popular adaptive method in machine learning is Adam (Kingma and Ba, 2014). Similarly to Adadelta, Adam keeps a running average of the estimates of the first (Z) and second (G) moments of the gradient:

$$G_t = \beta_2 G_{t-1} + (1-\beta_2)(\nabla^{\mathbf{W}^{\mathrm{PP}}})^2, \tag{2.17}$$

$$\mathbf{M}_t = \frac{\hat{Z}_t}{\sqrt{\hat{G}_t^*} + \epsilon}, \tag{2.18}$$

where $\hat{G}_t^* = \frac{G_t}{1-\beta_2^t}$ and $\hat{Z}_t = \frac{Z_t}{1-\beta_1^t}$. In addition, Adam includes a bias correction term to account for the initial estimate of these moments (β_1 and β_2) being inaccurate at the beginning of training. Adam also uses momentum to help smooth the updates and escape from poor local minima. As is the case with Adagrad, RMSprop and Adadelta, we consider more biologically plausible variants at the different levels of modulation. For this method, we assume that the first moment, $Z_t = \beta_1 Z_{t-1} + (1-\beta_1)\nabla^{\mathbf{W}^{\mathrm{PP}}}$, can be computed locally.

2.3.2.5 Yi-ChoCA

The Yi method is an adaption of the optimisation method proposed by Yi et al. (2021). It scales the learning rate by the current value of the cost function:

$$\mathbf{M}_t = \frac{1}{n}\sum_{i=1}^{n}(\mathbf{u}_N^{\mathrm{P}} - \mathbf{u}_N^{\mathrm{trgt}})^2, \tag{2.19}$$

where n is each training example in a batch.

2.3.2.6 Eve-ChoCA

The Eve method is an adaption of the optimisation method proposed by Hayashi et al. (2016). It scales the learning rate dependent on the history of values evaluated from the cost function:

$$\mathbf{C}_t = \frac{1}{n}\sum_{i=1}^{n}(\mathbf{u}^{\mathrm{P}}_{N,t} - \mathbf{u}^{\mathrm{trgt}}_{N,t})^2, \tag{2.20}$$

$$\mathbf{D}_t = \frac{|\mathbf{C}_t - \mathbf{C}_{t-1}|}{\min\{\mathbf{C}_t, \mathbf{C}_{t-1}\} - \mathbf{C}^*}, \tag{2.21}$$

$$\hat{\mathbf{D}}_t = \mathrm{clip}(\mathbf{D}_t, [1/c, c]), \tag{2.22}$$

$$\tilde{\mathbf{D}}_t = \beta\tilde{\mathbf{D}}_{t-1} + (1-\beta)\hat{\mathbf{D}}_t, \tag{2.23}$$

$$\mathbf{M}_t = \frac{1}{\tilde{\mathbf{D}}_t} \tag{2.24}$$

where $\mathbf{C}_t$ is the evaluated cost at timestep t, $\mathbf{C}^*$ is the value at its global minimum (i.e. 0) and n is each training example in a batch. c is a clipping constant, set to 10 in our experiments, to constrain the magnitude of the modulation factor to prevent updates that are too large or too small.

2.3.3 Related work

In machine learning, a few adaptive methods have been developed to address the issue of the memory requirements linearly scaling with the number of model parameters. Shazeer and Stern (2018) present Adafactor, an adaptation of Adam, that maintains per-row and per-column sums of the moving averages and uses these to obtain an estimate for the original full-size modulation matrix. Anil et al. (2019) build on this work and present the SM3 algorithm that estimates the per-synapse statistics from other measures such as the maximum of gradient magnitudes over larger subsets of synapses. From a biological perspective, these methods are less plausible as they require the computation and application of per-synapse modulation factors. Some recent work argues that adaptive methods are not that important in machine learning (Wilson et al., 2017). However, these studies do not mention that, despite the fact that SGD can achieve similar or even better performance in the long run, adaptive methods can often learn much quicker, which would be a critical property in maximising an animal's chances of survival.

Within a biologically motivated reinforcement learning framework, Doya (2002) proposes that cholinergic neuromodulation is responsible for controlling the learning rate. Other computational accounts have also demonstrated a role of neuromodulatory systems in modulating the learning rate (Jepma et al., 2018; Sales et al., 2019; Aljadeff et al., 2019). For example, recent work by Miconi et al. (2020) uses standard backpropagation to optimise standard Hebbian learning rules coupled with neuromodulated eligibility traces. Our model is different in that we start from biologically plausible backpropagation framework that is then modulated via the cholinergic system, which is adaptive but not optimised via gradient descent.

2.4 Experimental details

2.4.1 Parameter requirements for each optimiser

Table 2.1 shows the number of optimiser parameters required to perform each optimisation method. The table highlights the large reduction in parameters when applying the adaptive methods at broader levels of modulation compared to their synapse level counterparts. For example, from the synapse to neuron level with Adagrad, RMSprop, Adadelta and Adam methods, there is a 99.76%, 99.76%, 99.76% and 49.87% reduction in parameters, respectively.

Method	No. Parameters (MNIST & F-MNIST)	
SGD	0	-
sAdagrad	203530	(100%)
nAdagrad	532	(0.26%)
lAdagrad	4	(0.0020%)
fAdagrad	1	(0.00049%)
sRMSprop	203530	(100%)
nRMSprop	532	(0.26%)
lRMSprop	4	(0.0020%)
fRMSprop	1	(0.00049%)
sAdadelta	407060	(100%)
nAdadelta	1064	(0.26%)
lAdadelta	8	(0.0020%)
fAdadelta	2	(0.00049%)
sAdam	407060	(100%)
nAdam	204062	(50.13%)
lAdam	203534	(50%)
fAdam	203531	(50%)

TABLE 2.1: **Number of additional parameters required by each adaptive method.** The percentage values reflect the relative number of additional parameters required compared with the synapse level methods.

2.4.2 Hyperparameter Search

A hyperparameter search was performed for the learning rate of the MNIST and F-MNIST image classification tasks. The range of values evaluated for SGD, Eve and all Adagrad and Adadelta methods was $\eta \in \{1.0, 0.1, 0.01, 0.001, 0.0001\}$. For all RMSprop and Adam methods, the ranges were $\eta \in \{0.1, 0.01, 0.001, 0.0001, 0.00001\}$ and $\eta \in \{0.001, 0.0001, 0.00001, 0.000001, 0.0000001\}$, respectively. The range for the Yi was $\eta \in \{10, 5, 2, 1, 0.1\}$. Table 2.2 shows the best parameters found and used in the two tasks. Bayesian hyperparameter optimisation was performed (using Weights & Biases (Biewald, 2020)) for the standard ANN model used on the Spiral dataset and the 3 hidden layer EDN models used on F-MNIST (Tables 2.3-2.4).

Method	MNIST [ChoCA]	F-MNIST [ChoCA]
SGD	0.1	0.1
Yi	5	5
Eve	0.1	0.1
sAdagrad	0.01	0.01
nAdagrad	0.01	0.01
lAdagrad	0.01	0.01
fAdagrad	0.001	0.01
sRMSprop	0.0001	0.0001
nRMSprop	0.0001	0.0001
lRMSprop	0.0001	0.0001
fRMSprop	0.0001	0.0001
sAdadelta	0.1	0.1
nAdadelta	0.1	0.1
lAdadelta	0.1	0.1
fAdadelta	0.1	0.1
sAdam	0.0001	0.0001
nAdam	0.0001	0.0001
lAdam	0.0001	0.0001
fAdam	0.0001	0.0001

TABLE 2.2: **Learning rates obtained using hyperparameter search for the EDN model trained on MNIST and F-MNIST.** These values correspond to the learning rate of the pyramidal-to-pyramidal feedforward weights, η^{PP}, with interneuron learning rates, $\eta^{\mathrm{IP}} = \eta^{\mathrm{PI}} = 2\eta^{\mathrm{PP}}$.

Method	Learning rate (η)
SGD	0.889
Yi	0.993
Eve	0.945

TABLE 2.3: **Learning rates obtained using Bayesian hyperparameter optimisation for the standard ANN models trained on the Spiral dataset.**

2.4.3 Learning tasks

In this section, we give details of the different tasks and model architectures that were used.

Method	**Learning rate (η)**	**η^{IP} scalar**	**η^{PI} scalar**	**Beta**
SGD	0.00979	1.80	0.755	-
Eve	0.00402	1.88	0.480	0.828
Yi	0.0319	3.39	3.70	-

TABLE 2.4: **Hyperparameters obtained using Bayesian hyperparameter optimisation for the 3 hidden layer EDN models trained on F-MNIST.** Beta is the decay constant used for the modulation factor stored with the Eve method.

2.4.3.1 MNIST and F-MNIST

For training on the MNIST and F-MNIST classification tasks, the same network structure was used for both tasks: 784 input neurons (28x28 greyscale images), a single hidden layer with 256 neurons and 10 output neurons. Sigmoid activation functions were used for the hidden layer neurons in all cases. An additional network structure of 784-256-256-256-10 was used to compare the performance of the Yi and Eve optimisers against SGD on F-MNIST. The network contained hidden layer interneurons that matched the number of pyramidal neurons in the layer above to allow for the self-prediction of next layer pyramidal neuron output activities. Sigmoid activation functions were also used for these interneurons. The feedforward weights and biases were initialised from $U(-0.1, 0.1)$, the interneuron weights were initialised in a self-predicting configuration and the feedback weights were tied to the feedforward weights such that $\mathbf{W}^{\text{PP}}_{k,k+1} = (\mathbf{W}^{\text{PP}}_{k+1,k})^T$. These parameters were updated according to the synaptic plasticity rules outlined in Section 2.3.1 with learning rates obtained through the hyperparameter search, given in Tables 2.2 & 2.4, with a batch size of 10. All nudging factors, λ_N, λ_k (for all k) and λ_{I}, were set to 0.1.

2.4.3.2 Spiral

The Spiral dataset consists of two spiral planars. The first spiral was generated by randomly sampling datapoints, x, from $2\sqrt{x}cos(4\pi\sqrt{x})$ and the second spiral from $2\sqrt{x}sin(4\pi\sqrt{x})$. Datapoints from the first and second spirals were assigned a label of 0 and 1, respectively (Fig. 2.2).

For training on the Spiral dataset, we used an ANN model with a network architecture of 2 inputs, 3 hidden layers with 20 units and 2 output units. Also, we used a sigmoid for the feedforward activation functions and a batch size of 50.

2.4.3.3 Catastrophic forgetting tasks

To evaluate the effect of each method on mitigating catastrophic forgetting, the split MNIST benchmark and variants of this task were used (Zenke et al., 2017). For all tasks, we used a one hidden layer EDN model with the same network architecture as mentioned for the MNIST

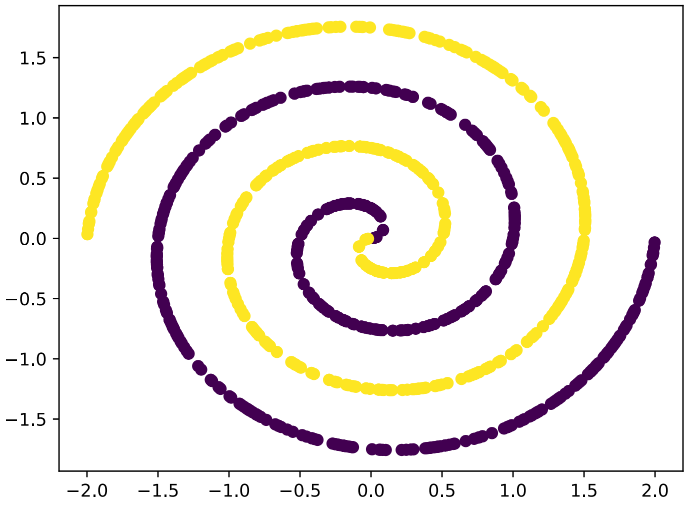

FIGURE 2.2: **The Spiral dataset.** Yellow and purple spirals form the two classes that the model must learn to classify.

and F-MNIST tasks (Section 2.4.3.1). We used a multi-head approach in which the loss was only computed for digits present in the current subset trained on. Unlike Zenke et al. (2017), the optimiser state was not reset after training each task.

MNIST-split-five and F-MNIST-split-five

MNIST-split-five and F-MNIST-split-five are variants of the split MNIST benchmark (Zenke et al., 2017), in which each of the MNIST and F-MNIST datasets were split into two subsets of five classes (Kemker et al., 2017). During the initial phase of the split-five tasks, the model was trained with a batch size of 10 on the first subset. Once this base set was learned, the model was trained with a batch size of 1 on the remaining 5 classes. In both phases, examples were sampled i.i.d. from the corresponding training set. We evaluated the mean classification accuracy on the base set after training at regular intervals to assess the model's memory of the classes learned in the first phase.

Split MNIST and split F-MNIST

Split MNIST (Zenke et al., 2017) is a catastrophic forgetting benchmark and split F-MNIST is a slightly harder variant using the F-MNIST dataset. These tasks involve dividing the entire MNIST and F-MNIST training datasets into five subsets, each consisting of consecutive digits. A model is trained on these five subsets to differentiate between two consecutive digits in the range of 0 to 9.

2.5 Results

To empirically evaluate our adaptive methods at different levels of modulation, we used the MNIST and F-MNIST image classification tasks. The Spiral dataset was also used to further test the Yi and Eve optimisation methods. We refer to the adaptive methods as {s,n,l,f}method where s = synapse, n = neuron, l = layer and f = full network.

2.5.1 Synapse-specific neuromodulation is not required for efficient adaptive learning

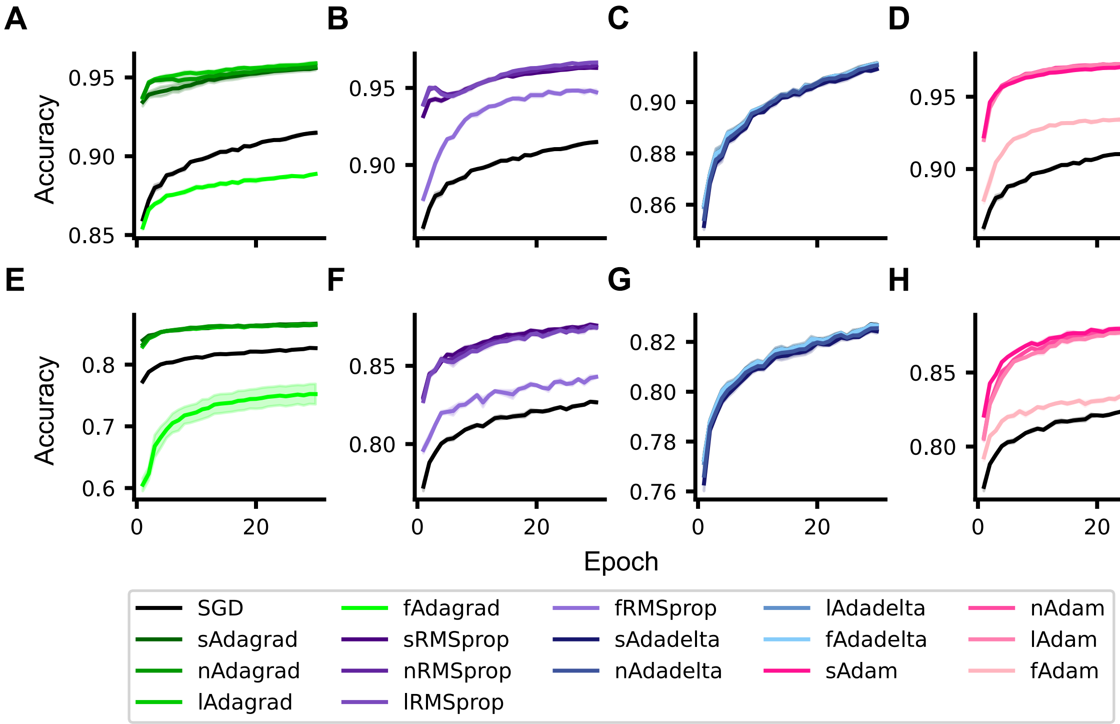

FIGURE 2.3: **Efficient adaptive learning without synaptic-specific neuromodulation.** Test set accuracy of each adaptive method for (**A-D**) MNIST and (**E-H**) F-MNIST. Data represented as mean and standard error (n = 5).

First, we evaluated whether our gradient-based methods (Adagrad-ChoCA, RMSprop-ChoCA, Adadelta-ChoCA and Adam-ChoCA) could be used to train the EDN model and if methods operating with less fine-grained modulation could retain the desirable properties of their synapse-level counterparts. A key property of standard adaptive learning methods is the increased performance they can achieve. To demonstrate that our methods maintain this property, we evaluated the test accuracy for each task throughout training (Fig. 2.3). Interestingly, when modulation was applied to the weight changes (Eqs. 2.2-2.3), the self-predicting state remained stable. In all cases, training with neuron- and layer-level modulation achieved similar accuracies to synapse-level modulation. These results suggest that the majority of the methods do not suffer from significant reductions in performance despite using much less gradient information (see Table 2.1 for memory usage comparison) and having much coarser-grained control. When modulation was applied at levels lower than the full network, all ChoCA methods performed better than SGD with the exception of Adadelta for which the performance was comparable (Fig. 2.3C, G). These results indicate that it is beneficial to adaptively control synaptic plasticity throughout training and are consistent with a wide range of experimental observations showing an important function of the cholinergic system in learning (Rasmusson, 2000; Hasselmo, 2006; Ballinger et al., 2016).

2.5.2 The effect of removing cholinergic modulation on interneuron synaptic plasticity and learning

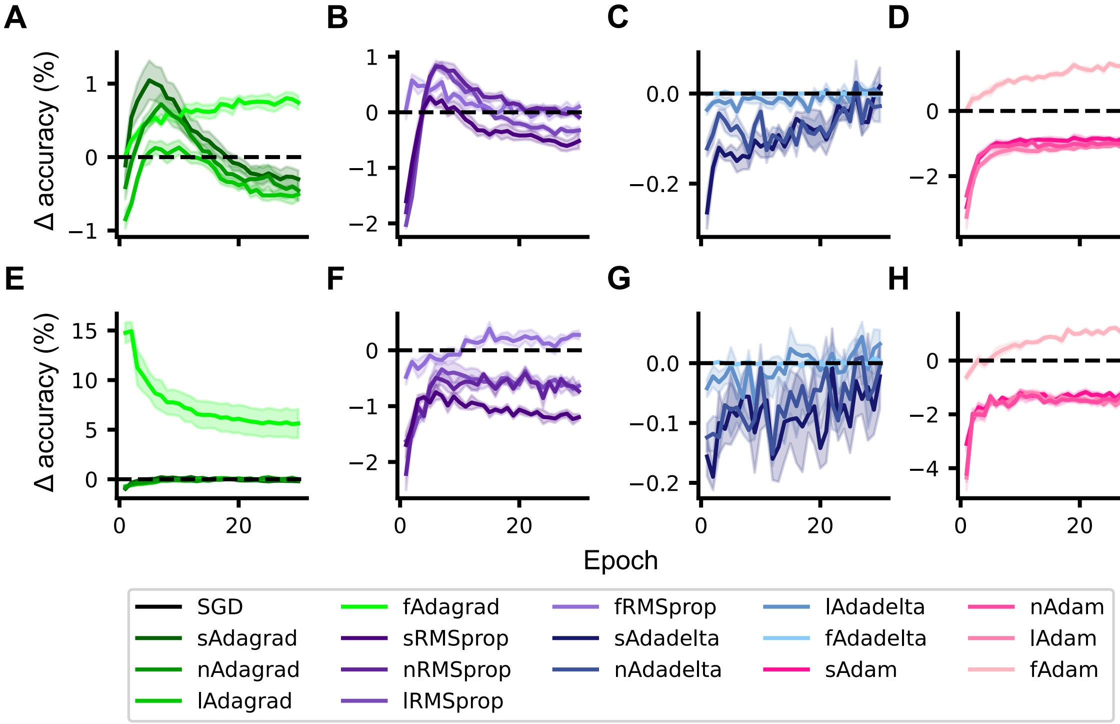

FIGURE 2.4: **Impact of removing cholinergic modulation of interneuron synapses on adaptive learning.** Difference in performance between the EDN model without interneuron modulation and a model with standard microcircuit neuromodulation, for both (**A-D**) MNIST and (**E-H**) F-MNIST. Both models are initialised in the optimal self-predicting state. Data represented as mean and standard error (n = 5).

To test the importance of microcircuit modulation by the cholinergic system, we investigated the difference in performance between modulating all connections (Eqs. 2.2-2.3) and only the feedforward weights. For all methods, excluding network-level modulation, the results show that the application of interneuron modulation increases overall performance (Fig. 2.4). This suggests that modulating interneuron synaptic plasticity is important for efficient adaptive learning.

2.5.3 Adaptive learning increases sparsity and selectivity

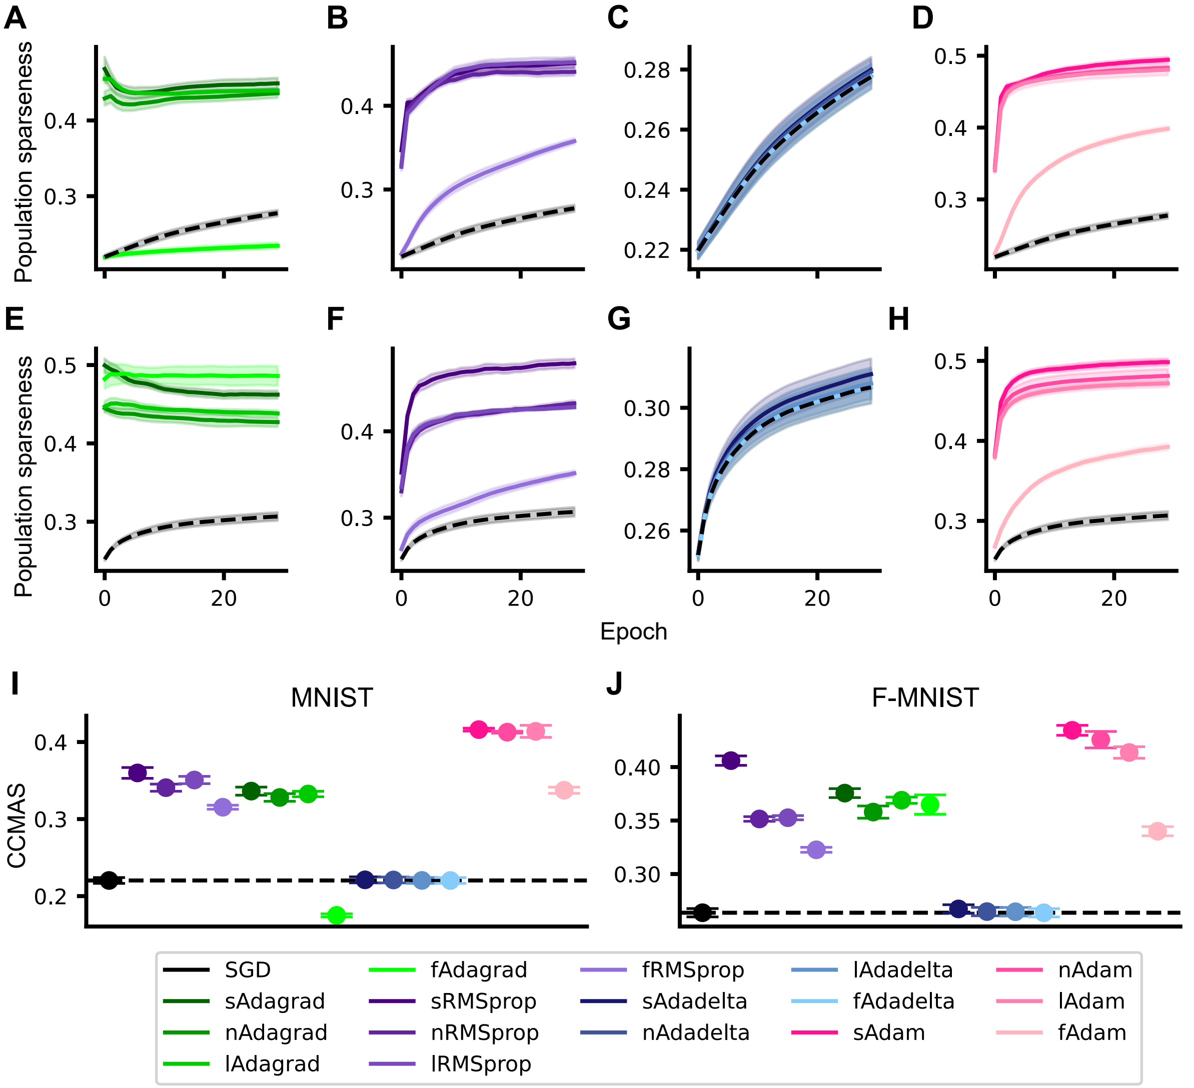

FIGURE 2.5: **Adaptive learning yields a better use of network representation capacity by increasing sparsity and selectivity.** The sparsity of activations present in each method are shown across epochs of training for (**A-D**) MNIST and (**E-H**) F-MNIST. (**B**) The selectivities of neurons to particular classes in the (**I**) MNIST and F-MNIST (**J**) datasets across methods are displayed using the CCMAS metric with a black dashed line indicating SGD as a reference for baseline selectivity. Both the sparsity and selectivity measures were calculated from neurons in the hidden layer. Data represented as mean and standard error (n = 5).

Adaptive learning methods implicitly control which parts of neural networks change over training, which should encourage more distributed neuronal representations in comparison to non-adaptive methods (e.g. SGD). We evaluated sparsity following Cayco-Gajic et al. (2017)[2] of each method in both the MNIST and F-MNIST tasks and found that, with the exception of Adadelta-ChoCA and full network Adagrad, the methods produced sparser representations compared to SGD (Fig. 2.5A-H). Sparsity is thought to be an important property in encouraging neurons to have feature selective responses (Sharkey and Sharkey, 1995; Murre, 2014). We investigated whether the increased sparsity seen in these methods would result in increased selectivity for specific classes as quantified by the class-conditional mean activity selectivity (CCMAS) metric (Morcos et al., 2018). Consistent with the population sparseness results shown for MNIST and F-MNIST, the same methods displayed an increase in selectivity for distinct classes in the datasets (Fig. 2.5I-J).

---

[2] Sparsity is measured as $\frac{N-\frac{(\sum_i x_i)^2}{\sum_i x_i^2}}{N-1}$, where N is the number of neurons and x_i is the neuron activity (Cayco-Gajic

2.5.4 The effect of adaptive learning on catastrophic forgetting

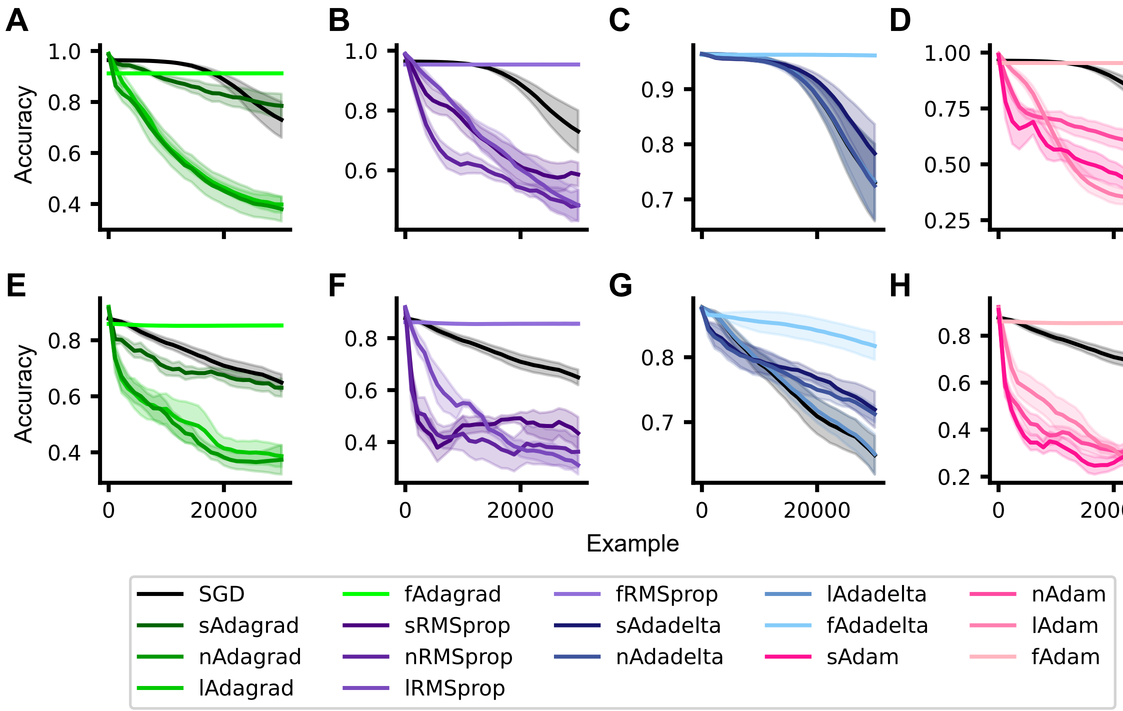

FIGURE 2.6: **The impact of adaptive learning on catastrophic forgetting for the MNIST- and F-MNIST-split-five tasks.** Classification accuracy for (**A-D**) MNIST and (**E-H**) F-MNIST variants of the split MNIST benchmark to quantify the degree of catastrophic forgetting as novel examples from unseen classes are presented after training on five classes. All methods are displayed against SGD as a reference for baseline levels of forgetting. Data represented as mean and standard error (n = 5).

The learning rate modulation of these adaptive methods diminishes updates in synapses that have experienced large amounts of change. This property of the dynamics could provide a mechanism to consolidate synapses important for features required to solve the task at hand. Therefore, as these methods are exposed to future tasks, the modulation parameters could act as a form of memory for the past tasks, thus, mitigating catastrophic forgetting. Additionally,

---

et al., 2017).

we were interested in whether the increased sparsity and selectivity properties observed with the adaptive methods would affect catastrophic forgetting. This is because sparse activations can promote generalisation and make more efficient use of network resources which helps to prevent overfitting, making sparse networks less prone to catastrophic forgetting (Srivastava et al., 2014). To examine the effect of the adaptive methods on catastrophic forgetting, we first evaluated the methods on MNIST-split-five and F-MNIST-split-five, variants of the split MNIST benchmark (Zenke et al., 2017), in which each dataset was split into two subsets of five classes (Kemker et al., 2017). The model is initially trained on the first subset and the task is to continue to remember this whilst training on the remaining 5 classes. We used a multi-head approach in which the loss was only computed for digits present in the current subset trained on (see Section 2.4.3.3 for complete experiment details). We evaluated the mean classification accuracy on the base set after regular intervals of training on the remaining classes to assess the model's memory of the classes learned in the first phase. Our results show that the majority of the adaptive methods exhibit equal or greater catastrophic forgetting when compared to to SGD (Fig. 2.6). However, network-level modulation consistently demonstrates lower catastrophic forgetting. In the network-level setup, as all gradients are collapsed into one modulation factor, the gradients of each layer in the network affects the updates within the same layer and the other layers. Therefore, if the magnitude of the output layer gradients is much bigger than the other layers, then updates in the other layers are suppressed more relative to the output layer. Due to the multihead approach freezing the output heads for the initial classes, the representations for them are unaffected when training occurs for the remaining classes which means that the overall output for the initial classes remains the same resulting in minimal catastrophic forgetting.

In addition to the MNIST-split-five and F-MNIST-split-five tasks, we evaluated the performance of each adaptive method on the split MNIST benchmark (Zenke et al., 2017) (Fig. 2.7) and a slightly harder variant using the F-MNIST dataset (Fig. 2.8). For the split MNIST and split F-MNIST benchmarks, we directly followed the experimental methodology provided by Zenke et al. (2017) using our previously mentioned models. The optimiser state was not reset after training each task as we were evaluating how the dynamics of these adaptive methods would impact catastrophic forgetting. The performance of each adaptive method was evaluated on the mean classification accuracy on all previous tasks as a function of number of tasks trained. In general, the adaptive methods were not as good as SGD at remembering previous tasks (Fig. 2.7-2.8).

2.5.5 Neuromodulation of synaptic plasticity directly using cost information improves learning

Based on the findings from gradient-based methods, which indicate that coarser-grained modulation can be as effective as synapse-level modulation for improving learning, we sought to investigate the characteristics of new adaptive learning methods that are more aligned with the current understanding of how the cholinergic system functions. Therefore, inspired by Yi et al.

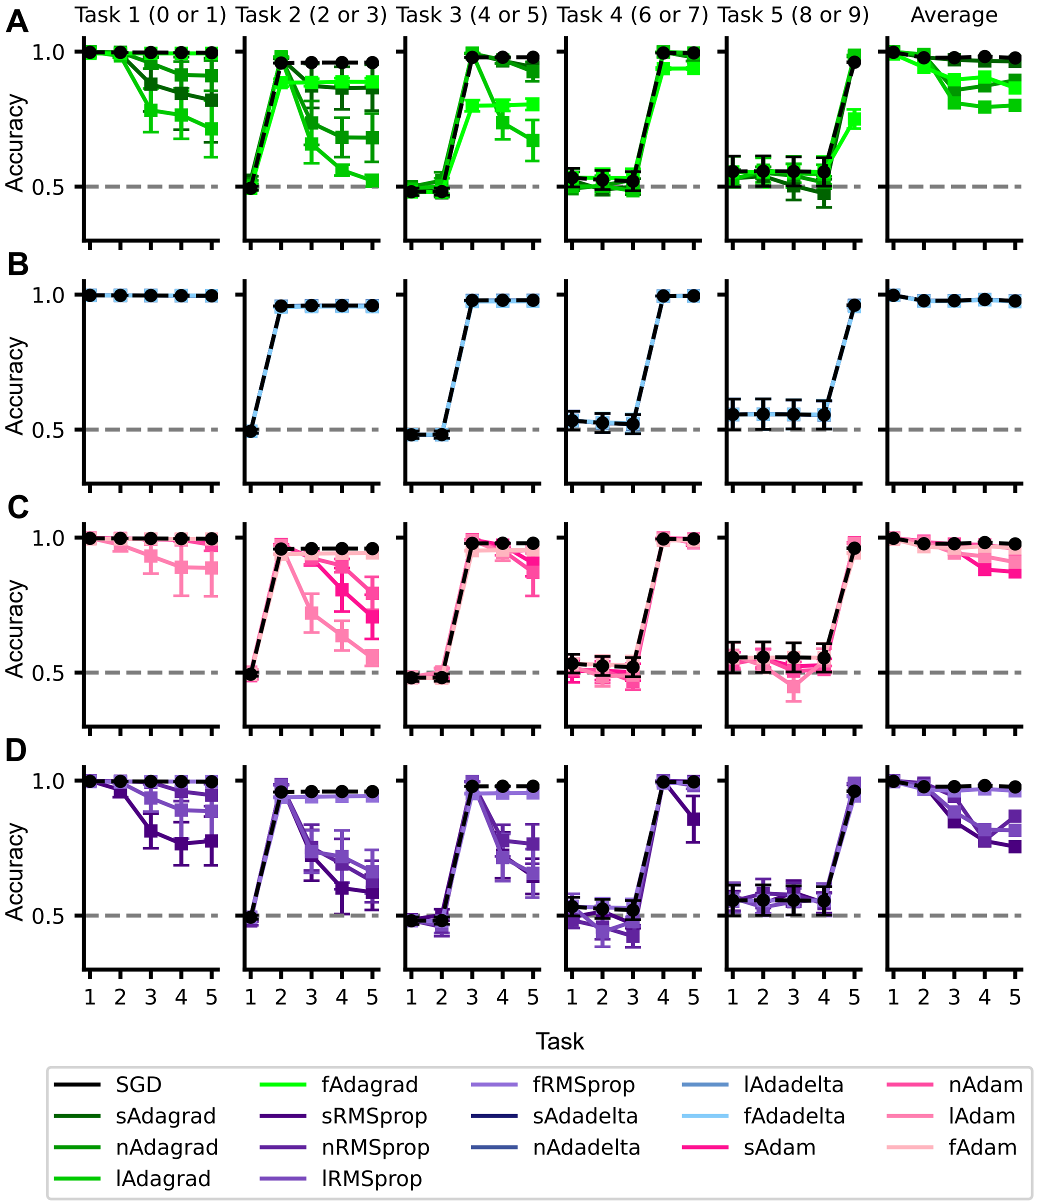

FIGURE 2.7: **Performance of adaptive methods on split MNIST.** The first five panels show classification accuracies on the five two-digit tasks after training on each task (x-axis). The panel furthest to the right shows the mean classification accuracy as a function of the number of tasks. All methods of (**A**) Adagrad, (**B**) RMSprop, (**C**) Adadelta and (**D**) Adam are displayed against SGD as a reference for baseline levels of forgetting. In this setup, five binary readout heads are used so an accuracy of 0.5 corresponds to chance represented by the grey dashed line. Data represented as mean and standard error (n = 5).

(2021) and Hayashi et al. (2016), we developed two novel optimisation methods, called Yi-ChoCA and Eve-ChoCA, which modulates the learning rate using the error information from the cost function. Unlike the gradient-based methods, these methods strictly have a global, homogeneous effect and increases modulatory output to error signals to facilitates synaptic plasticity, consistent with how the cholinergic system is thought to behave (Hasselmo, 2006; Obermayer et al.,

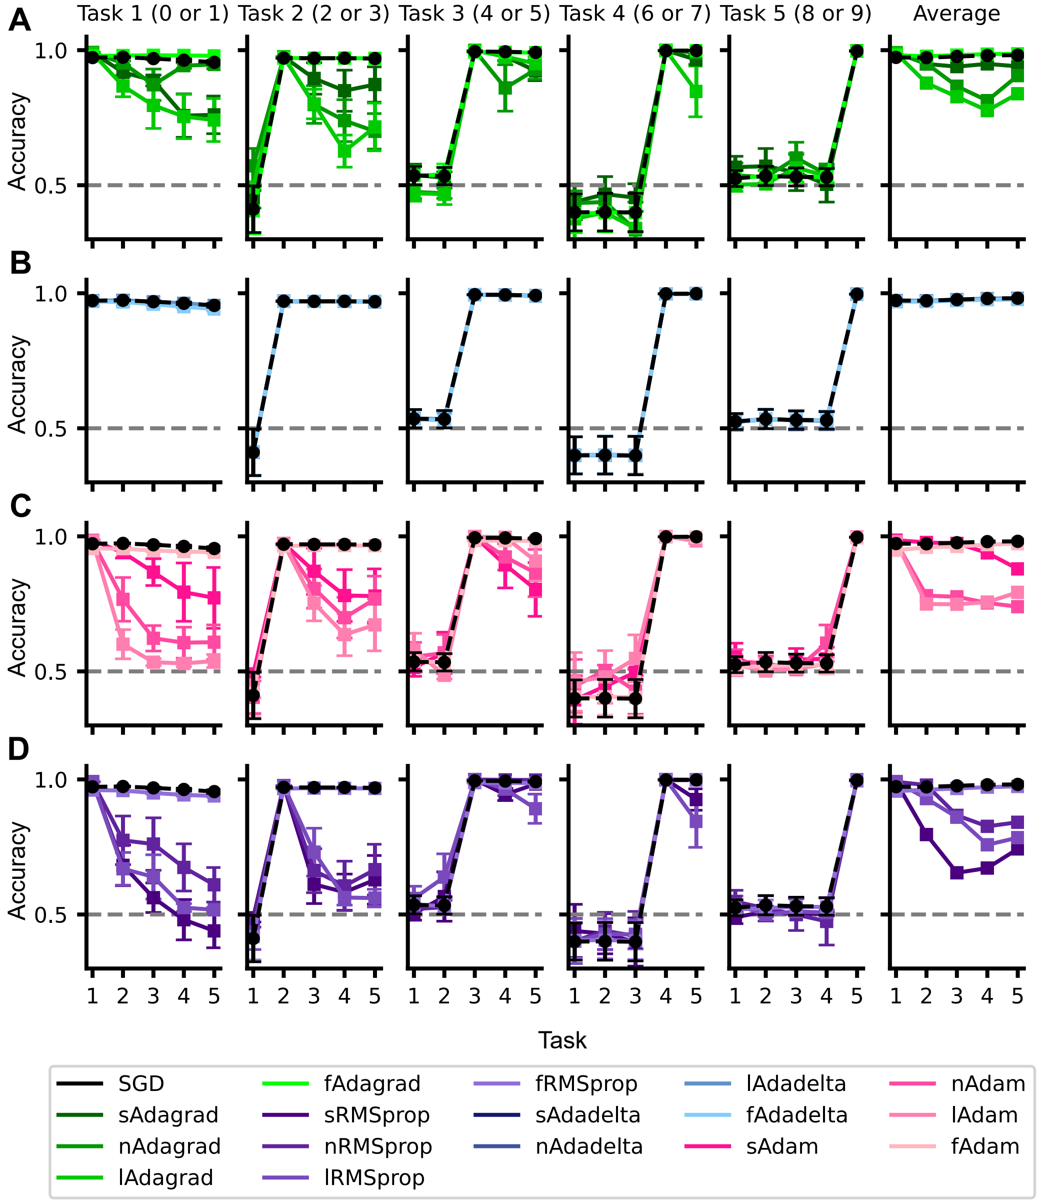

FIGURE 2.8: **Performance of adaptive methods on split F-MNIST.** The first five panels show classification accuracies on the five two-digit tasks after training on each task (x-axis). The panel furthest to the right shows the mean classification accuracy as a function of the number of tasks. All methods of (**A**) Adagrad, (**B**) RMSprop, (**C**) Adadelta and (**D**) Adam are displayed against SGD as a reference for baseline levels of forgetting. In this setup, five binary readout heads are used so an accuracy of 0.5 corresponds to chance represented by the grey dashed line. Data represented as mean and standard error (n = 5).

2017; Záborszky et al., 2018; Disney and Higley, 2020). In addition to the MNIST and F-MNIST datasets, we tested the properties of these two cost-based methods on the Spirals dataset.

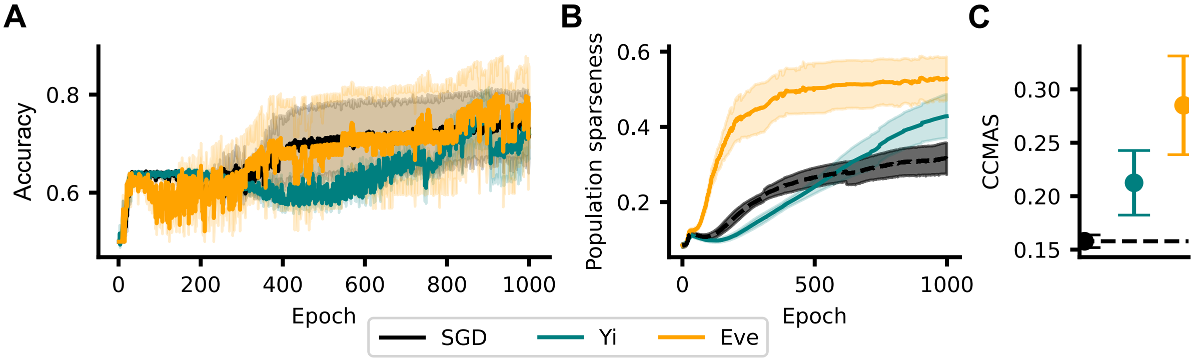

FIGURE 2.9: **Performance of cost-based methods on the Spirals dataset in standard deep neural networks.** (**A**) Test set accuracy of each adaptive method for the Spirals dataset. (**B**) The sparsity of activations present in each method are shown across epochs of training. (**C**) The selectivities of neurons to particular classes across methods are displayed using the CCMAS metric with a black dashed line indicating SGD as a reference for baseline selectivity. Data represented as mean and standard error (n = 5).

2.5.5.1 Spirals

Initially, both cost-based adaptive methods were tested on the Spirals dataset using a standard 3 hidden layer ANN. We found that the learning performance of our methods were similar to that of SGD (Fig. 2.9A). They displayed increased sparsity and selectivity as seen with the gradient-based adaptive methods (Fig. 2.9B-C).

2.5.5.2 MNIST

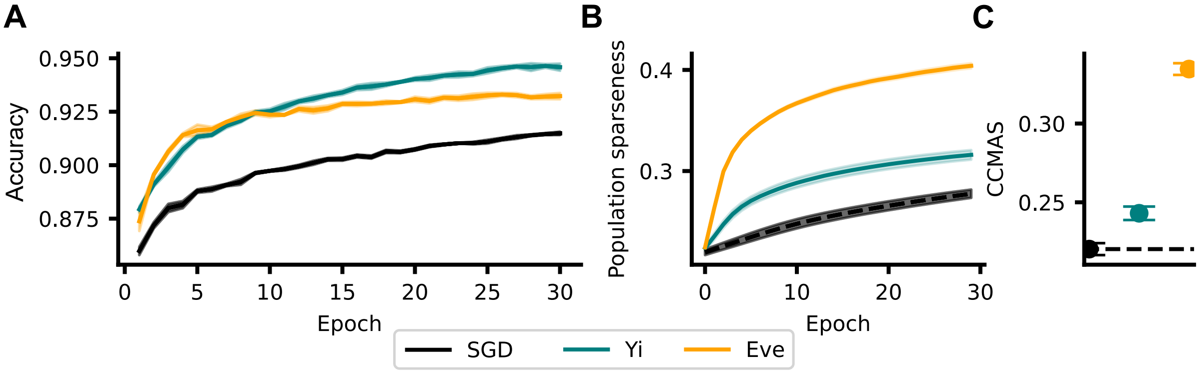

FIGURE 2.10: **Performance of cost-based adaptive methods on the MNIST dataset in the EDN model.** (**A**) Test set accuracy of each adaptive method for the MNIST dataset. (**B**) The sparsity of activations present in each method are shown across epochs of training. (**C**) The selectivities of neurons to particular classes across methods are displayed using the CCMAS metric with a black dashed line indicating SGD as a reference for baseline selectivity. Data represented as mean and standard error (n = 5).

When these cost-based adaptive methods were applied in a one hidden layer EDN model on the MNIST image classification task, we found an increase in performance compared to SGD (Fig. 2.10A). Similar to the properties seen on the Spirals dataset in ANNs, these methods also demonstrate increased sparsity and selectivity (Fig. 2.10B-C).

2.5.5.3 F-MNIST

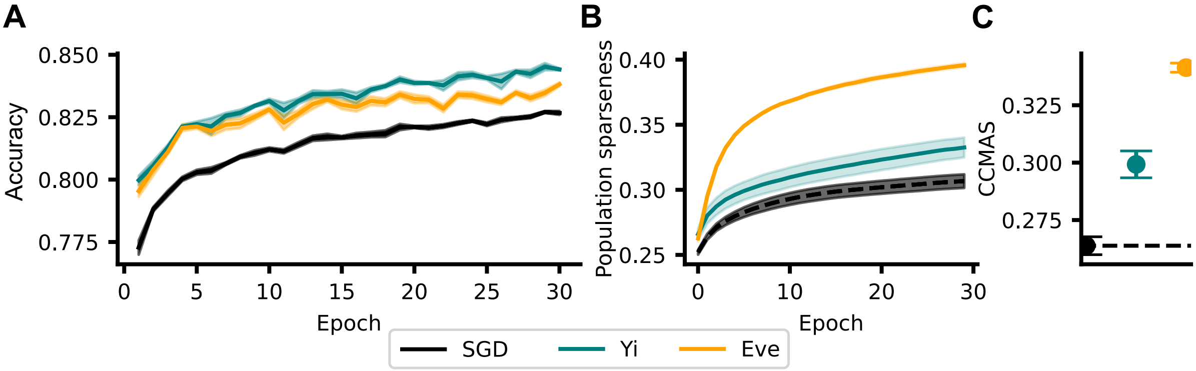

FIGURE 2.11: **Performance of cost-based adaptive methods on the F-MNIST dataset in the EDN model.** (**A**) Test set accuracy of each adaptive method for the F-MNIST dataset. (**B**) The sparsity of activations present in each method are shown across epochs of training. (**C**) The selectivities of neurons to particular classes across methods are displayed using the CCMAS metric with a black dashed line indicating SGD as a reference for baseline selectivity. Data represented as mean and standard error (n = 5).

Similar to the results found with the MNIST dataset, cost-based adaptive learning in a one hidden layer EDN model on F-MNIST improved learning performance compared to SGD, promoted sparse activations and increased selectivity of neurons to distinct classes (Fig. 2.11). Next, we wanted to test whether these properties changed with a deeper network. Therefore, these adaptive methods were applied in a three hidden layer EDN model and evaluated on the F-MNIST task. With this deeper network, Eve-ChoCA still performs slightly better and displays faster learning compared to SGD (Fig. 2.12A). However, Yi-ChoCA performs significantly worse than SGD and also loses the increased sparsity and selectivity properties previously shown. The Eve-ChoCA method maintains the selectivity property for all hidden layers and the increased sparsity seen before, remain in the earlier layers (Fig. 2.12B-C).

Together, these results suggest that, similar to the gradient-based adaptive methods, both cost-based adaptive methods improve learning performance and display a consistent increase in sparsity and selectivity across different tasks in the EDN model with shallower network architectures. However, the Yi-ChoCA method loses its increased performance, sparsity and selectivity properties when it is applied in deeper networks whereas the Eve-ChoCA is more robust to changes in network architecture.

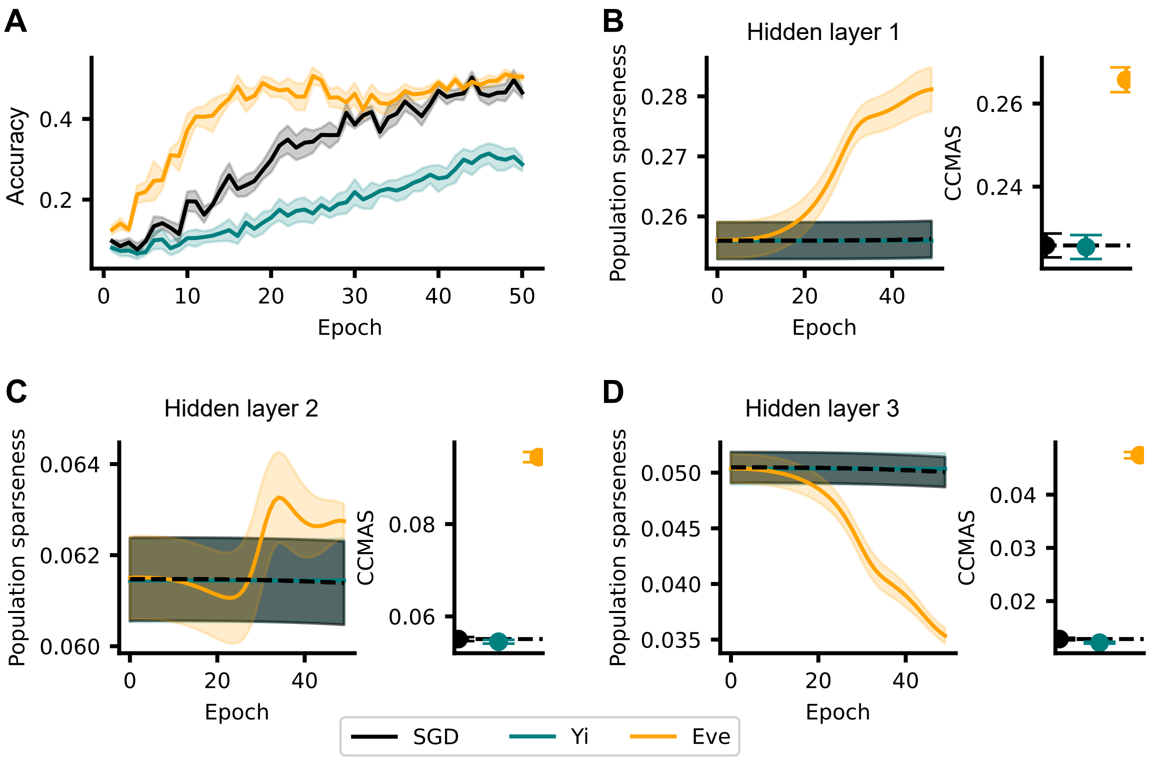

FIGURE 2.12: **Performance of cost-based adaptive methods on the F-MNIST dataset in a 3 hidden layer EDN model.** (**A**) Test set accuracy of each adaptive method for the F-MNIST dataset. (**B-D**; left) The sparsity of activations present in each method are shown across epochs of training. (**B-D**; right) The selectivities of neurons to particular classes across methods are displayed using the CCMAS metric with a black dashed line indicating SGD as a reference for baseline selectivity. Both the sparsity and selectivity measures were calculate for all hidden layers. Data represented as mean and standard error (n = 5).

2.5.6 The effect of cost-based adaptive modulation of interneurons on learning

Similar to Section 2.5.2, we wanted to examine the effect of removing the modulation, provided by the cost-based adaptive methods, of interneuron synaptic plasticity in the EDN model. In the one hidden layer EDN model, the results demonstrate that interneuron modulation is critical for learning with the Yi-ChoCA method as removal leads to a significant decrease in performance on the MNIST and F-MNIST tasks. However, an opposite effect is seen with the Eve-ChoCA method as the removal of interneuron modulation results in an increase in performance on the two image classification tasks (Fig. 2.13A-B). In the three hidden layer EDN model, both Yi-ChoCA and Eve-ChoCA methods suffer a reduction in performance with the removal of interneuron modulation (Fig. 2.13C). These data suggest that interneuron modulation is important for the Yi-ChoCA method but has inconsistent effects for the Eve-ChoCA method depending on the network architecture.

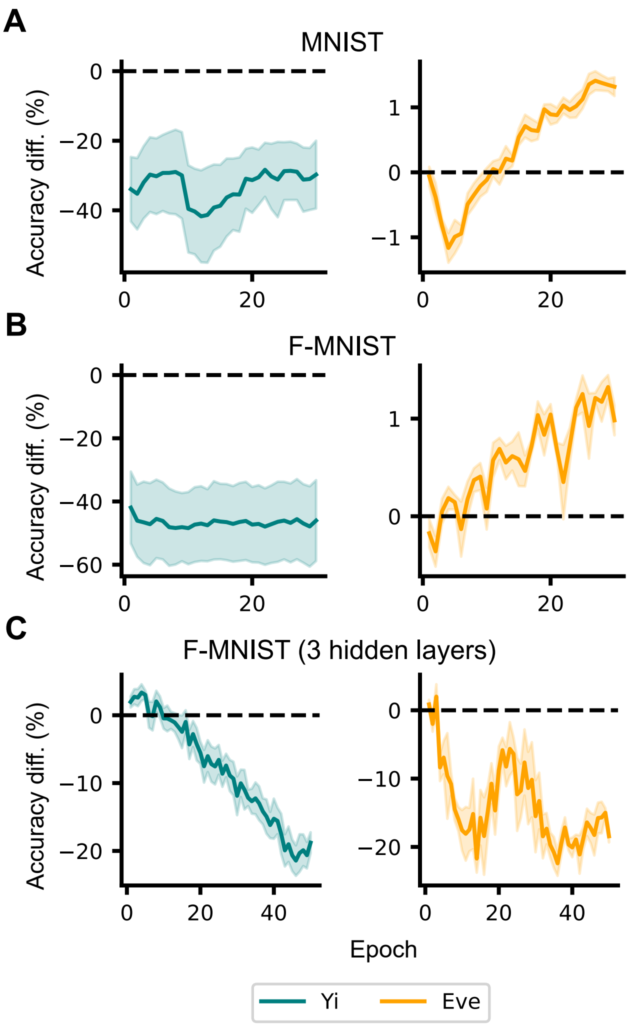

FIGURE 2.13: **Impact of removing cost-based cholinergic modulation of interneuron synaptic plasticity on adaptive learning.** Difference in performance between the EDN model without interneuron modulation and a model with standard microcircuit neuromodulation for (**A**) MNIST, (**B**) F-MNIST and (**C**) F-MNIST with a deeper network architecture of 3 hidden layers instead of one. Both models are initialised in the optimal self-predicting state. Data represented as mean and standard error (n = 5).

2.5.7 The effect of cost-based adaptive learning on catastrophic forgetting

The same tasks in Section 2.5.4 were used to examine the effect of the Yi-ChoCA and Eve-ChoCA methods on catastrophic forgetting. For the MNIST-split-five and F-MNIST-split-five tasks, both cost-based adaptive methods demonstrated greater level of catastrophic forgetting in comparison to SGD (Fig. 2.14). This result was consistent for Yi-ChoCA on split MNIST but similar performances to SGD were seen with Eve-ChoCA on split MNIST and both methods on split F-MNIST (Fig. 2.15). These results indicate that both cost-based methods do not mitigate catastrophic forgetting.

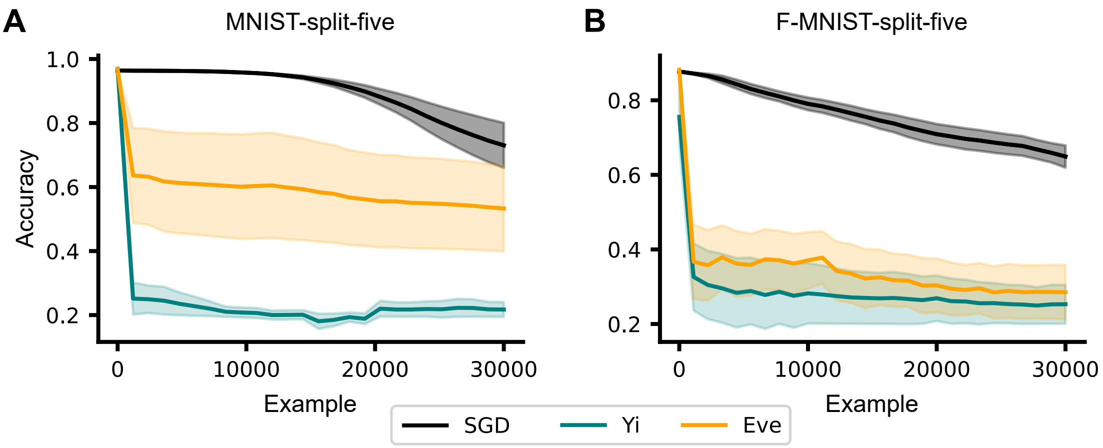

FIGURE 2.14: **The impact of cost-based adaptive learning on catastrophic forgetting for the MNIST- and F-MNIST-split-five tasks.** Classification accuracy for (**A**) MNIST and (**B**) F-MNIST variants of the split MNIST benchmark to quantify the degree of catastrophic forgetting as novel examples from unseen classes are presented after training on five classes. All methods are displayed against SGD as a reference for baseline levels of forgetting. Data represented as mean and standard error (n = 5).

2.6 Discussion

Here, we have proposed a framework for the cholinergic system as an adaptive learning module that integrates error information to modulate synaptic plasticity throughout the brain. In our model, we have specifically explored the idea of acetylcholine as a neuromodulator that controls the learning rate of a system via cholinergic modulation of synaptic credit assignment akin to adaptive learning algorithms used in machine learning (Doya, 2002; Duchi et al., 2011; Tieleman et al., 2012; Zeiler, 2012; Kingma and Ba, 2014; Shazeer and Stern, 2018; Anil et al., 2019; Yi et al., 2021; Hayashi et al., 2016). Our results show that despite its diffuse nature, cholinergic neuromodulation in the brain can enhance the efficiency of learning when this is coupled with synaptic credit assignment. We have shown across different tasks that, with the exception of network-level Adagrad-ChoCA and Adadelta-ChoCA, the adaptive modulation of synaptic plasticity in cortical microcircuits provides significant improvements to learning. There has been a longstanding debate about whether the cholinergic system acts as a global modulator, influencing brain activity in a non-specific, homogeneous manner, or whether it acts more specifically to modulate particular neural circuits or regions (Disney and Higley, 2020). We have shown with the gradient-based adaptive methods that more heterogeneous models of cholinergic modulation are better at enhancing learning. A recent study by Lohani et al. (2020) using dual-color mesoscopic imaging of both acetylcholine and calcium activity across the neocortex demonstrates that different behavioural states are associated with distinct spatiotemporal patterns of cholinergic signalling. This suggests that the cholinergic system can provide dynamic, heterogeneous fluctuations in acetylcholine depending on the current state of the brain. This does not mean that the cholinergic system does not act as a diffuse,

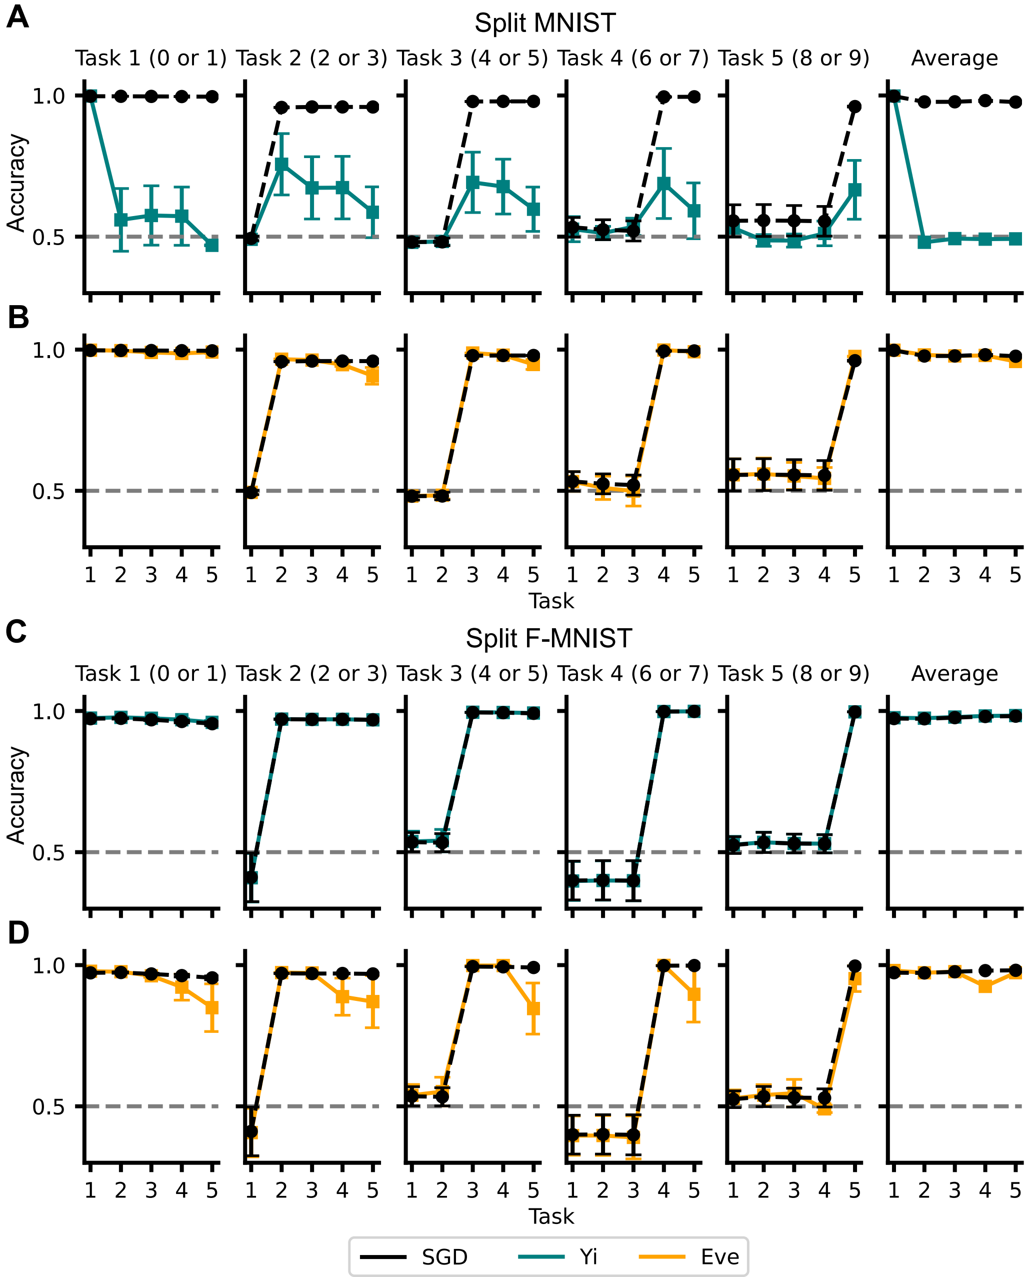

FIGURE 2.15: **Performance of cost-based adaptive methods on split MNIST and split F-MNIST.** The first five panels show classification accuracies on the five two-digit tasks after training on each task (x-axis). The panel furthest to the right shows the mean classification accuracy as a function of the number of tasks. Both methods were tested on (**A-B**) split MNIST and (**C-D**) split F-MNIST. They are displayed against SGD as a reference for baseline levels of forgetting. In this setup, five binary readout heads are used so an accuracy of 0.5 corresponds to chance represented by the grey dashed line. Data represented as mean and standard error (n = 5).

global neuromodulator but implies that it is also capable of more targeted effects. We have also shown that the homogeneous, global effect of the cost-based adaptive methods, which are more consistent with experimental observations of the cholinergic system as a network that encodes valence-free prediction error signals, facilitates learning, albeit to a lesser degree than the finer-resolution gradient-based adaptive methods (Hangya et al., 2015; Sturgill et al., 2020). However, the Eve-ChoCA method was better at mitigating catastrophic forgetting in the split MNIST and split F-MNIST tasks compared to the gradient-based adaptive methods that improved learning on the image classification tasks. In addition, the network-level gradient-based adaptive methods consistently performed better than their less diffuse counterparts on the forgetting tasks. Collectively, this indicates that there is a context-dependent function for both homogeneous and heterogeneous cholinergic signals. In addition, our use of adaptive methods removes the need for hand-tuned layer-specific learning rates as used in the EDN model by Sacramento et al. (2018).

In the current implementation of adaptive learning in our model, the cholinergic modulation directly alters the learning rate of the plasticity rules. We can interpret this modulation as a change in the excitability of the distal apical dendrite where the error signals are computed (Williams and Fletcher, 2019). For example, a large modulation factor indicates a high concentration of acetylcholine release which increases the excitability of the distal dendrite; thus, increasing the potential and scaling the error signal to promote plasticity of the bottom-up synapses. However, if this altered error signal is directly propagated down the network then earlier layers in the network would experience updates that are misaligned with backprop-derived gradients. This would likely cause the model to not correctly follow gradient descent and ultimately harm learning in the network. Therefore, a correction would need to be made for this modulation in the outgoing signals. If we increase the excitability of the distal dendrite with cholinergic modulation as shown by Williams and Fletcher (2019) then we would need to depress the weight of the feedback excitatory and inhibitory connections. The reduction in weight of the feedback excitatory pathway is consistent with experimental observations by Palacios-Filardo et al. (2021). They show that when you patch onto a CA1 pyramidal cell there is a reduction in the excitatory response from the temporammonic feedback pathway in the presence of acetylcholine. Another prediction would be that the SST-to-PC connection must also be down-regulated by acetylcholine which is examined in Chapter 4.

In the EDN model, the distal dendrite-targeting interneurons play a critical role in the generation error signals that are used for the update of feedforward synapses. We found that inhibiting cholinergic modulation of interneuron synaptic plasticity in the non-network-level gradient-based adaptive methods and cost-based adaptive methods in deep networks significantly reduced learning capabilities. Without this modulation, the feedforward weights are likely changing at a rate that the interneuron weights cannot follow resulting in the EDN losing its self-predicting state; thus, generating error signals that begin to not align with backprop-derived gradients. If the brain performs credit assignment in a similar manner to the EDN model, this suggests that cholinergic modulation of these interneurons is important for learning.

Acetylcholine has been shown to play a role in regulating the activity of neurons in the brain. Some studies suggest that acetylcholine release in specific brain regions or circuits leads to increased sparsity of neuronal firing patterns, which may enhance neural coding efficiency and improve cognitive function (Pinto et al., 2013; Goard and Dan, 2009). Consistent with these findings, we have demonstrated that the adaptive methods capable of facilitating learning produce sparse activations and increase the selectivity of neurons in the network. Although we thought the increased sparsity property would help mitigate catastrophic forgetting, a contrasting effect was seen. Contrary to our initial beliefs, sparse activations can also lead to a network becoming more prone to catastrophic forgetting as neurons becoming specialised for specific inputs can result in reduced generalisation and increased sensitivity to changes in the input distribution. This can make it harder for the network to adapt to new tasks without forgetting previous ones. Overall, the relationship between sparse activations and catastrophic forgetting is complex and depends on various factors which means we cannot be certain of the effects of this property.

We acknowledge that there is a prominent role for the cholinergic system in memory and we demonstrate that our implementation of the cholinergic system does not yield properties of reduced forgetting which we consider important for memory consolidation (Hasselmo, 2006). However, we are not implying that the cholinergic system is strictly participating in this form of adaptive learning. We are simply stating that this could be one of many functions of the cholinergic system. In fact, the framework we have proposed in which the cholinergic system receives and integrates error information from cortical networks enables it to freely process this information to achieve various functions. It would be interesting to explore a meta-learning approach in which the cholinergic system can be explicitly modelled as a separate neural network (Beaulieu et al., 2020). This provides an unbiased approach for optimising the required strategy to achieve known functions of the cholinergic system including efficient learning and mitigating catastrophic forgetting. For example, a similar approach to Andrychowicz et al. (2016) could be used to explicitly model the cholinergic system as a recurrent neural network that meta-learns to produce modulation factors which affect the learning rate as a function of the error information it receives. Unlike, the current adaptive methods, an architectural constraint can be applied to the network so that the modulation is diffuse. By adding physiological constraints to these models, the meta-learning process can act as the evolutionary process of the cholinergic system which can provide us with a theoretical understanding of how the cholinergic system may be acting to produce its effects.

2.6.1 Limitations

Although we have proposed a framework in which error signals are transmitted to the cholinergic system, it is unclear exactly how this could be achieved with the current EDN model. Additionally, the EDN model requires the SST interneurons to be one-to-one with the number of pyramidal cells in the next layer which is not equal to the relative proportions between interneurons and pyramidal cells observed in the brain. We attempt to address these issues with our new model in Chapter 3 which uses burst ensemble multiplexing to simultaneously communicate inference and error signals within spiking activity patterns (Naud and Sprekeler, 2018).

Our biological interpretation of how the cholinergic modulation can alter the learning rate of plasticity in our current model applies well to the cholinergic system as a global modulator. This is because a homogeneous signal allows the entire network to have access to the same modulatory information. With modulation at a finer resolution, the interneurons in one brain region would need access to the modulation factor in another brain region. This would imply that different parts of the microcircuit in one area receives different cholinergic signals and that acetylcholine release onto interneurons is identical to that of pyramidal cells in a upstream brain area.

2.6.2 Closing remarks

In conclusion, we have introduced a brain-wide deep learning model of cholinergic neuromodulation as an adaptive learning system that integrates error information to dynamically modulate synaptic plasticity. We have shown that this neuromodulation produces efficient learning when coupled with a biologically plausible credit assignment mechanism and applied in a similar manner to adaptive learning algorithms found in machine learning.

CHAPTER 3

SINGLE-PHASE DEEP LEARNING IN CORTICO-CORTICAL NETWORKS

3.1 Statement of contributions

This work was done in collaboration with a PhD student, Will Greedy, from the Department of Computer Science at the University of Bristol. A large amount of the content in this chapter is directly taken from our co-first author published paper "Single-phase deep learning in cortico-cortical networks" (Greedy et al., 2022). The spiking and continuous model data was predominantly produced by Will Greedy. The reinforcement learning data was predominantly produced by me. The analytical proof of BurstCNN approximating the backpropagation algorithm was mainly done by Will Greedy and Joseph Pemberton, another PhD student from the Department of Computer Science at the University of Bristol, with minor input from me. Will and I contributed equally to the rest of this project.

3.2 Introduction

For effective learning, synaptic modifications throughout the brain should result in improved behavioural function. This requires a process by which credit should be assigned to synapses given their contribution to behavioural output (Roelfsema and Holtmaat, 2018; Richards et al., 2019; Lillicrap et al., 2020). In multilayer networks, credit assignment is particularly challenging as the impact of changing a synaptic connection depends on its downstream brain areas. Classical local Hebbian plasticity rules, even when coupled with global neuromodulatory factors, are unable to communicate enough information for effective credit assignment through multiple layers of processing (Lillicrap et al., 2020). In machine learning, the error-backpropagation (backprop) algorithm is the most successful solution to the credit assignment problem. However, it relies on a number of biologically implausible assumptions to compute gradient information used for synaptic updates. Previous work has attempted to address these implausibilities but important issues remain open when mapping backprop to the neuronal physiology. Earlier attempts relied on using single-compartment neuron models (Lillicrap et al., 2016; Richards and Lillicrap, 2019) but this poses a problem as single-compartment neurons are unable to simultaneously store the necessary inference and credit assignment signals. One solution is to model neurons with an apical dendritic compartment that separately stores credit information (Sacramento et al., 2018; Richards and Lillicrap, 2019), supported by the electrotonic separation of the soma and apical dendrites (Williams and Stuart, 2002). These distal credit signals can then be communicated to the soma through non-linear dendritic events that trigger bursting at the soma (Xu et al., 2012), thereby inducing long-term synaptic plasticity (Sjostrom et al., 2008). In particular, two recent approaches, Error-encoding Dendritic Networks (EDNs) (Sacramento et al., 2018) and Burstprop (Payeur et al., 2021), have demonstrated how such multi-compartment neuron models can be used to construct networks capable of backprop-like credit assignment. EDNs encode credit signals at apical dendrites resulting from the mismatch between dendritic-targeting interneuron activity and downstream activity. Burstprop uses bursting, controlled by dendritic excitability, as a mechanism to communicate credit signals. However, these models still have major issues, such as the inability to effectively backpropagate error signals through many layers (EDNs) and the requirement for a multi-phase learning process (Burstprop).

Here, we propose a new model called the Bursting Cortico-Cortical Network (BurstCCN) as a solution to the credit assignment problem which addresses several outstanding issues of current biologically plausible backprop research. Our model builds upon prior multi-compartment neuron models (Sacramento et al., 2018; Payeur et al., 2021): it encodes credit signals in distal dendritic compartments which trigger bursting activity at the soma to drive backprop-like synaptic updates. We demonstrate that combining well-established properties of cortical neurons such as bursting activity, short-term plasticity (STP) and dendrite-targeting interneurons provides a biologically plausible mechanism for performing credit assignment. In contrast to previous models, BurstCCN is highly effective at backpropagating credit signals in multi-layer architectures while only requiring a single-phase learning process. We implement multiple versions of the BurstCCN at different levels of abstraction in order to demonstrate some of its key properties and to empirically confirm our theoretically motivated claims.

First, we use a spike-based implementation of the BurstCCN to demonstrate its ability to learn without the need for multiple phases. We further show the importance of this single-phase learning by training a continuous-time rate-based version of the BurstCCN on a continuous-time non-linear regression task. Next, we show both empirically and analytically that our model's dynamics result in learning that approximately follows backprop-derived gradients. Then, we use a simplified discrete-time BurstCCN implementation to demonstrate that the model achieves good performance on non-trivial image classification tasks (MNIST and CIFAR-10), even in the

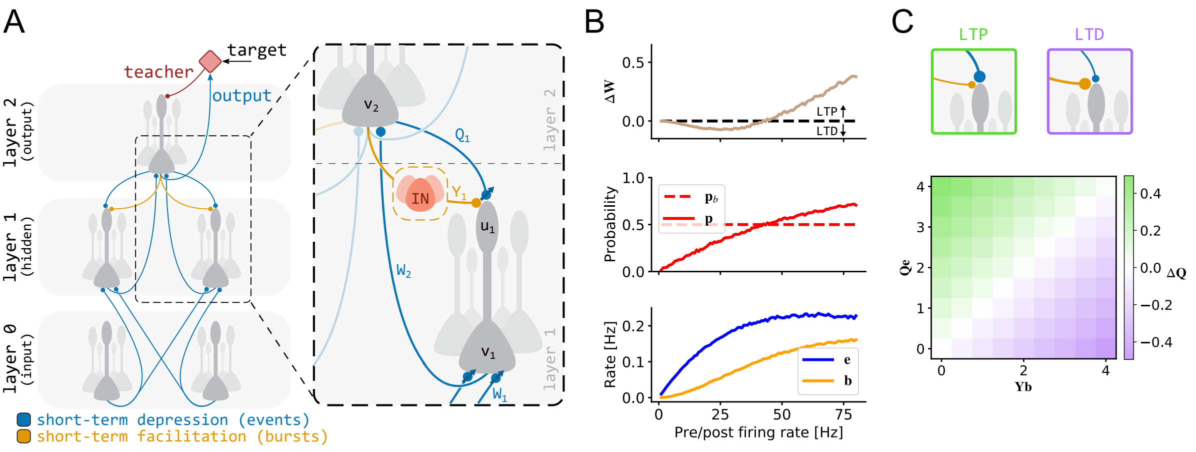

FIGURE 3.1: **Bursting cortico-cortical networks (BurstCCN) for credit assignment through bursting activity.** (**A**) Network schematic consisting of neuron ensembles and connection-type-specific STP. Events from the input are propagated forward through short-term depressing (STD) connections, **W**. Output event rates are compared to a target value which generates a teaching signal that is presented to the output layer apical dendrites. This acts as an error signal and appears as a deflection in the dendritic potential from its resting potential which causes changes to bursting activity from its baseline. The error-carrying bursting signals are propagated back through short-term facilitating connections, **Y**, which we interpret as being communicated by populations of dendrite-targeting interneurons. Events are also propagated backwards via STD connections, **Q**, to provide a means of cancelling baseline bursting activity. The difference in activity from these two feedback connections results in changes to dendritic excitability that lead to burst-dependent synaptic plasticity. (**B**) Burst-dependent plasticity rule. Simple setup of a single connection between a pre- and post-synaptic cell that are both modelled with Poisson spike trains with equal rates. As the firing rates increase, (top) plasticity of the synaptic weight switches from long-term depression (LTD) to long-term potentiation (LTP) (middle) when the burst probability increases above the baseline value. (bottom) The magnitude of the weight change is scaled by the event rate. (**C**) Homeostatic plasticity rule for **Q** weights. The difference between the signals through **Q** and **Y** dictates the direction and magnitude of synaptic plasticity.

presence of random feedback synaptic weights. We also demonstrate that the BurstCCN can be applied to a reinforcement learning (RL) context and performs similarly to artificial neural networks in the CartPole-v0 RL task. Finally, we demonstrate that by incorporating various types of interneurons and their respective STP connectivity, it is possible to adjust our model to adhere to Dale's law while retaining the capacity to learn MNIST.

3.3 Bursting Cortico-Cortical Networks

3.3.1 Burst Ensemble Multiplexing

Burst Ensemble Multiplexing (BEM) (Naud and Sprekeler, 2018) refers to the idea that ensembles of cortical neurons are capable of simultaneously representing multiple distinct signals within the patterns of their spiking activity. Typically, pyramidal cells receive top-down and bottom-up signals into their apical and basal dendrites, respectively. Bottom-up basal inputs affect the rate of spiking and top-down apical inputs convert these somatically induced spikes into high-frequency bursts. Postsynaptic populations can then use STP to decode these distinct signals from the overall spiking activity.

The BurstCCN uses the concept of BEM in a similar way to Burstprop (Payeur et al., 2021) in which ensembles of cells encode both feedforward inference signals and feedback error signals. The model encodes these signals as the rates of *events* and *bursts*, respectively, across the ensembles. Here, the specific definition of a burst is a collection of spikes with interspike intervals less than 16 ms and an event is either a burst or a single isolated spike (i.e. a spike not followed or preceded by another within 16ms). The burst probability of an ensemble is defined as the probability that an event at a given time is a burst and is computed as a ratio of the event rate ($\mathbf{e}$) and burst rate ($\mathbf{b}$): $\mathbf{p} = \mathbf{b}/\mathbf{e}$.

3.3.2 Rate-based BurstCCN

In our discrete-time implementation of the rate-based BurstCCN, example input-output pairs are processed independently in discrete timesteps. For each example, the event rates of the input layer, $\mathbf{e}_0$, encode the input stimulus. The model then computes each subsequent layer's activities, equivalent to that of a standard feedforward artificial neural network (Fig. 3.1A). Specifically, somatic potentials are computed by integrating basal input as $\mathbf{v}_l = \mathbf{W}_l \mathbf{e}_{l-1}$ where $\mathbf{W}_l$ are short-term depressing (STD) feedforward weights from layer $l-1$ to layer l. The STD nature of these weights ensures that only event rate information propagates forwards. Each layer's event rates are then computed by applying a non-linear activation function, f, to the somatic potentials, $\mathbf{e}_l = f(\mathbf{v}_l)$. These linear and nonlinear operations are repeated for each layer in the network to ultimately obtain the output layer event rates, $\mathbf{e}_L$, where L denotes the total number of layers.

The desired target output of the network, $\mathbf{e}_{target}$, is compared to the output layer event rates to produce a signed error, $\mathbf{e}_{target} - \mathbf{e}_L$, which is used as a teaching signal. This error information is then propagated backwards through each layer in the network by altering the apical dendritic compartment potential and, as a result, the burst probability of each pyramidal ensemble. At the output layer, the burst probability is computed directly as $\mathbf{p}_L = \mathbf{p}^b_L + \mathbf{p}^b_L \odot (\mathbf{e}_{target} - \mathbf{e}_L) \odot h(\mathbf{e}_L)$ where $\odot$ denotes the element-wise product, $\mathbf{p}^b_L$ represents the baseline burst probability in the absence of any teaching signal and $h(\mathbf{e}_l) = f'(\mathbf{v}_l) \odot \mathbf{e}_l^{-1}$. These burst probabilities are used at the output layer ($l = L$) to compute the burst rates as $\mathbf{b}_l = \mathbf{e}_l \odot \mathbf{p}_l$ which are decoded and sent backwards to layer $l-1$ apical dendrites by a set of short-term facilitating (STF) feedback weights, $\mathbf{Y}_{l-1}$. The STF feedback weights and STD feedforward weights are similarly used in Burstprop. However, the BurstCCN additionally includes a novel set of apical dendrite-targeting STD feedback weights, $\mathbf{Q}_{l-1}$, which send event rates backwards. We interpret the STF feedback connections as being provided via a type of dendrite-targeting interneuron and STD feedback as direct connections in line with recent experimental studies (Lee et al., 2013; Kinnischtzke et al., 2014; Petrof et al., 2015; Zolnik et al., 2020; Naskar et al., 2021; Martinetti et al., 2022). The signals through both sets of feedback weights lead to the apical potentials in the previous layer, $\mathbf{u}_{l-1} = \mathbf{Q}_{l-1}\mathbf{e}_l - \mathbf{Y}_{l-1}\mathbf{b}_l$. These determine the layer's burst probabilities which are computed as $\mathbf{p}_{l-1} = \bar{\sigma}(\mathbf{u}_{l-1} \odot h(\mathbf{e}_{l-1}))$ where $\bar{\sigma}$ denotes the sigmoid function, σ, with scaling and offset parameters, $\bar{\sigma}(x) = \sigma(\alpha x + \beta)$ (Payeur et al. (2021); see SM of Greedy et al. (2022)). The same process is repeated backwards for each layer until the input layer to obtain their dendritic potentials and burst probabilities. Note that for all experiments, we set $\alpha = 4$ (and $\beta = 0$) to prevent this function from implicitly scaling down the errors propagating backwards through each layer by a factor of 4 (since $\frac{d\sigma}{dx} \approx \frac{1}{4}$ around $x = 0$).

After the error information has been propagated backwards, feedforward synaptic weight changes are computed using a burst-dependent synaptic plasticity rule:

$$\Delta \mathbf{W}_l = \eta_l^{(\mathbf{W})}((\mathbf{p}_l - \mathbf{p}_l^b) \odot \mathbf{e}_l)\mathbf{e}_{l-1}^T \tag{3.1}$$

where $\eta^{(\mathbf{W})}$ is a learning rate and $\cdot^T$ is the transpose operation. Importantly, the learning rule depends on the change in burst probability from the predefined layer-wise baseline burst probability, $\mathbf{p}_l^b = p_l^b(1,\ldots,1)^T$, which represents the signed error signal required for backprop-like learning. Consequently, when we make both pre- and postsynaptic cells fire following Poisson statistics we obtain long-term depression and long-term potentiation for low and high firing rates, respectively (Fig. 3.1B). This is in line with a large number of experimental studies of cortical synapses (Bienenstock et al., 1982; Sjöström et al., 2001). It can be shown that the updates produced by this learning rule approximate those obtained by the backpropagation algorithm in the weak-feedback case (see Section 3.5.3.1 and SM of Greedy et al. (2022)).

In the absence of a teaching signal, it is important for pyramidal ensembles to produce a baseline level of bursting such that no weight changes occur (cf. Eq. 3.1). This holds true for the output layer as there are no other inputs onto the apical dendrites. However, for the hidden layers the event rate signals through $\mathbf{Q}$ and the burst rate signals through $\mathbf{Y}$ need to exactly cancel each other out such that the apical dendritic potentials are at rest (i.e. $\mathbf{u} = 0$). For any $\mathbf{Y}$ weights, there is always an optimal set of $\mathbf{Q}$ weights that will produce this exact cancellation regardless of the event rates propagating through the network. Specifically, they must be set as $\mathbf{Q}_l = p_l^b \mathbf{Y}_l$ which we refer to as the weights being in a *Q-Y symmetric* state. However, it is not biologically plausible for the $\mathbf{Q}$ synapses to have direct knowledge of $\mathbf{Y}$. Therefore, inspired by earlier work (Vogels et al., 2011; Sacramento et al., 2018), we use a learning rule for $\mathbf{Q}$ to provide this cancellation:

$$\Delta \mathbf{Q}_l = -\eta_l^{(\mathbf{Q})}\mathbf{u}_l\,\mathbf{e}_{l+1}^T \tag{3.2}$$

which explicitly aims to silence the apical potentials (Fig. 3.1C). In the absence of a teaching signal at the output layer, all $\mathbf{Q}$ weights will eventually converge to their optimal values and achieve a symmetric state under reasonable assumptions (see SM of Greedy et al. (2022)). Note that we can have add this learning rule on the **Y** feedback weights to cancel the activity through the **Q** weights, which produces similar results (Fig. 3.5).

When teaching signals are applied at the output layer, it is important to note that only the bursting activity propagated through the **Y** connections changes because the event rates through **Q** are unaffected by the dendritic activity. This enables single-phase learning as the symmetry in the two feedback connection types (**Q** and **Y**) can be exploited to directly compare *without teacher* signals (i.e. at baseline) to *with teacher* signals.

3.3.3 Continuous-time BurstCCN

In contrast to the discrete-time implementation, the continuous-time BurstCCN does not process each input example independently from one another at discretised timesteps. The network dynamics instead evolve through a continuous-time simulation incrementing with timesteps of $\Delta t = 0.1s$, where there is a memory of the prior network state at each timestep and no parallel processing of mini-batches. The signals into the input layer are now given by a time-varying function, $\mathbf{x}(t)$, along with target signals to the output layer, $\mathbf{y}(t)$. The input layer event rate is set instantaneously as $\mathbf{e}_0(t) = \mathbf{x}(t)$. All somatic potentials and other layer event rates evolve with:

$$\frac{d\mathbf{v}_l(t)}{dt} = \frac{1}{\tau_v}\Big(-\mathbf{v}_l(t) + \mathbf{W}_l(t)\mathbf{e}_{l-1}(t)\Big) \tag{3.3}$$

$$\mathbf{e}_l(t) = f(\mathbf{v}_l(t)) \tag{3.4}$$

where τ_v is the membrane leak time constant of the soma. The output layer burst probabilities are also set instantaneously as $\mathbf{p}_L(t) = \mathbf{p}_L^b + \mathbf{p}_L^b \odot (\mathbf{y}(t) - \mathbf{e}_L(t)) \odot h(\mathbf{e}_l(t))$. The hidden layer dendritic potentials, burst probabilities and burst rates evolve with:

$$\frac{d\mathbf{u}_l(t)}{dt} = \frac{1}{\tau_u}\Big(-\mathbf{u}_l(t) + h(\mathbf{e}_l(t)) \odot \Big[\mathbf{Q}_l(t)\mathbf{e}_{l+1}(t) - \mathbf{Y}_l(t)\mathbf{b}_{l+1}(t)\Big]\Big) \tag{3.5}$$

$$\mathbf{p}_l(t) = \sigma(\mathbf{u}_l(t)) \tag{3.6}$$

$$\mathbf{b}_l(t) = \mathbf{p}_l(t) \odot \mathbf{e}_l(t) \tag{3.7}$$

where τ_u is the membrane leak time constant of the apical dendrite. Finally, the feedforward weights change over time following:

$$\frac{d\mathbf{W}_l(t)}{dt} = \frac{1}{\tau_W}\Big(\Big[\mathbf{p}_l(t) - \mathbf{p}_l^b\Big] \odot \mathbf{e}_l(t)\Big)\mathbf{e}_{l-1}^T(t) \tag{3.8}$$

where τ_W is a time constant that determines the learning rate.

3.3.4 Spiking BurstCCN

Previous backprop-like learning models rely on computations across distinct phases to obtain the necessary error signals to learn (Guerguiev et al., 2017; Payeur et al., 2021). For example, Burstprop, proposed by Payeur et al. (2021), computes the difference between the level of bursting between two phases: an initial phase with no teaching signal and a second with a teaching signal present. The burst-dependent plasticity rule used in the spiking implementation of Burstprop requires neuron ensembles to compute a moving average of their burst probability over time. This quantity needs to reach stability in order to generate useful error signals. Therefore, there must be a period with no plasticity before a teaching signal is presented. This requires global plasticity switches between the two phases which is problematic because a high level of coordinated plasticity is needed across the whole network. To address these issues, we build upon Burstprop's plasticity rule and propose a spiking implementation of our rate-based burst-dependent plasticity rule (Eq. 3.1):

$$\frac{dw_{ij}}{dt} = \eta\Big[B_i(t) - P_i^b E_i(t)\Big]\tilde{E}_j(t) \tag{3.9}$$

where P^b is the baseline burst probability, $\tilde{E}$ is an eligibility trace of presynaptic activity (with time constant τ_{pre}), η is a learning rate and B_i and E_i are burst and event trains, respectively. In contrast to Burstprop, we use this constant baseline burst probability (P^b) instead of a moving average of the burst probability to control the relative magnitude of positive and negative weight changes. This makes it unnecessary to have a separate phase to compute its value since it is no longer time-dependent. However, to make single phase learning possible, we have also introduced a set of STD feedback weights that cancel baseline bursting feedback activity into apical compartments such that only deviations from the baseline are communicated. With just STF feedback weights, as is the case for Burstprop, cells will not consistently burst at baseline in the absence of a teaching signal. This is due to the bursting signals they communicate having a dependence on the feedforward activity (i.e. events) and necessarily varying across input stimuli.

We include this spike-based burst-dependent synaptic plasticity rule in a spiking implementation of the BurstCCN. Unlike the two rate-based implementations, the spiking BurstCCN more accurately models the internal neuron spiking dynamics instead of abstracting these details away and only considering the ensemble-level behaviour. Pyramidal neurons are modelled with two compartments corresponding to the soma and apical dendrites. The remaining details of this model along with the hyperparameters used are described by Payeur et al. (2021) and the simulation of all neurons was carried out using the Auryn simulator (Zenke and Gerstner, 2014).

3.3.5 Dalean BurstCCN

For the Dalean implementation of the rate-based BurstCCN, each connection type (**W**, **Q** and **Y**) was divided into separate excitatory and inhibitory pathways. For the feedforward **W** weights, the excitatory pathway was implemented as a direct excitatory connection, **A**, between pyramidal cells (PCs) in subsequent layers, whereas the inhibitory pathway, **B**, was implemented using a feedforward inhibitory interneuron population which we interpreted as being parvalbumin (PV) interneurons. Initially, the effective weight from neuron i to j, $\mathbf{W}_{ij}$, was initialised using a normal distribution and was split into the excitatory pathway and inhibitory pathway weights as follows:

$$\mathbf{A}_{ij} = \begin{cases} \mathbf{W}_{ij} + c & \mathbf{W}_{ij} > 0 \\ c & \mathbf{W}_{ij} \leq 0 \end{cases} \tag{3.10}$$

$$\mathbf{B}_{ij} = \begin{cases} -\mathbf{W}_{ij} + c & \mathbf{W}_{ij} < 0 \\ c & \mathbf{W}_{ij} \geq 0 \end{cases} \tag{3.11}$$

where c is a constant added to all weights to prevent initial starting weights of 0. The sign of excitatory, $\mathbf{A}_{ij}$, and inhibitory, $\mathbf{B}_{ij}$, weights are kept constant. Together, the combination of the excitatory and inhibitory weight pathways form the effective weight between two neurons, $\mathbf{W}_{ij} = \mathbf{A}_{ij} - \mathbf{B}_{ij}$. The feedforward interneuron population in the inhibitory pathway received event rates from PCs in the previous layer, $\mathbf{v}_l^{PV} = \mathbf{e}_{l-1}$ (the PC-to-PV weights were interpreted as the identity matrix), and inhibited current layer PCs by transmitting the negated form of this information, $-\mathbf{B}_l\mathbf{v}_l^{PV} = -\mathbf{B}_l\mathbf{e}_{l-1}$. The direct excitatory pathway transmitted event rates from one layer to the next as $\mathbf{A}_l\mathbf{e}_{l-1}$. Therefore, the somatic potentials of the PCs are computed by integrating the excitatory and inhibitory inputs as $\mathbf{v}_l = \mathbf{A}_l\mathbf{e}_{l-1} - \mathbf{B}_l\mathbf{v}_l^{PV} = \mathbf{A}_l\mathbf{e}_{l-1} - \mathbf{B}_l\mathbf{e}_{l-1} = \mathbf{W}_l\mathbf{e}_{l-1}$.

A similar approach to separate the excitatory and inhibitory pathways was undertaken for the feedback $\mathbf{Q}$ and $\mathbf{Y}$ weights. For the $\mathbf{Q}$ weights, a direct excitatory connection, $\mathbf{C}$ (defined similarly to $\mathbf{A}$ in Equation 3.10 but $\mathbf{W}_{ij}$ replaced with $\mathbf{Q}_{ij}$), was implemented to communicate event rates backwards to previous layer PCs. A feedforward inhibitory interneuron population which we interpreted as being neuron derived neurotrophic factor (NDNF) interneurons was used for the inhibitory pathway, $\mathbf{D}$ (defined similarly to $\mathbf{B}$ in Equation 3.11 but $\mathbf{W}_{ij}$ replaced with $\mathbf{Q}_{ij}$). The feedforward interneuron population in the inhibitory pathway received event rates from PCs in the next layer, $\mathbf{v}_l^{NDNF} = \mathbf{e}_{l+1}$ (the PC-to-NDNF weights were interpreted as the identity matrix), and inhibited current layer PCs by transmitting the negated form of this information, $-\mathbf{D}_l\mathbf{v}_l^{NDNF} = -\mathbf{D}_l\mathbf{e}_{l+1}$. The direct excitatory pathway transmitted event rates backwards to the current layer as $\mathbf{C}_l\mathbf{e}_{l+1}$. Therefore, the apical potential contribution are computed by integrating the excitatory and inhibitory inputs as $\mathbf{C}_l\mathbf{e}_{l+1} - \mathbf{D}_l\mathbf{v}_l^{NDNF} = \mathbf{C}_l\mathbf{e}_{l+1} - \mathbf{D}_l\mathbf{e}_{l+1} = \mathbf{Q}_l\mathbf{e}_{l+1}$.

For the $\mathbf{Y}$ weights, the excitatory pathway, $\mathbf{E}$, was implemented by the disinhibition of an inhibitory interneuron population which we interpreted as vasoactive intestinal peptide-expressing (VIP) interneurons disinhibiting somatostatin (SST) interneurons. The inhibitory pathway, $\mathbf{F}$, was implemented by a feedforward inhibitory interneuron population which we interpreted as a different population of SST interneurons (SST2). $\mathbf{E}$ and $\mathbf{F}$ weights are defined similarly to $\mathbf{A}$ and $\mathbf{B}$ in Equations 3.10-3.11 but $\mathbf{W}_{ij}$ replaced with $\mathbf{Y}_{ij}$. The feedforward interneuron population in the inhibitory pathway received burst rates from PCs in the next layer, $\mathbf{v}_l^{SST2} = \mathbf{b}_{l+1}$ (the PC-to-SST2 weights were interpreted as the identity matrix), and inhibited current layer PCs by transmitting the negated form of this information, $-\mathbf{F}_l\mathbf{v}_l^{SST2} = -\mathbf{F}_l\mathbf{b}_{l+1}$. The excitatory pathway consisted of VIP interneurons that received bursting information from PCs in the next layer, $\mathbf{v}_l^{VIP} = \mathbf{b}_{l+1}$, inhibiting downstream SST interneurons, $\mathbf{v}_l^{SST} = \mathbf{1} - \mathbf{v}_l^{VIP}$ (the constant of 1 is added as a tonic level of activation to provide SST activity for the VIP interneurons to inhibit). This SST interneuron population then inhibits the apical dendrite and the overall apical contribution is computed by integrating the disinhibitory and inhibitory inputs as:

$$\begin{aligned}
-\mathbf{E}_l\mathbf{v}_l^{SST} - \mathbf{F}_l\mathbf{v}_l^{SST2} &= -\mathbf{E}_l\left(\mathbf{1} - \mathbf{v}_l^{VIP}\right) - \mathbf{F}_l\mathbf{v}_l^{SST2} \\
&= -\mathbf{E}_l\left(\mathbf{1} - \mathbf{b}_{l+1}\right) - \mathbf{F}_l\mathbf{b}_{l+1} \\
&= \mathbf{E}_l\mathbf{b}_{l+1} - \mathbf{F}_l\mathbf{b}_{l+1} - \mathbf{E}_l\mathbf{1} \\
&= \mathbf{Y}_l\mathbf{b}_{l+1} - \mathbf{E}_l\mathbf{1}.
\end{aligned}$$

which is equivalent to the contribution of the $\mathbf{Y}$ feedback pathway in Section 3.3.2 with an additional term $-\mathbf{E}_l\mathbf{1}$. In this ideal case, we introduced a bias term, $\mathbf{E}_l\mathbf{1}$, that exactly cancels this extra term. We believe that in theory this could be learned by a separate cell population, although this was not explicitly modelled.

Together, the signals through the feedback weights lead to the apical potentials in the current layer, $\mathbf{u}_l = \mathbf{Q}_l\mathbf{e}_{l+1} - \mathbf{Y}_l\mathbf{b}_{l+1}$, which determine the layer's burst probabilities.

The learning rules for the excitatory pathway weights are the same as Equations 3.1-3.2 and are negated for the inhibitory pathway weights.

3.3.6 Related work

As previously mentioned, BurstCCN takes inspiration from two prior models: EDNs (Sacramento et al., 2018) and Burstprop (Payeur et al., 2021). Similar to these models, the BurstCCN uses a separate apical dendritic compartment to represent an error signal. To silence this apical compartment and maintain correct error signals, the EDN uses a homeostatic plasticity rule from local interneurons to cancel the signals received from a separate feedback pathway. In the BurstCCN, we use the same principle by adapting this plasticity rule for learning of the novel **Q** weights. Unlike the EDN, we use a similar idea to Burstprop in which error signals are encoded as bursts in the neural activity and decoded by STP dynamics.

Within each layer, Burstprop includes a set of recurrent connections onto the apical compartments which aim to maintain the dendritic potential in the linear regime of the feedback non-linearity. Updating the weights of these connections requires separate learning phases and it is unclear how the plasticity rule can be justified. In contrast, the BurstCCN does not require these connections. Instead, the novel set of STD feedback connections (**Q**) onto the apical dendrites provide a mechanism for single-phase learning and perform a similar role of linearising the feedback. Additionally, burst-dependent plasticity in our model relies on a constant baseline burst probability instead of using a moving average of the burst probability (Section 3.3.4).

3.4 Experimental details

In the following section, we provide additional details on the experimental setups used to obtain the results.

3.4.1 Spiking XOR task

In the spiking XOR task, we used networks consisting of 5 distinct neuron populations: 2 populations that encoded each input, 2 hidden layer populations and a single population that encoded the output. Each population consisted of 500 individual neurons. These neurons were sparsely connected in the feedforward and feedback directions with a connection probability of 0.05. All neurons received balanced excitatory and inhibitory Poisson noise into both their somatic and dendritic compartments.

In the two-phase setup, learning rates for the weights and biases were set to $\eta_W = 0.004$ and $\eta_B = 0.0001$, respectively. We reduced learning rates to $\eta_W = 0.0004$ and $\eta_B = 0.00001$ in the one-phase setup due to the increased duration where plasticity was on.

Each simulation was carried out for $16000s$ (2000 examples × $8s$ per example) with timesteps of $dt = 0.1s$. Learning of the feedforward weights was carried out by each model's burst-dependent synaptic plasticity rule with the time constant of the pre-synaptic input eligibility trace $\tau_{pre} = 0.1s$. The moving average time constant for Burstprop, τ_{ma}, was set to $2s$. Baseline burst probabilities in the BurstCCN were set to $P^b_{hidden} = 0.18$ and $P^b_{out} = 0.401$ in the hidden and output layers, respectively. The remaining task details can be found in (Payeur et al., 2021).

3.4.2 Continuous-time input-output task

This task was carried out using a 2-layer continuous-time BurstCCN with 3 inputs, 50 hidden units and a single output unit. The simulation was ran for 10^6 timesteps with $dt = 0.1s$ to give a total simulation time of $10^5 s$. During the first $100s$ of the simulation, plasticity and teaching signals were switched off to provide a means of comparing the initialised network to the trained network (Fig. 3.3C-D). The time constants used in the simulation for the somatic potentials, dendritic potentials and the synaptic weights were set to $\tau_v = 0.1s$, $\tau_u = 0.1s$ and $\tau_W = 100.0s$, respectively.

3.4.3 CartPole-v0 reinforcement learning benchmark

To train on the CartPole-v0 RL task from the OpenAI Gym environment (Brockman et al., 2016), we used a standard deep Q-network (DQN) model and a BurstCCN equivalent which involved substituting the conventional ANN in the DQN model with the BurstCCN (Mnih et al., 2013). The same network structure was used for both model types: 4 inputs, 2 hidden layers with 50 units and 2 output units. We used a sigmoid for the feedforward activation functions and a batch size of 512. The models were trained with momentum and Table 3.1 shows the best hyperparameters found for each model. To enhance training stability, we employed a large experience replay buffer of 50000 samples and updated the target network every 100 steps (Mnih et al., 2015). Once the experience replay buffer captured 50000 samples, the main network started updating every step.

Due to the use of sigmoidal activations, the output range was restricted to 0 and 1. Therefore, to fit the rewards into the range of a sigmoid, we scaled them such that the maximum total episodic reward was 0.6. The network was trained to produce values in the range of 0.2 and 0.8 which corresponded to the predicted value with an offset of 0.2.

Model type	**Learning rate (η)**	**Momentum**
ANN	1.2	0.99
BurstCCN	0.008	0.99

TABLE 3.1: CartPole-v0 hyperparameters

3.4.4 Rate-based model experiments

3.4.4.1 Hyperparameter search

Bayesian hyperparameter optimisation was performed (using Weights & Biases (Biewald, 2020)) for the shared hyperparameters of each rate-based model shown in Tables 3.2-3.3 and 3.5-3.6. This was also performed for the model-specific parameters that are stated in the following relevant sections.

3.4.4.2 Feedback plasticity on Q weights

As described in Section 3.5.3.2, this task involved fixing the **W** and **Y** weights and allowing the **Q** weights to learn following the plasticity rule given in Equation 3.2. Throughout training, independent Gaussian noise was added to the event rates propagating forwards to give somatic potentials, $\mathbf{v}_l = \mathbf{W}_l(\mathbf{e}_{l-1} + \xi)$ where $\xi \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$. This noise was introduced to decorrelate neural activities in each layer to facilitate the alignment of **Q** to **Y** (see alignment proof in SM of Greedy et al. (2022)). For all configurations, we set $\sigma = 0.1$ during training but removed the noise when evaluating the apical potential magnitudes (Fig. 3.4B,F) and alignment angles (Fig.3.4C-D,G-H). Similarly, in the without-teacher setups, teaching signals were added only for the evaluation of the alignment angles to backprop or feedback alignment. Table 3.2 shows the best hyperparameters for each setup optimised for **Q**-**Y** alignment.

3.4.4.3 MNIST

For MNIST, we used standard ANN, BurstCCN, Burstprop and EDN models. The same network structure was used for all model types: 784 inputs (28×28 grayscale images), n number of hidden

Feedback mode	Teacher	Q learning rate ($\eta^{(\mathbf{Q})}$)	Y init. scale ($\sigma_{\mathbf{Y}}$)	Q init. scale ($\sigma_{\mathbf{Q}}$)
Random (FA)	No	0.00520	0.5	0.0148
	Yes	0.00380	0.5	0.0488
Symmetric	No	0.00337	N/A	0.0075
	Yes	0.00677	N/A	0.0735

TABLE 3.2: Q-weight learning hyperparameters

layers with 500 units and 10 output units. In all cases, we used a sigmoid for the feedforward activation functions and a batch size of 32. Table 3.3 shows the best hyperparameters found for the 5-layer networks of each model which were shared across all network sizes.

BurstCCN The additional $\mathbf{Q}$ weights were trained with a learning rate of 3.5×10^{-5} and the baseline burst probability $\mathbf{p}_l^b = 0.5$ for all layers.

Burstprop The additional recurrent weights were initialised from a Gaussian distribution (with $\mu = 0$ and $\sigma = 1.3 \times 10^{-4}$) and trained with a learning rate of 2.6×10^{-4}. The baseline burst probability at the output layer was set to 0.2.

EDN Learning rates for the interneuron-to-pyramidal ($\eta_{l,l}^{\mathrm{PI}}$) and pyramidal-to-interneuron ($\eta_{l,l}^{\mathrm{IP}}$) weights were set proportionally to the feedforward learning rates ($\eta_{l+1,l}^{\mathrm{PP}}$) as $\eta_{l,l}^{\mathrm{PI}} = \eta_{l,l}^{\mathrm{IP}} = \alpha\eta_{l+1,l}^{\mathrm{PP}}$ with the proportionality constant $\alpha = 0.643$.

Model Type	Learning rate (η)	Y init. scale ($\sigma_{\mathbf{Y}}$)	Momentum	Weight decay
ANN	0.201	1.49	0.474	1.09×10^{-9}
BurstCCN	0.0246	0.638	0.836	4.01×10^{-10}
BurstCCN (**Q**-**Y** sym)	0.0246	0.638	0.836	4.01×10^{-10}
Burstprop	0.481	1.71	0.347	7.59×10^{-6}
EDN	0.00578	2.60	0.198	2.01×10^{-7}

TABLE 3.3: MNIST hyperparameters

3.4.4.4 CIFAR-10

For CIFAR-10, we used convolutional ANN, BurstCCN and Burstprop models with three convolutional layers followed by two fully-connected layers (Table 3.4). In all cases, we used a sigmoid for the feedforward activation functions and a batch size of 32. Table 3.5 and Table 3.6 show the best hyperparameters shared across the different models in the random fixed **Y** and **W**-**Y** symmetric feedback weight regimes, respectively.

BurstCCN The additional **Q** weights were trained with a learning rate of 2.21×10^{-3} and the baseline burst probability $\mathbf{p}_l^b = 0.5$ for all layers.

Burstprop Under the random fixed **Y** feedback weight regime, the additional recurrent weights were initialised from a Gaussian distribution (with $\mu = 0$ and $\sigma = 2.0 \times 10^{-4}$) and trained with a learning rate of 2.0×10^{-4}. Under the **W**-**Y** symmetric weight regime, the additional recurrent weights were initialised from a Gaussian distribution (with $\mu = 0$ and $\sigma = 2.5 \times 10^{-2}$) and trained with a learning rate of 2.2×10^{-5}. For both weight regimes, the baseline burst probability at the output layer was set to 0.2.

Layer number	**Layer type**	**Size**
0	Input	$32 \times 32 \times 3$
1	Convolutional	5×5, 64, stride=2
2	Convolutional	5×5, 128, stride=2
3	Convolutional	3×3, 256, stride=1
4	Fully-connected	1480 neurons
5	Fully-connected	10 neurons

TABLE 3.4: CIFAR-10 architecture

Model Type	**Learning rate (η)**	**Y init. scale ($\sigma_{\mathbf{Y}}$)**	**Momentum**	**Weight decay**
ANN	0.00335	0.363	0.502	5.96×10^{-8}
BurstCCN	0.0343	0.424	0.0142	1.01×10^{-10}
BurstCCN (**Q**-**Y** sym)	0.0132	0.290	0.578	2.65×10^{-9}
Burstprop	0.102	0.981	0.698	3.91×10^{-9}

TABLE 3.5: CIFAR-10 hyperparameters (random fixed **Y** feedback weight regime)

Model Type	**Learning rate (η)**	**Momentum**	**Weight decay**
ANN	0.0814	0.447	1.60×10^{-6}
BurstCCN (**Q**-**Y** sym)	0.155	0.496	1.50×10^{-8}
Burstprop	0.146	0.869	2.53×10^{-8}

TABLE 3.6: CIFAR-10 hyperparameters (**W**-**Y** symmetric weight regime)

3.4.5 MNIST with Dalean model

Bayesian hyperparameter optimisation was performed (using Weights & Biases Biewald (2020)) and best hyperparameters are shown in Table 3.7. We used the Dalean BurstCCN network architecture of 784 inputs (28×28 grayscale images), 6 hidden layers with 500 units and 10 output units. Also, we used a sigmoid for the feedforward activation functions and a batch size of 32.

Hidden layer learning rate (η)	Output layer learning rate (η)	Q learning rate ($\eta^{(\mathbf{Q})}$)	Feedback scale
0.185	0.189	0.000731	0.644

TABLE 3.7: Dalean BurstCCN hyperparameters for MNIST

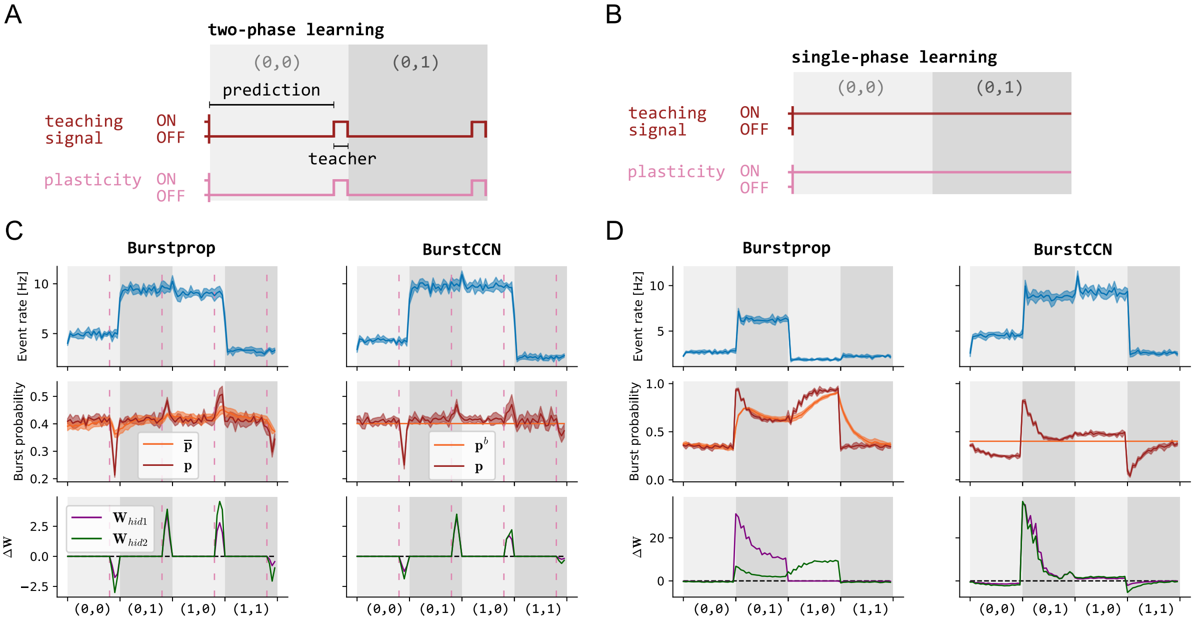

FIGURE 3.2: **Spiking BurstCCN does not require multi-phase learning to solve the XOR classification task.** Schematic of the (A) two-phase and (B) single-phase learning settings. (**A**) For each input during two-phase learning, networks are given a 7.2s prediction period during which teaching signals and plasticity are turned OFF, followed by a 0.8s learning period where both teaching signals and plasticity are turned back ON. (**B**) During single-phase learning, both the teaching signals and plasticity remain ON throughout training. (**C, D**) Top: event rate (**e**) of the output layer. Middle: burst probability (**p**) for the output layer and the baseline or moving average of the burst probability ($\mathbf{p}^b$ or $\bar{\mathbf{p}}$) for BurstCCN and Burstprop, respectively. Bottom: the resulting weight updates for connections from hidden layer neurons. Model results represent mean ± standard error (n = 5).

3.5 Results

3.5.1 BurstCCN can learn with a single learning phase

A key motivation for developing the BurstCCN was to design a model capable of learning without the need for separate learning phases, while being consistent with a range of cortical features across multiple levels. To demonstrate that our model can perform single-phase learning, we trained the spiking version of our model on the XOR classification task and contrasted it with Burstprop, which requires a two-phase learning process (Fig. 3.2). In both single- and two-phase learning regimes, the input stimulus is presented for a total of 8s before the next example is shown. The two-phase learning regime has an initial prediction phase, lasting 7.2s for each input presentation, where plasticity is switched off throughout the network and the output neurons do not receive any teaching signals (Fig. 3.2A). This is followed by a teacher phase for the remaining 0.8s where plasticity is restored and teaching signals are delivered at the output. The single-phase regime removes the initial prediction phase and extends the teacher phase to the full duration of the input stimulus (Fig. 3.2B).

Our results show that both models were capable of successfully learning the task in the two-phase regime as indicated by the high output event rates in response to the $(0,1)$ and $(1,0)$ inputs and low event rates for the $(0,0)$ and $(1,1)$ inputs (Fig. 3.2C). However, when training in the single-phase regime, only BurstCCN was able to learn the task (Fig. 3.2D). The inability of Burstprop to learn the task can be explained by comparing the moving average of the burst probability ($\overline{\mathbf{p}}$) with the actual burst probability ($\mathbf{p}$) which determines the sign of synaptic weight updates (Fig. 3.2D). Burstprop failed to learn in the single-phase learning setup due to the teaching signal remaining on and preventing $\overline{\mathbf{p}}$ from being able to provide a stable representation of the without-teacher burst probability.

3.5.2 BurstCCN can learn with dynamic input-output

Typically, studies that have attempted to solve the credit assignment problem with biologically plausible implementations of backprop make an implicit assumption that during learning there is a period where the continuous-time input stream is fixed (Sacramento et al., 2018; Payeur et al., 2021). This is required in most cases to allow the network to stabilise its activities before learning can take place. With single-phase learning, we can relax this assumption to enable learning in conditions where the inputs and their corresponding teaching signals are dynamically changing over time. We assessed this ability by training the continuous-time BurstCCN (Section 3.3.3) on an online non-linear regression task (Fig. 3.3). This task consisted of three sinusoidal inputs, $x_i(t) = sin(\alpha_i t + \beta_i)$, with random frequencies $\alpha_i \sim U(0, \frac{\pi}{2})$ and phase offsets $\beta_i \sim U(0, 2\pi)$ (Fig. 3.3A). The network had a single output unit for which a non-trivial target was obtained by passing the same inputs to a 3-25-1 artificial neural network (ANN). This approximates a setting in which a given cortical area learns to regress its input onto the activity of another cortical area. The ANN weights were randomly initialised with $w^1_{ij} \sim U(-\sqrt{3}, \sqrt{3})$ for the first layer and $w^2_{ij} \sim U(-0.6, 0.6)$ for the second layer. Despite the BurstCCN initially producing outputs that were significantly different to the target (Fig. 3.3C), the results show that over training it learned to produce output patterns that closely matched the non-linear and dynamic target (Fig. 3.3B,D). This highlights that the BurstCCN is capable of adequately backpropagating useful error signals when both inputs and teaching signals are constantly changing.

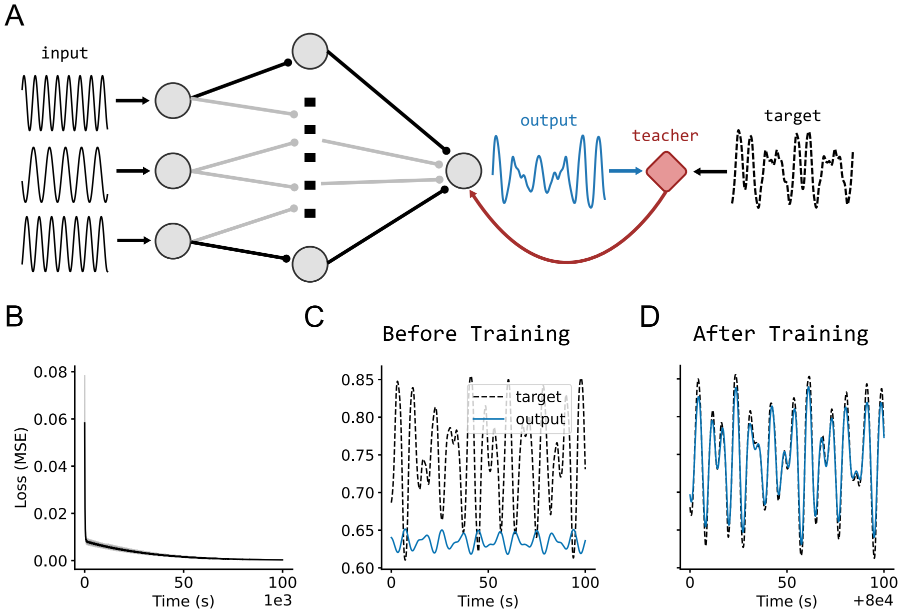

FIGURE 3.3: **BurstCCN can learn a dynamic non-linear regression task.** (**A**) Schematic of the task. Three sinusoidal waves with random frequencies are given as inputs. The task is to learn to match the target pattern which is obtained by passing the same inputs through a fixed, randomly initialised ANN. (**B**) Learning curve for the (continuous-time) BurstCCN. (**C, D**) Example output traces for (C) before and (D) after training. Model results represent mean ± standard error (n = 5).

3.5.3 Feedback plasticity rule facilitates alignment to backprop updates

Next, we wanted to understand how well our model approximates backprop. As stated above, the purpose of the learning rule for the feedback STD **Q** connections (Eq. 3.2) is to silence the apical compartments in every ensemble by cancelling activity through the feedback STF **Y** connections. When a teaching signal is applied, this becomes important for computing the correct local error signal that is used for learning and backpropagated to previous layers. Here, we show both analytically and empirically using the discrete version of the model how the computed errors relate to backprop.

3.5.3.1 BurstCCN with weak feedback approximates backpropagation algorithm

Under some small assumptions, we analytically show that the feedback pathway of BurstCCN is approximately communicating the same error gradients that are computed by backprop. Specifically, we assume that the feedback weights are optimally aligned (i.e. $\mathbf{Q}_l = p_l^b \mathbf{Y}_l$) and focus on the change in burst rate, $\delta \mathbf{b}_l := (\mathbf{p}_l - \mathbf{p}_l^b) \odot \mathbf{e}_l$. If we let $E^{\text{task}} = ||\mathbf{e}_L - \mathbf{e}_{target}||^2$ define the task error then, by construction, the change in burst rate at the output layer is equivalent to the negative error gradient, $\delta \mathbf{b}_L = -\frac{\partial E^{\text{task}}}{\partial \mathbf{v}_L}$. For the hidden layers, we derive the following iterative relationship (see SM of Greedy et al. (2022) for further details):

$$\delta \mathbf{b}_l = f'(\mathbf{v}_l) \odot (-\mathbf{Y}_l) \delta \mathbf{b}_{l+1} + \mathcal{O}(\mathbf{u}_l^3). \tag{3.12}$$

This approximates the same relationship present in backprop up to a third-order[1] term with respect to the apical potentials $\mathbf{u}_l$ if the feedback weights are set to be symmetric with the feed-forward weights (i.e. $\mathbf{Y}_l = -\mathbf{W}_{l+1}^T$). We refer to this as the *W-Y symmetric* state. The link between weight updates from simply performing gradient descent with backprop and the BurstCCN can be seen clearly:

$$\Delta \mathbf{W}_l^{\text{BurstCNN}} = \eta_l^{(\mathbf{W})} \delta \mathbf{b}_l \, \mathbf{e}_{l-1}^T \tag{3.13}$$

$$\Delta \mathbf{W}_l^{\text{backprop}} = -\eta_l^{(\mathbf{W})} \frac{\partial E^{\text{task}}}{\partial \mathbf{v}_l} \mathbf{e}_{l-1}^T \tag{3.14}$$

It remains to be shown that the apical potentials, $\mathbf{u}_l$, of every layer are indeed appropriately small (so that the approximation error, $||\mathbf{u}_l^3||$, is small). Under the assumption $\mathbf{u}_{l+1}$ is small, we can derive the recursive relationship $\mathbf{u}_l \approx f'(\mathbf{v}_{l+1}) \odot (-\mathbf{Y}_l)\mathbf{u}_{l+1}$ (see SM of Greedy et al. (2022) for further details). We show that if f' is bounded (as is the case for sigmoid and many activation functions) and the weights $\mathbf{Y}_l$ are reasonably small then $||\mathbf{u}_l^3|| \leq ||\mathbf{u}_{l+1}^3||$. This means that if the error gradient at the output layer, $\frac{\partial E^{\text{task}}}{\partial \mathbf{e}_N}$, is small then, by induction, $||\mathbf{u}_l^3||$ is small for every layer and $\Delta \mathbf{W}_l^{\text{BurstCNN}} \approx \Delta \mathbf{W}_l^{\text{backprop}}$.

3.5.3.2 Learning Q feedback connections better approximates backprop-derived gradients

We empirically evaluated our feedback plasticity rule by updating *only* the **Q** weights of a randomly initialised 5-layer discrete-time BurstCCN with all other weight types (**W** and **Y**) fixed. We used multiple initialisations and training regimes to understand how the plasticity rule behaves in different scenarios. The network was either initialised in the **W**-**Y** symmetric state or with random feedback weights (where $\mathbf{Y}_l \neq -\mathbf{W}_{l+1}^T$). We computed the angle between the update that would have been made by the feedforward plasticity rule (Eq. 3.1) and either backprop or feedback alignment (Lillicrap et al., 2016) for the symmetric and random configurations, respectively. We examined both cases: in the theoretically ideal case for learning **Q** where no teaching signal is present (Fig. 3.4A-D) and with a teaching signal at the output layer (Fig. 3.4E-H).

In all cases, as the alignment between the **Q** and **Y** connections improved (Fig. 3.4A,E), the apical potential decreased (Fig. 3.4B,F) and this resulted in updates that more closely aligned to backprop (Fig. 3.4C,G) and feedback alignment (Fig. 3.4D,H). In the absence of a teaching signal, this alignment angle to both backprop and feedback alignment eventually became very small which supports our analytical results that show our model approximates these methods (Fig. 3.4C-D). Despite producing less aligned feedforward updates in the presence of a teaching signal, the updates computed were still informative since they were consistently well below 90° of the direction of steepest descent (Fig. 3.4G).

---

[1] Here we use abuse of notation $\mathbf{u}_l^3 = \left(u_{l,1}^3, u_{l,2}^3, \dots\right)^T$ to represent the element-wise cubic of $\mathbf{u}_l$

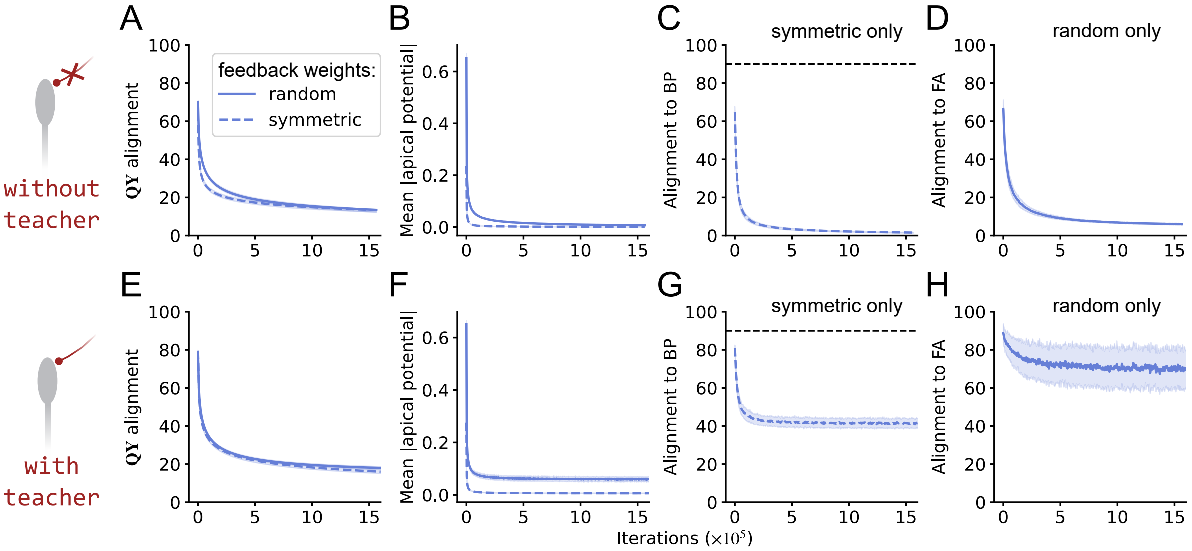

FIGURE 3.4: **Feedback learning rule enables a close alignment with backprop and feedback alignment.** The network is a randomly initialised 5-layer discrete-time BurstCCN with random (solid line) or symmetric (dashed line), fixed **W** and **Y** weights. The **Q** weights are updated in the presence of **(A-D)** no teaching signal or **(E-H)** a teaching signal. (A,E) Alignment between **Q** and **Y** weights, (B,F) the mean absolute value of the apical potentials, (C,G) the alignment to backprop (BP) and (D,H) feedback alignment (FA) as the **Q** weights learn to silence apical dendrite potential. Updates below 90° marked by the black dashed line are considered useful as they still follow the direction of backprop on average. Model results represent mean ± standard error (n = 5).

3.5.3.3 Feedback plasticity on Y-weights

The feedback plasticity rule used on the **Q** weights could likewise be applied to the **Y** weights to produce a similar effect. This would be in line with long-term synaptic plasticity observations of dendrite-targeting interneurons (Chiu et al., 2018).

$$\Delta \mathbf{Y}_l = \eta_l^{(\mathbf{Y})} \mathbf{u}_l \, \mathbf{e}_{l+1}^T \tag{3.15}$$

In an identical experimental setup to the one used for investigating the feedback plasticity rule on the **Q** weights (Section 3.5.3.2), we examined the impact on the alignment to backprop when the feedback plasticity rule was instead applied onto the **Y** weights. In all cases, as the alignment between the **Q** and **Y** connections improved (Fig. 3.5A,E), the apical potential decreased (Fig. 3.5B,F) and this resulted in updates that more closely aligned to backprop (Fig. 3.5C,G) and feedback alignment (Fig. 3.5D,H). In the absence of a teaching signal, this alignment angle to both backprop and feedback alignment eventually became very small (Fig. 3.5C-D). Despite producing less aligned feedforward updates in the presence of a teaching signal, the updates computed were still informative since they were consistently well below 90° of the direction of steepest descent (Fig. 3.5G). Overall, this suggests that the feedback plasticity rule on the **Q** weights is equally effective when applied to the **Y** weights.

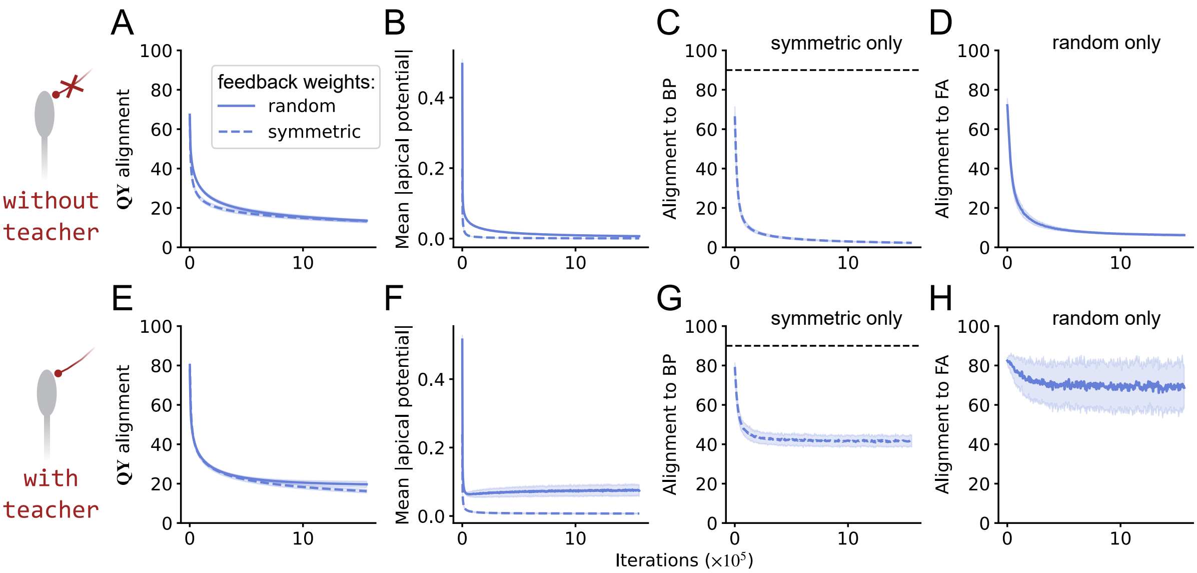

FIGURE 3.5: **Feedback plasticity rule on Y weights facilitates updates to align with backprop and feedback alignment.** The network is a randomly initialised 5-layer discrete-time BurstCCN with random (solid line) or symmetric (dashed line), fixed **W** and **Y** weights. The **Y** weights are updated in the presence of (**A-D**) no teaching signal or (**E-H**) a teaching signal. (A,E) Alignment between **Q** and **Y** connections, (B,F) the mean absolute value of the apical potentials, (C,G) the alignment to backprop (BP) and (D,H) feedback alignment (FA) as **Y** weights learn to silence apical dendrite potential. Updates below 90° marked by the black dashed line are considered useful as they still follow the direction to backprop on average. Model results represent mean ± standard error (n = 5).

3.5.4 BurstCCN learns image classification tasks with multiple hidden layers

3.5.4.1 MNIST

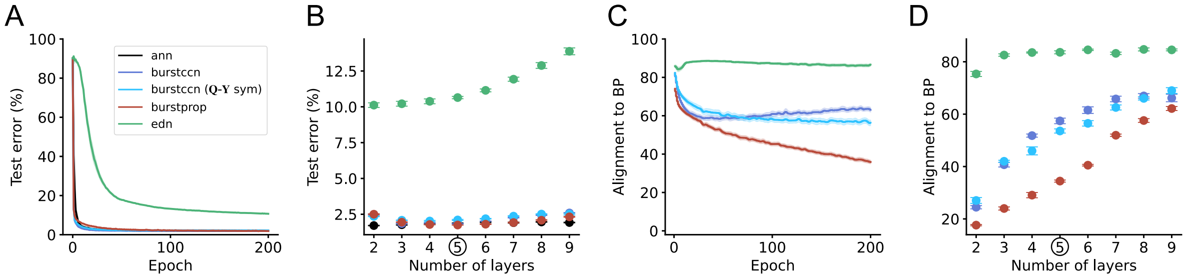

FIGURE 3.6: **BurstCCN learns to classify handwritten digits (MNIST) with deep networks.** (**A**) Learning curve of 5-layer ANN (black), BurstCCN (blue), BurstCCN ($\eta^{(\mathbf{Q})} = 0$) (light blue), Burstprop (red) and EDN (green). (**B**) Test error with different numbers of hidden layers for all models. (**C**) Alignment to backprop (BP) over time for all 5-layer models. (**D**) Alignment to backprop with different numbers of hidden layers for all models. The black circle indicates that the hyperparameters for each model were optimised for 5-layer networks. Model results represent mean ± standard error (n = 5).

Next, to test whether our model can indeed perform backprop-like deep learning, we trained a number of (discrete-time) BurstCCN architectures on the MNIST handwritten digit classification task (LeCun and Cortes, 2010). We compared the BurstCCN with Burstprop (Payeur et al., 2021) and EDNs (Sacramento et al., 2018) using similar architectures (Section 3.4.4.3). We focused on the more biologically plausible case of using random fixed feedback weights (i.e. feedback alignment (Lillicrap et al., 2016); see Section 3.5.5 for symmetric feedback weight case) with the remaining connection types of the different models updated using their respective plasticity rules. We also tested the BurstCCN in its idealised case where the feedback STD weights (**Q**) were fixed in the **Q**-**Y** symmetric state (see Section 3.3.2). We denote this model as "BurstCCN (**Q**-**Y** sym)".

Using 5-layer networks, the BurstCCN obtained a test error of 1.84±0.01%, comparable to that of Burstprop with 1.75±0.01% and significantly outperforming the EDN with 10.65±0.09% (Fig. 3.6A). As the network depth was increased, both BurstCCN and Burstprop retained high performances but the EDN showed a substantial decay in performance with deeper networks (Fig. 3.6B). We then compared the alignment between the models and backprop. For the 5-layer networks, Burstprop's updates were most closely aligned to backprop, followed by the two BurstCCN models which all vastly outperformed the EDN (Fig. 3.6C). As expected, the BurstCCN with **Q**-**Y** symmetry could better propagate error signals. By increasing the network depth, we demonstrate that it was more difficult to produce updates that were closely aligned to backprop. However, we show that the BurstCCN was still capable of backpropagating useful error signals in relatively deep networks (Fig. 3.6D).

3.5.5 Symmetric MNIST

Here, we trained the same models used in Section 3.5.4.1 under the **W**-**Y** symmetric weight regime. This regime is implausible due to sharing of weights but we investigated this to isolate the ability of each network to backpropagate errors in the ideal setting that most resembles backprop.

Using 5-layer networks, BurstCCN (**Q**-**Y** sym) obtained a test error of 1.86±0.03% similar to that of an ANN (1.97±0.04%) and outperforming both Burstprop (2.79±0.03%) and EDN (12.28±0.26%) (Fig. 3.7A). However, the BurstCCN with **Q** weight learning achieved a performance of 5.27±0.07%. This was lower than the random fixed **Y** weight regime (Fig. 3.6) due to the increased difficulty of aligning with the rapidly moving **Y** weights as the network learns. As the network depth was increased, both BurstCCN and Burstprop retained high performances but the EDN showed a substantial decay in performance with deeper networks (Fig. 3.7B). We then compared the alignment between the models and backprop. For the 5-layer networks, BurstCCN consistently produced updates that aligned with backprop at around 50° (Fig. 3.7C). As expected, the BurstCCN with **Q**-**Y** symmetry was significantly better at propagating error signals. With increased network depths, we demonstrate that it was more difficult to produce updates that were closely aligned to backprop. However, we show that the BurstCCN was still capable of backpropagating useful error signals in relatively deep networks (Fig. 3.7D).

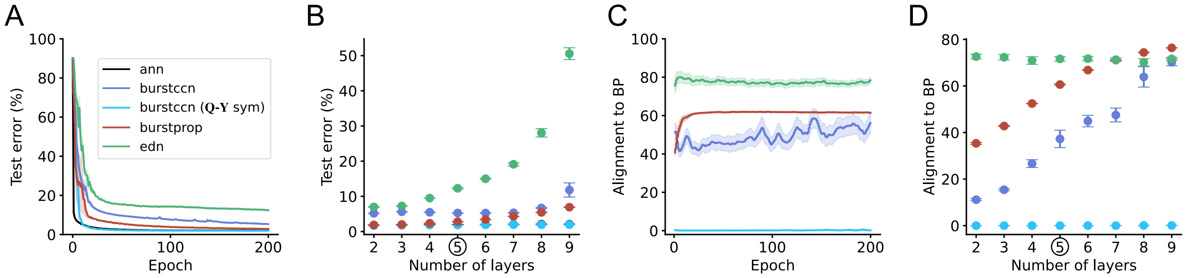

FIGURE 3.7: **BurstCCN with W-Y symmetric weights.** (**A**) Learning curve and (**C**) alignment to backprop of 5-layer BurstCCN (blue), BurstCCN ($\eta^{(\mathbf{Q})} = 0$) (light blue), Burstprop (red) and EDN (green). (**B**) Different number of hidden layers across all models. (**D**) Alignment to backprop (BP) across number of hidden layers. The black circle indicates that the hyperparameters for each model were optimised for 5-layer networks. Model results represent mean ± standard error (n = 5).

3.5.5.1 CIFAR-10

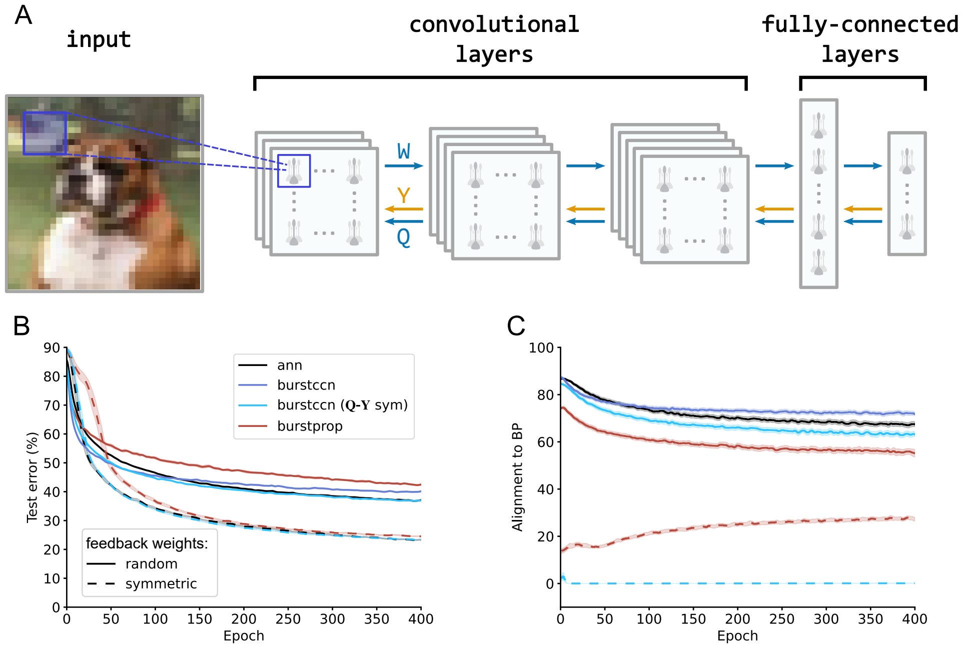

FIGURE 3.8: **BurstCCN with convolutional layers learns to solve natural image classification task (CIFAR-10).** (**A**) Schematic of BurstCCN architecture consisting of an input layer, three convolutional layers, a fully-connected hidden layer and output layer. For the BurstCCN, each layer was connected with a set of feedforward weights, **W**, and feedback weights, **Y** and **Q**. (**B**) Learning curve and (**C**) alignment to backprop of the different models with random (solid lines) and symmetric (dashed lines) feedback weight regimes. Model results represent mean ± standard error (n = 5).

Finally, we wanted to investigate the capabilities of the BurstCCN on more challenging tasks that are commonly tested in deep learning. We constructed a deep network consisting of three convolutional layers followed by a fully-connected hidden layer and output layer (Fig. 3.8A). We trained ANN, BurstCCN and Burstprop models using this network architecture on the CIFAR-10 image classification task (Krizhevsky et al., 2014). BurstCCN (**Q**-**Y** sym) was trained in the **Q**-**Y** symmetric regime whereas BurstCCN was initialised in this state and **Q** weights were then updated using the corresponding plasticity rule. All model types were tested with two feedback weight regimes: **W**-**Y** symmetric and random fixed **Y** feedback weights (i.e. feedback alignment).

After training in the random feedback weight regime, we observed a test error of 38.99±0.18% for BurstCCN, similar to performances achieved by an ANN (36.30±0.16%) and Burstprop (41.32±0.14%) (Fig. 3.8B). For the **W**-**Y** symmetric regime which most resembles backprop, BurstCCN (22.92±0.03%) performed significantly better than all random feedback setups and, once again, obtained a similar error to the symmetric ANN (22.62±0.10%) and Burstprop (24.15±0.17%) models. In the symmetric setups, there was a large improvement in the alignment angles to backprop compared to the random feedback setup (Fig. 3.8C). This suggests that they were backpropagating errors more effectively which likely explains the increase in performance. However, as seen within the random feedback setups, an improvement in this alignment does not guarantee an improvement to performance. This is because each model will traverse a different learning trajectory and converge to a different local minimum but the alignment angle remains a good indicator of expected performance.

3.5.6 Deep reinforcement learning with BurstCCN

Thus far, we have solely presented results for BurstCCN regarding supervised machine learning tasks, which prompts the question of whether explicit supervised learning signals occur in the brain. Although complex spikes found in the cerebellum provide evidence for the existence of these signals, it is still uncertain if they are present in other regions of the brain (Hull, 2020). While we do not directly answer this question, our objective is to exhibit that supervised learning signals are not necessarily required for effective learning in the BurstCCN. To achieve this, we incorporate the BurstCCN into a reinforcement learning (RL) scenario that only necessitates globally available reward signals. We integrate the BurstCCN into a Q-learning framework by substituting a conventional ANN with the BurstCCN in a standard deep Q-network (DQN) approach (Mnih et al., 2013). Q-learning is a popular method in RL that uses the concept of the temporal difference (TD) error to update their estimates of the state-action value function (Q-function), which is used to evaluate the expected future rewards of taking an action in a given state. The TD error measures the error between the predicted and actual outcomes in a sequence of states and actions. The dopamine reward system in the brain is believed to behave in a similar manner by using reward prediction errors to signal discrepancies between expected and actual rewards received. The dopamine reward prediction error can be seen as a biological implementation of the TD error, providing a plausible mechanism for how the brain learns from rewards (Daw and Tobler, 2014; Starkweather and Uchida, 2021; Tsutsui-Kimura et al., 2020;

Glimcher, 2011).

To investigate whether the BurstCCN can learn in an RL context with global reward signals, we chose to evaluate its performance on the CartPole-v0 RL benchmark from the OpenAI Gym environment (Fig. 3.9A) (Brockman et al., 2016). The goal of the task is to balance a pole upright on a cart by moving the cart left or right along a one-dimensional track. An episode ends when the agent succeeds in preventing the pole from falling over for 200 timesteps or fails if it can no longer keep it upright. The task is solved when the average total timesteps for the episode reaches 195 over 100 consecutive trials. For a fair comparison, we used a 3-layer network architecture comprising of sigmoidal units to train both an ANN and the BurstCCN. Both models received the state of the environment which consists of four continuous variables: the position and velocity of the cart, and the angle and angular velocity of the pole. Our results demonstrate that, although it takes longer than an ANN, the BurstCCN is capable of solving the CartPole-v0 task (Fig. 3.9B-C).

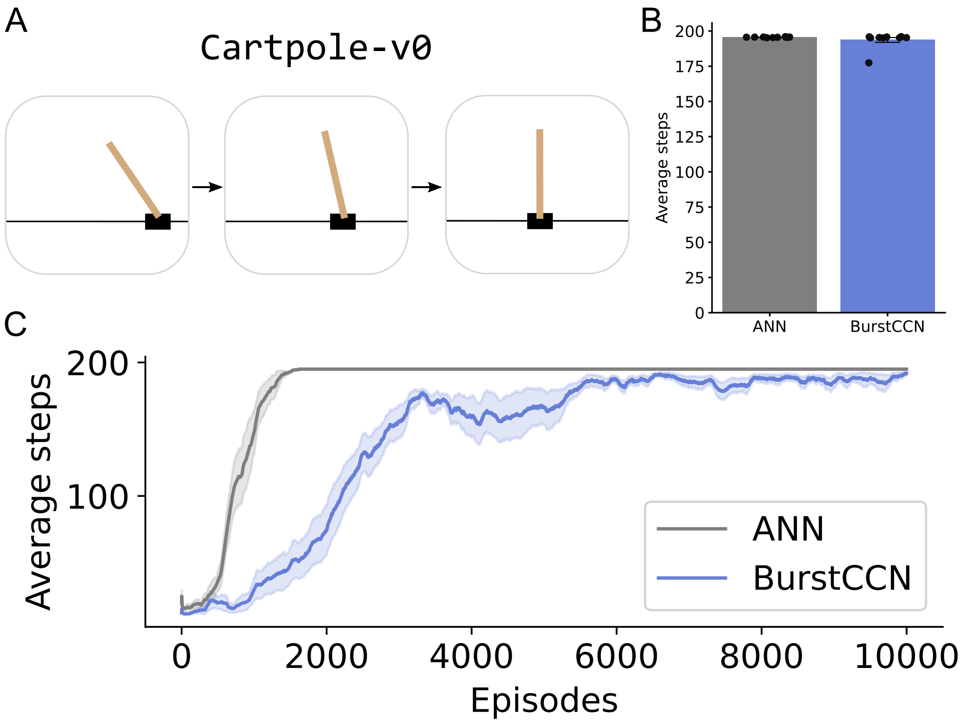

FIGURE 3.9: **BurstCCN learns a policy to perform well on the CartPole-v0 reinforcement learning benchmark.** (**A**) Schematic of the CartPole-v0 task: the agent is the cart moving left and right to keep the pole upright. A reward is given at every timestep the pole remains upright for a maximum of 200 timesteps. (**B**) The moving average of the number of timesteps the agent keeps the pole upright after 10000 episodes. A moving average of 195 timesteps is considered solving the task. (**C**) The learning timecourse for the BurstCCN (blue) and a standard artificial neural network (grey). Model results represent mean ± standard error (n = 10).

3.5.7 Dalean BurstCCN learns MNIST

In neuroscience, Dale's law is a principle that states a neuron releases the same neurotransmitters at all of its postsynaptic connections regardless of the nature of the target neuron (Strata et al., 1999). According to this, a presynaptic neuron can only exert either an excitatory or an inhibitory effect on its postsynaptic partners, but not both simultaneously. Most ANNs and all previously described BurstCCN implementations violate this architectural constraint as it impairs the learning performance of these models. In these networks, neurons consist of a mixture of positive and negative outgoing synaptic weights which can change their sign after initialisation. This is one of the most notable disconnects between the observed physiological architecture of neural circuits and the design of ANNs in the field of deep learning (Cornford et al., 2020). Recent efforts have been made to reconcile this disparity by demonstrating it is possible to modify neural networks to comply with Dale's law without undermining their learning potential (Payeur et al., 2021). Here, we have made comparable alterations to the architectural connectivity of the BurstCCN, introducing distinct inhibitory units into both the feedforward and feedback pathways (Fig. 3.10A). We have proposed the presence of certain interneuron subtypes in each pathway based on experimental observations of their cortical STP connectivity (Lee et al., 2013; Kinnischtzke et al., 2014; Petrof et al., 2015; Zolnik et al., 2020; Naskar et al., 2021; Martinetti et al., 2022). Our results show that this model is capable of producing updates that consistently align within 90° of backprop-derived gradients and after training, it learns to classify MNIST handwritten digits with a small error of $4.23 \pm 0.558\%$ (Fig. 3.10B-C).

3.6 Discussion

We have introduced a new model to address the credit assignment problem by integrating known properties of cortical networks to capture the essence of backprop. We have shown that by combining the biological mechanisms of namely bursting, STP and dendrite-targeting inhibition it is possible to construct a model that effectively learns in a continuous, online setting reminiscent of learning in the brain. We have demonstrated that such a model is capable of learning complex image classification tasks from machine learning that many current models of approximate backprop struggle to solve. In addition, we have shown that the model does not require explicit supervised learning signals by integrating it into a RL framework. Furthermore, even when we impose architectural constraints to comply with Dale's law, the BurstCCN is still capable of learning to perform well on MNIST. The ability of BurstCCN to perform single-phase learning has largely been made possible by the introduction of an additional set of $\mathbf{Q}$ feedback connections. These connections removed the requirement for temporally non-local computations by providing a method of representing non-teacher bursting activity for the generation of error signals which approximate backprop-derived gradients.

Our model proposes specific STP dynamics on the feedforward and feedback connections. It

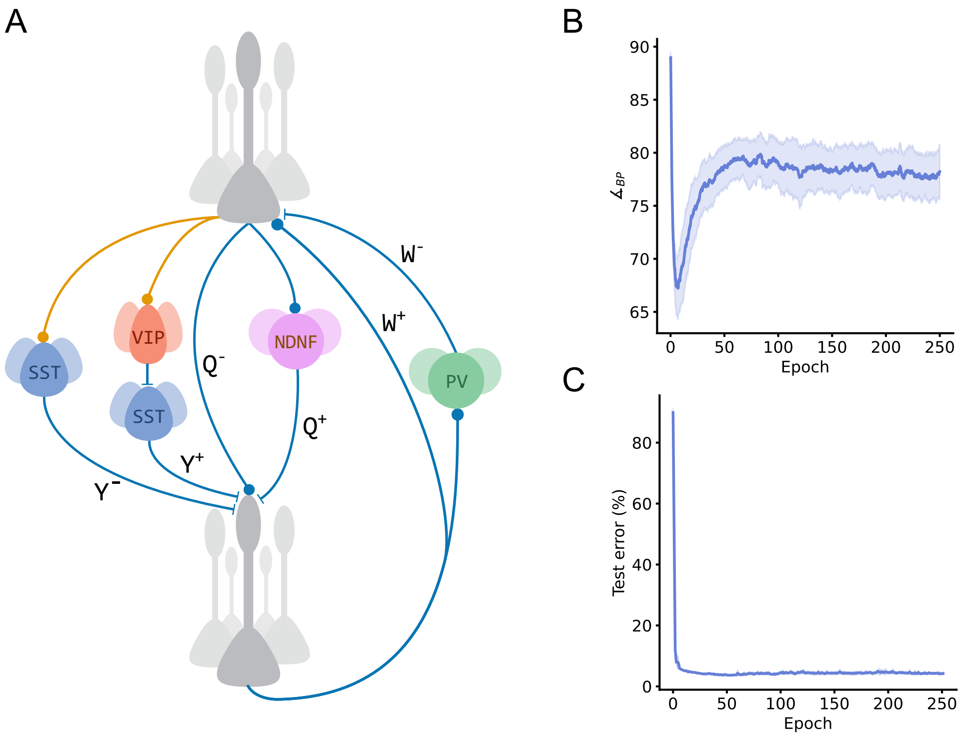

FIGURE 3.10: **BurstCCN with Dalean constraints learns to classify handwritten digits (MNIST).** (**A**) Schematic of the Dalean BurstCCN with proposed interneuron populations: parvalbumin (PV), somatostatin (SOM), vasoactive intestinal peptide (VIP) and neuron-derived neurotrophic factor (NDNF). (**B**) Alignment to backprop (BP) over time. (**C**) The learning timecourse of the model. Model results represent mean ± standard error (n = 5).

requires STD on cortico-cortical projections onto pyramidal cells in line with experimental evidence (Lee et al., 2013; Kinnischtzke et al., 2014; Petrof et al., 2015; Zolnik et al., 2020; Naskar et al., 2021). In addition, it suggests a key role for dendrite-targeting interneurons such as SST-positive Martinotti cells in the feedback pathway. There is evidence that these interneurons receive STF top-down connections whereas top-down projections onto pyramidal cells exhibit STD dynamics as required by our model (Lee et al., 2013; Kinnischtzke et al., 2014; Petrof et al., 2015; Zolnik et al., 2020; Naskar et al., 2021; Martinetti et al., 2022). Although we have introduced independent inhibitory units into our Dalean implementation of the BurstCCN, the specific interneuron subtypes proposed are not explicitly modelled. Therefore, it would be interesting to see the impact of adding this extra complexity on the ability of the model to learn and to further increase biological plausibility.

A prediction from our model is that manipulations of interneurons with STF connections would lead to disruptions in burst decoding from the layer (brain area) above thereby obstructing learning in the brain area below. Additionally, as error signals alter the level of bursting in the network, the model predicts that the variance in bursting activity and the distal dendritic potentials would correlate with the severity of errors made by the network during learning.

The current model does not include a proposed function for the commonly observed local feedback element of dendrite-targeting SST interneurons. In the EDN, this feature was critical for the generation and backpropagation of error signals but here we have incorporated the SST interneurons in a different mechanism. As this part of the circuitry is currently not modelled, it provides an opportunity to explore the possible benefits it could functionally contribute to the model. As these SST interneurons will be receiving activity and feeding this back onto local neurons, they have the necessary information and connectivity to be able to implement a concept from machine learning known as layer normalisation (Ba et al., 2016). Layer normalisation ensures that all neurons in a given layer have a consistent activity distribution across all features for a particular input. In deep ANNs, this technique has been shown to be effective in improving training and generalisation. In terms of biological plausibility, layer normalisation can be thought of as a homeostatic process to stabilise and normalise the range of network activities which is typically found in the brain to prevent the onset of neurological disorders such as epileptogenesis. For the BurstCCN, this could provide network stability by discouraging the saturation of neurons which is detrimental to error backpropagation.

As we have established here, the BurstCCN provides improvements on previous biologically plausible models of backprop including the model used in the previous chapter. Therefore, it would be interesting to examine the benefits of integrating the BurstCCN with the ChoCA framework from Chapter 2.

3.6.1 Limitations

Although our model captures a wide range of biological features, some biological implausibilities remain. Currently, we use feedback alignment to provide a solution to the *weight transport problem* (Crick, 1989) but this has a substantial impact on performance, particularly in more challenging tasks. Therefore, it would be important to explore some of the recently introduced plausible feedback learning rules (Lee et al., 2015; Akrout et al., 2019; Ahmad et al., 2020; Kunin et al., 2020) which could be used in conjunction with our proposed learning rules to outperform feedback alignment (Lillicrap et al., 2016).

Sigmoidal units are used in the BurstCCN to represent the physiological firing of neurons but this becomes an issue for regression tasks with output targets outside the range of the sigmoid activation function (0 and 1). Therefore, the model is restricted to RL tasks with a finite time horizon to allow for a bounded total episodic reward. We can then scale this bounded total reward to fit within the range of the sigmoid activation function otherwise our network saturates.

Another limitation in the current RL setup is the DQN approach. Although it has proved to be a powerful tool for RL, it relies on several assumptions that are not consistent with physiology of the brain. One assumption is that the learning process requires two neural networks: a main network and a target network. These two networks have the same architecture and different weights but after a certain amount of time, the weights from the main network are copied to the target network. It is unclear in the brain where such a target network would exist. Additionally, the networks do not update in an online fashion after each time step due to stability: it waits and updates its weights based on samples from a memory buffer that stores and replays previous experiences. In contrast, the brain continuously updates and integrates new experiences into existing knowledge.

3.6.2 Closing remarks

Our model provides a novel solution to the credit assignment problem that has demonstrated improvements from previous models alike including Burstprop and the EDN (the model used in Chapter 2). Overall, this work suggests that a range of cortical features from sub-cellular to the systems level jointly underlie single-phase, efficient deep learning in the brain.

CHAPTER 4

CHOLINERGIC MODULATION OF OLM INTERNEURONS AND THEIR REGULATION OF LTP

4.1 Introduction

With frequent exposure to novelty, it is imperative that the brain is capable of rapidly adapting in order to execute appropriate responses to novel situations. Neuromodulators are released under these conditions to provide the necessary regulation of network activities for synaptic plasticity to occur. Acetylcholine (ACh) is one of the major neuromodulators released during novelty that evokes functional reconfiguration of neural networks to create a window for learning (Palacios-Filardo et al., 2021). The importance of ACh in learning and memory has been widely examined (Hasselmo, 2006; Teles-Grilo Ruivo and Mellor, 2013; Giocomo and Hasselmo, 2007). Physical and chemical lesions of the cholinergic neuromodulatory system impairs learning in a myriad of memory tasks (Hasselmo, 2006). ACh facilitates the induction of synaptic plasticity via the activation of nicotinic and muscarinic ACh receptors (nAChRs and mAChRs, respectively). The intracellular signalling pathways downstream of these receptor subtypes act to up- or down-regulate the activity of both GABAergic and glutamatergic neurons (Picciotto et al., 2012). In the hippocampus, the activation of muscarinic M1 AChRs directly excites CA1 pyramidal cells (PCs) and induces long-term potentiation (LTP) of their glutamatergic inputs (Dennis et al., 2016; Buchanan et al., 2010; Dasari and Gulledge, 2011). In addition, information processing in the hippocampal circuits is regulated by the differential inhibitory actions of ACh primarily on feedforward parvalbumin (PV) and cholecystokinin (CCK) interneurons in the Schaffer collateral (SC) and temporammonic (TA) pathways, respectively (Palacios-Filardo et al., 2021; Cea-del Rio et al., 2010, 2011; Szabó et al., 2010; Raza et al., 2017).

Another class of interneurons in the hippocampus affected by cholinergic neuromodulation are the feedback, dendrite-targeting interneurons commonly identified by the expression of the neuropeptide somatostatin (SST) (Pelkey et al., 2017). A study by Lovett-Barron et al. (2014) demonstrated that the activity of the SST interneurons was crucial for contextual fear conditioning and pharmacological intervention revealed that cholinergic inputs were responsible for driving their activity. The SST molecular marker encompasses a diverse population of interneurons; the SST interneurons Lovett-Barron et al. (2014) eluded to were the dendrite-targeting Oriens-Lacunosum Moleculare (OLM) cells. This major subclass of hippocampal interneurons is heterogeneous to other hippocampal SST interneuron populations primarily due to their exclusive expression of the $\alpha 2$ subunit of nAChRs (Chrna2) (Mikulovic et al., 2015). OLM interneurons have been shown to control synaptic plasticity in the hippocampus by disinhibiting and inhibiting the proximal and distal dendritic compartments of CA1 PCs, respectively (Leão et al., 2012). These cells receive direct cholinergic input and it is known that mAChR modulation increases their excitability (Pelkey et al., 2017). However, there is no direct evidence to suggest that this increase in excitability translates to an increase in inhibition at the SST/OLM-to-PC synapse. Furthermore, the involvement of the inhibitory output from OLM interneurons in regulating associative long-term potentiation (aLTP) remains unclear.

Therefore, investigating the impact of ACh on the functional output of these dendrite-targeting interneurons is crucial for understanding how cholinergic modulation regulates synaptic plasticity. To examine the effects of ACh on the SST/OLM-to-PC synapse and the regulation of plasticity by the OLM interneurons, *ex vivo* whole-cell patch clamp electrophysiology experiments were carried out in acute hippocampal slices. Carbachol (CCh), an acetylcholine (ACh) receptor agonist, was found to reduce the inhibitory activity of OLM and SST interneurons onto CA1 pyramidal cells despite increasing their intrinsic excitability. Furthermore, OLM interneurons inhibited the induction of aLTP in the SC pathway and facilitated long-term depression (LTD) in the TA pathway.

4.2 Materials and methods

4.2.1 Materials

All reagents were purchased from Fisher Scientific unless otherwise stated.

4.2.2 Animals

Animal care and use was in accordance with the UK Animals (Scientific Procedures) Act of 1986 and were approved by the University of Bristol ethical review committee. Experiments were performed on male and female mice aged 4 to 15 weeks. C57BL/6J (Charles River) mice were used as the background strain. Cre reporter allele mice (The Jackson Laboratory) were used to tag specific neuronal populations: somatostatin (SST) interneurons (SST-Cre; MGI: 4838416) and oriens-lacunosum moleculare (OLM) interneurons (Chrna2-Cre; MGI: 5490768). These Cre reporter mice were crossed with homozygous Ai32 (MGI: 5577173) to generate litters of heterozygous offspring expressing channelrhodopsin (ChR2). Chrna2-Cre reporter mice were also crossed with homozygous Ai14 (MGI: 3817869) to generate litters of heterozygous offspring expressing TdTomato. Mice were housed under regulated temperature (22 ± 2°C) in a normal 12 h light/dark cycle with ad libitum access to food and water.

4.2.3 Hippocampal slice preparation

Following cervical dislocation and decapitation, mouse brains were removed and dissected in ice-cold cutting solution (in mM: 205 sucrose, 10 glucose, 26 $NaHCO_3$, 2.5 KCl, 1.25 NaH_2PO_4, 0.5 $CaCl_2$, 5 $MgSO_4$) saturated with 95% O_2 and 5% CO_2. 400 µm thick parasagittal slices were prepared with a Leica VT1200 vibratome (Leica, Germany) and slices were incubated in artificial cerebrospinal fluid (aCSF; in mM: 124 NaCl, 10 glucose, 24 $NaHCO_3$, 3 KCl, 1.25 NaH_2PO_4, 2.5 $CaCl_2$, 1.3 $MgSO_4$) at 34°C for 25 min and then stored at room temperature for at least 30 min before use.

4.2.4 Electrophysiology

In all electrophysiology experiments, slices were transferred to a submerged recording chamber perfused with aCSF at a flow rate of 2.5 ml/min and kept at 32°C via an inline heating system (TC-324, Warner Instruments). The constant flow and removal of aCSF was maintained by a 101U/R peristaltic pump (Watson Marlow) and a Dymax 5 vacuum pump (Charles Austin), respectively. Target cells were visualised with infrared differential interference contrast optics on an Olympus BX-50WI microscope. Patch electrodes with a resistance of 3-7 MΩ were pulled from borosilicate filamented glass capillaries (1.5 OD x 0.86 ID x 100 L mm, Harvard Apparatus) with a horizontal puller (P-87, Sutter Instrument Co., UK).

Recordings from hippocampal neurons were made with a MultiClamp 700A amplifier (Molecular Devices, USA), filtered at 2.4 kHz and digitised at 10 kHz with a CED Power1401 data acquisition board and Signal 5.12 acquisition software (CED, Cambridge, UK). Cells were voltage clamped at 0 or –70 mV (not corrected for junction potential). Series resistance was monitored throughout the experiments and cells that showed >40% change were discarded from subsequent analysis. Recordings were also rejected from analysis if the series resistance was greater than 35 MΩ.

Excitatory synaptic responses were evoked in independent pathways with two 2 ms square voltage steps applied at 20 Hz through bipolar stimulating electrodes located in the Schaffer collateral (SC) pathway from CA3 and the temporammonic (TA) pathway at the stratum lacunosum-moleculare (SLM). Whilst, inhibitory synaptic responses were evoked optogenetically with five

2 ms square light steps applied at 10 Hz via a mounted 470 nm LED illuminator (M470L2 with a LEDD1B driver, Thorlabs, US) through a 40x objective lens.

4.2.4.1 Carbachol dose-response

In ChR2-expressing Chrna2-Cre (Chrna2-Cre x Ai32) and SST-Cre (SST-Cre x Ai32) mice, whole-cell patch-clamp recordings were made in CA1 pyramidal neurons with patch pipettes filled with an intracellular solution containing (in mM) 130 $CsMeSO_3$, 4 NaCl, 10 HEPES, 0.5 EGTA, 10 TEA, 1 QX-314 chloride, 2 Mg-ATP, 0.5 Na-GTP. Cells were voltage clamped at 0 mV, far from the chloride reversal potential, to maximise the amplitude of inhibitory postsynaptic current (IPSC) recordings. A stable baseline period of 10 min was achieved before the consecutive application of carbachol (CCh) at 0.1, 0.3, 1, 3, 10 and 30 μM for 10 min periods each. The 0.1 μM dose was excluded in experiments targetting the OLM interneurons. The rise time of the IPSC was calculated as the time it took for the response to reach 80% of its maximum amplitude from 20%. The tau decay time constant was calculated by fitting a monoexponential curve to the waveform after the max amplitude had been reached. The dose-response curved was fit using a four parameter logistic curve. The paired pulse ratio (PPR) was calculated by dividing the ampltiude of the second peak by the amplitude of the first peak (P2/P1). Data points and bar plots represent the mean ± standard error of the last 3 mins of each drug concentration.

4.2.4.2 Somatostatin interneuron excitability

In slices prepared from SST-Cre x Ai32 animals, current clamp recordings were made from SST cells located in the stratum oriens (SO) with patch pipettes filled with an intracellular solution containing (in mM) 135 K-gluconate, 10 HEPES, 8 NaCl, 2 Na-ATP, 0.3 Mg-GTP, 0.2 EGTA and adjusted to pH 7.4, 390-395 mOsm. Once a cell was patched, SST cells were identified by their step-like response to a 100 ms blue light pulse. The resting membrane potential of the cell had to be at least –55 mV or lower for the experiment to begin. Action potentials were evoked in the cell with 1.5 s square current steps that started at –50 pA and increased in +10/20 pA increments up to a maximum of +250 pA. However, to prevent inducing excessive damage to the cell, if the current step caused the membrane potential of the cell to exceed –30 mV then no further steps were delivered to the cell. These recordings were made without CCh, in the presence of 10 μM CCh and if the resting membrane potential increased after drug application, the cell was repolarised back to baseline and another set of recordings were made in the presence of 10 μM CCh. Spikes were counted in Python 3.9 using the SciPy package and results were presented as the average number of spikes elicited using a current step that produced approximately half of the maximum number of spikes during the baseline (without drug) period.

4.2.4.3 Associative LTP

In slices prepared from Chrna2-Cre x Ai32 animals, voltage clamp recordings were made from CA1 pyramidal neurons with patch pipettes filled with an intracellular solution containing (in mM) 120 KMeSO3, 10 HEPES, 0.2 EGTA, 4 Mg-ATP, 0.3 Na-GTP, 8 NaCl, 10 KCl and adjusted to pH 7.4, 280-285 mOsm. These experiments were carried out in parasagittal slices to maximise the density and connectivity of temporammonic fibres in the SLM. Cells were voltage clamped at –70 mV to obtain excitatory currents. Postsynaptic currents were recorded in three independent pathways: SC test pathway, TA test pathway and SC control pathway. A steady 5 min baseline was achieved before attempting to induce plasticity via the application of a theta burst protocol (TBP) to the two test pathways. The TBP consisted of a train of 10 bursts where each burst contained 5 pulses at 100 Hz with the frequency of bursts set at 5 Hz. The TBP was applied through the two test pathway stimulating electrodes either in combination with light pulses following the same pattern, to activate the ChR2-expressing OLM interneurons, or in the absence of light. During the TBP, cells were in the current-clamp configuration and were injected with a current that initially held their resting membrane potentials at -70 mV before the TBP. The waveforms from the activities of these cells during this period were also recorded. The protocol was performed within 10 min of reaching the whole-cell configuration to prevent washout of plasticity. After TBP, responses from both pathways were recorded for a further 25 min. Consecutive traces were averaged to produce a mean response every 30 seconds. The mean amplitude response of the baseline period was used to normalise the responses of each pathway. Plasticity was assessed by comparing the average ESPC amplitude during the last 5 min of the experiment between the control and test pathways.

The waveforms during the TBP were analysed for number of spikes, number of bursts and the total area underneath the waveform. Spikes were counted in Python 3.9 using the SciPy package. A burst was counted for each theta-burst period with 2 or more spikes. The total area underneath the waveform was calculated first by detecting and removing spikes in the waveform; fitting a biexponential curve to capture the fast and slow components of an afterhyperpolaristion in order to estimate the baseline; then taking the area between the fitted baseline curve and the waveform itself.

4.2.5 Immunohistochemistry

Prior to immunohistochemical staining of target neurons, mice brains were fixed via perfusion fixation. First, mice received an intraperitoneal injection of sodium pentobarbital at a lethal dose of 100 mg/kg. After the cessation of breathing and reflexive responses from foot pinching, the abdomen was opened and the rib cage was separated to expose the heart. A small incision was made in the right atrium to allow blood and old solution to drain whilst 20 ml of 0.01 M phosphate-buffered saline (PBS) was being injected into the left ventricle followed by another

20 ml of 4% paraformaldehyde in 0.01 M PBS (PFA). The mice were then decapitated, the skulls were cut open following the midline sutures (sagittal and interfrontal sutures) and the brains were extracted. Before slicing, the brains were submerged in 4% PFA for 24 hours and then transferred to 30% sucrose in 0.01 M PBS until they sank to the bottom of the container which took approximately 48 hours. 40 µm thick parasagittal slices were obtained using a freezing microtome and stored in 0.01 M PBS with 0.01% sodium azide.

For immunohistochemistry, free-floating slices were washed with gentle agitation in 0.01 M PBS for 10 min and three additional times (10 min each) in 0.01 M PBS with 0.2% Triton X-100 (0.2% PBST). Slices were then incubated in a blocking solution containing 5% donkey serum (Sigma-Aldrich D9663-10ML) in 0.3% PBST for 1 hour at room temperature. Slices were subsequently incubated at 4°C for 24 hours in blocking solution containing either anti-SST (1:1000 Santa Cruz SC-7819) or anti-GFP (1:1000 Abcam ab13970) which are antibodies for visualising SST and ChR2, respectively. After 24 hours, slices were brought up to room temperature for 1 hour. Following three washes with 0.2% PBST (10 min each), slices were incubated in blocking solution containing either anti-goat Alexa Fluor 488 (1:500, Invitrogen Antibodies A-11055) or anti-goat Alexa Fluor 594 (1:500, Invitrogen Antibodies A-11058) and anti-chicken Alexa Fluor 488 (1:500, Jackson ImmunoResearch 703-545-155) for 2.5 hours at room temperature. Finally, slices were washed twice in 0.2% PBST and once in 0.01 M PBS (10 min each), before being stained with DAPI (1:10000, Invitrogen D1306), mounted on microscope slides and covered with rectangular glass cover slips using Fluorsave (Merck Millipore).

4.2.6 Microscopy

Sections were examined using a Leica DMI6000B microscope with a HC PL FLUOTAR 10x/0.30 DRY objective lens. The DAPI (excitation 350/50 nm, dichroic 400 nm, emission 460/50 nm), GFP (excitation 480/40 nm, dichroic 505 nm, emission 527/30 nm) and Tx Red (excitation 560/40 nm, dichroic 595 nm, emission 645/76 nm) filter cubes were used. Images were captured using a Photometrics Prime 95B sCMOS camera and processed using Leica Application Suite X (Leica, Germany).

4.2.7 Imaging analysis

All images were taken in sections typical of hippocampal slices used for ex vivo electrophysiological recordings. Hippocampal layers within the CA1 were identified based on density of DAPI-stained nuclei, 250 x 600 µm sections of CA1 (covering all layers) were processed using ImageJ software and fluorescence intensity analysis of each channel (fluorophore of interest) was performed using MATLAB (R2022b). The normalised intensity profile was calculated by averaging the fluorescence intensity across the width of the image to provide representative intensities from the top to the bottom of the CA1. The maximum fluorescence intensity was used to normalise the trace and the normalised trace was averaged across slices. The fluorescence contribution per layer was calculated as an average across the layer.

4.2.8 Statistical analysis

The experimental unit was defined as a cell with only one cell recorded per slice. Data are reported as means ± SEM or medians and were analysed for statistical significance using SPSS (IBM Software). The normality of the data was assessed using the Shapiro-Wilk test in order to determine the use of parametric or non-parametric testing. Student's paired t-test (two-tailed) was used to examine statistical differences in fluorescence intensities in the different layers of the hippocampus, kinetics between OLM and SST interneurons, SST interneuron changes in excitability and associative LTP with interneuron activation. Wilcoxon matched pairs signed ranks tests (two-tailed) was used on the associative LTP without interneuron activation data. The Friedman test was used for PPR and input resistance analysis. Further Bonferroni-corrected Wilcoxon matched pairs signed ranks tests (two-tailed) were performed on the PPR and input resistance data. Pearson correlation and Spearman's rank correlation was used to examine relationships between features in the waveforms recorded and plasticity during the TBP with and without interneuron activation, respectively.

4.3 Results

4.3.1 Characterisation of Chrna2-cre mouse line

To selectively activate dendrite-targeting interneurons in hippocampal CA1, mouse lines that expressed genes for this interneuron type were required. Somatostatin (SST) is a common gene found in cortical interneurons that target the distal apical dendrites of pyramidal cells (Tremblay et al., 2016; Udakis et al., 2020). Additionally, Chrna2 is a gene for the α subunit of nicotinic acetylcholine receptors (nAChR) that has been shown to selectively express in a subpopulation of dendrite-targeting SST interneurons, known as OLM cells, located in the stratum oriens of the hippocampus (Leão et al., 2012; Freund and Buzsáki, 1996). Therefore, mice expressing Cre-recombinase under the promoter for the SST (SST-Cre) and Chrna2 (Chrna2-Cre) genes were bred with homozygous Ai32 mice, which have floxed-ChR2 genetics, to selectively express EGFP-tagged channelrhodopsin and enable optogenetic control of these interneuron populations. Subsequent histological analysis of SST and Chrna2 gene expression was performed to investigate whether there were differences in expression patterns. Immunohistochemistry of hippocampi from mice with Cre-mediated TdTomato (red fluorescent protein) expression under the control of the Chrna2 promoter (Chrna2-Cre x TdTomato mice) confirmed that TdTomato-positive somata and innervations in CA1 almost exclusively resided in stratum oriens (SO) and stratum lacunosum moleculare (SLM), respectively (Fig. 4.1A). A repeated measures ANOVA with a Greenhouse-Geisser correction for departure from sphericity revealed that mean normalised fluorescence intensity differed significantly between different layers of the hippocampus (*F*(1.253, 10.023) = 1259.74, $P < 0.001$). Simple effects analysis, using the Bonferroni correction for multiple comparisons, determined that the fluorescence found in the SP (3.18 ± 0.85%) and SR (4.03 ± 0.81%) were not significantly different (-0.85 (95%, -2.19 to 0.49) %, $p = 0.354$) and both layers were significantly lower compared to the SO and SLM (SO-SP: 8.87 (95%, -7.59 to 10.14) %, $p < 0.001$; SO-SR: 8.02 (95%, 6.23 to 9.81) %, $p < 0.001$; SP-SLM: -54.16 (95%, -59.18 to -49.15) %, $p < 0.001$; SR-SLM: -53.32 (95%, -58.08 to -48.55) %, $p < 0.001$). A very similar result was seen in Chrna2-Cre mice crossed with Ai32 mice (Chrna2-Cre x Ai32) (Fig. 4.1B). A repeated measures ANOVA with a Greenhouse-Geisser correction for departure from sphericity, once again, revealed that mean normalised fluorescence intensity differed significantly between different layers of the hippocampus (*F*(2.083, 16.666) = 60.578, $P < 0.001$). Post hoc pairwise comparisons using the Bonferroni correction showed that, unlike Chrna2-Cre x TdTomato mice, there was no significant difference in fluorescence between SO and SP (6.06 (95%, -7.89 to 20.01) %, $p = 1.000$) but both had significantly greater fluorescence compared to SR (SO-SR: 22.18 (95%, 13.37 to 30.99) %, $p < 0.001$; SP-SR: 16.12 (95%, 5.24 to 26.99) %, $p = 0.005$). The fluorescence measured at the SLM was significantly the highest compared to the other layers (SLM-SO: 19.90 (95%, 7.61 to 32.19) %, $p = 0.003$; SLM-SP: 25.96 (95%, 13.40 to 38.53) %, $p < 0.001$; SLM-SR: 42.08 (95%, 37.01 to 47.15) %, $p < 0.001$). Histological analysis of SST expression indicated that these cells and their innervations, although were predominantly located in the same locations as OLM interneurons, were also partially found in the stratum pyramidal (SP) and stratum radiatum (SR) (Fig. 4.1A). A repeated measures ANOVA determined that mean normalised fluorescence intensity differed significantly between different layers of the hippocampus (*F*(3, 24) = 787.787, $P < 0.001$). Post hoc analysis with Bonferroni adjustment revealed that there was no significance difference in fluorescence between SO and SR (0.55 (95%, -3.74 to 4.85) %, $p = 1.000$), and both layers had significantly higher fluorescence than SP (SO-SP: 7.29 (95%, 3.25 to 11.33) %, $p < 0.001$; SR-SP: 6.74 (95%, 3.32 to 10.15) %, $p < 0.001$). Similar to the expression of the Chrna2-Cre dependent markers, the fluorescence measured at the SLM was significantly the highest compared to the other layers (SLM-SO: 45.64 (95%, 40.39 to 50.90) %, $p = 0.003$; SLM-SP: 52.93 (95%, 48.38 to 57.48) %, $p < 0.001$; SLM-SR: 46.20 (95%, 42.32 to 50.08) %, $p < 0.001$). Taken together, these results suggest that Chrna2-expressing cells are highly specific to the OLM interneurons found in SO but SST-expressing cells, although primarily confined to SO, are also found or innervate the SR. This is supported by observations from Leão et al. (2012) and Freund and Buzsáki (1996). Furthermore, SST and Chrna2 cells heavily project to the SLM of CA1 indicating that both cell populations are suitable for mediating the activity found at the distal apical dendrite of CA1 pyramidal cells.

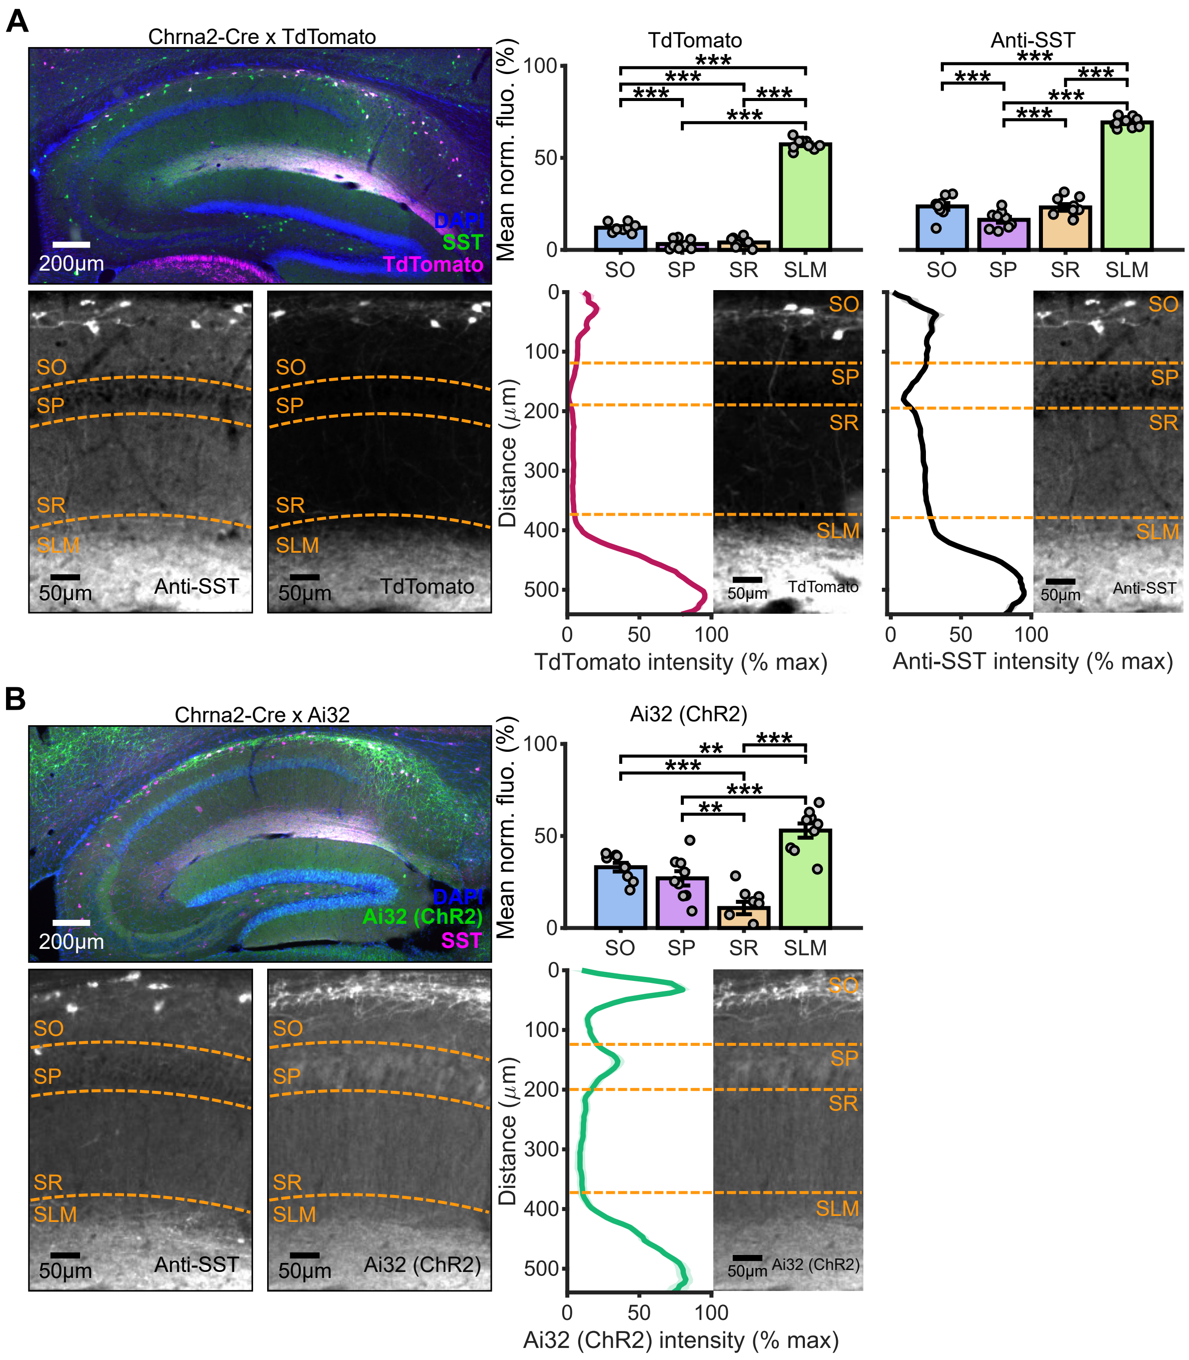
A
Chrna2-Cre x TdTomato
DAPI
SST
TdTomato
200µm
TdTomato
Anti-SST
Mean norm. fluo. (%)
100
50
0
***
SO
SP
SR
SLM
50µm
Anti-SST
TdTomato
Distance (µm)
0
100
200
300
400
500
TdTomato intensity (% max)
Anti-SST intensity (% max)
B
Chrna2-Cre x Ai32
DAPI
Ai32 (ChR2)
SST
Ai32 (ChR2)
**
Ai32 (ChR2) intensity (% max)

FIGURE 4.1: **OLM interneurons reside in the stratum oriens and innervate the stratum lacunosum moleculare of the hippocampus.** (**A**; top left) Immunohistochemical staining in hippocampal slice of Chrna2-Cre x TdTomato animal showing expression of SST (green) and TdTomato in OLM interneurons (magenta). Sections of CA1 displaying expression patterns of SST (bottom far left) and TdTomato (bottom middle left) in the different hippocampal layers: stratum oriens (SO), stratum pyramidal (SP), stratum radiatum (SR) and stratum lacunosum moleculare (SLM). The normalised intensity profiles show the mean TdTomato (bottom middle right) and SST (bottom far right) fluorescence levels as a function of distance from the top (SO) to the bottom (SLM) of CA1. Bar plots demonstrate the mean fluorescence

within each layer for TdTomato (top middle right) and SST (top far right) (N = 9 slices from three animals). (**B**) Same as (**A**) but for Chrna2-Cre x Ai32 animals and the expression of ChR2. Data represented as mean ± SEM; ** $p < 0.01$, *** $p < 0.001$ (repeated measures ANOVA with Bonferroni-corrected simple effects analysis.

4.3.2 Effect of cholinergic modulation on the inhibitory synaptic transmission from OLM and SST interneurons

Electrophysiological data has shown that neocortical SST and hippocampal OLM cells are potently activated by ACh via a muscarinic ACh receptor (mAChR) mediated mechanism but the collective effect of ACh on the SST/OLM to CA1 pyramidal cell (PC) synapses remains unclear (Muñoz and Rudy, 2014; Xu et al., 2013; Lawrence et al., 2006; Lawrence, 2008). To investigate the cholinergic action on SST/OLM-to-PC synapses, whole-cell patch clamp recordings of inhibitory currents were made from CA1 PCs by photostimulating ChR2-expressing SST or OLM interneurons with blue (470 nm) light (Fig. 4.2A). In contrast to a previous study by Udakis et al. (2020), optogenetic activation of these interneurons prior to carbachol (CCh), AChR agonist, demonstrated that OLM-derived IPSCs expressed significantly slower rise and decay kinetics compared to SST-derived IPSCs (rise time: 4.60 ± 0.37 ms OLM versus 2.20 ± 0.58 ms SST, $t(12) = 3.453$, $p = 0.005$; decay time: 44.42 ± 2.62 ms OLM versus 22.13 ± 4.03 ms SST, $t(12) = 4.986$, $p < 0.001$) indicating further differences in the two cell populations (Fig. 4.2B). Inhibitory currents from both SST and OLM interneurons decreased in a dose-dependent manner after the application of increasing concentrations of CCh (OLM inhibition with IC_{50} of 0.859 μM; SST inhibition with IC_{50} of 4.52 μM) (Fig. 4.2C-F). To determine the mechanism behind the cholinergic-mediated inhibition of SST/OLM-to-PC synapses, the paired pulse ratio (PPR), defined as the ratio of the IPSC evoked by the second pulse of light to the first, and the input resistance were measured. Friedman tests determined a statistical difference in OLM PPR depending on the concentration of CCh ($\chi^2(5) = 11.987$, $p = 0.035$) but this interaction was not seen with OLM input resistance and CCh concentration ($\chi^2(5) = 3.032$, $p = 0.695$). Dunn-Bonferroni post hoc analysis of OLM PPR revealed there were no significant differences between any of the concentrations. Repeated measures ANOVAs with Greenhouse-Geisser corrections for departure from sphericity revealed no significant interaction between SST PPR or input resistance and the concentration of CCh (PPR: $F(1.994, 5.983) = 2.189$, $P = 0.193$; input resistance: $F(1.543, 4.630) = 0.735$, $P = 0.494$). These data suggest that ACh inhibits SST/OLM-to-PC synapses and the effect is mediated by a presynaptic mechanism in OLM-to-PC synapses but not in SST-to-PC synapses.

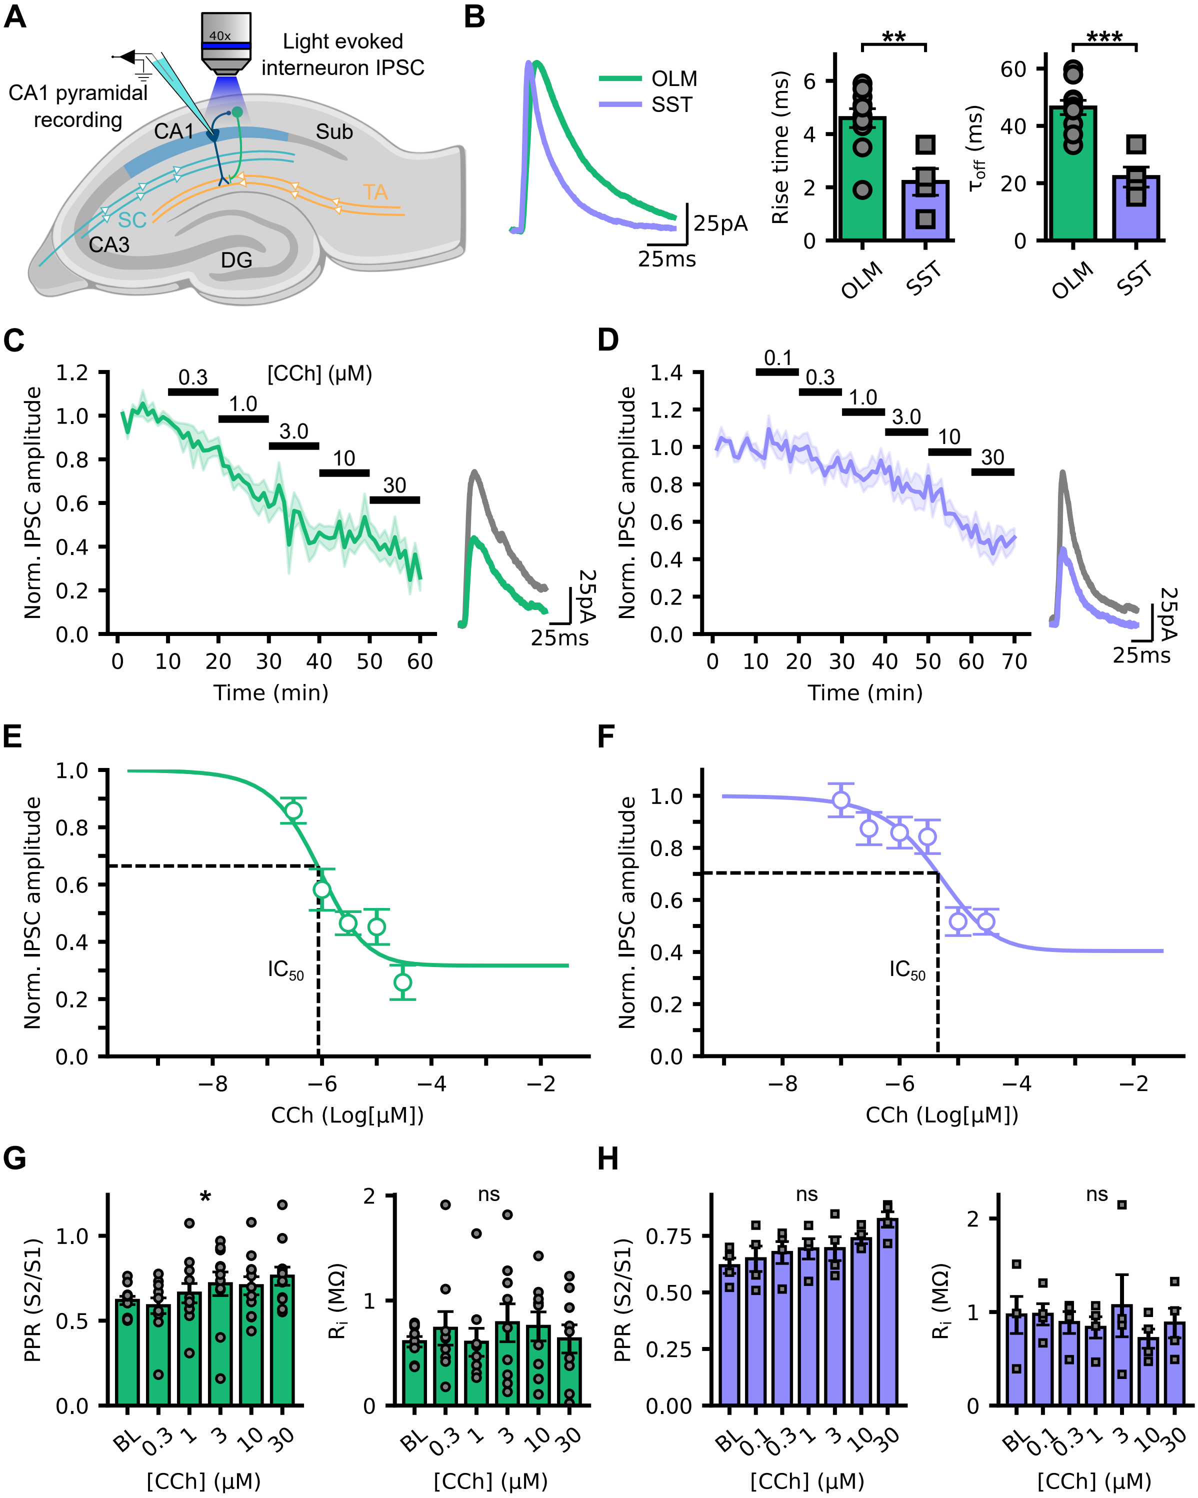

FIGURE 4.2: **CCh causes a dose-dependent decrease in OLM and SST interneuron inhibition.** (**A**) Schematic of the hippocampal slice recording setup: whole-cell patch-clamp recordings of IPSCs in CA1 pyramidal cells (dark blue) from photostimulated dendrite-targeting OLM interneurons (green). (**B**;

left) Example normalised IPSC traces from OLM (green) and SST (purple) interneurons highlighting representative differences in synapse kinetics. Summary of IPSC rise (middle) and decay (right) times for SST and OLM interneurons (unpaired t-test, two-tailed; N = 10, OLM cells; N = 4, SST cells). (**C**) A timecourse of the experiment showing a dose-dependent decrease in normalised IPSC responses from OLM interneurons with increasing concentrations of CCh. (**D**) Same as (**C**) but with SST interneurons. (**E**) A dose-response curve with the IC_{50} labelled for a clearer visualisation of this dose-dependent effect in OLM interneurons. (**F**) Same as (**E**) but for SST interneurons. (**G**; left) The PPR at varying concentrations of CCh calculated by dividing the amplitude of the first pulse by the second (S2/S1) (Friedman test; N = 11 cells). (right) The input resistance across the different concentrations of CCh (Friedman test; N = 9). (**H**) Same as (**G**) but for SST interneurons (repeated measure ANOVA; N = 4). BL, baseline period prior to bath application of CCh. Data represented as mean ± SEM; * $p < 0.05$.

4.3.3 Effect of cholinergic modulation on the intrinsic excitability of stratum oriens SST interneurons

After establishing the profound inhibition of the SST/OLM-to-PC synapses by ACh and an indication that this effect is potentially via a presynaptic mechanism in OLM-to-PC synapses, SO SST interneurons were patched to examine whether this inhibition was due to a cholinergic-mediated reduction in their intrinsic excitability (Fig. 4.3A). Current-clamp recordings were made from these cells to assess their resting membrane potential (RMP) and firing properties in response to increasingly more depolarised current steps in the presence of 10 μM CCh (Fig. 4.3B). After the application of CCh, there was a significant increase in the RMP from -57.16 ± 1.87 mV ,during baseline, to -50.65 ± 2.74 mV ($t(3) = -3.413$, $p = 0.042$); a depolarisation of 6.51 ± 1.91 mV. A repeated measures ANOVA also determined that there was a significant main effect of drug condition on the average spike firing ($F(2, 6) = 9.005$, $P = 0.016$). Despite a mean increase of 8.42 ± 1.96 spikes from pre- (8.96 ± 4.57 spikes) to post-drug (17.38 ± 5.54 spikes) application, post-hoc simple effect analysis, using a Bonferroni correction for multiple comparisons, revealed that there were no significant interactions between the different experimental conditions (Baseline-Drug: -8.42 (95%, -17.93 to 1.10) spikes, $p = 0.070$; Baseline-Repolarised: -3.54 (95%, -9.74 to 2.65) spikes, $p = 0.208$; Drug-Repolarised: 4.88 (95%, -7.44 to 17.19) spikes, $p = 0.451$) (Fig. 4.3C-D). These results suggest that ACh increases the intrinsic excitability of stratum oriens SST interneurons by depolarising their RMPs.

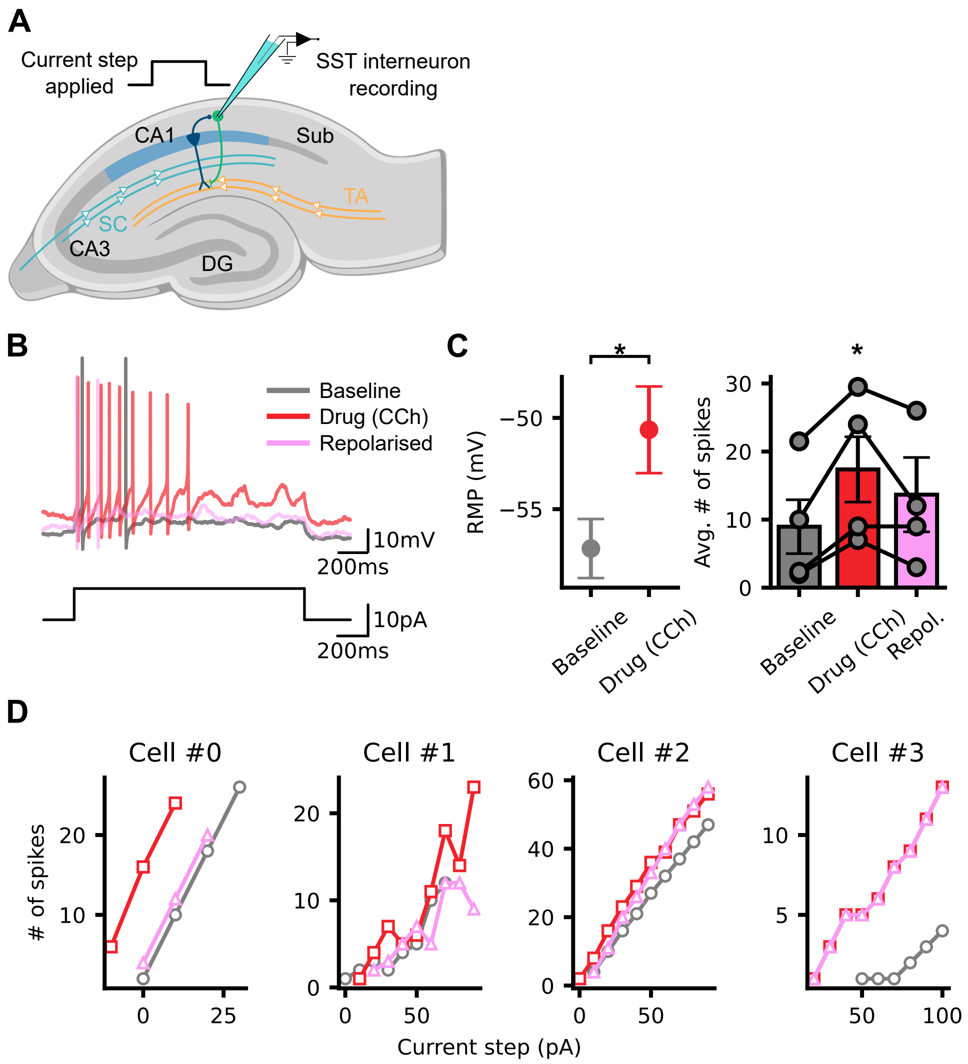

FIGURE 4.3: **CCh increases the excitability of SST interneurons.** (**A**) Schematic of the hippocampal slice recording setup: whole-cell patch-clamp recordings of action potentials in SST interneurons evoked by increasingly higher amplitude current steps. (**B**) Example traces of OLM interneuron activity after the application of a 10 pA current step during baseline (grey), in the presence of CCh (red) and after repolarising the resting membrane potential (RMP; pink) of the cell back to the baseline state. (**C**) Summary of RMP before and after CCh application (left; Student's paired t-test, two-tailed; N = 4) and the average number of spikes, elicited by a current step that evokes half the number of maximum spikes during the baseline period, in each experiment state (right; repeated measures ANOVA; N = 4). (**D**) The excitability profiles of each SST interneuron that contributed to the summary statistics (**C**). Data represented as mean ± SEM; * $p < 0.05$.

4.3.4 OLM interneurons regulate aLTP in CA1

The CA1 region of the hippocampus is crucial for the formation of new memories. This process is thought to occur through the integration of spatiotemporally coherent inputs from two sources:

the entorhinal cortex and the CA3 region of the hippocampus via the temporammonic (TA) and Schaffer collateral (SC) pathways, respectively. The coincident timing of these two afferent inputs interact supra-linearly at the distal apical dendrites of CA1 PCs which can initiate dendritic plateau potentials and complex burst firing that drive associative LTP (aLTP) (Bittner et al., 2015; Sjöström and Häusser, 2006). OLM interneurons target the distal apical dendritic regions that mediate this nonlinear dendritic integration, thus they should have the capability to control synaptic plasticity (Leão et al., 2012; Freund and Buzsáki, 1996). To investigate whether OLM interneurons can regulate aLTP, whole-cell patch clamp recordings were made from CA1 PCs and aLTP was assessed using a theta-burst stimulation protocol (TBS) consisting of simultaneous theta-burst stimulations of the TA and SC pathways. Activity from another independent SC pathway was also recorded as a negative control (Fig. 4.4A). Wilcoxon signed rank tests showed that the TBS caused a significant robust expression of aLTP in the SC pathway but not in the TA pathway (Control-SC: Z = -2.366, p = 0.018; Control-TA: Z = -0.676, p = 0.499) (Fig. 4.4B). After establishing the ability of the TBS to induce plasticity, these experiments were repeated but with optogenetic activation of OLM interneurons during the induction protocol (Fig. 4.4C). Student's paired t-tests revealed a lack of aLTP in the SC pathway but a significant depression in activity was seen in the TA pathway (Control-SC: 0.004 ± 0.159, $t(6)$ = 0.024, p = 0.982; Control-TA: 0.431 ± 0.093, $t(6)$ = 4.626, p = 0.002) (Fig. 4.4D). Together, these results demonstrate that OLM interneuron activity can regulate the induction of aLTP.

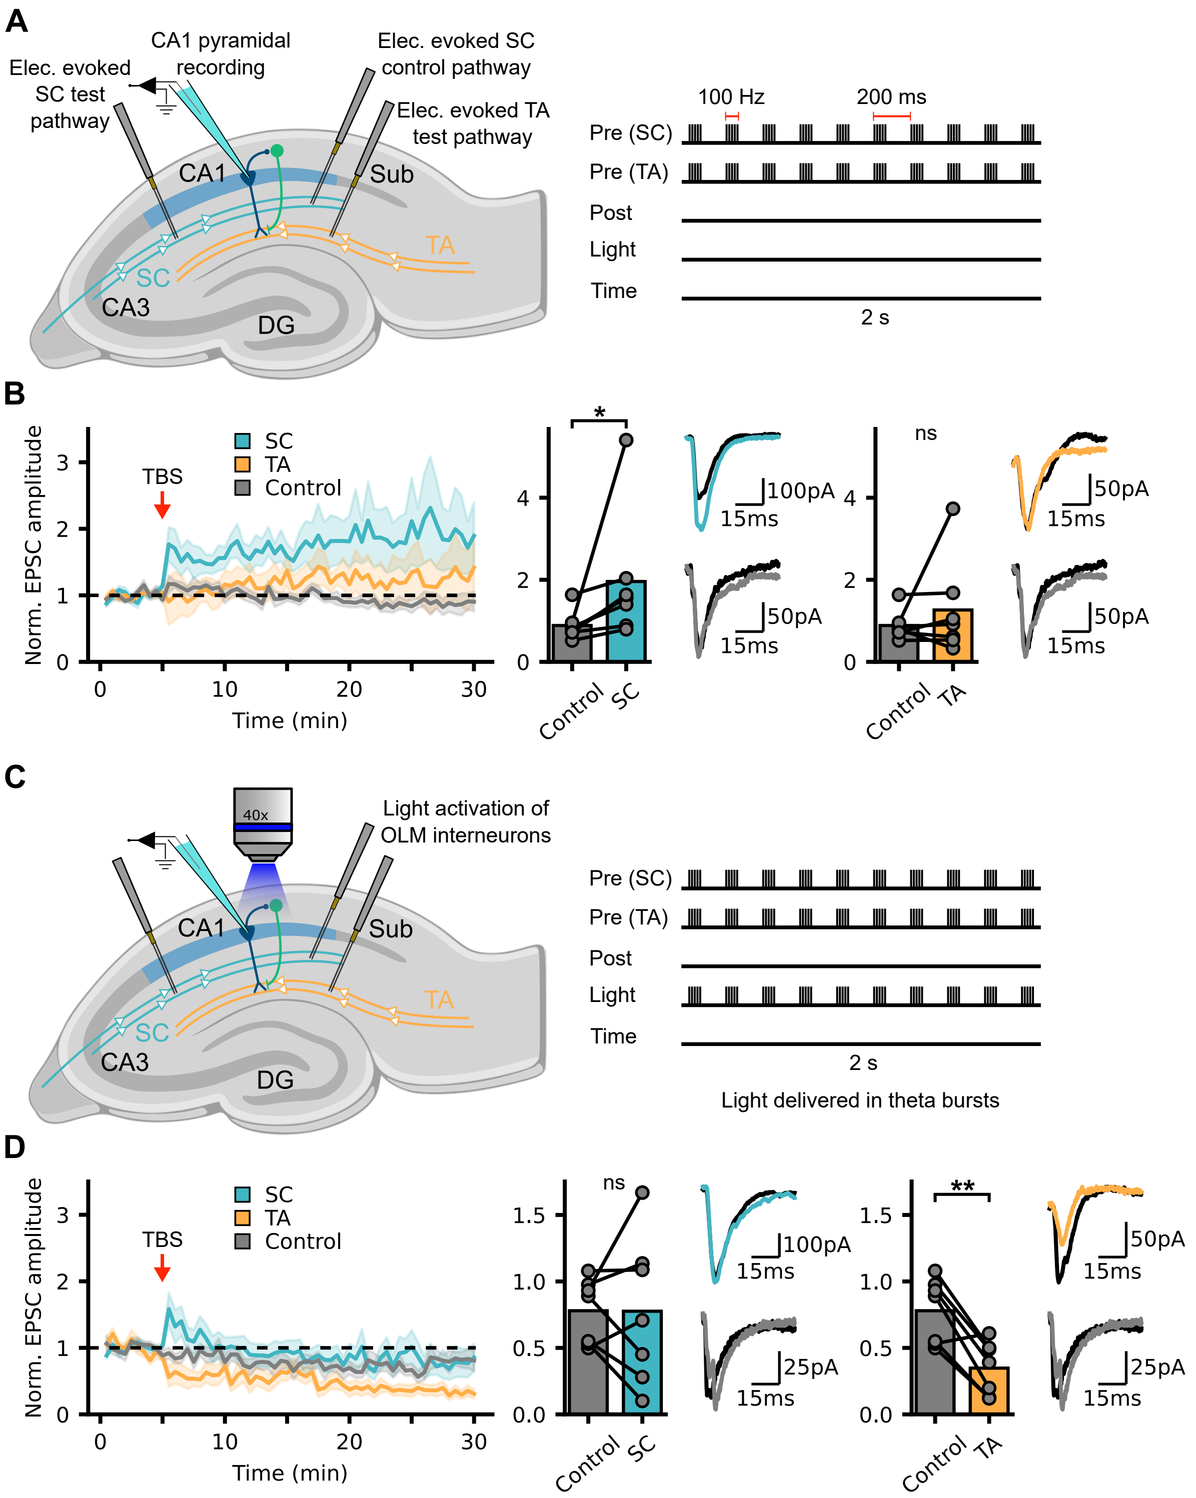

FIGURE 4.4: **OLM interneuron activity prevents theta-burst-induced associative LTP in the SC pathway.** (**A**; left) Schematic of the hippocampal slice recording setup. Whole-cell patch-clamp recordings of EPSCs in CA1 pyramidal cells evoked by stimulation of three independent pathways: SC test, TA test and SC control. (right) A 2 s aLTP theta-burst induction protocol (TBS) illustrating stimulation of the

SC test and TA test pathways in the absence of somatic depolarisation and photostimulation of OLM interneurons. (**B**; left) A timecouse of the TBP-induced aLTP at SC synapses. Summary of the mean plasticity at the control, SC (middle) and TA (right) pathways in the last 5 min of the experiment with example traces before (black) and after (coloured) the TBS (Wilcoxon matched pairs signed ranks test, two-tailed; N = 7). (**C**; left) Schematic of experimental setup same as (**A**; left) but with photostimulation of OLM interneurons during the TBS (right). (**D**; left) A timecouse of attenuated aLTP at SC synapses and aLTD at TA synapses. Summary of the mean plasticity at the control, SC (middle) and TA (right) pathways in the last 5 min of the experiment with example traces before (black) and after (coloured) the TBS with concurrent OLM interneuron activation (Student's paired t-test, two-tailed; N = 7). Data represented as mean ± SEM; * p < 0.05, ** p < 0.01.

4.3.5 Effects of spiking, bursting and amount of depolarisation on aLTP

Active mechanisms in dendrites, such as backpropagating action potentials and dendritic calcium plateau potentials, drive complex bursting activity in the soma of cells resulting in synaptic plasticity (Xu et al., 2012; Harnett et al., 2013; Larkum et al., 1999; Larkum, 2013; Takahashi and Magee, 2009; Tsay et al., 2007; Jarsky et al., 2005; Kamondi et al., 1998; Sjöström and Häusser, 2006; Sjostrom et al., 2008; Bittner et al., 2015). Therefore, to gain a better understanding of how spiking, bursting and depolarisation affects aLTP and how these relationships are altered by OLM interneuron intervention, exploratory analysis was performed on the induction traces from the previous plasticity experiments. In induction traces without the optogenetic activation of OLM interneurons, there were non-significant positive correlations in both the SC and TA pathways of number of spikes (SC Spearman's rank correlation: $r_s(7) = 0.143$, $p = 0.760$; TA Pearson correlation: $r(7) = 0.736$, $p = 0.059$) (Fig. 4.5A-B), bursts (SC Spearman's rank correlation: $r_s(7) = 0.143$, $p = 0.760$; TA Pearson correlation: $r(7) = 0.634$, $p = 0.126$) (Fig. 4.5D-E) and total depolarisation (AUC) (SC Spearman's rank correlation: $r_s(7) = 0.312$, $p = 0.482$; TA Pearson correlation: $r(7) = 0.637$, $p = 0.098$) (Fig. 4.5G-H) with strength of aLTP. In contrast, when there was optogenetically evoked inhibition during the induction, there were non-significant negative correlations in both the SC and TA pathways of number of spikes (SC Pearson correlation: $r_s(7) = -0.297$, $p = 0.518$; TA Pearson correlation: $r(7) = -0.007$, $p = 0.988$) (Fig. 4.5C), bursts (SC Pearson correlation: $r_s(7) = -0.395$, $p = 0.380$; TA Pearson correlation: $r(7) = -0.274$, $p = 0.552$) (Fig. 4.5F) and total depolarisation (SC Pearson correlation: $r_s(7) = -0.243$, $p = 0.600$; TA Pearson correlation: $r(7) = -0.046$, $p = 0.921$) (Fig. 4.5I) (Fig. 4.5G-H) with strength of aLTP. These data suggest that there are no significant relationships of number of spikes, bursts and depolarisation with synaptic plasticity.

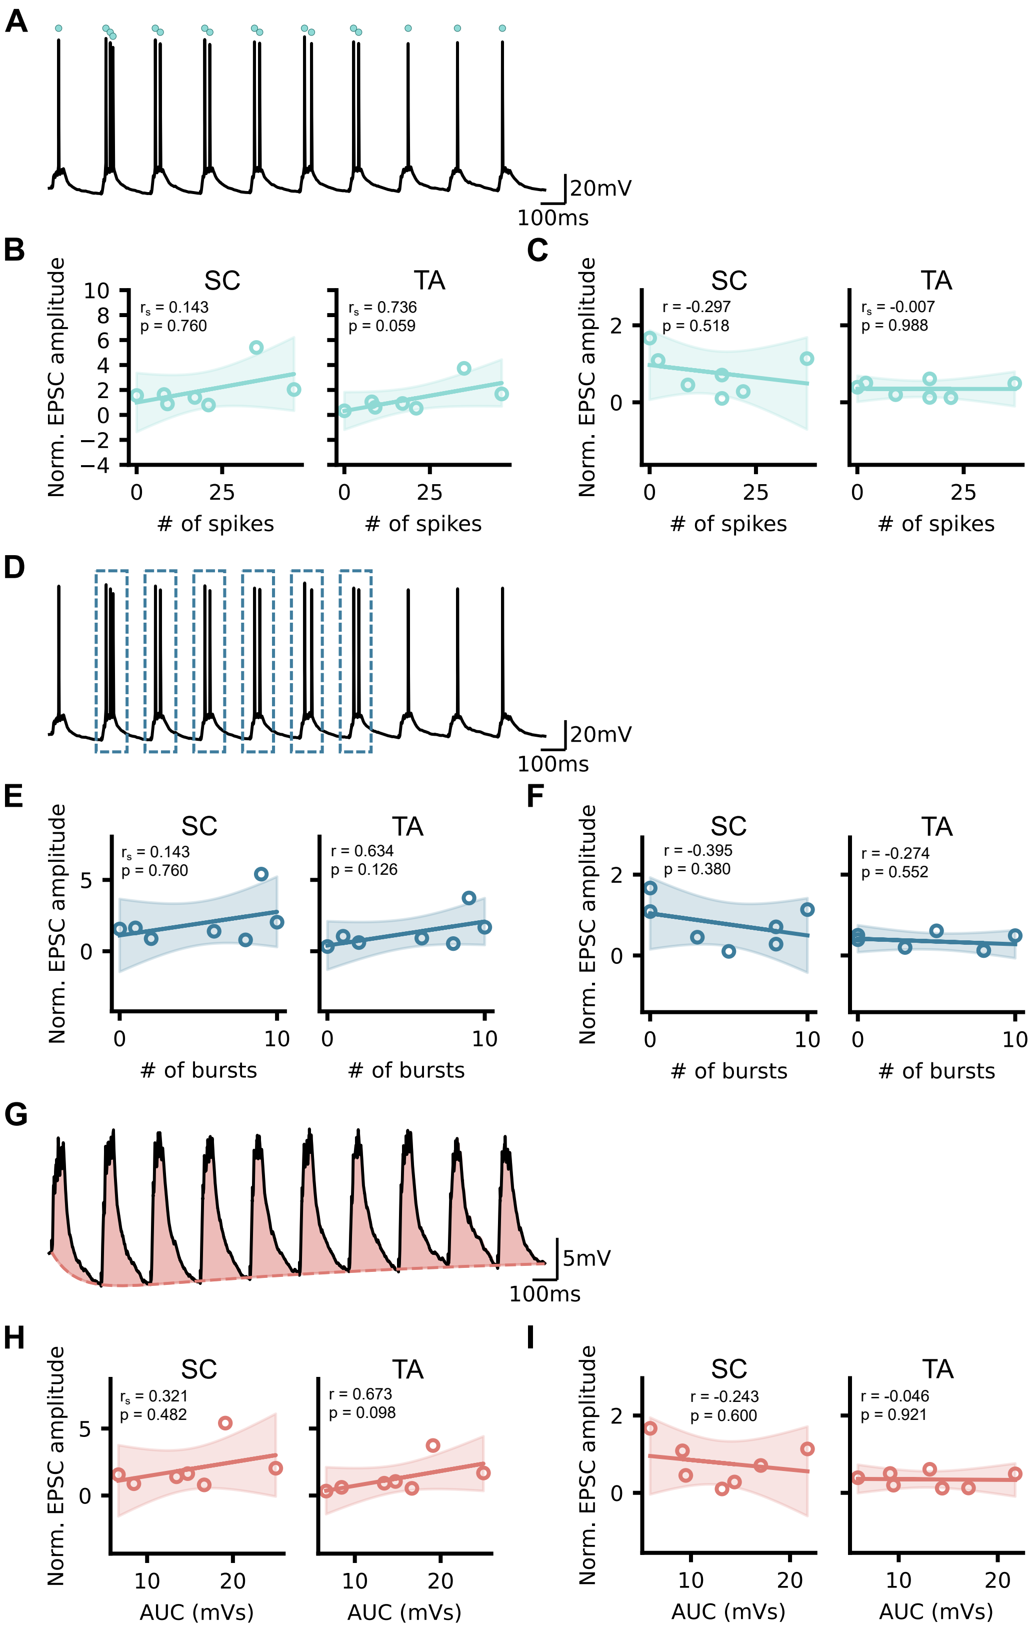

FIGURE 4.5: **The number of spikes, bursts and amount of depolarisation during in-**

duction does not correlate with the expression of plasticity. (**A**) Example trace of CA1 pyramidal cell activity during TBS with cyan dots indicating the count of a spike. Correlation between number of spikes, during the TBS in the absence (**B**) and presence (**C**) of simultaneous photostimulation of OLM interneurons, and plasticity in the SC and TA pathways. (**D**) Example trace of CA1 pyramidal cell activity during TBS with dashed blue box indicating the count of a burst. Correlation between number of bursts, during the TBS in the absence (**E**) and presence (**F**) of simultaneous photostimulation of OLM interneurons, and plasticity in the SC and TA pathways. (**G**) Example trace of CA1 pyramidal cell activity during TBS with a dashed, pink biexponential curve fitted as the baseline and shaded area representing amount of depolarisation. Correlation between amount of depolarisation, during the TBS in the absence (**E**) and presence (**F**) of simultaneous photostimulation of OLM interneurons, and plasticity in the SC and TA pathways. Data represented as values from each experiment, line of best fit and shaded 95% confidence intervals. Spearman's rank correlations performed on data correlating to TBS-induced plasticity in SC. Pearson correlations were performed on the rest of the data (N = 7).

4.4 Discussion

In the hippocampus, a substantial portion of the GABAergic interneurons are composed of SST cells. However, the SST molecular marker is found in a diverse population of interneurons comprising of Chrna2-positive OLM cells, Chrna2-negative OLM cells, bistratified cells and radiatum cells (Honoré et al., 2021). Previous behavioural and electrophysiological studies have shown that the activity of these interneurons is important for memory formation and synaptic plasticity, and it is thought that they are mediated by the cholinergic projections from the medial septum (Lovett-Barron et al., 2014; Honoré et al., 2021; Leão et al., 2012). Despite extensive research on the effects of cholinergic neuromodulation on the intrinsic excitability of these cells, how ACh affects their eventual inhibitory transmission onto CA1 PCs is not well understood (Pelkey et al., 2017; Zheng et al., 2011; Lawrence et al., 2006; Lawrence, 2008). This chapter addresses this question and further investigates the role of the OLM interneurons in synaptic plasticity. First, the immunohistochemical experiments performed here have confirmed the variety in the SST interneuron population and the specificity of the Chrna2-positive OLM cells as a dendrite-targeting interneuron population. Analysis of their synapse kinetics onto CA1 PCs further highlighted the differences of the SST and OLM interneuron populations. The faster rise and decay kinetics of light-evoked SST IPSCs in comparison to OLM IPSCs suggested that a proportion of the SST interneurons, likely constituting the bistratified cells and radiatum cells, synapse to dendritic regions more proximal to the soma.

In both cell populations, the application of increasing concentrations of CCh caused a dose-dependent reduction in inhibitory transmission onto CA1 PCs. The differences in the potency of CCh were likely due to the SST cell population comprising of additional interneuron types. The mechanism by which CCh acts to cause this depression in inhibitory output remains inconclusive. Initial PPR analysis indicated that for the OLM cells, this cholinergic-mediated effect was due to a presynaptic mechanism as there was a significant main effect of CCh but further post-hoc analysis did not provide any significant differences between the various CCh concentrations. This could be due to the Bonferroni correction being too conservative or there being no significant differences in the PPRs. The lack of a significant difference in the PPRs, as seen with SST cells, suggested that this could be due to a postsynaptic mechanism, potentially via the down-regulation of postsynaptic GABA receptors. However, it is important to not exclude the potential of a presynaptic mechanism as a study by Burke et al. (2018) has shown that there are non-canonical forms of presynaptic neuromodulation that result in no changes to PPR. As expected, both SST and OLM cells displayed no changes in input resistance across the CCh concentrations: the use of a caesium-based internal blocked potassium channels, thus it was not possible to see the changes in potassium conductance of the cell from the downstream effects of mAChR activation.

Although CCh hindered the inhibitory output of SST and OLM cells onto CA1 PCs, it was found to increase the intrinsic excitability of these interneurons by depolarising their resting membrane potentials. This up-regulatory effect of activity has been previously shown in a number of studies, suggesting that the cholinergic modulation of inhibitory transmission is functionally decoupled from changes in membrane excitability (Pelkey et al., 2017; Lawrence et al., 2006; Lawrence, 2008).

In the aLTP experiments, a TBS was optimised to produce sufficient LTP in the SC pathway. The intervention of the TBS by simultaneous optogenetic activation of OLM interneurons prevented the induction of LTP in the SC pathway and caused LTD in the TA pathway. This demonstrated that OLM interneurons are capable of modulating hippocampal synaptic plasticity. It is known that the coincident activity of SC and TA inputs is important for the generation of plateau potentials that induce LTP at glutamatergic synapses of CA1 PCs (Takahashi and Magee, 2009; Dudman et al., 2007). Recent studies have demonstrated that these plateau potentials are important for the generation and feature selectivity of CA1 place cells (Bittner et al., 2015, 2017). As the OLM interneurons target the distal dendritic compartments of the CA1 PC, a potential mechanism to regulate synaptic plasticity is by mediating the pairing of the proximal and distal CA1 inputs which are required to evoke plateau potentials. OLM-mediated inhibition of the TA pathway can gate information flow into the hippocampus by disrupting the necessary external information from the entorhinal cortex (EC) reaching the CA1. This means the CA1 circuit cannot associate the local contextual information provided by CA3, via the SC pathway, with the spatial information from the EC to encode context-dependent spatial maps (Bittner et al., 2015). Cholinergic neuromodulation can therefore be a higher level control of hippocampal synaptic plasticity by controlling the inhibitory transmission of OLM interneurons. ACh release can act as a novelty detection signal to facilitate memory formation in the hippocampus by inhibiting the OLM-to-PC synapse and glutamatergic synapses onto OLM interneurons (Zheng et al., 2011).

The exploratory analysis of the induction traces from the aLTP experiments determined, contrary to previous studies, that there were no significant relationships of number of spikes, bursts and depolarisation with synaptic plasticity (Xu et al., 2012; Harnett et al., 2013; Larkum et al., 1999; Larkum, 2013; Takahashi and Magee, 2009; Tsay et al., 2007; Jarsky et al., 2005; Kamondi et al., 1998; Sjöström and Häusser, 2006; Sjostrom et al., 2008; Bittner et al., 2015). However, the small sample size indicates that the lack of significance could be due to a lack of power in the statistical analysis. In general, there was a trend showing a positive correlation of these properties with synaptic plasticity in experiments without OLM interneuron activation which became anti-correlated when the TBS included optogenetic stimulation of OLM interneurons. The attenuation of LTP with concomitant OLM interneuron activation, despite similar numbers of spikes, bursts and depolarisation, could be potentially explained by a theoretical burst-dependent plasticity rule proposed by Payeur et al. (2021). The plasticity rule suggests that there is a baseline level of bursting that a cell exhibits and subsequent increases or decreases in burst activity result in LTP or LTD, respectively. In this case, the TBS is a strong enough stimulus to guarantee an increase in bursting from the cells' baseline burst rates, thus inducing LTP. However, inhibitory transmission from OLM interneurons could be weakening the stimulation enough to reduce the number of bursts below the cells' baseline burst rate resulting in LTD.

4.4.1 Limitations

For the SST excitability experiments, SST-Cre x Ai32 mice were used to locate SO dendrite-targeting interneurons. As previously mentioned, a diverse population of cells constitutes the SST interneuron population including bistratified and radiatum cells that do not target the distal dendritic compartments of CA1 PCs (Honoré et al., 2021). This means that the changes in intrinsic excitability of the cells recorded could include cholinergic effects on cell types that are not of interest. Additionally, during the experiment, there were no drugs applied to block the influence of spontaneous network activity on the patched cell. CCh effects on other cells in the hippocampal circuit could externally be affecting the excitability of the interneuron being recorded from. Therefore, an ideal experiment would be to use the Chrna2-Cre x Ai32 mouse line to identify and patch OLM interneurons with the bath application of NBQX, DAPV and picrotoxin to isolate the cell by blocking potential regenerative dendritic events, glutamatergic and GABAergic inputs, respectively.

The sample sizes of the dose-response experiments could ideally be larger to distinguish where the differences in OLM PPR occur and whether the trend that PPR increased with increasing concentrations of CCh seen in the SST cells is definitively significant. Similarly, a greater sample size for the aLTP experiments would help to verify whether the trends seen in the correlations of number of spikes, bursts and depolarisation with synaptic plasticity were real and more interestingly, if OLM interneuron activation decouples these properties and the induction of long-term plasticity.

4.4.2 Future experiments

It is known that coincident timing of activities from the SC and TA pathways evoke plateau potentials that induce synaptic plasticity (Bittner et al., 2015; Takahashi and Magee, 2009; Dudman et al., 2007). It would be interesting to see whether the TBS used in these experiments to induce aLTP is also following a similar mechanism via plateau potential generation. If so, to understand the mechanism by which OLM interneurons mediate synaptic plasticity, it would be important to investigate their influence on the generation of plateau potentials. A potential approach could be to use two-photon calcium imaging: visualising the calcium events at the the distal apical dendritic regions of CA1 PCs whilst applying the TBS and also with the activation of the OLM interneurons.

Finally, it would be crucial to test the impact of ACh on OLM-mediated learning and memory *in vivo* at the behavioural level. Previous studies have shown that the activity of OLM interneurons can modulate fear-related memory encoding and learning in spatial memory tasks but none have investigated the impact of hippocampal cholinergic modulation on OLM interneuron-dependent learning (Lovett-Barron et al., 2014; Siwani et al., 2018; Haam et al., 2018). In addition, spatial memory is highly dependent on the hippocampus via the formation of place cells. If OLM interneurons are important in spatial memory tasks, it is still not clear how these OLM interneurons affect CA1 place cell generation. Therefore, it would be beneficial to understand their involvement and whether cholinergic modulation regulates this process.

4.4.3 Closing remarks

In conclusion, the optogenetic activation of OLM interneurons prevents the induction of LTP and the application of CCh inhibits the OLM-to-PC synapse despite increasing the intrinsic excitability of the OLM cells. This suggests that the cholinergic modulation of OLM interneuron transmission could potentially be a gating mechanism for CA1 synaptic plasticity.

CHAPTER

5

INDEPENDENT OPTICAL EXCITATION OF CHAT- AND CHRNA2-EXPRESSING NEURAL POPULATIONS

5.1 Introduction

Dual optogenetics involves the use of two different types of opsins (light-sensitive proteins) to control the activity of specific neurons in the brain. These opsins can be activated by different wavelengths of light and can be used to either excite or inhibit neurons. However, there is a risk of cross-talk between the two opsins, which can lead to unintended activation or inhibition of neurons.

One potential source of cross-talk is spectral overlap between the two opsins. This occurs when the wavelengths of light used to activate one opsin also activates the other. For example, if one opsin is activated by blue light and the other opsin is activated by green light, there may be some overlap in the spectrum of light that activates both opsins. This is quite common with the earlier developed excitatory opsins such as channelrhodopsin-2 (ChR2) and red-shifted variants including VChR1 and C1V1, which all have fairly broad action spectra that overlap with each other (Zhang et al., 2008; Yizhar et al., 2011). This can result in unintended activation or inhibition of neurons expressing one or both opsins and is one of the main complications when attempting to optogenetically stimulate individual neural populations.

As the optogenetic toolbox continues to develop and expand its repertoire of available opsin variants, steps are gradually being taken to address this long-standing issue of spectral cross-talk. One approach to tackling this problem is by creating long wavelength-sensitive opsins with narrower action spectra to avoid impeding onto the action spectra of blue light-sensitive opsins. One recent development is the advent of Chrimson, a red-shifted excitatory opsin that has a spectral peak at 590 nm which is 45 nm more red-shifted than many existing red-shifted opsins (Klapoetke et al., 2014). Another approach is to shift the peak excitation to a shorter wavelength. The issue that exists for current opsins with this feature is that they either produce weaker or similar photocurrents compared to ChR2 which, when used in combination with red-shifted opsins, cannot offer sufficient excitation of neurons without cross-talk (Mattis et al., 2012). Therefore, efforts have also gone into making more sensitive blue light-sensitive opsins such as ChRGer which produce photocurrents four times greater than ChR2 and reach their maximum response with much lower blue light irradiances (Bedbrook et al., 2019).

In the previous chapter it was shown that CCh hinders the inhibitory synaptic transmission of OLM interneurons onto CA1 PCs. This form of cholinergic modulation could suggest a potential mechanism for ACh to gate hippocampal synaptic plasticity. However, it remains unclear whether the concentrations of CCh used are reflective of physiological ACh release. Therefore, a dual optogenetics approach was undertaken in order to independently stimulate endogenous release of ACh and examine its effects on the optogenetically evoked IPSCs at the OLM-to-PC synapses.

To validate this dual optogenetic approach, the spectral cross-talk of the ChRger and Chrimson opsins was evaluated and the optimal blue light irradiance for spectral separability was obtained. This irradiance proved capable of photostimulating the release of endogenous ACh in hippocampal slices from mice with ChR2-expressing cholinergic neurons. An attempt was made to measure the concentration of this physiological ACh with the use of the ACh3.0 sensor (Jing et al., 2020). Two approaches were examined to selectively express ChRger in the cholinergic fibres and Chrimson in the OLM interneurons. The Cre and FlpO recombinase gene expression systems are widely used genetic tools based on site-specific DNA recombinases that facilitate precise DNA modification in living organisms. Cre and FlpO recombinase recognises specific DNA sequences called loxP and FRT sites, respectively. Both enzymes catalyse recombination between those sites to enable precise genetic modification. Therefore, one approach was a pure transgenic approach by creating a novel mouse line which expressed Cre recombinase in Chrna2-expressing cells such as the OLM interneurons and FlpO recombinase in the cholinergic cells. Theoretically, subsequent Cre- and FlpO-dependent viral expression of the desired opsins should be targeted to the chosen cell types. Another approach involved the use of a recently developed virus that expresses Cre recombinase under the choline acetyltransferase (ChAT) promoter in Chrna2-Cre mice. To prevent non-specific viral expression of opsins intended for cholinergic neurons, this approach was validated to ensure that Chrna2-expressing cells in the medial septum, which all expressed Cre recombinase, were not projecting to the hippocampus. This approach demonstrated greater expression specificity and was used to independently excite the OLM interneuron population and cholinergic fibres in hippocampal CA1. Activation of septo-hippocampal projections was found to produce a non-significant reduction in OLM inhibitory transmission onto CA1 PCs.

5.2 Materials and methods

5.2.1 Materials

All reagents were purchased from Fisher Scientific unless otherwise stated.

5.2.2 Animals

Animal care and use was in accordance with the UK Animals (Scientific Procedures) Act of 1986 and were approved by the University of Bristol ethical review committee. Experiments were performed on male and female mice aged 4 to 15 weeks. C57BL/6J (Charles River) mice were used as the background strain. Cre and FlpO reporter allele mice (The Jackson Laboratory) were used to tag specific neuronal populations: oriens-lacunosum moleculare (OLM) interneurons (Chrna2-Cre, MGI: 5490768), cholinergic neurons (ChAT-Cre, MGI: 3699161; ChAT-FlpO, MGI: 6387001). The Cre reporter mice were crossed with homozygous Ai32 (MGI: 5577173) to generate litters of heterozygous offspring expressing channelrhodopsin (ChR2). These experiments employed a novel mouse line created by crossing Chrna2-Cre mice with ChAT-FlpO mice (Chrna2-Cre x ChAT-FlpO). Mice were housed under regulated temperature (22 ± 2°C) in a normal 12 h light/dark cycle with ad libitum access to food and water.

5.2.3 Surgical procedures

For stereotaxic injections of adeno-associated viruses (AAVs), mice were acclimatised to the animal unit for 5 days prior to surgery. Mice were anaesthetised with isoflurane (1.5-4%) and buprenorphine (Vetergesic) was administered at 0.05 mg/kg for pain relief. The heads of the mice were shaved and the mice were placed in a stereotaxic frame. Chlorhexidine was applied to the scalp and a midline skin incision was made over the dorsal surface of the skull to expose bregma and lambda which acted as landmarks to direct the injection. The connective tissue was carefully removed and the head was leveled by adjusting the dorsal-ventral (DV) measurement of bregma to the reference measurement made at lambda. Lacri-lube (ocular ointment to protect the eyes from drying out) was applied on the eyes of the mice and mice were maintained on a heat pad throughout surgery. A 1 mm diameter burr hole was made over each site of injection. The injections for the medial septum (MS) were made at a medial-lateral angulation of 8° and the coordinates relative to bregma (mm) were anterior-posterior (AP): +0.7 and +0.9, medial-lateral (ML): ±0.6 and DV: -4.5 and 4.0. The bilateral injection coordinates for the dorsal striatum (DS) were AP: +1.1, medial-lateral (ML): ±1.5 and DV: -3.2 and for the hippocampus were AP: -2.0, medial-lateral (ML): ±1.6 and DV: -2.0. A glass micropipette pulled with a horizontal puller (P-87, Sutter Instrument Co., UK) and bevelled with a 25° tip was primed with double distilled water (ddH$_2$O), a 200 nL air bubble to monitor volume injected and back filled with the appropriate AAVs. The The AAVs were infused with a volume of 200 nL per site into the MS and 300 nL per site into the DS and hippocampus at 50 µL/min. After infusion, the micropipette was left at the site for 10 min to allow for the virus to diffuse. Following surgery, the incision was sutured and mice received a 300 µL subcutaneous injection of saline into the flank. Mice were monitored for the next 3-4 weeks, to ensure sufficient virus expression, until they were used in experiments.

5.2.3.1 AAV vectors

A list of the AAV vectors used:

- AAV2.9-CAG-DIO-ChRger2-TS-YFP (titre: 3.49 x 10^{13} viral genomes/mL, VectorBuilder with plasmids from Addgene Cat#127239)
- AAV5-Syn-FLEX-rc[ChrimsonR-tdTomato] (titre: 8.50 x 10^{12} viral genomes/mL, Addgene Cat#62723-AAV5)
- AAV1-CAG-FLEXFRT-ChR2(H134R)-mCherry (titre: 1.50 x 10^{13} viral genomes/mL, Addgene Cat#75470-AAV1)
- AAV9-hSyn-GRAB-ACh3.0 (titre: 2.03 x 10^{13} viral genomes/mL, WZ Biosciences)
- AAV2.1-EF1a-fDIO-EYFP (titre: 2.10 x 10^{13} viral genomes/mL, Addgene Cat#55641-AAV1)
- AAV2.9-ChAT-Cre-P2A-EGFP (titre: 2.59 x 10^{12} viral genomes/mL, BrainVTA Cat#PT-0652)

5.2.4 Hippocampal slice preparation

Following cervical dislocation and decapitation, mouse brains were removed and dissected in ice-cold cutting solution (in mM: 205 sucrose, 10 glucose, 26 $NaHCO_3$, 2.5 KCl, 1.25 NaH_2PO_4, 0.5 $CaCl_2$, 5 $MgSO_4$) saturated with 95% O_2 and 5% CO_2. 400 µm thick parasagittal slices were prepared with a Leica VT1200 vibratome (Leica, Germany) and slices were incubated in artificial cerebrospinal fluid (aCSF; in mM: 124 NaCl, 10 glucose, 24 $NaHCO_3$, 3 KCl, 1.25 NaH_2PO_4, 2.5 $CaCl_2$, 1.3 $MgSO_4$) at 34°C for 25 min and then stored at room temperature for at least 30 min before use.

5.2.5 Electrophysiology

In all electrophysiology experiments, slices were transferred to a submerged recording chamber perfused with aCSF at a flow rate of 2.5 ml/min and kept at 32°C via an inline heating system (TC-324, Warner Instruments). The constant flow and removal of aCSF was maintained by a 101U/R peristaltic pump (Watson Marlow) and a Dymax 5 vacuum pump (Charles Austin), respectively. Target cells were visualised with infrared differential interference contrast optics on an Olympus BX-50WI microscope. Patch electrodes with a resistance of 3-7 MΩ were pulled from borosilicate filamented glass capillaries (1.5 OD x 0.86 ID x 100 L mm, Harvard Apparatus) with a horizontal puller (P-87, Sutter Instrument Co., UK).

Recordings from hippocampal neurons were made with a MultiClamp 700A amplifier (Molecular Devices, USA), filtered at 2.4 kHz and digitised at 10 kHz with a CED Power1401 data acquisition board and Signal 5.12 acquisition software (CED, Cambridge, UK). Cells were voltage clamped at 0 (not corrected for junction potential). Series resistance was monitored throughout the experiments and cells that showed >40% change were discarded from subsequent analysis. Recordings were also rejected from analysis if the series resistance was greater than 35 MΩ.

Inhibitory synaptic responses were evoked optogenetically with square light steps applied via mounted LED illuminators (M395L5, M470L2 and M590L2 with a LEDD1B driver, Thorlabs, US) through 5x or 10x objective lenses.

5.2.5.1 Independent optical excitation of OLM interneurons

To characterise the kinetics and the optical separability of the opsins (ChR2, ChRger, Chrimson), Chrna2-Cre mice were either crossed with Ai32 (Chrna2-Cre x Ai32) or virally injected with AAV5-Syn-FLEX-rc[ChrimsonR-tdTomato] or AAV2.9-CAG-DIO-ChRger2-TS-YFP in the hippocampus. Whole-cell patch-clamp recordings were made in CA1 pyramidal neurons with patch pipettes filled with an intracellular solution containing (in mM) 130 $CsMeSO_3$, 4 NaCl, 10 HEPES, 0.5 EGTA, 10 TEA, 1 QX-314 chloride, 2 Mg-ATP, 0.5 Na-GTP. Cells were voltage clamped at 0 mV, far from the chloride reversal potential, to maximise the amplitude of inhibitory postsynaptic current (IPSC) recordings.

For the characterisation of channel kinetics, patch-clamped cells were photostimulated with 5 ms light pulses through a 10x objective lens from a 470 nm LED (M470L2, Thorlabs, US) at irradiance of 14.29 mW/mm^2 for ChR2 and ChRger, and by a 590 nm LED (M590L2, Thorlabs, US) at irradiance of 2.11 mW/mm^2 for Chrimson. Channel closing rates (τ_{off}) were calculated by fitting a monoexponential curve to the photocurrent after the max amplitude had been reached. The time to 90% of the peak of the response was measured after the beginning of light onset.

Optical excitation separability was examined by photostimulating patch-clamped cells with 2 or 5 ms light steps at increasingly higher irradiances and measuring the photocurrents produced by each opsin. By observing the amplitude of the photocurrents produced by the activation of opsins under different light wavelengths and irradiances, valuable insights can be gained regarding the opsins' sensitivity to specific ranges of light wavelengths and intensities. The different output light intensities were achieved by sending increasingly larger amplitude voltage steps from the CED Power1401 data acquisition board, via Signal 5.12 acquisition software, to the LEDD1B driver whilst on "Modulation (MOD)" mode. Exact voltage outputs from the CED acquisition board and resultant irradiances recorded using an optical power meter (PM100D power meter console with a S120C photodiode power sensor, Thorlabs) are presented in (Table 5.1). Experiments were repeated in the presence of 1 μM TTX, 0.1 mM 4-AP and 4 mM Ca^{2+} to isolate terminal release.

Voltage signals (mV)	395 nm mW/mm^2)	470 nm (mW/mm^2)	590 nm (mW/mm^2)
1	0.085	0.022	0.000
3	0.787	0.309	0.016
5	1.515	0.638	0.040
10	3.345	1.440	0.127
20	6.976	3.038	0.342
30	10.464	4.626	0.575
50	16.703	7.692	1.054
60	19.378	9.170	1.280
80	24.028	11.853	1.714
100	27.710	14.621	2.090
120	27.710	14.289	2.114

TABLE 5.1: **The voltage outputs from the CED acquisition board and resultant irradiances for each LED through a 10x objective lens.**

Voltage signals (mV)	470 nm (mW/mm^2)
1	0.005
3	0.056
5	0.108
10	0.234
20	0.468
30	0.683
50	1.070
60	1.241
80	1.581
100	1.880
120	1.880

TABLE 5.2: **The voltage outputs from the CED acquisition board and resultant irradiances for a 470 nm LED through a 5x objective lens.**

5.2.5.2 Endogenous acetylcholine release in ChAT-cre x Ai32

In ChR2-expressing Chat-Cre animals (ChAT-Cre x Ai32), whole-cell patch-clamp recordings were made in OLM interneurons with patch pipettes filled with an intracellular solution containing (in mM) 130 $CsMeSO_3$, 4 NaCl, 10 HEPES, 0.5 EGTA, 10 TEA, 1 QX-314 chloride, 2 Mg-ATP, 0.5 Na-GTP. 5 ms light steps at different irradiances were used to elicit responses in OLM interneurons mediated by nicotinic receptors (Table 5.2).

5.2.5.3 Effect of endogenous acetylcholine on dendritic inhibition

Chrna2-Cre animals were injected with AAV2.9-ChAT-Cre-P2A-EGFP and AAV2.9-CAG-DIO-ChRger2-TS-YFP in the MS, and AAV5-Syn-FLEX-rc[ChrimsonR-tdTomato] in the hippocampus.

Slices were prepared after 3-4 weeks of viral transfection and whole-cell patch-clamp recordings were made in CA1 pyramidal neurons with patch pipettes filled with an intracellular solution containing (in mM) 130 $CsMeSO_3$, 4 NaCl, 10 HEPES, 0.5 EGTA, 10 TEA, 1 QX-314 chloride, 2 Mg-ATP, 0.5 Na-GTP. Cells were voltage clamped at 0 mV to obtain inhibitory currents. OLM interneuron IPSCs were evoked by two 5 ms amber light pulses at 10 Hz every 10 s through a 10x objective lens from a 590 nm LED (M590L2, Thorlabs, US) at irradiance of 2.11 mW/mm^2. After a stable 5 min baseline period was achieved, 5 ms blue light pulses at 2 Hz from a 470 nm LED (M470L2, Thorlabs, US) at irradiance of 0.64 mW/mm^2 were added for 5 min to evoke physiologically maximal acetylcholine release (Palacios-Filardo et al., 2021; Jing et al., 2018). Then, the blue light photostimulation was switched off for the remaining 20 min of the experiment. If the blue light photostimulation affected the inhibitory activity of the OLM interneurons, then the experiment was repeated in the presence of atropine (25 μM) and mecamylamine (50 μM). Consecutive traces were averaged to produce a mean response every 30 seconds.

After the experiment, cells were voltage clamped at -70 mV; both amber and blue light pulses were used to examine off-target expression of the viruses. If there were any responses, first atropine (25 μM) and mecamylamine (50 μM) were applied to check whether it was cholinergic and then NBQX (20 μM) and DAPV (25 μM) were applied to check if it was excitatory.

OLM interneurons in the stratum oriens were also patch-clamped and voltage clamped at -70 mV to obtain excitatory cholinergic responses mediated by nicotinic acetylcholine receptors.

5.2.6 Two-photon imaging characterisation of ACh3.0 sensor

A two-photon microscope (Scientifica) equipped with a tuneable Ti:Sapphire pulsed laser (100 fs pulsewidth, MaiTai HP, SpectraPhysics) and ScanImage MATLAB acquisition software (MBF Bioscience) was used for these imaging experiments. The laser was tuned to an excitation wavelength of 920 nm and the power was typically set at 50-65 mW. The hippocampi of C57BL/6J (Charles River) or ChAT-Cre x Ai32 mice were injected with AAV9-hSyn-GRAB-ACh3.0. After 3-4 weeks time for virus expression, parasagittal hippocampal slices were prepared from these animals. Slices were transferred to a submerged recording chamber perfused with aCSF at a flow rate of 2.5 ml/min and kept at 32°C via an inline heating system (TC-324, Warner Instruments). The constant flow of aCSF was maintained by a MINIPULS 3 peristaltic pump (GILSON). Images were acquired with a 60x water immersion lens at resolutions of either 256 x 256 pixels or 512 x 512 pixels. A CED Power1401 data acquisition board and Signal 5.12 acquisition software (CED, Cambridge, UK) was used to trigger the onset of scanning, light stimulation from a 470 nm laser-diode fibre (light source, Doric lenses LDFLS_473/070; fibre optic cannulae, Thorlabs CFMLC22L20) and the opening of a shutter to the PMTs. Images were processed in Python 3.9 and $\Delta F/F_0$ was calculated by subtracting and dividing by the mean baseline image for all images. The mean baseline image was generated by averaging per pixel across the baseline period. Each datapoint in the time course traces represent an average of 5 consecutive images. The datapoints on the bar plots and example raw and normalised images at each experiment state are an average of 60 images within each period.

5.2.6.1 Carbachol dose-response

Once a location with expression of the ACh3.0 sensor in CA1 of slices prepared from C57BL/6J mice was identified, 256 x 256 pixels or 512 x 512 pixels images were acquired at 3.42 Hz or 0.86 Hz, respectively. After the acquisition of 60 images for the baseline period, 100 μM carbachol (CCh) was bath applied for a further 600 images before being washed out of the system. The numbers stated relate to the acquisition of 256 x 256 pixels images; half the number of images were acquired at each experiment state for 512 x 512 pixels images.

To examine the fluorescence of the ACh3.0 sensor with different concentrations of CCh, 512 x 512 pixels images were acquired at 3.42 Hz. After the acquisition of 30 images for the baseline period, increasingly higher concentrations of CCh were bath applied: 0.3, 1, 3, 10 and 100 μM for 300 images each.

5.2.6.2 Detecting photostimulated release of endogenous acetylcholine

After locating ACh3.0 expression in CA1 of slices prepared from ChAT-Cre x Ai32 animals, 256 x 256 pixels images were acquired at 1 Hz. A baseline period of 60 images was acquired before the onset of 5 ms light pulses at 1 Hz to optogenetically stimulate the release of endogenous acetylcholine (Jing et al., 2018). During these light pulses, the shutter to the PMTs was closed to prevent the light emitted by the laser-diode fibre burning the PMTs. Image acquisition resumed 40 ms post-photostimulation. The photostimulation period lasted 20 s (29 images) before a quiescent period of 40 s. This pattern was repeated two more times to make a total of three photostimulation periods. Once the photostimulation period was over, 100 μM CCh was applied as a reference to compare the change in fluorescence from the endogenous release of acetylcholine. These images were acquired at 0.1 Hz for 10 min followed by an additional 10 min wash out period. The same protocol, without the drug application period, was performed in slices prepared from C57BL/6J mice and also in the absence of a slice in the recording chamber to act as controls.

5.2.7 Retrograde tracing of septo-hippocampal innervations

Cholera Toxin Subunit B (CTB) with Alexa Fluor 594 Conjugate (Invitrogen C34777) was injected into the hippocampus of Chrna2-Cre x Ai32 mice. Animals were left for 3-4 weeks to allow sufficient time for the tracer to retrogradely transport from the hippocampal terminals to the cells in the MS and then animals were sacrificed for immunohistochemistry.

5.2.8 Characterisation of Chrna2-Cre x ChAT-FlpO mouse line

AAV1-CAG-FLEXFRT-ChR2(H134R)-mCherry was injected into the MS of Chrna2-Cre x ChAT-FlpO animals. AAV1-CAG-FLEXFRT-ChR2(H134R)-mCherry and AAV2.1-EF1a-fDIO-EYFP were also separately injected into the DS of these animals. After 3-4 weeks of virus expression, animals were sacrificed for immunohistochemistry.

5.2.9 Characterisation of AAV2.9-ChAT-Cre-P2A-EGFP

AAV2.9-ChAT-Cre-P2A-EGFP was injected into the MS of C57BL/6J mice animals. After 3-4 weeks of virus expression, animals were sacrificed for immunohistochemistry.

5.2.10 Immunohistochemistry

Prior to immunohistochemical staining of target neurons, mice brains were fixed via perfusion fixation. First, mice received an intraperitoneal injection of sodium pentobarbital at a lethal dose of 100 mg/kg. After the cessation of breathing and reflexive responses from foot pinching, the abdomen was opened and the rib cage was separated to expose the heart. A small incision was made in the right atrium to allow blood and old solution to drain whilst 20 ml of 0.01 M phosphate-buffered saline (PBS) was being injected into the left ventricle followed by another 20 ml of 4% paraformaldehyde in 0.01 M PBS (PFA). The mice were then decapitated, the skulls were cut open following the midline sutures (sagittal and interfrontal sutures) and the brains were extracted. Before slicing, the brains were submerged in 4% PFA for 24 hours and then transferred to 30% sucrose in 0.01 M PBS until they sank to the bottom of the container which took approximately 48 hours. 40 µm thick parasagittal slices were obtained using a freezing microtome and stored in 0.01 M PBS with 0.01% sodium azide.

For immunohistochemistry, free-floating slices were washed with gentle agitation in 0.01 M PBS for 10 min and three additional times (10 min each) in 0.01 M PBS with 0.2% Triton X-100 (0.2% PBST). Slices were then incubated in a blocking solution containing 5% donkey serum (Sigma-Aldrich D9663-10ML) in 0.3% PBST for 1 hour at room temperature. Slices were subsequently incubated at 4°C for 24 hours in blocking solution containing primary antibodies for visualising target neurons. After 24 hours, slices were brought up to room temperature for 1 hour. Following three washes with 0.2% PBST (10 min each), slices were incubated in blocking solution containing secondary antibodies for 2.5 hours at room temperature. Primary and secondary antibodies used are listed below. Finally, slices were washed twice in 0.2% PBST and once in 0.01 M PBS (10 min each), before being stained with DAPI (1:10000, Invitrogen D1306), mounted on microscope slides and covered with rectangular glass cover slips using Fluorsave (Merck Millipore).

A different protocol was used for visualising Chrimson expression of virally injected Chrna2-Cre animals. Free-floating slices were washed three times with gentle agitation in 0.01 M PBS (10 min each). Afterwards, slices were incubated for 10 min in 1% H_2O_2 in PBS followed by three 10 min washes in PBST. Slices were then incubated in a blocking solution containing 5% normal goat serum and 2.5% bovine serum albumin in 0.2% PBST for 1 hour at room temperature. Slices were subsequently incubated at 4°C for 48 hours in blocking solution containing anti-mCherry (1:1000 Abcam ab205402) and anti-SST (1:1000 Abcam ab111912) for visualising Chrimson and SST interneurons, respectively. After 48 hours, slices were brought up to room temperature for 1 hour. Following three washes with 0.2% PBST (10 min each), slices were incubated in blocking solution containing anti-chicken Alexa Fluor 594 (1:500 Abcam ab150172) and anti-rabbit Alexa Fluor 488 (1:250 Vector Laboratories DI-1488-1.5) for 24 hours at 4°C. Finally, slices were washed four times in 0.2% PBST, before being stained with DAPI (1:10000, Invitrogen D1306), mounted on microscope slides and covered with rectangular glass cover slips using Fluorsave (Merck Millipore).

A list of the primary antibodies used:

- Anti-SST (1:1000 Santa Cruz SC-7819) for visualising SST interneurons.
- Anti-SST (1:1000 Abcam ab111912) for visualising SST interneurons.
- Anti-mCherry (1:1000 Abcam ab205402) for visualising Chrimson.
- Anti-mCherry (1:1000 Abcam ab167453) for visualising FlpO-dependent ChR2.
- Anti-GFP (1:1000 Abcam ab13970) for visualising FlpO-dependent GFP and ChR2 (Ai32).
- Anti-CTB (1:1000 Sigma-Aldrich C3062) for visualising CTB.
- Anti-ChAT (1:250 Merck Millipore AB144P) for visualising ChAT neurons.

A list of the secondary antibodies used:

- Anti-rabbit Alexa Fluor 488 (1:250 Vector Laboratories DI-1488-1.5)
- Anti-chicken Alexa Fluor 594 (1:500 Abcam ab150172)
- Anti-rabbit Alexa Fluor 594 (1:500 Invitrogen A-21207)
- Anti-goat Alexa Fluor 488 (1:500 Invitrogen A-11055)
- Anti-goat Alexa Fluor 594 (1:500 Invitrogen A-11058)
- Anti-goat Alexa Fluor 647 (1:400 Invitrogen A-21447)
- Anti-chicken Alexa Fluor 488 (1:500 Jackson ImmunoResearch 703-545-15)

5.2.11 Microscopy

Hippocampal slice sections containing ChRger and Chrimson were examined using a Leica DMI6000B microscope with a HC PL FLUOTAR 10x/0.30 DRY objective lens. The DAPI (excitation 350/50 nm, dichroic 400 nm, emission 460/50 nm), GFP (excitation 480/40 nm, dichroic 505 nm, emission 527/30 nm) and Tx Red (excitation 560/40 nm, dichroic 595 nm, emission 645/76 nm) filter cubes were used. Images were captured using a Photometrics Prime 95B sCMOS camera and processed using Leica Application Suite X (Leica, Germany).

MS and DS sections were examined using a Leica SP8 AOBS confocal laser scanning microscope with a HC PL APO CS2 20x/0.75 DRY objective lens. A 65 mW Ar laser, 20 mW DPSS yellow laser (561 nm), 10 mW Red He/Ne (633 nm) and 50 mW 405 nm diode laser were used to visualise GFP, mCherry, Alexa Fluor 647 and DAPI, respectively. Hybrid detectors were tuned to the emission spectrum of each fluorophore and ensured no overlap between spectra. Images were processed using Leica Application Suite X (Leica, Germany).

5.2.12 Imaging analysis

All hippocampal images were taken in slices typically used for ex vivo electrophysiological recordings. Hippocampal layers within the CA1 were identified based on density of DAPI-stained nuclei, 250 x 600 µm sections of CA1 (covering all layers) were processed using ImageJ software and fluorescence intensity analysis of each channel (fluorophore of interest) was performed using MATLAB (R2022b). The normalised intensity profile was calculated by averaging the fluorescence intensity across the width of the image to provide representative intensities from the top to the bottom of the CA1. The maximum fluorescence intensity was used to normalise the trace and the normalised trace was averaged across slices. The fluorescence contribution per layer was calculated as an average across the layer.

400 x 400 µm representative sections of the MS and DS were processed in ImageJ software for cell counting. Cell counting was performed using the ModularImageAnalysis (MIA) plugin in ImageJ software with a custom workflow developed by Stephen Cross (Wolfson Imaging Facility, University of Bristol).

5.2.13 Statistical analysis

The experimental unit was defined as a slice or a cell with only one cell recorded per slice. Data are reported as means ± SEM or medians and were analysed for statistical significance using SPSS (IBM Software). The normality of the data was assessed using the Shapiro-Wilk test in order to determine the appropriate parametric or non-parametric test.

5.3 Results

5.3.1 Characterisation of ChRger and Chrimson expression in OLM interneurons

To investigate the viability of ChRger and Chrimson as opsins for the optical excitation of independent neural populations, the expression selectivities of the viral constructs were first examined to determine the likelihood of non-specific activation. The hippocampi of Chrna2-Cre mice were stereotaxically injected with the Cre-dependent AAV2.9-CAG-DIO-ChRger2-TS-YFP or AAV5-Syn-FLEX-rc[ChrimsonR-tdTomato] vectors to selectively express ChRger or Chrimson in OLM interneurons. Subsequent histological analysis of Chrimson-injected mice (Chrna2-Cre x Chrimson) showed Chrimson-positive somata and innervations were almost exclusively located in the stratum oriens (SO) and stratum lacunosum moleculare (SLM), respectively (Fig. 5.1A). A repeated measures ANOVA determined a main effect of hippocampal layer on the expression of Chrimson fluorescence intensity ($F(3, 12) = 368.663$, $P < 0.001$). Post hoc pairwise comparisons using the Bonferroni correction revealed that there were no significant differences in fluorescence intensity between the stratum pyramidale (SP; 11.85 ± 2.62%) and stratum radiatum (SR; 14.59 ± 2.19%) layers (-2.74 (95%, -6.29 to 0.81) %, $p = 0.120$), and both layers had significantly lower fluorescences compared to the SO and SLM (SO-SP: 15.86 (95%, 9.87 to 21.84) %, $p = 0.001$; SO-SR: 13.11 (95%, 7.38 to 18.84) %, $p = 0.002$; SP-SLM: -53.71 (95%, -66.22 to -41.20) %, $p < 0.001$; SR-SLM: -50.97 (95%, -63.28 to -38.65) %, $p < 0.001$). The SLM displayed significantly the highest fluorescence amongst all layers (SLM-SO: 37.85 (95%, 28.94 to 46.77) %, $p < 0.001$). Immunohistochemistry of hippocampi from mice injected with the ChRger adeno-associated virus (AAV) demonstrated a similar expression profile (Fig. 5.1B). A repeated measures ANOVA with a Greenhouse-Geisser correction for departure from sphericity revealed that mean normalised fluorescence intensity differed significantly across the hippocampal layers ($F(1.256, 10.050) = 48.685$, $P < 0.001$). Simple effects analysis, using the Bonferroni correction for multiple comparisons, revealed the same significant differences as seen with the Chrimson expression (SO-SP: 11.03 (95%, 1.04 to 21.02) %, $p = 0.030$; SO-SR: 12.42 (95%, 0.69 to 24.15) %, $p = 0.037$; SO-SLM: -29.72 (95%, 0.01 to -7.50) %, $p = 0.010$; SP-SR: 1.39 (95%, -2.94 to 5.72) %, $p = 1.000$; SP-SLM: -40.75 (95%, -56.00 to -25.50) %, $p < 0.001$; SR-SLM: -42.14 (95%, -54.80 to -29.48) %, $p < 0.001$). Immunohistochemical staining for somatostatin (SST) also showed a large overlap of Chrimson and ChRger expression with SST which further supports the specificity of these AAVs as the OLM interneurons are a subpopulation of SST interneurons (Fig. 5.1). Together, these results suggest that both viral constructs are highly selective for Cre-expressing cells as the expression profiles of the opsins heavily reflect the distribution of OLM interneurons and their innervations in hippocampal CA1.

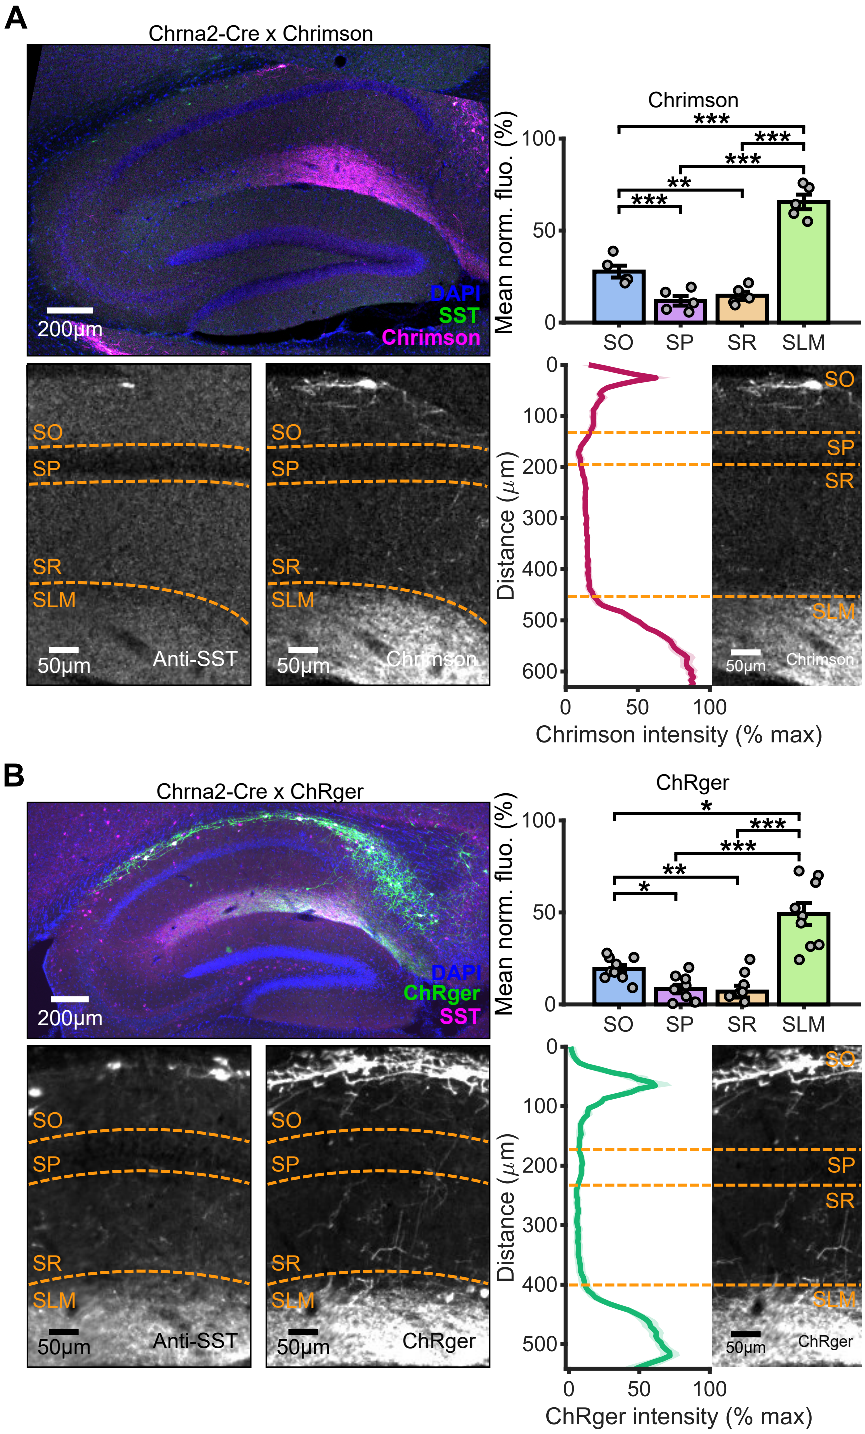
A
Chrna2-Cre x Chrimson
DAPI
SST
Chrimson
200µm
Chrimson
***
***
***
**
***
Mean norm. fluo. (%)
100
50
0
SO
SP
SR
SLM
SO
SP
SR
SLM
50µm
Anti-SST
SO
SP
SR
SLM
50µm
Chrimson
Distance (µm)
0
100
200
300
400
500
600
0
50
100
Chrimson intensity (% max)
SO
SP
SR
SLM
50µm
Chrimson
B
Chrna2-Cre x ChRger
DAPI
ChRger
SST
200µm
ChRger
*
***
***
**
*
Mean norm. fluo. (%)
100
50
0
SO
SP
SR
SLM
SO
SP
SR
SLM
50µm
Anti-SST
SO
SP
SR
SLM
50µm
ChRger
Distance (µm)
0
100
200
300
400
500
0
50
100
ChRger intensity (% max)
SO
SP
SR
SLM
50µm
ChRger

FIGURE 5.1: **Chrimson and ChRger selectively express in stratum oriens OLM interneurons.** (**A**; top left) Immunohistochemical staining in hippocampal slice of Chrna2-Cre animals with viral expression of Chrimson (Chrna2-Cre x Chrimson) showing expression of SST (green) and Chrimson in OLM interneurons (magenta). Sections of CA1 displaying expression patterns of SST (bottom left) and Chrimson (bottom middle) in the different hippocampal layers: stratum oriens (SO), stratum pyramidal (SP), stratum radiatum (SR) and stratum lacunosum moleculare (SLM). The normalised intensity profile shows the

mean Chrimson (bottom right) fluorescence level as a function of distance from the top (SO) to the bottom (SLM) of CA1. Bar plot demonstrates the mean fluorescence within each layer for Chrimson (top right) (N = 9 slices from three animals). (**B**) Same as (**A**) but looking at the expression of ChRger in virally-transfected Chrna2-Cre animals (Chrna2-Cre x ChRger). Data represented as mean ± SEM; * p < 0.05, ** p < 0.01, *** p < 0.001 (repeated measures ANOVA with Bonferroni-corrected simple effects analysis.

5.3.2 Independent optical excitation of OLM interneurons

Strong efforts have been made to develop opsins with red-shifted excitation spectra (e.g. VChR1, C1V1, ReaChR and Chrimson) to avoid spectral overlap with blue light-sensitive opsins when they are used in combination to optically excite different neural populations (Zhang et al., 2008; Yizhar et al., 2011; Klapoetke et al., 2014; Lin et al., 2013). However, action spectra data collated from Klapoetke et al. (2014) and Bedbrook et al. (2019) highlights that residual absorption of blue light remains a fundamental issue with these red-shifted opsins (Fig. 5.2A). Therefore, after establishing the expression selectivity of ChRger and Chrimson, it was critical to examine the spectral overlap of the opsins to determine if there was significant cross-talk between the opsins that would violate independent optical stimulation of specific neural populations. To assess the wavelength sensitivity, photocurrent amplitude and channel kinetics of Chrimson, ChRger and ChR2 (the control blue light-sensitive opsin), voltage-clamp recordings at 0 mV were made from CA1 PCs after photostimulating OLM interneurons independently expressing each opsin with 5 ms pulses of violet (395 nm), blue (470 nm) and amber (590 nm) light (Fig. 5.2B). Channel kinetics were measured from IPSCs optogenetically evoked at each opsins' respective peak wavelengths: amber 590 nm light (irradiance of 2.114 mW/mm^2) for Chrimson and blue 470 nm light for ChR2 and ChRger (irradiance of 14.289 mW/mm^2). A Kruskal-Wallis H test determined a statistically significant difference in off-kinetics between the different opsins ($\chi^2(2) = 8.549$, $p = 0.014$) with a mean decay time of 8.22 ms for ChR2, 14.44 ms for ChRger and 3.00 ms for Chrimson. Simple effects analysis, using the Bonferroni correction for multiple comparisons, revealed that the difference was between Chrimson and ChRger (11.44 ± 4.63 ms, $p = 0.040$) (Fig. 5.2C-D). There were no statistically significant differences between the off-kinetics of ChR2 and ChRger (-6.22 ± 2.79 ms, $p = 0.077$), as well as Chrimson and ChR2 (5.22 ± 1.13 ms, $p = 0.776$). A one-way ANOVA also showed that there was a significant differences in on-kinetics between the different opsins ($F(2,17) = 7.274$, $p = 0.005$) with a mean rise time of 11.70 ms for ChR2, 8.92 ms for ChRger and 15.45 ms for Chrimson. A Tukey HSD post hoc test revealed that the difference was between Chrimson and ChRger (6.53 ± 1.86 ms, $p = 0.007$) (Fig. 5.2E). There were no statistically significant differences between the off-kinetics of ChR2 and ChRger (2.78 ± 1.12 ms, $p = 0.059$), as well as Chrimson and ChR2 (3.75 ± 1.86 ms, $p = 0.138$). This data suggests that there are no significant differences in channel kinetics between ChR2 and ChRger but ChRger has significantly slower off-kinetics and faster on-kinetics compared to Chrimson.

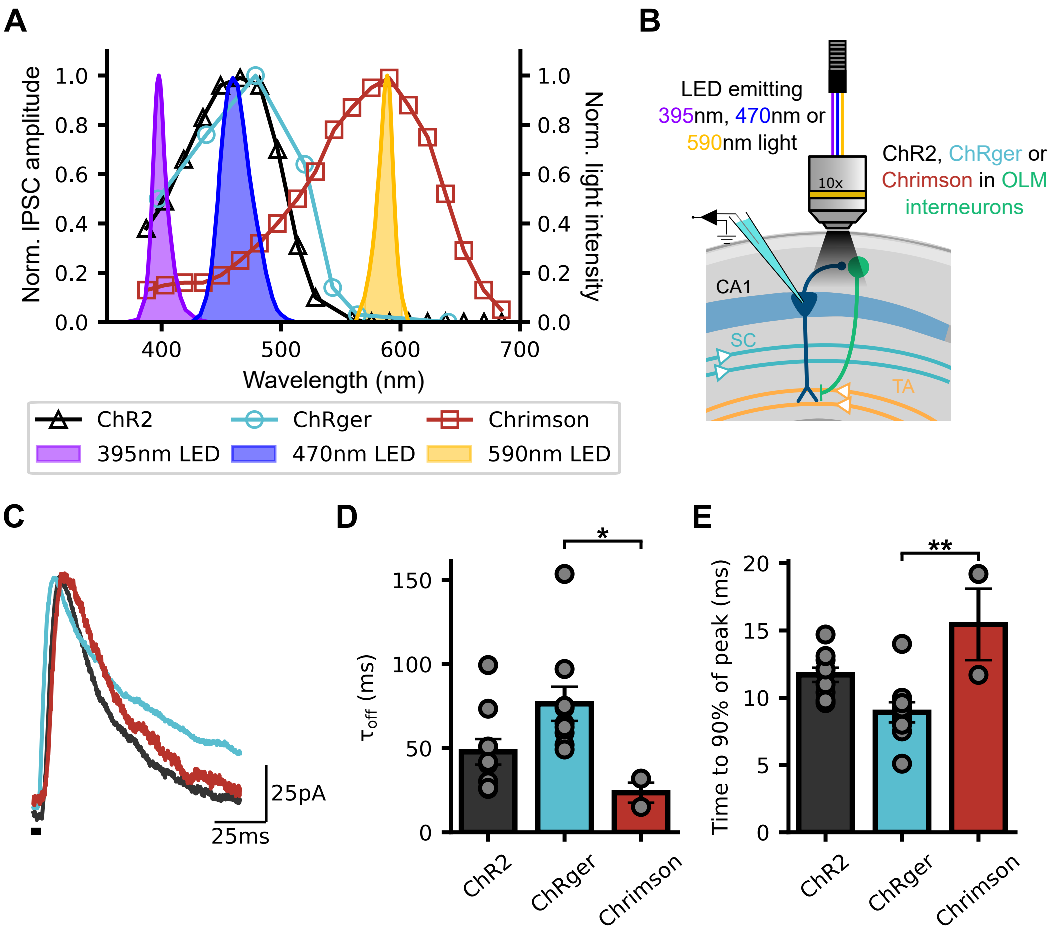

FIGURE 5.2: **The channel kinetic properties of ChR2, ChRger and Chrimson.** (**A**) Action spectra of each opsin (data collated from Klapoetke et al. (2014) and Bedbrook et al. (2019)) and the light intensity spectra of the LEDs used in the experiments. (**B**) Schematic of the hippocampal slice recording setup: whole-cell patch-clamp recordings of IPSCs in CA1 pyramidal cells (dark blue) from photostimulated dendrite-targeting OLM interneurons (green) with different wavelength LEDs. (**C-E**) Channel kinetic properties measured from 5 ms light pulses. All opsins were illuminated near their respective peak wavelength: amber 590 nm light for Chrimson and blue 470 nm light for ChR2 and ChRger. (**C**) Example normalised IPSC traces of each opsin highlighting representative differences in synapse kinetics. Summary of channel off-kinetics (τ_{off}) (**D**) calculated from a monoexponential fit of the IPSC decay (independent-samples Kruskal-Wallis test with Bonferroni-corrected simple effects analysis; N = 9, ChR2; N = 9, ChRger; N = 2, Chrimson) and on-kinetics (**E**) as the time taken to reach 90% of the peak response post-illumination (one-way ANOVA with Tukey HSD post-hoc multiple comparisons; N = 9, ChR2; N = 9, ChRger; N = 2, Chrimson). Data represented as mean ± SEM; * $p < 0.05$, ** $p < 0.01$.

Next, the wavelength sensitivity and photocurrent amplitudes of each opsin were measured via photostimulation using 5 ms pulses of violet (395 nm), blue (470 nm) and amber (590 nm) light with increasingly higher light intensities. Under amber light conditions, Chrimson began to produce responses of 30.14 ± 8.60 pA at an irradiance of 0.575 mW/mm^2 and at the maximum irradiance of 2.114 mW/mm^2 it produced an IPSC amplitude of 86.45 ± 22.44 pA. ChR2 and ChRger displayed no responses across all amber light irradiances. Under blue light conditions, ChRger produced the largest photocurrents (182.61 ± 28.93 pA) at the maximum irradiance of 14.289 mW/mm^2, 37% and 171% greater than responses produced by ChR2 (133.18 ± 14.39 pA) and Chrimson (67.46 ± 15.06 pA), respectively. At a blue light irradiance of 1.440 mW/mm^2, Chrimson did not produce any responses (photocurrents started at 3.038 mW/mm^2) but ChRger was already producing IPSCs of 139.89 ± 32.98 pA which were over 3.9 times larger than ChR2 currents (35.73 ± 13.37 pA). Even with 2 ms blue light pulses, ChRger was producing similar responses to those seen at 5 ms (163.11 ± 26.55 pA). Under violet light conditions, Chrimson was unresponsive until after an irradiance of 16.703 mW/mm^2 but at 16.703 mW/mm^2, ChRger was producing responses of 138.54 ± 43.12 pA (Fig. 5.3A-B). To directly examine the terminal release evoked by the opsins, these experiments were repeated in the presence of 1 μM TTX, 0.1 mM 4-AP and 4 mM Ca^{2+}. Responses recorded under these conditions were used as an underestimation of how excitable opsin-expressing fibres would be to photostimulation. Under blue light conditions, ChR2 produced no responses whereas ChRger at an irradiance of 1.440 mW/mm^2, a blue light intensity that did not activate Chrimson, produced IPSCs of 114.50 ± 24.40 pA. At a violet light irradiance (16.703 mW/mm^2) that did not evoke Chrimson activity in the drug-free state, ChRger currents were 173.02 ± 54.50 pA (Fig. 5.3C). Together, these results indicate that ChRger is considerable more sensitive than ChR2: ChRger produces larger photocurrents and also at much lower light intensities. It also suggests that with the combination of the Chrimson and Chrger opsins, under certain light conditions (blue light irradiance of 1.440 mW/mm^2 or below), it is possible to achieve optical activation of distinct neural populations without cross-talk.

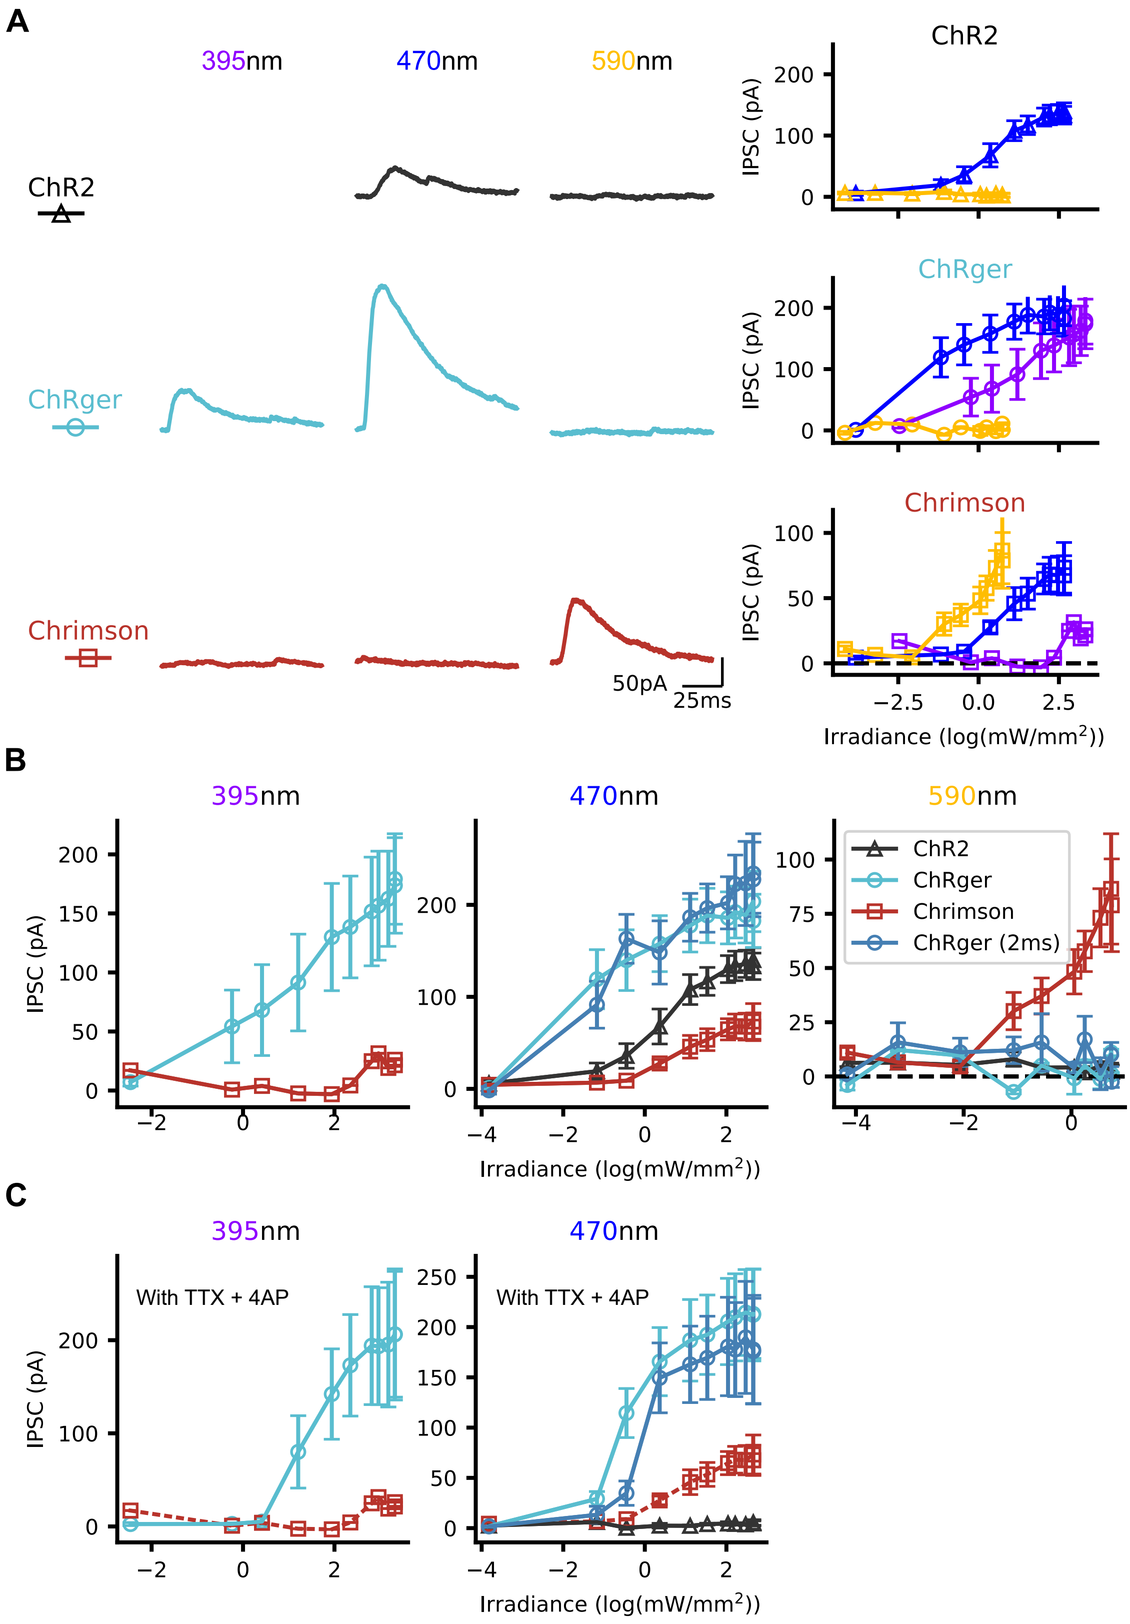

FIGURE 5.3: **Characterisation of opsin violet (395 nm), blue (470 nm) and amber (590 nm) light**

sensitivities for two-colour excitation in acute hippocampal slices. (**A**; left) Representative voltage-clamp traces of photocurrents produced by each opsin under 5 ms violet, blue and amber light pulses at irradiances of 10.5 mW/mm^2, 0.638 mW/mm^2 and 2.11 mW/mm^2, respectively. (right) The photocurrent amplitudes of each opsin evoked with different optical wavelengths and irradiances. (**B**) Comparison of IPSC amplitudes across different opsins obtained under the same illumination conditions. (**C**) Same as (**B**) but cells were in the presence of 1 µM TTX, 0.1 mM 4-AP and 4 mM Ca^{2+}. Dashed red line is a reference of Chrimson responses without the presence of drug. Data represented as mean ± SEM.

5.3.3 Photostimulating endogenous release of ACh in ChAT-Cre x Ai32 mice

After demonstrating that it is possible to achieve optical excitation of Chrimson and Chrger without spectral cross-talk, it was important to examine whether the light conditions required would be suitable for the optogenetic stimulation of ACh release from cholinergic fibres. To test this, OLM interneurons, identified by their morphology and location in the stratum oriens, were patched in voltage clamp mode at -70 mV in hippocampal slices prepared from mice with Cre-mediated ChR2 expression under the control of the ChAT promoter (ChAT-Cre x Ai32) (Fig. 5.4A). Nicotinic ACh receptor (nAChR) mediated EPSC recordings were made from these cells during photostimulation of ACh release with increasing blue light irradiances. Photostimulated responses began appearing with blue light irradiances of at least 0.683 mW/mm^2 which is much lower than the irradiance required to elicit Chrimson photocurrents (3.038 mW/mm^2). At an irradiance of 1.241 mW/mm^2, the EPSC amplitude produced was 25.47 pA. A Wilcoxon signed-rank test showed that there was no significant difference in EPSC amplitude when comparing resting state responses with no light stimulation (2.05 ± 1.88 pA) and responses produced by the optogenetic stimulation (33.29 ± 25.51 pA) of ACh release ($Z = 1.342$, $p = 0.180$) (Fig. 5.4B). However, a statistical test performed with such low sample size numbers (N = 2) is unreliable. Despite demonstrating that a blue light irradiance of 1.440 mW/mm^2 is capable of activating cholinergic fibres expressing a much less sensitive blue light-sensitive opsin (ChR2), there is not a sufficiently large enough sample size to draw definitive conclusions.

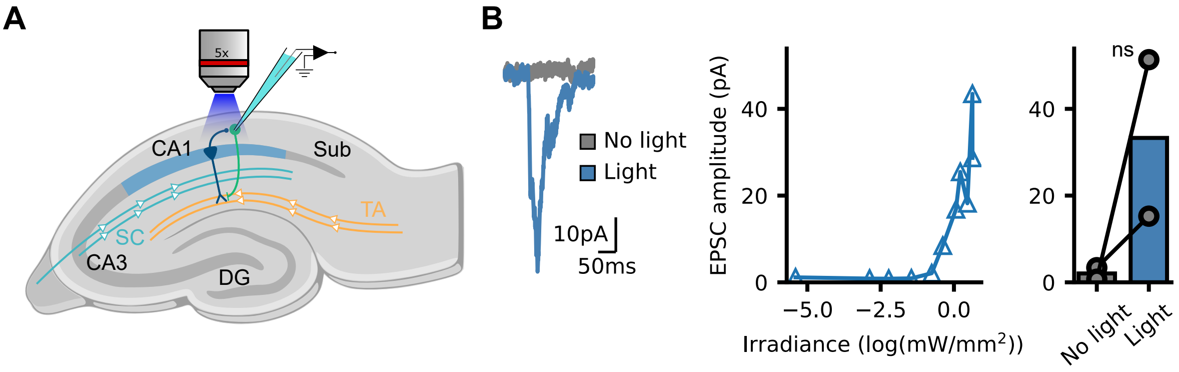

FIGURE 5.4: **Photostimulation of cholinergic innervations onto OLM interneurons.** (**A**) Schematic of the hippocampal slice recording setup: whole-cell patch-clamp recordings of EPSCs in

OLM interneurons evoked by 2 ms blue (470 nm) light pulses. (**B**; left) Example voltage-clamp traces of EPSCs evoked from optogenetic stimulation of cholinergic fibres and (right) the amplitude of responses at different irradiances. (**C**) Summary of mean EPSC amplitudes with (blue) and without (grey) photostimulation at maximum irradiance of 1.88 mW/mm^2.

5.3.4 Characterisation of ACh3.0 sensor

While nAChRs are present on OLM interneurons and can modulate their activity, successfully evoking nAChR-mediated responses in these interneurons is not guaranteed and these responses can vary depending on various factors such as the specific subtype of nAChR involved, the level of cholinergic innervations preserved in the hippocampal slice, thus the concentration of ACh available, and their proximity to the OLM interneurons (Leão et al., 2012; Son and Winzer-Serhan, 2008). Additionally, ACh is rapidly metabolised and reuptaken once released so using nAChR-mediated excitatory currents seen in OLM interneurons as a proxy for endogenous ACh release and quantifying its physiological concentration is quite unreliable. The ACh3.0 sensor is a biosensor used to detect physiological concentrations of ACh: it combines the type 3 muscarinic ACh receptor with a circularly permutated GFP so that upon ACh binding, a conformational change on the muscarinic receptor induces a sensitive fluorescence response (Jing et al., 2020). Therefore, this biosensor was expressed in CA1 neurons via stereotaxic hippocampal injections with AAV9-hSyn-GRAB-ACh3.0 to determine whether it was an effective tool for precise spatio-temporal detection of cholinergic transmission (Fig. 5.5A). Two-photon imaging of acute hippocampal slices showed that bath application of 100 μM CCh evoked a robust increase in fluorescence which decreased as the drug was removed from the system (Fig. 5.5B). To measure the sensitivity of the ACh3.0 sensor, progressively higher concentrations of CCh, from 0.3 μM to 100 μM, were bath-applied. When examining the fluorescence of the entire image, there was an initial increase in fluorescence in response to the first drug application of 0.3 μM but this rapidly decayed and further increases in CCh concentration up to 10 μM did not produce visible increases in fluorescence. However, at the peak concentration of 100 μM, there was fluorescence change similar to that of 0.3 μM (Fig. 5.6A-B). A repeated measures ANOVA with a Greenhouse-Geisser correction for departure from sphericity revealed that fluorescence did not differ significantly across the CCh concentrations ($F(1.269, 2.539) = 4.122$, $P < 0.155$) (Fig. 5.6C). This indicates that the biosensor is not very sensitive to varying concentrations of CCh but the large number of pixels with fluorescence intensities lower than baseline suggests that photobleaching has occurred. As an attempt to mitigate the impact of photobleaching, this analysis was repeated but on a smaller region of interest, within each slice, that displayed minimal photobleaching. Unlike analysis performed on the entire image, application of CCh caused an initial increase in fluorescence that was maintained up to a concentration of 10 μM. Subsequent bath application of 100 μM CCh induced a near two-fold increase in fluorescence from 10 μM CCh (Fig. 5.6D-E). However, a Friedman test determined that there was no significant main effect of CCh concentration on fluorescence ($\chi^2(6) = 8.286$, $p = 0.218$) (Fig. 5.6F). These results demonstrate that the ACh3.0 sensor is capable of detecting the presence of CCh, thus ACh, at concentrations as low as 0.3 µM but its sensitivity to concentrations between 0.3 µM and 10 µM is inseparable.

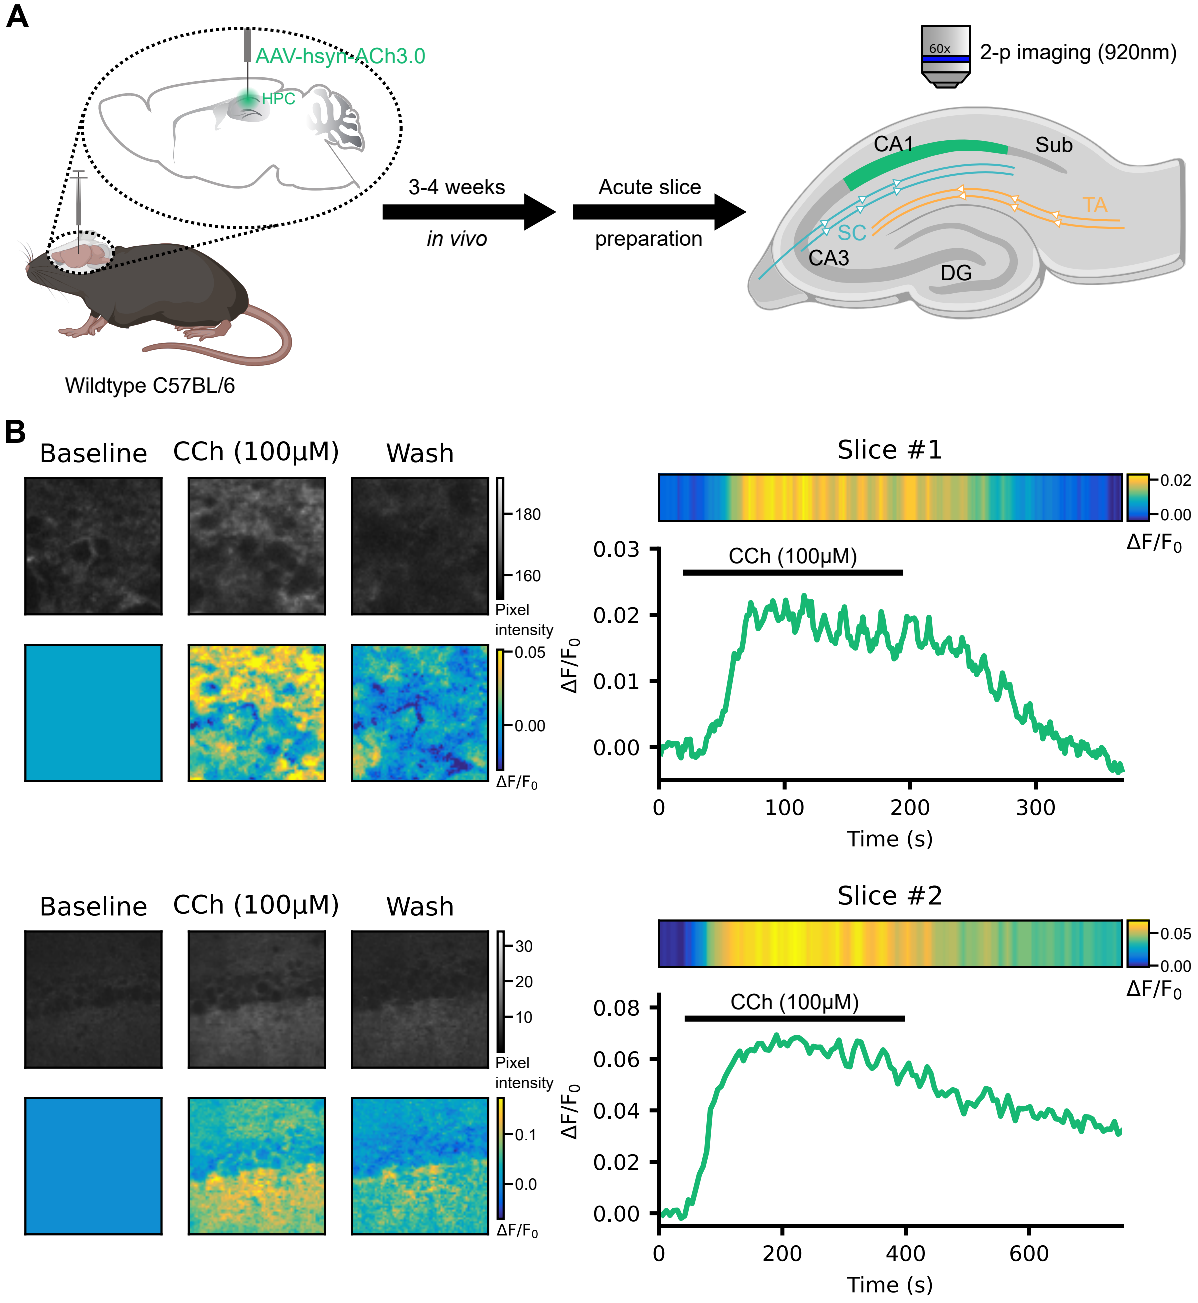

FIGURE 5.5: **Fluorescence response of ACh3.0-expressing cells in acute hippocampal brain slices to the bath application of 100 µM CCh.** (**A**) Schematic illustration of the experimental design used to measure the effects of bath applied CCh (100 µM) on fluorescence responses in ACh3.0-expressing CA1 hippocampal neurons of acute hippocampal brain slices (green). (**B**; left) Images display raw (top

left; grey) and $\Delta F/F_0$ (bottom left; pseudocoloured) ACh3.0 sensor fluorescence responses in CA1 to before (Baseline), during (CCh (100 μM)) and after (Wash) application of 100 μM CCh. (top right) Fluorescence responses over the course of the experiment. (bottom right) A timecourse of the $\Delta F/F_0$ showing the rise in fluorescence after CCh application. Top and bottom rows are regions of interest from different acute hippocampal slices.

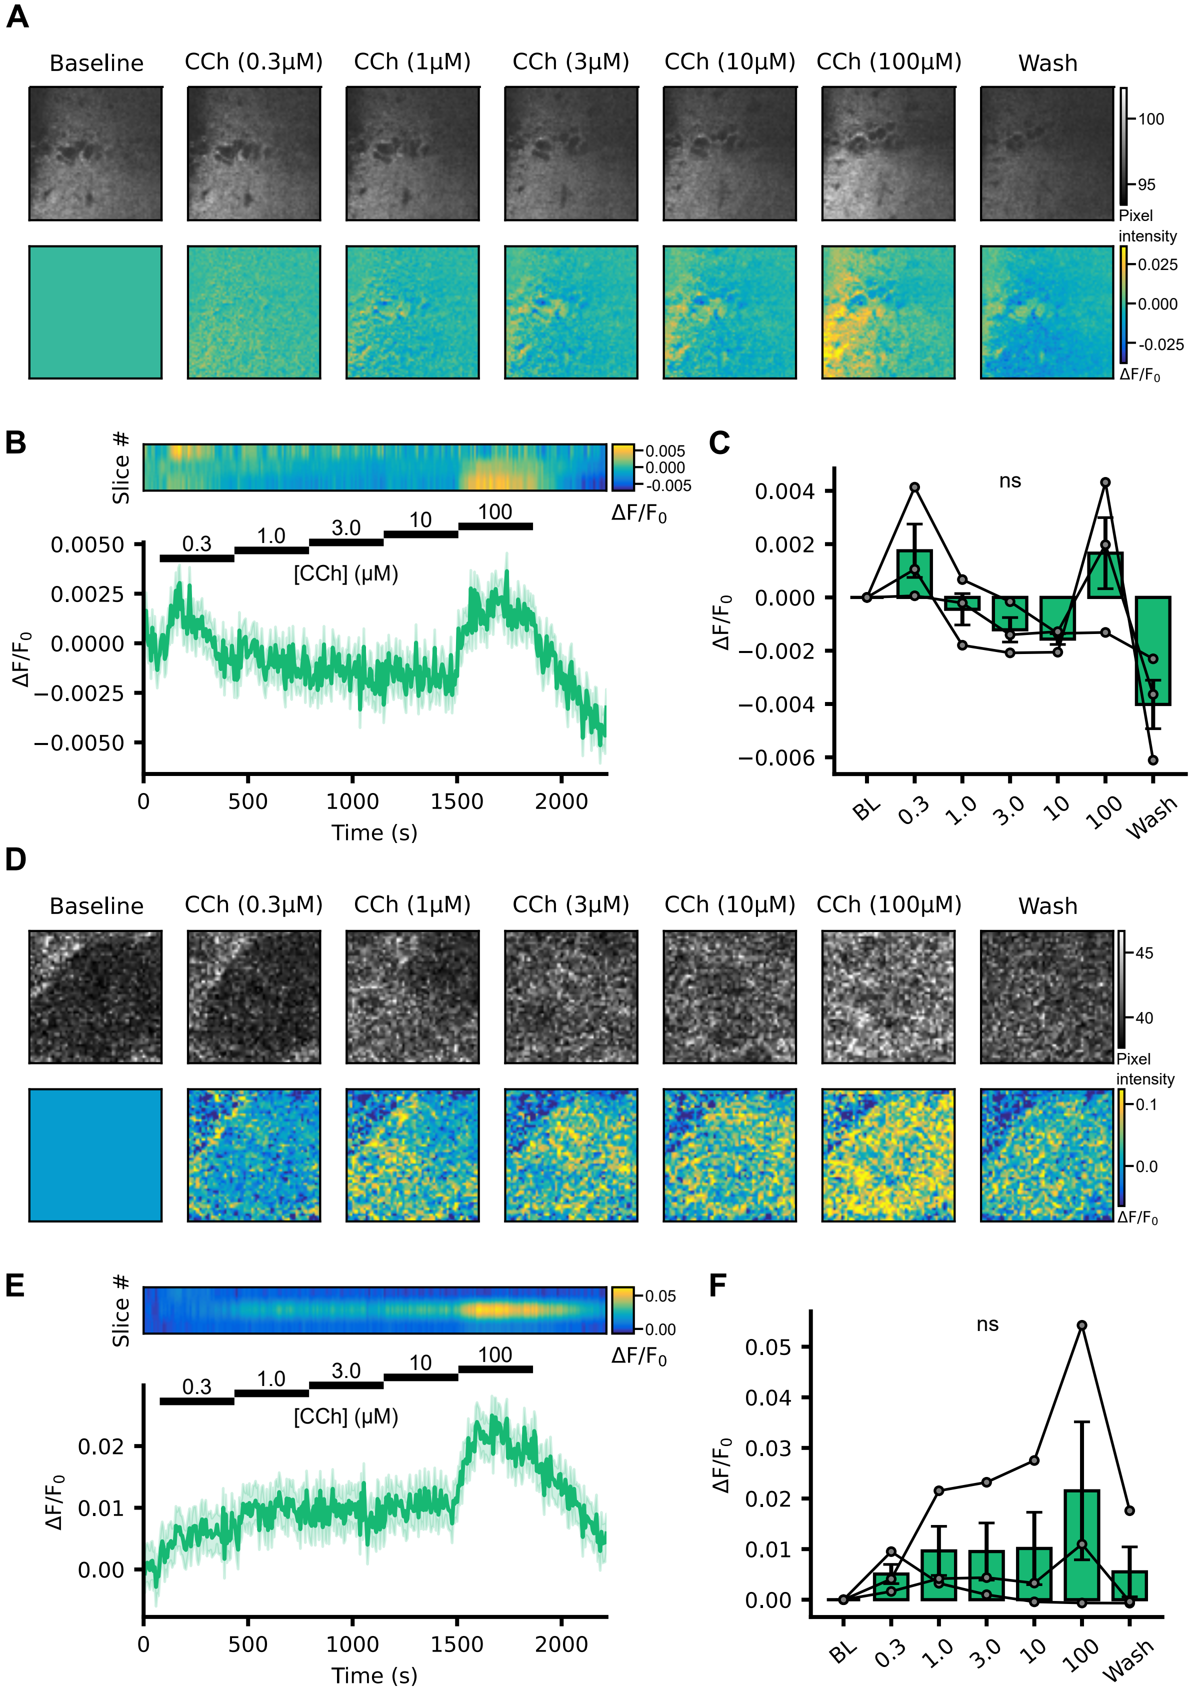
A
Baseline
CCh (0.3μM)
CCh (1μM)
CCh (3μM)
CCh (10μM)
CCh (100μM)
Wash
100
95
Pixel intensity
0.025
0.000
-0.025
ΔF/F₀
B
Slice #
0.005
0.000
-0.005
ΔF/F₀
0.3
1.0
3.0
10
100
[CCh] (μM)
0.0050
0.0025
0.0000
−0.0025
−0.0050
ΔF/F₀
0
500
1000
1500
2000
Time (s)
C
ns
0.004
0.002
0.000
−0.002
−0.004
−0.006
ΔF/F₀
BL
0.3
1.0
3.0
10
100
Wash
D
Baseline
CCh (0.3μM)
CCh (1μM)
CCh (3μM)
CCh (10μM)
CCh (100μM)
Wash
45
40
Pixel intensity
0.1
0.0
ΔF/F₀
E
Slice #
0.05
0.00
ΔF/F₀
0.3
1.0
3.0
10
100
[CCh] (μM)
0.02
0.01
0.00
ΔF/F₀
0
500
1000
1500
2000
Time (s)
F
ns
0.05
0.04
0.03
0.02
0.01
0.00
ΔF/F₀
BL
0.3
1.0
3.0
10
100
Wash

FIGURE 5.6: **Fluorescence response of ACh3.0-expressing cells in acute hippocampal brain slices to different concentrations of bath applied CCh.** (**A**) Images display raw (top; grey) and $\Delta F/F_0$ (bottom; pseudocoloured) ACh3.0 sensor fluorescence responses in CA1 across increasingly higher bath applied concentrations of CCh. (**B**; top) Fluorescence responses in different hippocampal slices over the course of the experiment. (bottom) The $\Delta F/F_0$ profile over time across increasing concentrations of bath applied CCh. (**C**) Summary of mean $\Delta F/F_0$ during each period of drug application (repeated measures ANOVA; N = 3). (**D-F**) Same as (**A-C**) but for a selected region within each hippocampal slice that expresses low levels of fluorophore photobleaching (Friedman test; N = 3). BL, baseline period prior to bath application of CCh; Wash, post-application of 100 μM CCh. Data represented as mean ± SEM.

After establishing the ability of the ACh3.0 sensor at detecting CCh, it was crucial to validate its behaviour to optogenetically evoked endogenous release of ACh. To evaluate the feasibility of combining optogenetic stimulation with optical imaging for simultaneous monitoring of cholinergic signals, CA1 neurons of ChAT-Cre x Ai32 mice were expressed with the ACh3.0 sensor. Cholinergic transmission was activated using single photon 5 ms LED pulses (470 nm) to optogenetically stimulate ChR2-expressing cholinergic fibres in acute hippocampal slices. Bath application of 100 μM CCh was used after to confirm ACh3.0 sensor expression in the slice. Light activation of cholinergic fibres produced an increase in fluorescence that was spatially different to the fluorescence change seen with 100 μM CCh (Fig. 5.7A). The fluorescence increase observed during the photostimulation period (0.00333 ± 0.00045) mirrored the duration of light being ON and was greater than the signal seen during CCh application (0.00230 ± 0.00110). A repeated measures ANOVA determined a main effect of experiment state on the fluorescence intensity captured ($F(3, 15) = 10.908$, $P < 0.001$). Post hoc pairwise comparisons using the Bonferroni correction revealed that during light stimulation there was a significant increase in fluorescence compared to BL and Wash (Light-BL: 0.003 (95%, 0.001 to 0.005) %, $p = 0.004$; Light-Wash: 0.004 (95%, 0.001 to 0.006) %, $p = 0.006$). There were no significant differences between the fluorescence signals captured during the optogenetic stimulation of ACh release (Light) and 100 μM CCh (0.001 (95%, -0.003 to 0.005) %, $p = 1.000$), baseline (BL) and post-application of 100 μM CCh (Wash) (0.000 (95%, -0.001 to 0.002) %, $p = 1.000$), as well as CCh and Wash (0.003 (95%, -0.001 to 0.007) %, $p = 0.215$) (Fig. 5.7B). However, a similar increase in fluorescence signal from baseline was observed in hippocampal slices from wildtype C57BL/6J mice without ACh3.0 expression during the same light stimulating protocol (Baseline-Light: -0.00704 ± 0.00049, $t(7) = -14.348$, $p < 0.001$) (Fig. 5.8A). There was no significant increase in fluorescence when light stimulation was triggered without a hippocampal slice in the recording chamber (Baseline-Light: -0.00092 ± 0.00052, $t(6) = -1.741$, $p = 0.132$) (Fig. 5.8B). These data suggests that the fluorescence signals detected during photostimulation could be an artefact from the biological tissue rather than ACh3.0 sensor fluorescence from binding endogenous ACh.

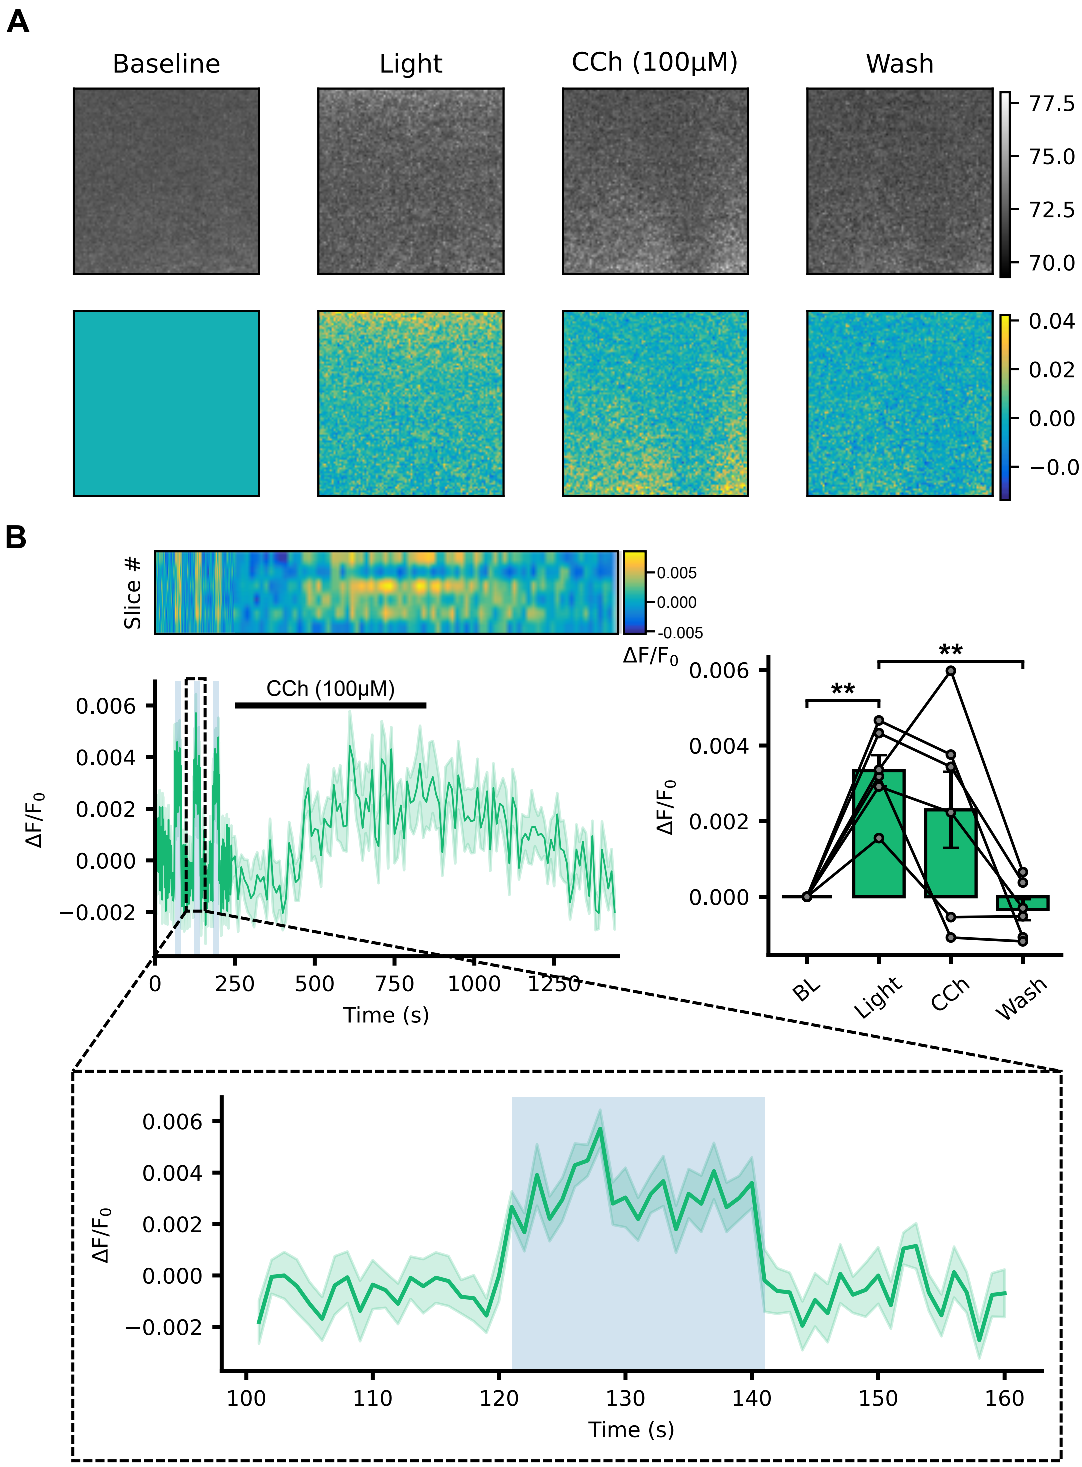

FIGURE 5.7: **Detecting optogenetically evoked endogenous ACh release with ACh3.0 sensor in ChAT-Cre x Ai32 acute hippocampal brain slices.** (**A**) Images display raw (top; grey) and $\Delta F/F_0$ (bottom; pseudocoloured) fluorescence detection in the absence (Baseline) and presence of (Light) 1 Hz blue light (470 nm) stimulation, and during (CCh (100 μM)) and after (Wash) bath application of 100 μM CCh in acute hippocampal brain slices containing ACh3.0-expressing neurons. (**B**; top left, upper) Fluorescence

responses in different hippocampal slices over the course of the experiment. (top left, lower) The $\Delta F/F_0$ profile over time in response to photostimulation of endogenous ACh release and bath application of 100 μM CCh. (top right) Summary of mean $\Delta F/F_0$ during each period the experiment (repeated measures ANOVA with Bonferroni-corrected simple effects analysis; N = 6). (bottom) Inset displaying a part of the experiment timecourse containing the fluorescence response to 1 Hz optogenetic stimulation. BL, baseline period prior to photostimulation; Wash, post-application of 100 μM CCh. Data represented as mean ± SEM; ; ** $p < 0.01$.

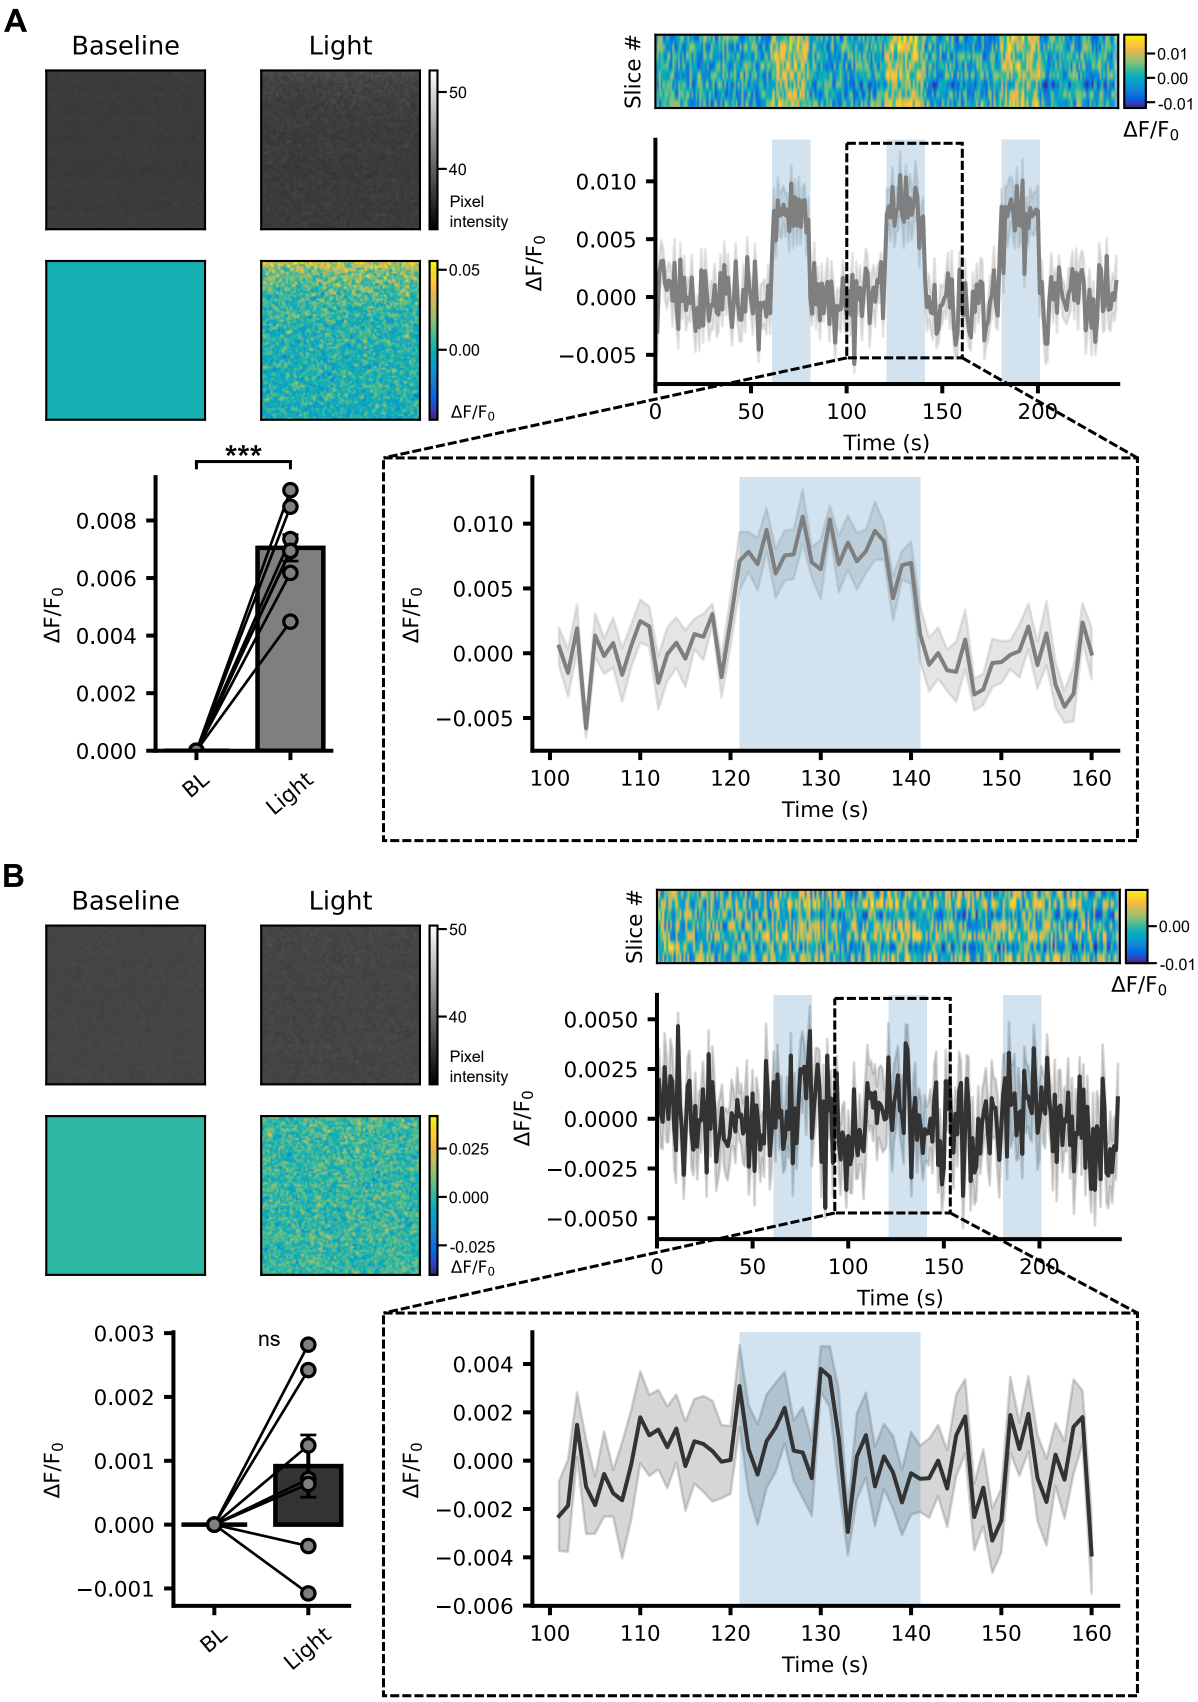
A
Baseline
Light
50
40
Pixel intensity
0.05
0.00
ΔF/F$_0$
Slice #
0.01
0.00
-0.01
ΔF/F$_0$
0.010
0.005
0.000
−0.005
ΔF/F$_0$
0
50
100
150
200
Time (s)
***
0.008
0.006
0.004
0.002
0.000
ΔF/F$_0$
BL
Light
0.010
0.005
0.000
−0.005
ΔF/F$_0$
100
110
120
130
140
150
160
Time (s)
B
Baseline
Light
50
40
Pixel intensity
0.025
0.000
-0.025
ΔF/F$_0$
Slice #
0.00
-0.01
ΔF/F$_0$
0.0050
0.0025
0.0000
−0.0025
−0.0050
ΔF/F$_0$
0
50
100
150
200
Time (s)
ns
0.003
0.002
0.001
0.000
−0.001
ΔF/F$_0$
BL
Light
0.004
0.002
0.000
−0.002
−0.004
−0.006
ΔF/F$_0$
100
110
120
130
140
150
160
Time (s)

FIGURE 5.8: **Fluorescence detection of 1 Hz blue light (470 nm) photostimulation in wildtype C57BL/6J acute hippocampal brain slices.** (**A**; top left) Images display raw (upper; grey) and $\Delta F/F_0$ (lower; pseudocoloured) fluorescence detection in the absence (Baseline) and presence of (Light) 1 Hz blue light (470 nm) stimulation in acute hippocampal brain slices without expression of ACh3.0 sensor. (top right, upper) Fluorescence responses in different hippocampal slices over the course of the experiment. (bottom right) Inset displaying a part of the experiment timecourse containing the fluorescence response to 1 Hz light stimulation. (bottom left) Summary of mean $\Delta F/F_0$ before (BL) and during (Light) light stimulation (Student's paired t-test, two-tailed; N = 8). (**B**) Same as ((**A**) but without a hippocampal slice (empty recording chamber) (Student's paired t-test, two-tailed; N = 7). Data represented as mean ± SEM; ; *** $p < 0.001$.

5.3.5 Characterisation of viral channelrhodopsin expression in Chrna2-Cre x FlpO mice

To explore the effect of cholinergic modulation on OLM interneuron inhibition of CA1 PCs, an intersectional approach employing the Cre and FlpO recombinase systems was initially examined. This involved breeding a mouse line in which Cre recombinase was targeted to the locus for the α subunit of nAChRs (Chrna2-Cre) with another mouse line in which FlpO recombinase was targeted to the locus for choline acetyltransferase (ChAT-FlpO). This generated a novel mouse line in which Cre and FlpO were selectively expressed in OLM interneurons and cholinergic neurons, respectively (Chrna2-Cre x ChAT-FlpO). To determine whether FlpO recombinase expression exhibited cholinergic neuron specificity in Chrna2-Cre x ChAT-FlpO mice, FlpO-dependent viruses for ChR2 and GFP were injected into the medial septum (MS) and dorsal striatum (DS) of these animals (Fig. 5.9A & Fig. 5.10A). Subsequent immunohistochemical staining with ChAT revealed that, in the MS, 1.75 ± 1.11% of ChAT-positive neurons were ChR2-expressing cells and 2.33 ± 1.57% of ChR2-expressing cells were ChAT-positive neurons (Fig. 5.9B-C). In the DS, 23.42 ± 5.14% of ChAT-positive neurons were ChR2-expressing cells and 5.27 ± 1.18% of ChR2-expressing cells were ChAT-positive neurons (Fig. 5.5B-C). Similar levels of specificity were observed with the FlpO-dependent GFP virus: 0.23 ± 0.21% of GFP-labelled cells were ChAT-positive neurons and 1.19 ± 1.09% of ChAT-positive neurons were GFP-labelled cells (Fig. 5.5D-E). These results suggest that there are large amounts of off-target expression with FlpO-dependent viruses in the novel Chrna2-Cre x ChAT-FlpO mouse line.

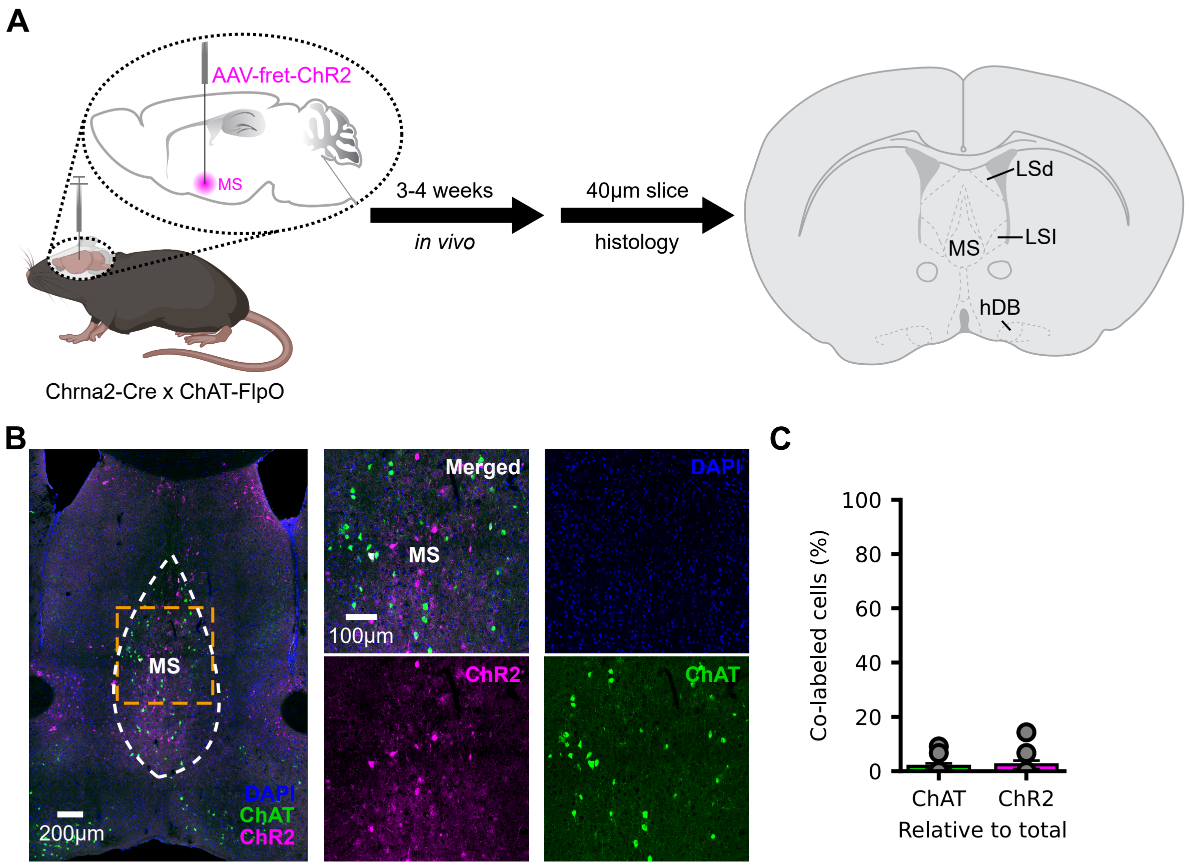
A
AAV-fret-ChR2
MS
Chrna2-Cre x ChAT-FlpO
3-4 weeks
in vivo
40µm slice
histology
LSd
LSI
MS
hDB
B
MS
200µm
DAPI
ChAT
ChR2
Merged
MS
100µm
DAPI
ChR2
ChAT
C
Co-labeled cells (%)
100
80
60
40
20
0
ChAT
ChR2
Relative to total

FIGURE 5.9: **AAV1-CAG-FLEXFRT-ChR2(H134R)-mCherry transfection of ChAT-expressing cells in the MS of ChAT-flp x Chrna2-cre mice.** (**A**) Schematic of AAV injection in MS to visualise expression of ChR2 in ChAT-expressing cells. (**B**; left) Representative image of ChR2 expression in MS. (right) Magnified image of target region (orange, dashed box) from representative image and associated labelling separated. (**C**) Percentage of cells coexpressing ChAT and ChR2 relative to total ChAT-labelled (green bar) and ChR2-labelled cells (magenta bar). Data represented as mean ± SEM.

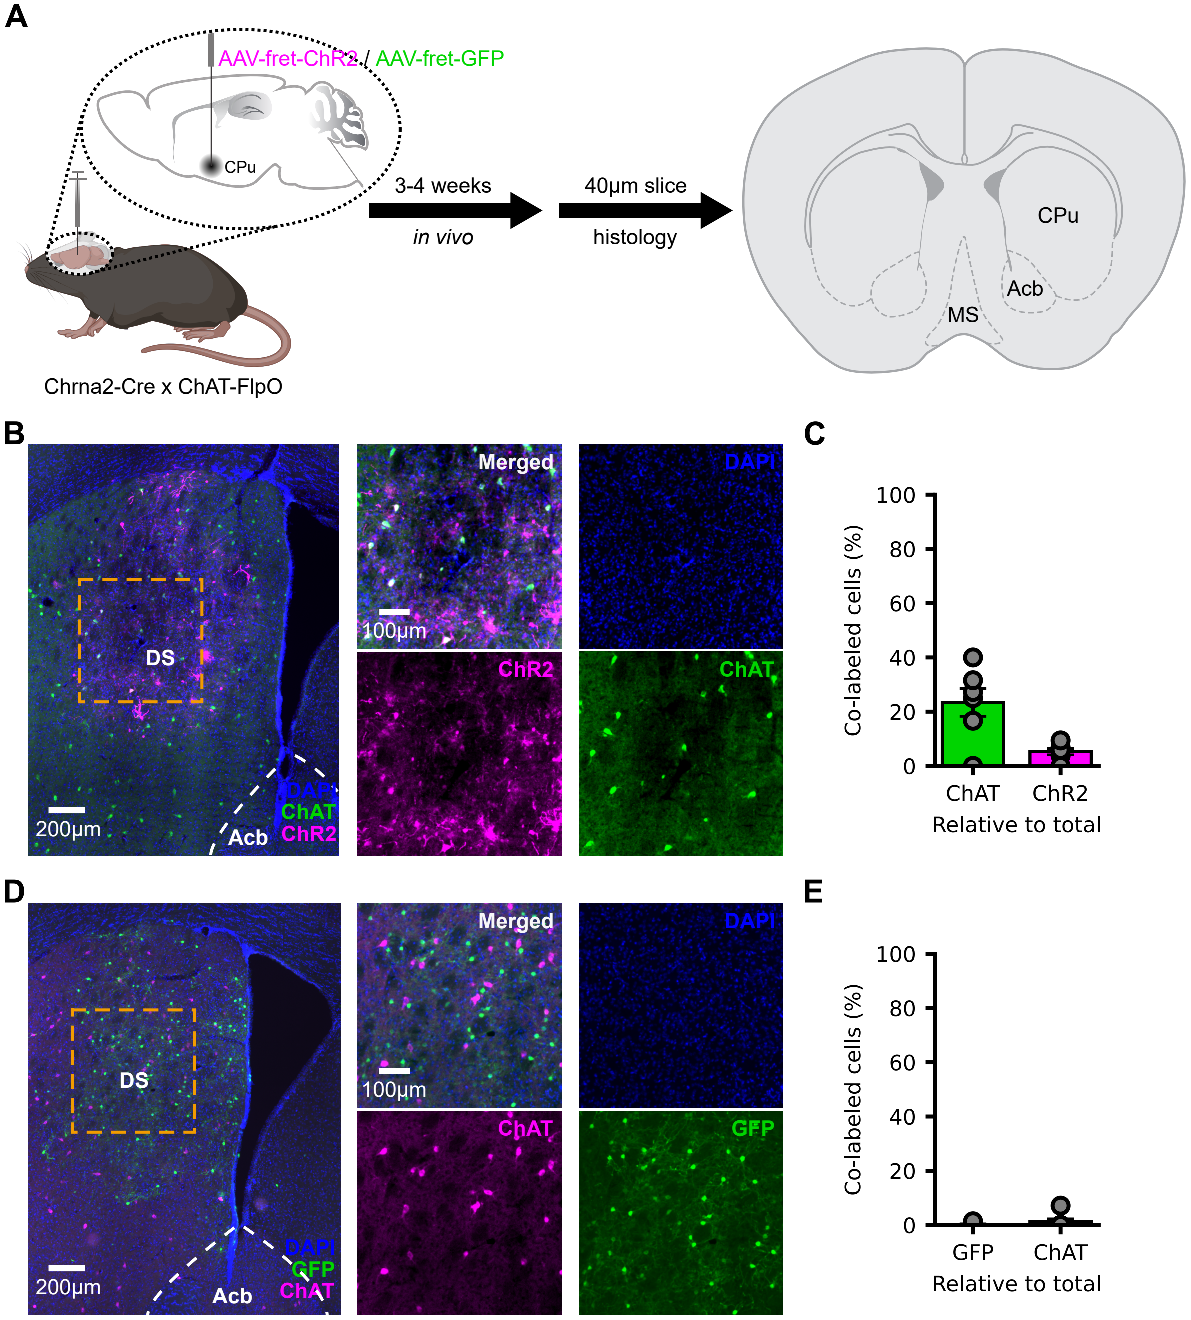
A
AAV-fret-ChR2 / AAV-fret-GFP
CPu
3-4 weeks
in vivo
40µm slice
histology
CPu
Acb
MS
Chrna2-Cre x ChAT-FlpO
B
DS
200µm
Acb
DAPI
ChAT
ChR2
Merged
100µm
DAPI
ChR2
ChAT
C
Co-labeled cells (%)
100
80
60
40
20
0
ChAT
ChR2
Relative to total
D
DS
200µm
Acb
DAPI
GFP
ChAT
Merged
100µm
DAPI
ChAT
GFP
E
Co-labeled cells (%)
100
80
60
40
20
0
GFP
ChAT
Relative to total

FIGURE 5.10: **FlpO-dependent ChR2 and GFP of ChAT-expressing cells in the dorsal striatum of ChAT-flp x Chrna2-cre mice.** (**A**) Schematic of AAV injection in the dorsal striatum (CPu/DS) to visualise expression of ChR2 or GFP in ChAT-expressing cells. (**B**; left) Representative image of ChR2 expression in DS. (right) Magnified image of target region (orange, dashed box) from representative image and associated labelling separated. (**C**) Percentage of cells coexpressing ChAT and ChR2 relative to total ChAT-labelled (green bar) and ChR2-labelled cells (magenta bar). (**D-E**) Same as (**A-B**). Data represented as mean ± SEM.

5.3.6 Characterisation of ChAT-cre virus

As the Cre/FlpO dual recombinase intersectional approach demonstrated poor specificity for cholinergic neurons, a double Cre recombinase approach using the AAV2.9-ChAT-Cre-P2A-EGFP vector in the MS of Chrna2-Cre x Ai32 mice was evaluated. First, the selectivity of this virus was examined by stereotaxically injecting it into the MS of wildtype C57BL/6J mice to express Cre-recombinase under the control of the promoter for ChAT (Fig. 5.11A). Immunohistochemistry showed that, in the MS, 8.04 ± 1.73% of EGFP-labelled cells were ChAT-positive neurons and 41.11 ± 5.99% of ChAT-positive neurons were EGFP-labelled cells (Fig. 5.11B-C). Additionally, dense networks of septo-hippocampal projections were observed in the hippocampus (Fig. 5.11D). These results suggest that the virus has low specificity in the MS but has robust expression in long range fibres.

5.3.7 Examining septo-hippocampal projections of Chrna2-expressing cells

It was necessary to further validate the double Cre recombinase approach by examining whether Chrna2-expressing cells in the MS also projected to the hippocampus in order to avoid non-specific opsin expression. To investigate this, the CTB retrograde tracer was stereotaxically injected into the hippocampi of Chrna2-Cre x Ai32 mice to track septo-hippocampal projecting neurons in the MS (Fig. 5.12A). The ChR2 of Chrna2-Cre x Ai32 mice was tagged with EYFP which enabled the visualisation of Chrna2-expressing cells. Histological analysis of the MS revealed that 0.01 ± 0.01% of ChAT-negative, EYFP-labelled cells were ChAT-negative, CTB-positive neurons and 0.04 ± 0.04% of ChAT-negative, CTB-positive neurons were ChAT-negative, EYFP-labelled cells (Fig. 5.12B-C). This indicates that a negligible proportion of septo-hippocampal projecting neurons express the Chrna2 gene.

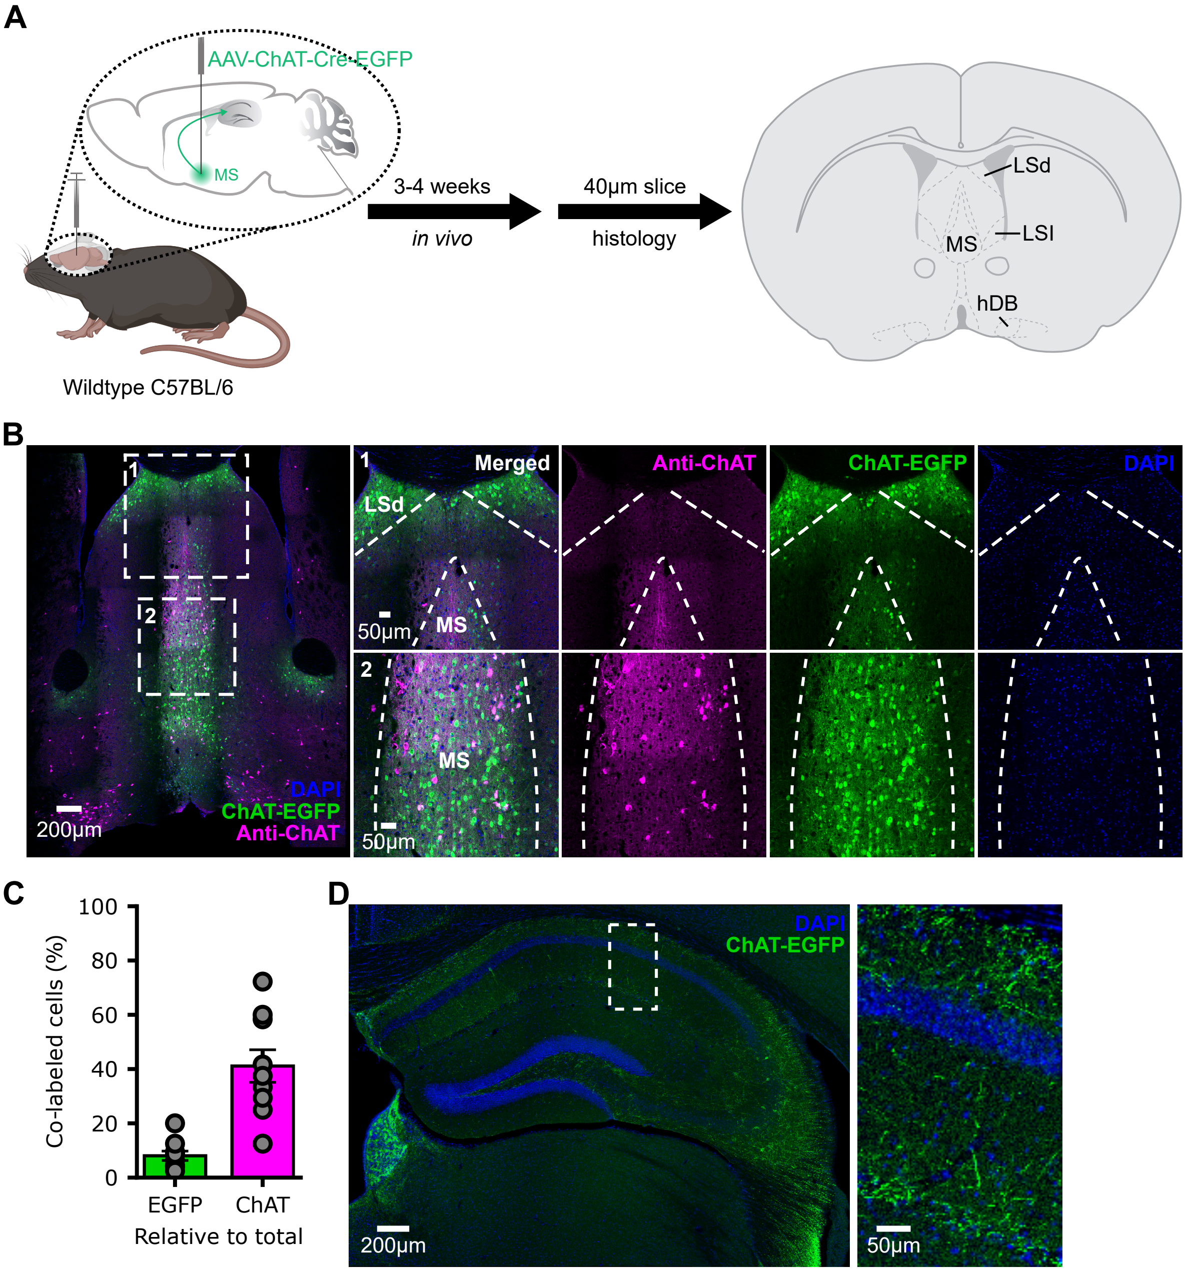
A
AAV-ChAT-Cre-EGFP
MS
3-4 weeks
in vivo
40μm slice
histology
LSd
LSI
MS
hDB
Wildtype C57BL/6
B
1
2
DAPI
ChAT-EGFP
Anti-ChAT
200μm
1
Merged
LSd
50μm
MS
Anti-ChAT
ChAT-EGFP
DAPI
2
MS
50μm
C
Co-labeled cells (%)
100
80
60
40
20
0
EGFP
ChAT
Relative to total
D
DAPI
ChAT-EGFP
200μm
50μm

FIGURE 5.11: **AAV2.9-ChAT-Cre-P2A-EGFP transfection of ChAT-expressing cells in the MS of C57BL/6J mice.** (**A**) Schematic of AAV injection in MS to characterise transfected cells and their projections to the hippocampus. (**B**; left) Representative image of ChAT-Cre viral expression in MS. (right) Magnified image of target regions (white, dashed box) from representative image and associated labelling separated. Percentage of cells coexpressing EGFP and ChAT relative to total EGFP-labelled (green bar) and ChAT-labelled cells (magenta bar). (**D**; left) Representative image of ChAT-Cre viral expression in septal projections to the hippocampus. (right) Inset, magnification of EGFP-labelled septo-hippocampal projections. Data represented as mean ± SEM.

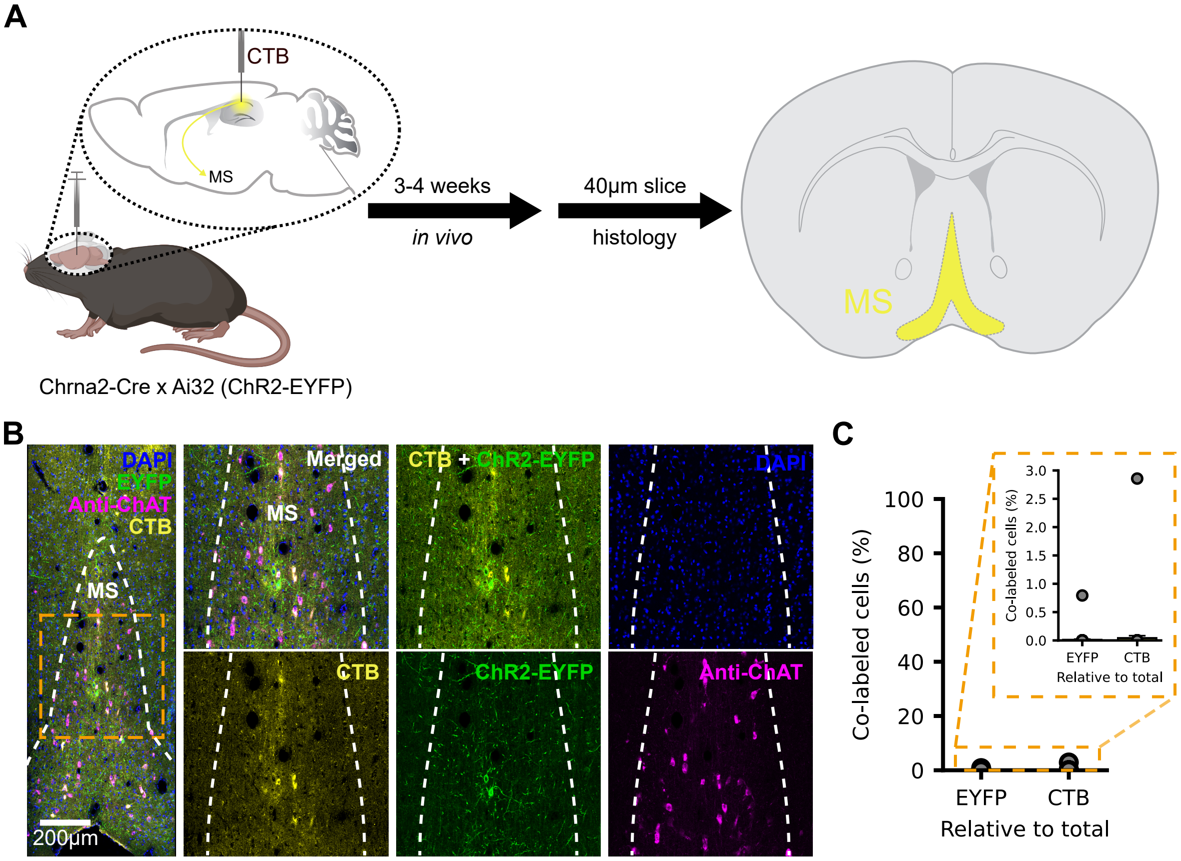
A
CTB
MS
3-4 weeks
in vivo
40µm slice
histology
MS
Chrna2-Cre x Ai32 (ChR2-EYFP)
B
DAPI
EYFP
Anti-ChAT
CTB
MS
200µm
Merged
MS
CTB + ChR2-EYFP
DAPI
CTB
ChR2-EYFP
Anti-ChAT
C
Co-labeled cells (%)
100
80
60
40
20
0
EYFP
CTB
Relative to total
3.0
2.5
2.0
1.5
1.0
0.5
0.0
Co-labeled cells (%)
EYFP
CTB
Relative to total

FIGURE 5.12: **Chrna2-expressing neurons in MS do not project to dorsal hippocampus.** (**A**) Schematic of CTB retrograde tracer injection into dorsal hippocampus of Chrna2-Cre x Ai32 mice to visualise septo-hippocampal projecting neurons in MS. (**B**; left) Representative image of CTB expression in MS. (right) Magnified image of target region (orange, dashed box) from representative image and associated labelling separated. (**C**) Percentage of cells coexpressing EYFP and CTB relative to total EYFP-labelled (Chrna2-Cre expressing; green bar) and CTB-labelled cells (yellow bar). Data represented as mean ± SEM.

5.3.8 The effect of endogenous acetylcholine on OLM inhibitory transmission

It has been shown that bath-applied CCh attenuates the inhibitory transmission of OLM interneurons onto CA1 PCs but whether endogenous ACh release would amount to the same effect remains unclear (Sec. 4.3.2). To investigate the effects of endogenous ACh on the OLM-to-PC synapse, a dual optogenetics strategy was undertaken to enable independent optical excitation of the cholinergic fibres and OLM interneurons. Despite the low selectivity of the AAV2.9-ChAT-Cre-P2A-EGFP vector, the double Cre recombinase approach remained the better option due to its higher transfection rate of ChAT-positive neurons in the MS compared to the FlpO-dependent ChR2 virus in Chrna2-Cre x ChAT-FlpO mice. Therefore, the AAV5-Syn-FLEX-rc[ChrimsonR-tdTomato] vector was injected into the hippocampi of Chrna2-Cre mice to selectively express Chrimson in OLM interneurons. Additionally, the AAV2.9-CAG-DIO-ChRger2-

TS-YFP and AAV2.9-ChAT-Cre-P2A-EGFP vectors were concurrently injected into the MS of these mice to express ChRger in cholinergic neurons (Fig. 5.13A). Voltage-clamp recordings at 0 mV were made from CA1 PCs during photostimulation of OLM interneurons with two 5 ms pulses of amber (590 nm) light. 2 Hz blue (470 nm) light stimulation at an irradiance of 0.64 mW/mm^2 (to prevent opsin cross-talk) was applied for 5 min to evoke physiologically maximal acetylcholine release (Palacios-Filardo et al., 2021; Jing et al., 2018) (Fig. 5.13B). Although optogenetic stimulation of cholinergic fibres caused a general reduction in OLM inhibition from baseline (baseline-opto: 0.102 ± 0.075 normalised IPSC amplitude), a Student's paired t-test determined that this was not a significant difference ($t(12) = 1.367$, $p = 0.197$). Immunohistochemistry confirmed that there was some coexpression of green fluorescence, the marker for both AAV2.9-CAG-DIO-ChRger2-TS-YFP and AAV2.9-ChAT-Cre-P2A-EGFP vectors, with ChAT staining in these mice (Fig. 5.13C). nAChR-mediated currents recorded from SO interneurons further supported the immunohistochemical characterisation of ChRger expression in cholinergic neurons (Fig. 5.13D). Photostimulation of the cholinergic fibres with blue light produced a fast current response in the OLM interneurons that was not seen with amber light (No light: 19.89 ± 1.35 pA; 470 nm light: 244.54 ± 156.94 pA; 590 nm light: 15.25 ± 8.20 pA). A Friedman test determined that light stimulation had no significant effect on the current response ($\chi^2(2) = 3.000$, $p = 0.223$). However, a statistical test performed with such low sample size numbers (N = 2) is unreliable. These results suggest that endogenous acetylcholine does not affect OLM interneuron inhibition in the same manner as bath application of CCh.

After taking a closer examination of the traces during the initial onset of blue light stimulation, large blue light-evoked inhibitory currents were observed in an subset of the experiments. This subset of experiments displayed a greater general suppression of OLM inhibition from baseline (baseline-opto: 0.220 ± 0.118 normalised IPSC amplitude) but a Student's paired t-test also deemed this a non-significant difference ($t(5) = 1.854$, $p = 0.123$) (Fig. 5.14A). Furthermore, in an experiment that showed a reduction in OLM inhibition during blue light optogenetic stimulation, atropine (25 μM) and mecamylamine (50 μM) was bath-applied afterwards to assess whether it was an ACh-dependent effect. Light stimulation in the presence of these cholinergic receptor antagonists did not prevent the attenuation of OLM inhibition (Fig. 5.14B). To determine whether there was off-target expression of ChRger and Chrimson on excitatory inputs of CA1 PCs, -70 mV recordings of voltage-clamped CA1 PCs, that demonstrated photostimulated depression in OLM-to-PC activity, were made in response to blue and amber light stimulation. Both blue and amber light evoked EPSCs that were prevented with NBQX (20 μM) and DAPV (25 μM) but not with atropine and mecamylamine (Fig. 5.14C). A repeated measures ANOVA determined that there was a significant main effect of drug on the normalised amplitude of blue light-evoked EPSCs ($F(2, 4) = 53.179$, $P < 0.001$). Simple effects analysis, using the Bonferroni correction, revealed that the main effect was due to a significant difference in amplitude between bath-applied atropine (Atro) with mecamylamine (Meca) and NBQX with DAPV (Atro & Meca - NBQX &

DAPV: 0.911 (95%, 0.140 to 1.682), $p = 0.036$). Despite a large reduction in the amplitude of amber light-evoked EPSCs, a Friedman test determined there was no significant main effect of drug on the excitatory currents ($\chi^2(2) = 5.600$, $p = 0.061$). Together, these results indicate there was off-target expression of ChRger in long-range excitatory and inhibitory connections from the MS, and potential off-target expression of Chrimson in excitatory inputs onto CA1 PCs. Additionally, these data suggests that septo-hippocampal innervations are capable of regulating OLM inhibiton on CA1 PCs via a cholinergic-independent mechanism but it is inconclusive whether endogenous ACh also has an effect.

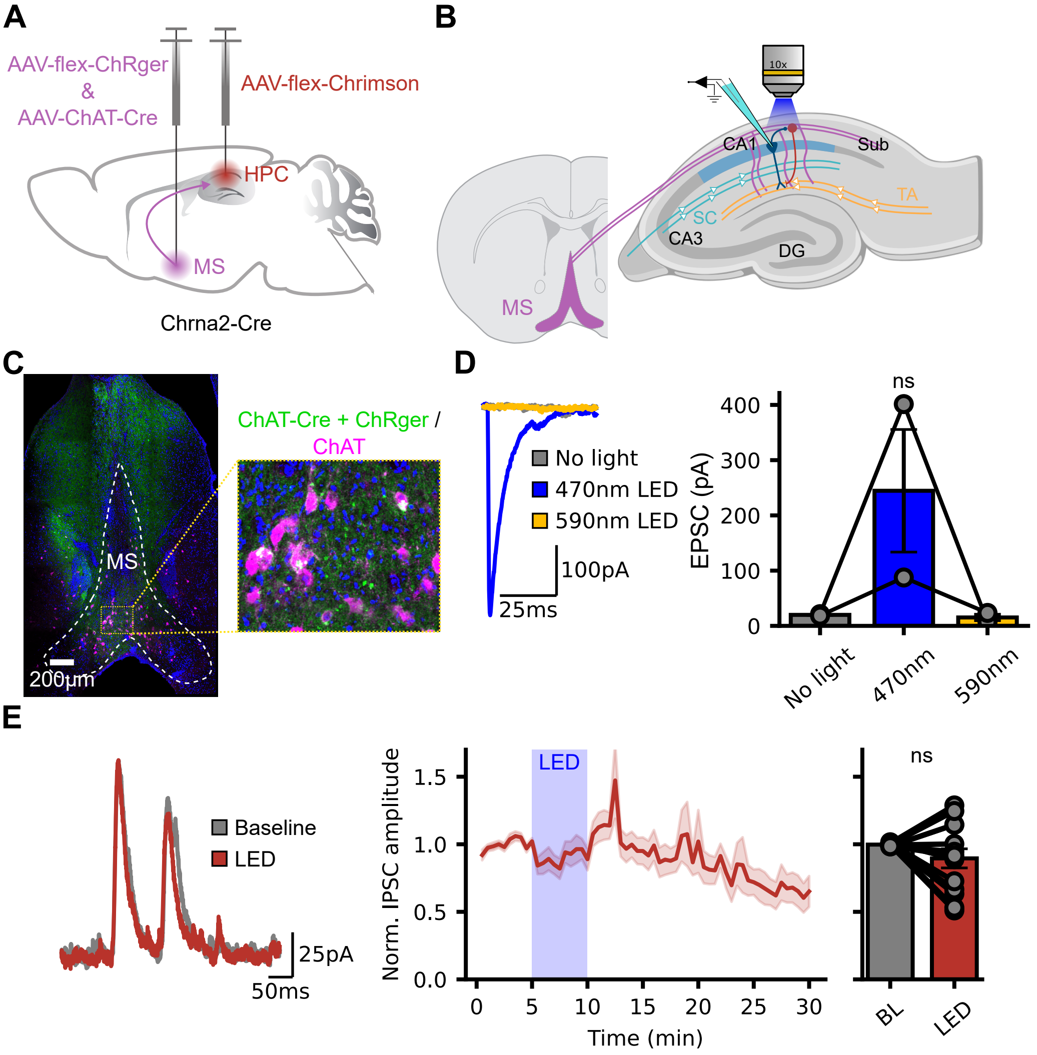

FIGURE 5.13: **The effect of optogenetically evoked ACh release on OLM interneuron inhibition.** (**A**) Schematic of AAV injection into MS and hippocampus (HPC) to enable independent optical activation of cholinergic fibres and OLM interneurons. (**B**) Schematic of the hippocampal slice recording setup: whole-cell patch-clamp recordings of IPSCs in CA1 pyramidal cells (dark blue) from photostimulated dendrite-targeting OLM interneurons (red) and blue light activated cholinergic fibres (purple). (**C**) Representative image of ChAT-Cre + ChRger coexpression (green) with ChAT (purple) in the MS. Inset, magnification

of target region (orange, dashed box). (**D**; left) Example traces of fast synaptic photocurrents recorded in OLM interneurons at -70 mV after blue and amber light stimulation of cholinergic fibres. (right) Summary of mean EPSC response under different light conditions (Friedman test; N = 2). (**E**; left) Example traces before (grey) and during (red) 2 Hz blue light photostimulation of cholinergic fibres. (middle) A timecourse of the OLM interneuron output with and without optogenetic activation of cholinergic fibres. (right) Summary of mean OLM interneuron output during baseline (BL) and photostimulation (LED) periods (Student's paired t-test, two-tailed; N = 13). Data represented as mean ± SEM.

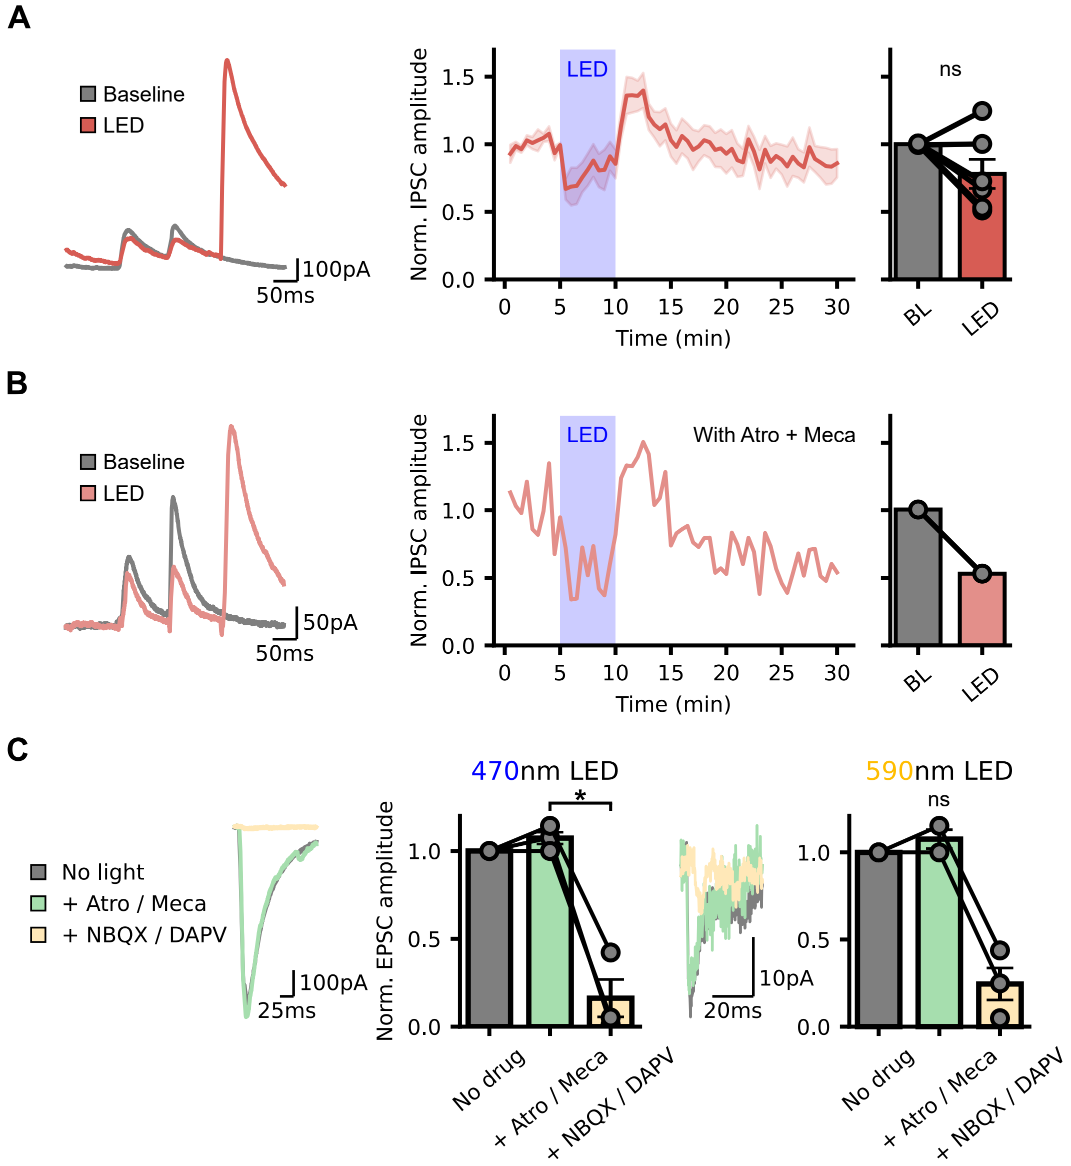

FIGURE 5.14: **Long-range inhibitory septo-hippocampal projections could potentially inhibit OLM interneurons.** (**A**; left) Example traces, recorded at 0 mV, before (grey) and during the onset of (peach) 2 Hz blue light photostimulation of septo-hippocampal projections. (middle) A timecourse from experiments expressing large inhibitory currents at the onset of light stimulation shows OLM interneuron output is inhibited by optogenetic activation of septo-hippocampal fibres. (right) Summary of mean OLM

interneuron output during baseline (BL) and photostimulation (LED) periods (Student's paired t-test, two-tailed; N = 6). (**B**) A repetition of an experiment in (**A**) that expressed inhibition from light stimulation but in the presence of atropine (25 μM) and mecamylamine (50 μM). (**C**; left) Representative traces of EPSCs recorded at 0 mV from experiments that expressed large inhibitory photocurrents during the onset of blue light stimulation. (middle) Summary of mean EPSC responses evoked by blue light at baseline (No drug), in the presence of atropine (25 μM) and mecamylamine (50 μM) (+ Atro / Meca), and also NBQX (20 μM) and DAPV (25 μM) (+ NBQX / DAPV) (repeated measures ANOVA with Bonferroni-corrected simple effects analysis; N = 3). (right) Same as (middle) but responses were evoked by amber light (Friedman test with posthoc Bonferroni-corrected Wilcoxon signed ranks tests; N = 3). Data represented as mean ± SEM.

5.4 Discussion

The experiments performed here have demonstrated that by using ChRger and Chrimson, there is a blue light irradiance range that can reliably evoke large ChRger-induced photocurrents without any Chrimson-induced photocurrents. At the blue light irradiance immediately before the onset of spectral cross-talk, ChRger-induced photocurrents were over three-fold greater than that of ChR2-induced photocurrents, highlighting the sensitivity of ChRger. The validation of these two opsins has presented a potential approach at optogenetically controlling distinct neural populations which has been an ongoing endeavour in neuroscience.

In ChAT-Cre x Ai32 mice, the blue light irradiance range used to avoid spectral cross-talk, obtained in the opsin validation experiments, was capable of stimulating endogenous ACh release. The ability to evoke cholinergic transmission with a much less sensitive opsin (ChR2) provided confidence that these light intensities would be sufficient at exciting cholinergic fibres expressing ChRger.

The ACh3.0 sensor was used to attempt to measure the physiological concentration of the optogenetically evoked release of ACh. These experiments showed that the biosensor was able to reliably detect fluorescence changes induced by the application of 100 μM CCh but it was less sensitive at distinguishing differences in fluorescence between lower concentrations ranging from 0.3-10 μM. Although non-significant, the ACh3.0 sensor seemed to be able to detect fluorescence changes at the initial 0.3 μM concentration which then either plateaued or decreased below baseline fluorescence until the application of 100 μM CCh. This inability to produce a dose-response curve is in contrast to the original study performed by (Jing et al., 2020). However, the experimental conditions were different as they measured the sensitivity of the biosensor by locally puffing ACh onto ACh3.0-expressing cultured HEK293T cells. This method of localised application enabled cells to rapidly experience the intended ACh concentrations which resulted in a much shorter experiment. In the case of bath application, the experiments here took much longer as time was required to ensure the recording chamber equilibrated to the new CCh concentration. This extra exposure to CCh could potentially result in desensitisation and internalisation of the M3 mAChRs that constitute the ACh3.0. In addition, the increased laser scanning could potentially cause photobleaching. Together, this could be a possible explanation for the lack of significant differences in fluorescence between the different CCh concentrations. When the biosensor was expressed in hippocampal CA1 cells of ChAT-Cre x Ai32 mice, photostimulation of cholinergic fibres caused a change in fluorescence similar in amplitude to the application of 100 μM CCh. However, a similar spatial and amplitude fluorescence change was found when light stimulation was applied to acute hippocampal slices without ACh3.0 expression and this fluorescence was not present when the hippocampal slice was removed. Therefore, it remains inconclusive whether the change in fluorescence is solely due to an artefact of the interaction between the light and the biological tissue or whether there was also a contribution from endogenous ACh-induced fluorescence of the ACh3.0 sensor.

It was found that the use of the AAV2.9-ChAT-Cre-P2A-EGF vector in Chrna2-Cre mice was a better system than the novel Chrna2-Cre x ChAT-FlpO mouse line to selectively express ChRger and Chrimson in cholinergic neurons and OLM interneurons, respectively. Therefore, this approach was undertaken to investigate the effects of endogenous ACh release on the OLM-to-PC synapse in CA1. Photostimulation of cholinergic fibres did not cause a significant reduction in OLM inhibitory transmission as previously seen with bath application of CCh in Chapter 4. However, there was a strong trend to suggest that activation of septo-hippocampal projections inhibits OLM inhibitory transmission. This was not a cholinergic neuromodulatory effect as the application of both nAChR and mAChR blockers did not attenuate this optogenetically evoked inhibitory effect. The presence of large inhibitory currents at the onset of blue light stimulation and the light-evoked excitatory currents indicated that the optogenetic inhibition of the OLM-to-PC synapse could potentially be via direct long-range inhibitory projections, long-range excitation of feedforward inhibition or both. This implies that there is a non-cholinergic but septo-hippocampal dependent mechanism to mediating the activity of CA1 OLM interneurons which could provide a different or complementary pathway for potentially gating hippocampal synaptic plasticity. This is supported by a study from Pang and Nocera (1999) which showed that selective cholinergic lesions of the medial septum with scopolamine did not impair spatial memory encoding as much as complete medial septal lesions with the addition of muscimol, a $GABA_A$ receptor agonist.

5.4.1 Limitations

In relation to the ChAT-Cre x Ai32 mice and the dual optogenetics experiments measuring the light-evoked EPSCs, it would be beneficial to increase their sample sizes for a more confident statistical conclusion. This would help determine whether the blue light irradiances used are truly capable of stimulating endogenous ACh release.

In the dose-response experiments for the characterisation of the ACh3.0 sensor, the inability to detect changes in fluorescence for the increasing concentrations of CCh, as mentioned previously, could be due to photobleaching or desensitisation of the biosensor. The effect of photobleaching on the data can be mitigated by detrending the data with the rate of photobleaching measured under the current experimental conditions. This could be done by scanning at the current sample rate once a supramaximal concentration of CCh has been bath applied and equilibrated to activate all ACh3.0 receptors. Additionally, to reduce the risks of desensitisation and photobleaching, localised application of CCh with a Picospritzer, similar to Jing et al. (2020), can be used instead of bath application.

Despite the use of the AAV2.9-ChAT-Cre-P2A-EGF vector in Chrna2-Cre mice being the best approach for the selective expression of the two opsins in independent neural populations, AAV2.9-ChAT-Cre-P2A-EGF characterisation displayed significant off-target expression in the MS and less than 50% expression in cholinergic neurons. Therefore, in the dual optogenetics experiments, the photostimulation in hippocampal CA1 could be of ChRger-expressing fibres from non-cholinergic neurons or there is insufficient cholinergic fibre activation. This could explain why there is no overall attenuation of OLM inhibitory transmission and when there is optogenetically evoked inhibition of the OLM-to-PC synapse, cholinergic blockers have no effect. Simultaneous two-photon imaging of ACh3.0 sensor fluorescence in these experiments during blue light stimulation would provide the best insight into whether there is ACh release. Also, post hoc immunohistochemistry of the used hippocampal slices could reveal if there was any ChRger expression in the cholinergic fibres. Whole-cell patch clamp experiments of light-induced nAChR-mediated EPSCs in mice stereotaxically injected with only the ChRger virus into the MS would further validate the ability to stimulate ACh release. To improve the selectivity of the AAV2.9-ChAT-Cre-P2A-EGF vector, different viral titres can be tested. Another approach would be to breed Chrna2-Cre mice with ChAT-Cre mice to produce a novel mouse line expressing Cre recombinase in both the cholinergic neurons and the OLM interneurons. The retrograde tracer experiments of the septo-hippocampal projections have shown that there is a negligible proportion of septo-hippocampal projecting neurons that express the Chrna2 gene. In addition, the ChAT-Cre mouse line has been heavily characterised to specifically target the ChAT-positive neurons, thus removing the issue of off-target expression in the MS. The issue that exists with this approach is that there are some ChAT-positive cells in the hippocampus that could potentially cause off-target expression of Chrimson (Granger et al., 2020). To test the suitability of this dual Cre recombinase approach, control experiments would need to be performed to examine the effect of these hippocampal ChAT-positive cells on OLM inhibitory transmission.

As cholinergic neurons in the MS are believed to fire at low frequencies, in the dual optogenetics experiments, cholinergic fibres were optogenetically stimulated at a frequency of 2 Hz to represent a typical high firing rate during wakefulness seen in rats (Simon et al., 2006; Zhang et al., 2010). This frequency has also been previously used in *ex vivo* whole-cell patch clamping experiments, to investigate cholinergic neuromodulation of information flow in CA1, to mimic a maximal physiological concentration of ACh in the hippocampus (Palacios-Filardo et al., 2021; Jing et al., 2018). However, a recent study by Ma et al. (2020) has shown that these cholinergic neurons exhibit higher firing rates during active exploration, ranging from 7.5-15 Hz, indicating that 2 Hz may not accurately reflect *in vivo* firing patterns corresponding with high ACh release in the hippocampus. This suggests that the stimulation frequency used in the dual optogenetics experiments may not have been high enough to release an amount of ACh representative of active animal behaviours which provides another explanation for the lack of cholinergic effect. Therefore, to help determine an appropriate stimulation frequency for a physiologically relevant concentration of ACh release, it would be beneficial to measure the amount of ACh released *in vivo* during active exploration using a cholinergic biosensor (Jing et al., 2020). This can then be compared with the concentration of ACh during *ex vivo* optogenetic stimulation of cholinergic fibres at varying frequencies.

5.4.2 Future experiments

The two-colour opsin validation experiments performed here have demonstrated a promising dual optogenetics approach. Since the current slice demonstration will not be identical to most experimental setups, it would be interesting to further validate this combination of opsins in other pathways and cell types. Whole-cell current clamp recordings of cells expressing these opsins during photostimulation with different wavelengths of light would determine whether there is unintended subthreshold depolarisation. In addition, short interval pairs of non-peak wavelength followed by peak wavelength light stimulation, e.g. blue light followed by amber light for Chrimson, would reveal whether the non-peak wavelength light stimulation alters synaptic transmission.

Despite slow off-kinetics from ChRger, it has been shown that it retains close to 100% spike fidelity in acute brain slices at stimulation frequencies up to 10 Hz which is within the range of physiological firing rates of cholinergic neurons in the MS (Bedbrook et al., 2019). However, it would be important to measure the spike fidelity of ChRger expressed specifically in cholinergic neurons as different experimental setups can yield different results.

5.4.3 Closing remarks

In conclusion, validation of two recently developed opsins, ChRger and Chrimson, has provided a promising dual optogenetics approach to controlling the activity of distinct neural populations. These opsins were expressed independently in the cholinergic neurons and OLM interneurons via the use of the AAV2.9-ChAT-Cre-P2A-EGF vector in Chrna2-Cre mice. Overall, photostimulation of the cholinergic fibres did not affect OLM inhibitory transmission onto CA1 PCs but the off-target expression of ChRger in the MS means that we cannot guarantee interference from the activity of other cell types.

CHAPTER 6

GENERAL DISCUSSION

How the brain accomplishes the process of credit assignment during learning is a longstanding debate in the field of neuroscience. In machine learning, this process is most commonly performed by the backprop algorithm and it produces results comparable to learning in humans. However, there are several biological inconsistencies in the direct translation of this algorithm into the brain. Although some of these concerns have been addressed in current models, there are still considerable biological implausibilities, and these models also ignore important factors that influence plasticity such as neuromodulation. Chapter 2 explored how cholinergic neuromodulation could be facilitating the process of learning with an implementation of plausible backprop-like learning in the brain using a typical cortical microcircuit network, consisting of pyramidal cells (PCs) and feedback SST interneurons. A framework of cholinergic neuromodulation was introduced whereby the cholinergic system integrates error information and releases acetylcholine to modulate synaptic plasticity across the brain. The cortical microcircuit model was specifically incorporated into this framework to examine the effects of cholinergic modulation on learning when coupled with a plausible credit assignment mechanism. Cholinergic modulation in this model was found to enhance learning in image classification tasks and excluding modulation of the SST interneurons in the microcircuit impaired learning. Consistent with experimental observations, the model of the cholinergic system is involved in promoting learning and the modulation of SST interneurons is important for this process (Hasselmo, 2006; Lovett-Barron et al., 2014).

Although the cholinergic system is widely assumed to be a global modulator, recent evidence of variability in the spatiotemporal dynamics of acetylcholine fluctuations across the brain brings this into question (Disney and Higley, 2020; Lohani et al., 2020). Testing neuromodulation at varying degrees of diffuseness in the model revealed that both homogeneous and heterogeneous cholinergic signals are beneficial for learning and that the advantages of each differ depending on the task. This implies that the cholinergic system may switch between distinct forms of signalling in order to perform a context-dependent function.

In this model, cholinergic modulation of the learning rate is interpreted as occurring via changes in dendritic excitability mediated by acetylcholine (Williams and Fletcher, 2019). However, these changes would alter the error signals that are propagated back through the network resulting in updates to synapses that would not approximate backprop gradients and thus impair learning. Therefore, a correction for the modulation is required on the outgoing signals which predicts that there needs to be a depression, proportional to the modulation factor, in the synaptic strength of the feedback pyramidal-pyramidal and lateral interneuron-pyramidal connections. Palacios-Filardo et al. (2021) demonstrated that acetylcholine reduces excitatory responses from the temporammonic (TA) feedback pathway onto CA1 pyramidal cells. In Chapter 4, *ex vivo* whole-cell patch clamp experiments showed that the application of carbachol, an acetylcholine receptor agonist, suppressed inhibitory synaptic transmission of OLM interneurons onto CA1 pyramidal cells. Furthermore, this cholinergic modulation of the OLM interneuron output was found to be functionally decoupled from its potentiating effects on membrane excitability. The likely locus of action is presynaptic although this is difficult to determine using PPR of optogenetically-evoked inhibitory transmission.

Additionally, in Chapter 4, the importance of OLM interneuron activity for synaptic plasticity was examined. *Ex vivo* whole-cell patch clamp recordings of CA1 pyramidal cells showed that simultaneous theta-burst stimulation of the Schaffer (SC) collateral and TA feedback pathways is capable of inducing SC LTP. When the OLM interneurons were optogenetically activated at the same time as the theta-burst stimulation of the two pathways, SC LTP was abolished and the TA pathway experienced LTD. This validates the EDN model used in Chapter 2 as it demonstrates that local feedback, dendrite-targeting interneurons are capable of modulating synaptic plasticity. When this is combined with the neuromodulatory effect of acetylcholine on the inhibitory output of these interneurons and its ability to control the formation of plateau potentials, it provides a potential mechanism for cholinergic regulation of synaptic plasticity (Williams and Fletcher, 2019; Griesius et al., 2022). As the OLM interneurons target the distal dendritic compartments of the CA1 PC, one possible mechanism for regulating synaptic plasticity is through mediating the coupling of the proximal and distal CA1 inputs required to elicit plateau potentials. Plateau potentials are important for the generation and feature selectivity of CA1 place cells (Bittner et al., 2015, 2017). Cholinergic neuromodulation, through modulating the inhibitory transmission of OLM interneurons, can thus represent a higher level regulation of hippocampus synaptic plasticity. Acetylcholine release can operate as a novelty detection signal in the hippocampus, facilitating memory formation by inhibiting the OLM-to-PC and glutamatergic synapses onto OLM interneurons (Zheng et al., 2011). Similar to the cholinergic framework in Chapter 2, cholinergic modulation is gating synaptic plasticity. This generates new testable hypotheses to investigate whether OLM interneurons and cholinergic neuromodulation contributes to the production of plateau potentials, place cell formation and remapping (Udakis et al., 2020; Hasselmo, 2006).

In Chapter 5, a dual optogenetic approach was undertaken to investigate whether the pharmacological effect of carbachol on the OLM-to-PC synapse seen in Chapter 4 remained consistent with endogenous acetylcholine. Validation of the ChRger and Chrimson opsins indicated that the combination of these two opsins could enable independent optical control of distinct neural populations. The ChRger and Chrimson opsins were expressed under the control of Cre-recombinase in Chrna-Cre mice stereotaxically injected in the medial septum (MS) with the AAV2.9-ChAT-Cre-P2A-EGFP vector as an attempt to provide selective expression localised to ChAT-positive MS neurons and Chrna2-expressing hippocampal OLM interneurons. Whole cell recordings of OLM interneuron transmission from CA1 pyramidal cells revealed that optogenetic stimulation of cholinergic fibres produced inconsistent effects on the inhibitory output of OLM interneurons onto CA1 pyramidal cells. This was likely due to non-specific optogenetic activation of non-cholinergic fibres from the MS, as the the AAV2.9-ChAT-Cre-P2A-EGFP vector displayed significant off-target expression of Cre-recombinase in ChAT-negative MS neurons. Therefore, a more selective approach to expressing the two opsins in the desired neural populations is required to determine more conclusively whether endogenous acetylcholine release has any effects on OLM inhibitory transmission.

In Chapter 2's framework of cholinergic neuromodulation, the error information must be communicated to the cholinergic system but there is no apparent mechanism in the error-encoding dendritic network (EDN) model used in Chapter 2. In Chapter 3, a novel model of credit assignment called the BurstCCN was proposed that was built upon the EDN and Burstprop models and employs the concept of burst ensemble multiplexing, which provides a means for communicating error information without interference from inference signals (Sacramento et al., 2018; Payeur et al., 2021; Naud and Sprekeler, 2018). By combining the biological mechanisms of bursting, STP, and dendrite-targeting inhibition, this model has been shown to effectively learn in a continuous, online setting reminiscent of learning in the brain. The BurstCCN has demonstrated its capability to learn complex image classification tasks from machine learning, which many current models of approximate backprop find difficult to solve. Furthermore, by including the model within a reinforcement learning framework, it has been demonstrated that it does not require explicit supervised learning signals. Additionally, even when architectural constraints are imposed to conform with Dale's law, the BurstCCN is still capable of learning to perform well on MNIST. The model also makes a prediction that the variance in bursting activity in the network would correlate with the the severity of the errors made during learning. Together, this new model improves on earlier biologically plausible backprop-like learning models and offers a mechanism to convey error information to the cholinergic system. Future research could explore the effects of applying the cholinergic neuromodulatory framework to this model.

To conclude, this research has shown that the cholinergic system as an adaptive learning module, when coupled with plausible backprop-like credit assignment, is capable of modulating synaptic plasticity and facilitating learning in the brain. *Ex vivo* slice electrophysiology experiments have demonstrated that dendrite-targeting interneurons, specifically the OLM interneurons of the hippocampus, can regulate synaptic plasticity and that cholinergic modulation of these interneurons have profound effects on their inhibitory output. These key findings suggest a potential mechanism for plateau potential generation and memory formation via the gating of synaptic plasticity through the cholinergic modulation of OLM inhibitory transmission.

BIBLIOGRAPHY

N. Ahmad, M. A. van Gerven, and L. Ambrogioni.
Gait-prop: A biologically plausible learning rule derived from backpropagation of error.
Advances in Neural Information Processing Systems, 33:10913–10923, 2020.

M. Akrout, C. Wilson, P. C. Humphreys, T. Lillicrap, and D. Tweed.
Using weight mirrors to improve feedback alignment.
arXiv preprint arXiv:1904.05391, 2019.

E. X. Albuquerque, E. F. Pereira, M. Alkondon, and S. W. Rogers.
Mammalian nicotinic acetylcholine receptors: from structure to function.
Physiological reviews, 2009.

J. Aljadeff, J. D'amour, R. E. Field, R. C. Froemke, and C. Clopath.
Cortical credit assignment by hebbian, neuromodulatory and inhibitory plasticity, 2019.

M. Alkondon and E. X. Albuquerque.
The nicotinic acetylcholine receptor subtypes and their function in the hippocampus and cerebral cortex.
Progress in brain research, 145:109–120, 2004.

M. Andrychowicz, M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, B. Shillingford, and N. De Freitas.
Learning to learn by gradient descent by gradient descent.
In *Advances in neural information processing systems*, pages 3981–3989, 2016.

R. Anil, V. Gupta, T. Koren, and Y. Singer.
Memory-Efficient Adaptive Optimization.
arXiv:1901.11150 [cs, math, stat], Sept. 2019.
URL `http://arxiv.org/abs/1901.11150`.
arXiv: 1901.11150.

S. D. Antic, W.-L. Zhou, A. R. Moore, S. M. Short, and K. D. Ikonomu.
The decade of the dendritic nmda spike.
Journal of neuroscience research, 88(14):2991–3001, 2010.

A. Asok, E. R. Kandel, and J. B. Rayman.
The neurobiology of fear generalization, 2019.
ISSN 16625153.

E. Augusto and F. Gambino.
Can nmda spikes dictate computations of local networks and behavior?
Frontiers in Molecular Neuroscience, 12:238, 2019.

J. L. Ba, J. R. Kiros, and G. E. Hinton.
Layer normalization.
arXiv preprint arXiv:1607.06450, 2016.

E. C. Ballinger, M. Ananth, D. A. Talmage, and L. W. Role.
Basal forebrain cholinergic circuits and signaling in cognition and cognitive decline.
Neuron, 91(6):1199–1218, 2016.

J. Basu and S. A. Siegelbaum.
The corticohippocampal circuit, synaptic plasticity, and memory.
Cold Spring Harbor perspectives in biology, 7(11):a021733, 2015.

S. Beaulieu, L. Frati, T. Miconi, J. Lehman, K. O. Stanley, J. Clune, and N. Cheney.
Learning to continually learn.
arXiv preprint arXiv:2002.09571, 2020.

C. N. Bedbrook, K. K. Yang, J. E. Robinson, E. D. Mackey, V. Gradinaru, and F. H. Arnold.
Machine learning-guided channelrhodopsin engineering enables minimally invasive optogenetics.
Nature methods, 16(11):1176–1184, 2019.

L. Benardo and D. Prince.
Cholinergic pharmacology of mammalian hippocampal pyramidal cells.
Neuroscience, 7(7):1703–1712, 1982.

G.-q. Bi and M.-m. Poo.
Synaptic modification by correlated activity: Hebb's postulate revisited.
Annual review of neuroscience, 24(1):139–166, 2001.

E. L. Bienenstock, L. N. Cooper, and P. W. Munro.
Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex.
Journal of Neuroscience, 2(1):32–48, 1982.

L. Biewald.

Experiment tracking with weights and biases, 2020.
URL `https://www.wandb.com/`.
Software available from wandb.com.

K. C. Bittner, C. Grienberger, S. P. Vaidya, A. D. Milstein, J. J. Macklin, J. Suh, S. Tonegawa, and J. C. Magee.
Conjunctive input processing drives feature selectivity in hippocampal ca1 neurons.
Nature neuroscience, 18(8):1133–1142, 2015.

K. C. Bittner, A. D. Milstein, C. Grienberger, S. Romani, and J. C. Magee.
Behavioral time scale synaptic plasticity underlies CA1 place fields.
Science, 2017.
ISSN 10959203.
doi: 10.1126/science.aan3846.

T. Bliss, G. Collingridge, and R. Morris.
Synaptic plasticity in health and disease: introduction and overview, 2014.

T. V. Bliss and G. L. Collingridge.
A synaptic model of memory: long-term potentiation in the hippocampus.
Nature, 361(6407):31–39, 1993.

T. V. Bliss and T. Lømo.
Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path.
The Journal of physiology, 232(2):331–356, 1973.

A. Blokland, W. Honig, and W. G. Raaijmakers.
Effects of intra-hippocampal scopolamine injections in a repeated spatial acquisition task in the rat.
Psychopharmacology, 109:373–376, 1992.

T. Branco and M. Häusser.
Synaptic integration gradients in single cortical pyramidal cell dendrites.
Neuron, 69(5):885–892, 2011.

E. Brazhnik, R. Muller, and S. Fox.
Muscarinic blockade slows and degrades the location-specific firing of hippocampal pyramidal cells.
Journal of Neuroscience, 23(2):611–621, 2003.

A. M. Brickman, Y. Stern, and S. A. Small.
Hippocampal subregions differentially associate with standardized memory tests.

Hippocampus, 21(9):923–928, 2011.

G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba.
Openai gym.
arXiv preprint arXiv:1606.01540, 2016.

Z. Brzosko, S. Zannone, W. Schultz, C. Clopath, and O. Paulsen.
Sequential neuromodulation of hebbian plasticity offers mechanism for effective reward-based navigation.
Elife, 6:e27756, 2017.

K. A. Buchanan, M. M. Petrovic, S. E. Chamberlain, N. V. Marrion, and J. R. Mellor.
Facilitation of long-term potentiation by muscarinic m1 receptors is mediated by inhibition of sk channels.
Neuron, 68(5):948–963, 2010.

J. Bures, A. A. Fenton, Y. Kaminsky, and L. Zinyuk.
Place cells and place navigation.
Proceedings of the National Academy of Sciences, 94(1):343–350, 1997.

T. S. Burger, M. E. Rule, and T. O'Leary.
Active dendrites enable robust spiking computations despite timing jitter.
bioRxiv, pages 2023–03, 2023.

K. J. Burke, C. M. Keeshen, and K. J. Bender.
Two forms of synaptic depression produced by differential neuromodulation of presynaptic calcium channels.
neuron, 99(5):969–984, 2018.

C. F. Cadieu, H. Hong, D. L. Yamins, N. Pinto, D. Ardila, E. A. Solomon, N. J. Majaj, and J. J. DiCarlo.
Deep neural networks rival the representation of primate it cortex for core visual object recognition.
PLoS computational biology, 10(12):e1003963, 2014.

L. Cai, R. B. Gibbs, and D. A. Johnson.
Recognition of novel objects and their location in rats with selective cholinergic lesion of the medial septum.
Neuroscience letters, 506(2):261–265, 2012.

N. A. Cayco-Gajic, C. Clopath, and R. A. Silver.
Sparse synaptic connectivity is required for decorrelation and pattern separation in feedforward networks.

Nature Communications, 8(1), 2017.
ISSN 20411723.
doi: 10.1038/s41467-017-01109-y.

C. A. Cea-del Rio, J. J. Lawrence, L. Tricoire, F. Erdelyi, G. Szabo, and C. J. McBain.
M3 muscarinic acetylcholine receptor expression confers differential cholinergic modulation to neurochemically distinct hippocampal basket cell subtypes.
Journal of Neuroscience, 30(17):6011–6024, 2010.

C. A. Cea-del Rio, J. J. Lawrence, F. Erdelyi, G. Szabo, and C. J. McBain.
Cholinergic modulation amplifies the intrinsic oscillatory properties of ca1 hippocampal cholecystokinin-positive interneurons.
The Journal of physiology, 589(3):609–627, 2011.

C. Q. Chiu, J. S. Martenson, M. Yamazaki, R. Natsume, K. Sakimura, S. Tomita, S. J. Tavalin, and M. J. Higley.
Input-specific nmdar-dependent potentiation of dendritic gabaergic inhibition.
Neuron, 97(2):368–377, 2018.

A. Citri and R. C. Malenka.
Synaptic plasticity: multiple forms, functions, and mechanisms.
Neuropsychopharmacology, 33(1):18–41, 2008.

R. E. Clark and L. R. Squire.
Similarity in form and function of the hippocampus in rodents, monkeys, and humans.
Proceedings of the National Academy of Sciences, 110(supplement_2):10365–10370, 2013.

J. S. Cohen, J. P. LaROCHE, and E. Beharry.
Response perseveration in the hippocampal lesioned rat.
Psychonomic Science, 23(3):221–223, 1971.

A. E. Cole and R. Nicoll.
Characterization of a slow cholinergic post-synaptic potential recorded in vitro from rat hippocampal pyramidal cells.
The Journal of physiology, 352(1):173–188, 1984.

M. W. Cole, P. Laurent, and A. Stocco.
Rapid instructed task learning: A new window into the human brain's unique capacity for flexible cognitive control.
Cognitive, affective & behavioral neuroscience, 13(1):1–22, Mar. 2013.
ISSN 1530-7026.
doi: 10.3758/s13415-012-0125-7.
URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3557598/`.

L. L. Colgin, E. I. Moser, and M.-B. Moser.
Understanding memory through hippocampal remapping.
Trends in neurosciences, 31(9):469–477, 2008.

S. H. Collin, B. Milivojevic, and C. F. Doeller.
Hippocampal hierarchical networks for space, time, and memory.
Current opinion in behavioral sciences, 17:71–76, 2017.

G. L. Collingridge, S. Peineau, J. G. Howland, and Y. T. Wang.
Long-term depression in the cns.
Nature reviews neuroscience, 11(7):459–473, 2010.

J. Cornford, D. Kalajdzievski, M. Leite, A. Lamarquette, D. M. Kullmann, and B. Richards.
Learning to live with dale's principle: Anns with separate excitatory and inhibitory units.
bioRxiv, pages 2020–11, 2020.

F. Crick.
The recent excitement about neural networks.
Nature, 337(6203):129–132, 1989.

H. Dannenberg, K. Young, and M. Hasselmo.
Modulation of hippocampal circuits by muscarinic and nicotinic receptors.
Frontiers in neural circuits, 11:102, 2017.

S. Dasari and A. T. Gulledge.
M1 and m4 receptors modulate hippocampal pyramidal neurons.
Journal of neurophysiology, 105(2):779–792, 2011.

N. D. Daw and P. N. Tobler.
Value learning through reinforcement: the basics of dopamine and reinforcement learning.
In *Neuroeconomics*, pages 283–298. Elsevier, 2014.

S. H. Dennis, F. Pasqui, E. M. Colvin, H. Sanger, A. J. Mogg, C. C. Felder, L. M. Broad, S. M. Fitzjohn, J. T. Isaac, and J. R. Mellor.
Activation of muscarinic m1 acetylcholine receptors induces long-term potentiation in the hippocampus.
Cerebral cortex, 26(1):414–426, 2016.

L. Descarries.
The hypothesis of an ambient level of acetylcholine in the central nervous system.
Journal of Physiology-Paris, 92(3-4):215–220, 1998.

A. A. Disney and M. J. Higley.

Diverse Spatiotemporal Scales of Cholinergic Signaling in the Neocortex.
The Journal of neuroscience : the official journal of the Society for Neuroscience, 40(4), 2020.
ISSN 15292401.
doi: 10.1523/JNEUROSCI.1306-19.2019.

K. Doya.
Metalearning and neuromodulation.
Neural Networks, 15(4-6), 2002.
ISSN 08936080.
doi: 10.1016/S0893-6080(02)00044-8.

J. Duchi, E. Hazan, and Y. Singer.
Adaptive subgradient methods for online learning and stochastic optimization.
Journal of machine learning research, 12(Jul):2121–2159, 2011.

J. T. Dudman, D. Tsay, and S. A. Siegelbaum.
A role for synaptic inputs at distal dendrites: instructive signals for hippocampal long-term plasticity.
Neuron, 56(5):866–879, 2007.

R. Eglen.
Muscarinic receptor subtypes in neuronal and non-neuronal cholinergic function.
Autonomic and Autacoid Pharmacology, 26(3):219–233, 2006.

E. Elvander, P. Schött, J. Sandin, B. Bjelke, J. Kehr, T. Yoshitake, and S. Ögren.
Intraseptal muscarinic ligands and galanin: influence on hippocampal acetylcholine and cognition.
Neuroscience, 126(3):541–557, 2004.

C. J. Frazier, Y. D. Rollins, C. R. Breese, S. Leonard, R. Freedman, and T. V. Dunwiddie.
Acetylcholine activates an α-bungarotoxin-sensitive nicotinic current in rat hippocampal interneurons, but not pyramidal cells.
Journal of Neuroscience, 18(4):1187–1195, 1998.

T. F. Freund and G. Buzsáki.
Interneurons of the hippocampus.
Hippocampus, 6(4):347–470, 1996.

F. Gambino, S. Pagès, V. Kehayas, D. Baptista, R. Tatti, A. Carleton, and A. Holtmaat.
Sensory-evoked ltp driven by dendritic plateau potentials in vivo.
Nature, 515(7525):116–119, 2014.

S. Ge and J. A. Dani.
Nicotinic acetylcholine receptors at glutamate synapses facilitate long-term depression or potentiation.
Journal of Neuroscience, 25(26):6084–6091, 2005.

W. Gerstner, M. Lehmann, V. Liakoni, D. Corneil, and J. Brea.
Eligibility traces and plasticity on behavioral time scales: experimental support of neohebbian three-factor learning rules.
Frontiers in neural circuits, 12:53, 2018.

K. P. Giese, N. B. Fedorov, R. K. Filipkowski, and A. J. Silva.
Autophosphorylation at thr286 of the α calcium-calmodulin kinase ii in ltp and learning.
Science, 279(5352):870–873, 1998.

L. M. Giocomo and M. E. Hasselmo.
Neuromodulation by glutamate and acetylcholine can change circuit dynamics by regulating the relative influence of afferent input and excitatory feedback.
Molecular neurobiology, 36:184–200, 2007.

P. W. Glimcher.
Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis.
Proceedings of the National Academy of Sciences, 108(supplement_3):15647–15654, 2011.

M. Goard and Y. Dan.
Basal forebrain activation enhances cortical coding of natural scenes.
Nature neuroscience, 12(11):1444–1449, 2009.

I. Goodfellow, Y. Bengio, and A. Courville.
Deep learning.
MIT press, 2016.

A. J. Granger, W. Wang, K. Robertson, M. El-Rifai, A. F. Zanello, K. Bistrong, A. Saunders, B. W. Chow, V. Nuñez, M. Turrero García, et al.
Cortical chat+ neurons co-transmit acetylcholine and gaba in a target-and brain-region-specific manner.
Elife, 9:e57749, 2020.

R. Gray, A. S. Rajan, K. A. Radcliffe, M. Yakehiro, and J. A. Dani.
Hippocampal synaptic transmission enhanced by low concentrations of nicotine.
Nature, 383(6602):713–716, 1996.

W. Greedy, H. W. Zhu, J. Pemberton, J. Mellor, and R. P. Costa.
Single-phase deep learning in cortico-cortical networks.
arXiv preprint arXiv:2206.11769, 2022.

C. Grienberger, X. Chen, and A. Konnerth.
Nmda receptor-dependent multidendrite ca2+ spikes required for hippocampal burst firing in vivo.
Neuron, 81(6):1274–1281, 2014.

S. Griesius, C. O'Donnell, S. Waldron, K. L. Thomas, D. M. Dwyer, L. S. Wilkinson, J. Hall, E. S. Robinson, and J. R. Mellor.
Reduced expression of the psychiatric risk gene dlg2 (psd93) impairs hippocampal synaptic integration and plasticity.
Neuropsychopharmacology, 47(7):1367–1378, 2022.

U. Güçlü and M. A. van Gerven.
Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream.
Journal of Neuroscience, 35(27):10005–10014, 2015.

J. Guerguiev, T. P. Lillicrap, and B. A. Richards.
Towards deep learning with segregated dendrites.
ELife, 6:e22901, 2017.

A. T. Gulledge and Y. Kawaguchi.
Phasic cholinergic signaling in the hippocampus: functional homology with the neocortex?
Hippocampus, 17(5):327–332, 2007.

J. Haam and J. L. Yakel.
Cholinergic modulation of the hippocampal region and memory function.
Journal of neurochemistry, 142:111–121, 2017.

J. Haam, J. Zhou, G. Cui, and J. L. Yakel.
Septal cholinergic neurons gate hippocampal output to entorhinal cortex via oriens lacunosum moleculare interneurons.
Proceedings of the National Academy of Sciences, 115(8):E1886–E1895, 2018.

B. Hangya, S. P. Ranade, M. Lorenc, and A. Kepecs.
Central cholinergic neurons are rapidly recruited by reinforcement feedback.
Cell, 162(5):1155–1168, 2015.

M. T. Harnett, N.-L. Xu, J. C. Magee, and S. R. Williams.

Potassium channels control the interaction between active dendritic integration compartments in layer 5 cortical pyramidal neurons.
Neuron, 79(3):516–529, 2013.

M. E. Hasselmo.
The role of acetylcholine in learning and memory, 2006.
ISSN 09594388.

H. Hayashi, J. Koushik, and G. Neubig.
Eve: A gradient based optimization method with locally and globally adaptive learning rates.
arXiv preprint arXiv:1611.01505, 2016.

N. Hiratani, Y. Mehta, T. Lillicrap, and P. E. Latham.
On the stability and scalability of node perturbation learning.
Advances in Neural Information Processing Systems, 35:31929–31941, 2022.

E. Honoré, A. Khlaifia, A. Bosson, and J.-C. Lacaille.
Hippocampal somatostatin interneurons, long-term synaptic plasticity and memory.
Frontiers in Neural Circuits, 15:687558, 2021.

C. Hull.
Prediction signals in the cerebellum: beyond supervised motor learning.
elife, 9:e54073, 2020.

S. Ikonen, R. McMahan, M. Gallagher, H. Eichenbaum, and H. Tanila.
Cholinergic system regulation of spatial representation by the hippocampus.
Hippocampus, 12(3):386–397, 2002.

T. Jarsky, A. Roxin, W. L. Kath, and N. Spruston.
Conditional dendritic spike propagation following distal synaptic activation of hippocampal ca1 pyramidal neurons.
Nature neuroscience, 8(12):1667–1676, 2005.

P. Jarzebowski, C. S. Tang, O. Paulsen, and Y. A. Hay.
Impaired spatial learning and suppression of sharp wave ripples by cholinergic activation at the goal location.
Elife, 10:e65998, 2021.

K. J. Jeffery, M. I. Anderson, R. Hayman, and S. Chakraborty.
A proposed architecture for the neural representation of spatial context.
Neuroscience & Biobehavioral Reviews, 28(2):201–218, 2004.

M. Jepma, S. B. Brown, P. R. Murphy, S. C. Koelewijn, B. De Vries, A. M. van den Maagdenberg, and S. Nieuwenhuis.
Noradrenergic and cholinergic modulation of belief updating.
Journal of Cognitive Neuroscience, 30(12), 2018.
ISSN 15308898.
doi: 10.1162/jocn_a_01317.

D. Ji, R. Lape, and J. A. Dani.
Timing and location of nicotinic activity enhances or depresses hippocampal synaptic plasticity.
Neuron, 31(1):131–141, 2001.

M. Jing, P. Zhang, G. Wang, J. Feng, L. Mesik, J. Zeng, H. Jiang, S. Wang, J. C. Looby, N. A. Guagliardo, et al.
A genetically encoded fluorescent acetylcholine indicator for in vitro and in vivo studies.
Nature biotechnology, 36(8):726–737, 2018.

M. Jing, Y. Li, J. Zeng, P. Huang, M. Skirzewski, O. Kljakic, W. Peng, T. Qian, K. Tan, J. Zou, et al.
An optimized acetylcholine sensor for monitoring in vivo cholinergic activity.
Nature methods, 17(11):1139–1146, 2020.

A. Kamondi, L. Acsady, and G. Buzsáki.
Dendritic spikes are enhanced by cooperative network activity in the intact hippocampus.
Journal of Neuroscience, 18(10):3919–3928, 1998.

A. M. Kaufman, T. Geiller, and A. Losonczy.
A role for the locus coeruleus in hippocampal ca1 place cell reorganization during spatial reward learning.
Neuron, 105(6):1018–1026, 2020.

R. Kemker, M. McClure, A. Abitino, T. Hayes, and C. Kanan.
Measuring catastrophic forgetting in neural networks, 2017.

N. S. Keskar and R. Socher.
Improving generalization performance by switching from adam to sgd.
arXiv preprint arXiv:1712.07628, 2017.

H. Kim, M. Beierlein, and B. W. Connors.
Inhibitory control of excitable dendrites in neocortex.
Journal of Neurophysiology, 74(4):1810–1814, 1995.

W. B. Kim and J.-H. Cho.
Encoding of contextual fear memory in hippocampal–amygdala circuit.
Nature communications, 11(1):1382, 2020.

D. P. Kingma and J. Ba.
Adam: A method for stochastic optimization.
arXiv preprint arXiv:1412.6980, 2014.

A. K. Kinnischtzke, D. J. Simons, and E. E. Fanselow.
Motor cortex broadly engages excitatory and inhibitory neurons in somatosensory barrel cortex.
Cerebral cortex, 24(8):2237–2248, 2014.

N. C. Klapoetke, Y. Murata, S. S. Kim, S. R. Pulver, A. Birdsey-Benson, Y. K. Cho, T. K. Morimoto, A. S. Chuong, E. J. Carpenter, Z. Tian, et al.
Independent optical excitation of distinct neural populations.
Nature methods, 11(3):338–346, 2014.

T. Klausberger and P. Somogyi.
Neuronal diversity and temporal dynamics: The unity of hippocampal circuit operations, 2008.
ISSN 00368075.

V. R. Konda and J. N. Tsitsiklis.
Onactor-critic algorithms.
SIAM journal on Control and Optimization, 42(4):1143–1166, 2003.

M. d. C. Krawczyk, J. Millán, M. G. Blake, and M. M. Boccia.
Role of prediction error and the cholinergic system on memory reconsolidation processes in mice.
Neurobiology of Learning and Memory, 185:107534, 2021.

A. Krizhevsky, V. Nair, and G. Hinton.
Cifar-10 (canadian institute for advanced research).
2014.
URL `http://www.cs.toronto.edu/~kriz/cifar.html`.

D. Kunin, A. Nayebi, J. Sagastuy-Brena, S. Ganguli, J. Bloom, and D. Yamins.
Two routes to scalable credit assignment without weight symmetry.
In *International Conference on Machine Learning*, pages 5511–5521. PMLR, 2020.

M. Larkum.
A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex.
Trends in neurosciences, 36(3):141–151, 2013.

M. E. Larkum and J. J. Zhu.
Signaling of layer 1 and whisker-evoked ca2+ and na+ action potentials in distal and terminal dendrites of rat neocortical pyramidal neurons in vitro and in vivo.
Journal of neuroscience, 22(16):6991–7005, 2002.

M. E. Larkum, J. J. Zhu, and B. Sakmann.
A new cellular mechanism for coupling inputs arriving at different cortical layers.
Nature, 398(6725):338–341, 1999.

M. E. Larkum, T. Nevian, M. Sandler, A. Polsky, and J. Schiller.
Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle.
Science, 325(5941):756–760, 2009.

M. E. Larkum, J. Wu, S. A. Duverdin, and A. Gidon.
The guide to dendritic spikes of the mammalian cortex in vitro and in vivo.
Neuroscience, 2022.

J. Larson and E. Munkácsy.
Theta-burst ltp.
Brain research, 1621:38–50, 2015.

J. Larson, D. Wong, and G. Lynch.
Patterned stimulation at the theta frequency is optimal for the induction of hippocampal long-term potentiation.
Brain research, 368(2):347–350, 1986.

J. J. Lawrence.
Cholinergic control of gaba release: emerging parallels between neocortex and hippocampus.
Trends in neurosciences, 31(7):317–327, 2008.

J. J. Lawrence, J. M. Statland, Z. M. Grinspan, and C. J. McBain.
Cell type-specific dependence of muscarinic signalling in mouse hippocampal stratum oriens interneurones.
The Journal of physiology, 570(3):595–610, 2006.

R. N. Leão, S. Mikulovic, K. E. Leão, H. Munguba, H. Gezelius, A. Enjin, K. Patra, A. Eriksson, L. M. Loew, A. B. Tort, et al.
Olm interneurons differentially modulate ca3 and entorhinal inputs to hippocampal ca1 neurons.
Nature neuroscience, 15(11):1524–1530, 2012.

Y. LeCun and C. Cortes.
MNIST handwritten digit database.
2010.
URL `http://yann.lecun.com/exdb/mnist/`.

D.-H. Lee, S. Zhang, A. Fischer, and Y. Bengio.
Difference target propagation.

In *Joint european conference on machine learning and knowledge discovery in databases*, pages 498–515. Springer, 2015.

S. Lee, I. Kruglikov, Z. J. Huang, G. Fishell, and B. Rudy.
A disinhibitory circuit mediates motor integration in the somatosensory cortex.
Nature neuroscience, 16(11):1662–1670, 2013.

S. Leutgeb, J. K. Leutgeb, C. A. Barnes, E. I. Moser, B. L. McNaughton, and M.-B. Moser.
Independent codes for spatial and episodic memory in hippocampal neuronal ensembles.
Science, 309(5734):619–623, 2005.

A. I. Levey, S. M. Edmunds, V. Koliatsos, R. G. Wiley, and C. J. Heilman.
Expression of m1-m4 muscarinic acetylcholine receptor proteins in rat hippocampus and regulation by cholinergic innervation.
Journal of Neuroscience, 15(5):4077–4092, 1995.

T. P. Lillicrap, D. Cownden, D. B. Tweed, and C. J. Akerman.
Random synaptic feedback weights support error backpropagation for deep learning.
Nature communications, 7(1):1–10, 2016.

T. P. Lillicrap, A. Santoro, L. Marris, C. J. Akerman, and G. Hinton.
Backpropagation and the brain.
Nature Reviews Neuroscience, 21(6):335–346, 2020.

J. Y. Lin, P. M. Knutsen, A. Muller, D. Kleinfeld, and R. Y. Tsien.
Reachr: a red-shifted variant of channelrhodopsin enables deep transcranial optogenetic excitation.
Nature neuroscience, 16(10):1499–1508, 2013.

L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han.
On the variance of the adaptive learning rate and beyond.
arXiv preprint arXiv:1908.03265, 2019.

S. Lohani, A. H. Moberly, H. Benisty, B. Landa, M. Jing, Y. Li, M. J. Higley, and J. A. Cardin.
Dual color mesoscopic imaging reveals spatiotemporally heterogeneous coordination of cholinergic and neocortical activity.
BioRXiv, pages 2020–12, 2020.

A. Losonczy, J. K. Makara, and J. C. Magee.
Compartmentalized dendritic plasticity and input feature storage in neurons.
Nature, 452(7186):436–441, 2008.

M. Lovett-Barron, G. F. Turi, P. Kaifosh, P. H. Lee, F. Bolze, X. H. Sun, J. F. Nicoud, B. V. Zemelman, S. M. Sternson, and A. Losonczy.
Regulation of neuronal input transformations by tunable dendritic inhibition.
Nature Neuroscience, 15(3), 2012.
ISSN 10976256.
doi: 10.1038/nn.3024.

M. Lovett-Barron, P. Kaifosh, M. A. Kheirbek, N. Danielson, J. D. Zaremba, T. R. Reardon, G. F. Turi, R. Hen, B. V. Zemelman, and A. Losonczy.
Dendritic inhibition in the hippocampus supports fear learning.
Science, 343(6173):857–863, 2014.

X. Ma, Y. Zhang, L. Wang, N. Li, E. Barkai, X. Zhang, L. Lin, and J. Xu.
The firing of theta state-related septal cholinergic neurons disrupt hippocampal ripple oscillations via muscarinic receptors.
Journal of Neuroscience, 40(18):3591–3603, 2020.

G. Major, M. E. Larkum, and J. Schiller.
Active properties of neocortical pyramidal neuron dendrites.
Annual review of neuroscience, 36:1–24, 2013.

R. C. Malenka and M. F. Bear.
Ltp and ltd: an embarrassment of riches.
Neuron, 44(1):5–21, 2004.

L. E. Martinetti, K. E. Bonekamp, D. M. Autio, H.-H. Kim, and S. R. Crandall.
Short-term facilitation of long-range corticocortical synapses revealed by selective optical stimulation.
Cerebral cortex, 32(9):1932–1949, 2022.

C. O. Martinez, V. H. Do, J. Martinez Jr, and B. E. Derrick.
Associative long-term potentiation (ltp) among extrinsic afferents of the hippocampal ca3 region in vivo.
Brain research, 940(1-2):86–94, 2002.

J. Mattis, K. M. Tye, E. A. Ferenczi, C. Ramakrishnan, D. J. O'shea, R. Prakash, L. A. Gunaydin, M. Hyun, L. E. Fenno, V. Gradinaru, et al.
Principles for applying optogenetic tools derived from direct comparative analysis of microbial opsins.
Nature methods, 9(2):159–172, 2012.

J. McGaughy, R. A. Koene, H. Eichenbaum, and M. E. Hasselmo.

Cholinergic deafferentation of the entorhinal cortex in rats impairs encoding of novel but not familiar stimuli in a delayed nonmatch-to-sample task.
Journal of Neuroscience, 25(44):10273–10281, 2005.

S. McKenzie, A. J. Frank, N. R. Kinsky, B. Porter, P. D. Rivière, and H. Eichenbaum.
Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas.
Neuron, 83(1):202–215, 2014.

T. Miconi, A. Rawal, J. Clune, and K. O. Stanley.
Backpropamine: training self-modifying neural networks with differentiable neuromodulated plasticity, 2020.

S. Mikulovic, C. E. Restrepo, M. M. Hilscher, K. Kullander, and R. N. Leão.
Novel markers for olm interneurons in the hippocampus, 2015.

M. Minsky.
Steps Toward Artificial Intelligence, 1961.
ISSN 00968390.

V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller.
Playing atari with deep reinforcement learning.
arXiv preprint arXiv:1312.5602, 2013.

V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al.
Human-level control through deep reinforcement learning.
nature, 518(7540):529–533, 2015.

A. S. Morcos, D. G. T. Barrett, N. C. Rabinowitz, and M. Botvinick.
On the importance of single directions for generalization, 2018.

R. Morris, E. Anderson, G. a. Lynch, and M. Baudry.
Selective impairment of learning and blockade of long-term potentiation by an n-methyl-d-aspartate receptor antagonist, ap5.
Nature, 319(6056):774–776, 1986.

R. G. Morris, P. Garrud, J. a. Rawlins, and J. O'Keefe.
Place navigation impaired in rats with hippocampal lesions.
Nature, 297(5868):681–683, 1982.

E. I. Moser, K. A. Krobert, M.-B. Moser, and R. G. Morris.
Impaired spatial learning after saturation of long-term potentiation.
Science, 281(5385):2038–2042, 1998.

M.-B. Moser, D. C. Rowland, and E. I. Moser.
Place cells, grid cells, and memory.
Cold Spring Harbor perspectives in biology, 7(2):a021808, 2015.

W. Muñoz and B. Rudy.
Spatiotemporal specificity in cholinergic control of neocortical function.
Current opinion in neurobiology, 26:149–160, 2014.

J. M. Murre.
Learning and categorization in modular neural networks.
Psychology Press, 2014.

S. Naskar, J. Qi, F. Pereira, C. R. Gerfen, and S. Lee.
Cell-type-specific recruitment of gabaergic interneurons in the primary somatosensory cortex by long-range inputs.
Cell reports, 34(8):108774, 2021.

R. Naud and H. Sprekeler.
Sparse bursts optimize information transmission in a multiplexed neural code.
Proceedings of the National Academy of Sciences, 115(27):E6329–E6338, 2018.

E. L. Newman, K. Gupta, J. R. Climer, C. K. Monaghan, and M. E. Hasselmo.
Cholinergic modulation of cognitive processing: insights drawn from computational models.
Frontiers in behavioral neuroscience, page 24, 2012.

R. A. Nicoll.
A brief history of long-term potentiation.
Neuron, 93(2):281–290, 2017.

J. Obermayer, M. B. Verhoog, A. Luchicchi, and H. D. Mansvelder.
Cholinergic modulation of cortical microcircuits is layer-specific: Evidence from rodent, monkey and human brain, 2017.
ISSN 16625110.

Y. Ohmura, T. Izumi, T. Yamaguchi, I. Tsutsui-Kimura, T. Yoshida, and M. Yoshioka.
The serotonergic projection from the median raphe nucleus to the ventral hippocampus is involved in the retrieval of fear memory through the corticotropin-releasing factor type 2 receptor.
Neuropsychopharmacology, 35(6):1271–1278, 2010.

J. O'Keefe.
Place units in the hippocampus of the freely moving rat.
Experimental neurology, 51(1):78–109, 1976.

J. O'Keefe and N. Burgess.
Geometric determinants of the place fields of hippocampal neurons.
Nature, 381(6581):425–428, 1996.

J. O'Keefe and J. Dostrovsky.
The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat.
Brain research, 1971.

J. O'Keefe and L. Nadel.
Précis of o'keefe & nadel's the hippocampus as a cognitive map.
Behavioral and Brain Sciences, 2(4):487–494, 1979.

J. O'Keefe, N. Burgess, J. G. Donnett, K. J. Jeffery, and E. A. Maguire.
Place cells, navigational accuracy, and the human hippocampus.
Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences, 353 (1373):1333–1340, 1998.

S. V. Ovsepian, R. Anwyl, and M. J. Rowan.
Endogenous acetylcholine lowers the threshold for long-term potentiation induction in the CA1 area through muscarinic receptor activation: In vivo study.
European Journal of Neuroscience, 2004.
ISSN 0953816X.
doi: 10.1111/j.1460-9568.2004.03582.x.

J. Palacios-Filardo and J. R. Mellor.
Neuromodulation of hippocampal long-term synaptic plasticity.
Current opinion in neurobiology, 54:37–43, 2019.

J. Palacios-Filardo, M. Udakis, G. A. Brown, B. G. Tehan, M. S. Congreve, P. J. Nathan, A. J. Brown, and J. R. Mellor.
Acetylcholine prioritises direct synaptic inputs from entorhinal cortex to ca1 by differential modulation of feedforward inhibitory circuits.
Nature communications, 12(1):5475, 2021.

K. C. Pang and R. Nocera.
Interactions between 192-igg saporin and intraseptal cholinergic and gabaergic drugs: role of cholinergic medial septal neurons in spatial working memory.
Behavioral neuroscience, 113(2):265, 1999.

A. Payeur, J. Guerguiev, F. Zenke, B. Richards, and R. Naud.
Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits.
bioRxiv, 2020.

A. Payeur, J. Guerguiev, F. Zenke, B. A. Richards, and R. Naud.
Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits.
Nature neuroscience, 24(7):1010–1019, 2021.

V. Pedrosa and C. Clopath.
The interplay between somatic and dendritic inhibition promotes the emergence and stabilization of place fields.
PLoS computational biology, 16(7):e1007955, 2020.

K. A. Pelkey, R. Chittajallu, M. T. Craig, L. Tricoire, J. C. Wester, and C. J. McBain.
Hippocampal gabaergic inhibitory interneurons.
Physiological reviews, 97(4):1619–1747, 2017.

I. Petrof, A. N. Viaene, and S. M. Sherman.
Properties of the primary somatosensory cortex projection to the primary motor cortex in the mouse.
Journal of neurophysiology, 113(7):2400–2407, 2015.

M. M. Petrovic, J. Nowacki, V. Olivo, K. Tsaneva-Atanasova, A. D. Randall, and J. R. Mellor.
Inhibition of post-synaptic kv7/kcnq/m channels facilitates long-term potentiation in the hippocampus.
PLoS One, 7(2):e30402, 2012.

M. R. Picciotto, M. J. Higley, and Y. S. Mineur.
Acetylcholine as a neuromodulator: cholinergic signaling shapes nervous system function and behavior.
Neuron, 76(1):116–129, 2012.

F. G. Pike, R. M. Meredith, A. W. Olding, and O. Paulsen.
Postsynaptic bursting is essential for 'hebbian'induction of associative long-term potentiation at excitatory synapses in rat hippocampus.
The Journal of physiology, 518(2):571–576, 1999.

L. Pinto, M. J. Goard, D. Estandian, M. Xu, A. C. Kwan, S.-H. Lee, T. C. Harrison, G. Feng, and Y. Dan.
Fast modulation of visual perception by basal forebrain cholinergic neurons.
Nature neuroscience, 16(12):1857–1863, 2013.

T. Pitler and B. Alger.
Cholinergic excitation of gabaergic interneurons in the rat hippocampal slice.
The Journal of Physiology, 450(1):127–142, 1992.

F. Pouille and M. Scanziani.
Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition.
Science, 293(5532):1159–1163, 2001.

F. Pouille, O. Watkinson, M. Scanziani, and A. J. Trevelyan.
The contribution of synaptic location to inhibitory gain control in pyramidal cells.
Physiological reports, 1(5), 2013.

L. Y. Prince, T. Bacon, R. Humphries, K. Tsaneva-Atanasova, C. Clopath, and J. R. Mellor.
Separable actions of acetylcholine and noradrenaline on neuronal ensemble formation in hippocampal ca3 circuits.
PLoS Computational Biology, 17(10):e1009435, 2021.

D. D. Rasmusson.
The role of acetylcholine in cortical synaptic plasticity.
Behavioural Brain Research, 115(2), 2000.
ISSN 01664328.
doi: 10.1016/S0166-4328(00)00259-X.

S. A. Raza, A. Albrecht, G. Çalışkan, B. Müller, Y. E. Demiray, S. Ludewig, S. Meis, N. Faber, R. Hartig, B. Schraven, et al.
Hipp neurons in the dentate gyrus mediate the cholinergic modulation of background context memory salience.
Nature communications, 8(1):189, 2017.

B. A. Richards and T. P. Lillicrap.
Dendritic solutions to the credit assignment problem.
Current opinion in neurobiology, 54:28–36, 2019.

B. A. Richards, T. P. Lillicrap, P. Beaudoin, Y. Bengio, R. Bogacz, A. Christensen, C. Clopath, R. P. Costa, A. de Berker, S. Ganguli, et al.
A deep learning framework for neuroscience.
Nature neuroscience, 22(11):1761–1770, 2019.

P. R. Roelfsema and A. Holtmaat.
Control of synaptic plasticity in deep cortical networks.
Nature Reviews Neuroscience, 19(3):166–180, 2018.

M. T. Rogan, U. V. Stäubli, and J. E. LeDoux.
Fear conditioning induces associative long-term potentiation in the amygdala.
Nature, 390(6660):604–607, 1997.

G. Rose and T. Dunwiddie.
Induction of hippocampal long-term potentiation using physiologically patterned stimulation.
Neuroscience letters, 69(3):244–248, 1986.

Z. B. Rosen, S. Cheung, and S. A. Siegelbaum.
Midbrain dopamine neurons bidirectionally regulate ca3-ca1 synaptic drive.
Nature neuroscience, 18(12):1763–1771, 2015.

S. Royer, B. V. Zemelman, A. Losonczy, J. Kim, F. Chance, J. C. Magee, and G. Buzsáki.
Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition.
Nature Neuroscience, 15(5), 2012.
ISSN 10976256.
doi: 10.1038/nn.3077.

L. M. T.-G. Ruivo, K. L. Baker, M. W. Conway, P. J. Kinsley, G. Gilmour, K. G. Phillips, J. T. Isaac, J. P. Lowry, and J. R. Mellor.
Coordinated acetylcholine release in prefrontal cortex and hippocampus is associated with arousal and reward on distinct timescales.
Cell reports, 18(4):905–917, 2017.

D. E. Rumelhart, G. E. Hinton, and R. J. Williams.
Learning representations by back-propagating errors.
Nature, 1986.
ISSN 00280836.
doi: 10.1038/323533a0.

J. Sacramento, R. P. Costa, Y. Bengio, and W. Senn.
Dendritic cortical microcircuits approximate the backpropagation algorithm.
In *Advances in Neural Information Processing Systems*, pages 8721–8732, 2018.

A. C. Sales, K. J. Friston, M. W. Jones, A. E. Pickering, and R. J. Moran.
Locus Coeruleus tracking of prediction errors optimises cognitive flexibility: An Active Inference model.
PLoS Computational Biology, 15(1), 2019.
ISSN 15537358.
doi: 10.1371/journal.pcbi.1006267.

S. J. Sara.
The locus coeruleus and noradrenergic modulation of cognition.
Nature reviews neuroscience, 10(3):211–223, 2009.

J. Schiller, G. Major, H. J. Koester, and Y. Schiller.

Nmda spikes in basal dendrites of cortical pyramidal neurons.
Nature, 404(6775):285–289, 2000.

L. C. Schmid, M. Mittag, S. Poll, J. Steffen, J. Wagner, H. R. Geis, I. Schwarz, B. Schmidt, M. K. Schwarz, S. Remy, and M. Fuhrmann.
Dysfunction of Somatostatin-Positive Interneurons Associated with Memory Deficits in an Alzheimer's Disease Model.
Neuron, 92(1), 2016.
ISSN 10974199.
doi: 10.1016/j.neuron.2016.08.034.

J.-W. Seok and C. Cheong.
Functional dissociation of hippocampal subregions corresponding to memory types and stages.
Journal of physiological anthropology, 39(1):1–11, 2020.

M. Shapiro, D. Simon, D. Olton, F. Gage III, O. Nilsson, and A. Björklund.
Intrahippocampal grafts of fetal basal forebrain tissue alter place fields in the hippocampus of rats with fimbria-fornix lesions.
Neuroscience, 32(1):1–18, 1989.

N. E. Sharkey and A. J. Sharkey.
An analysis of catastrophic interference.
Connection Science, 1995.

N. Shazeer and M. Stern.
Adafactor: Adaptive Learning Rates with Sublinear Memory Cost.
arXiv:1804.04235 [cs, stat], Apr. 2018.
URL `http://arxiv.org/abs/1804.04235`.
arXiv: 1804.04235.

M. E. Sheffield, M. D. Adoff, and D. A. Dombeck.
Increased prevalence of calcium transients across the dendritic arbor during place field formation.
Neuron, 96(2):490–504, 2017.

H. Z. Shouval, S. S.-H. Wang, and G. M. Wittenberg.
Spike timing dependent plasticity: a consequence of more fundamental learning rules.
Frontiers in computational neuroscience, 4:19, 2010.

A. P. Simon, F. Poindessous-Jazat, P. Dutar, J. Epelbaum, and M.-H. Bassant.
Firing properties of anatomically identified neurons in the medial septum of anesthetized and unanesthetized restrained rats.
Journal of Neuroscience, 26(35):9038–9046, 2006.

S. Siwani, A. S. França, S. Mikulovic, A. Reis, M. M. Hilscher, S. J. Edwards, R. N. Leão, A. B. Tort, and K. Kullander.
Olmα2 cells bidirectionally modulate learning.
Neuron, 99(2):404–412, 2018.

P. J. Sjöström and M. Häusser.
A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons.
Neuron, 51(2):227–238, 2006.

P. J. Sjöström, G. G. Turrigiano, and S. B. Nelson.
Rate, timing, and cooperativity jointly determine cortical synaptic plasticity.
Neuron, 32(6):1149–1164, 2001.

P. J. Sjostrom, E. A. Rancz, A. Roth, and M. Hausser.
Dendritic excitability and synaptic plasticity.
Physiological reviews, 88(2):769–840, 2008.

J.-H. Son and U. H. Winzer-Serhan.
Expression of neuronal nicotinic acetylcholine receptor subunit mrnas in rat hippocampal gabaergic interneurons.
Journal of Comparative Neurology, 511(2):286–299, 2008.

N. Spruston.
Pyramidal neurons: dendritic structure and synaptic integration.
Nature Reviews Neuroscience, 9(3):206–221, 2008.

N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov.
Dropout: a simple way to prevent neural networks from overfitting.
The journal of machine learning research, 15(1):1929–1958, 2014.

C. K. Starkweather and N. Uchida.
Dopamine signals as temporal difference errors: recent advances.
Current opinion in neurobiology, 67:95–105, 2021.

P. Strata, R. Harvey, et al.
Dale's principle.
Brain research bulletin, 50(5):349–350, 1999.

J. F. Sturgill, P. Hegedus, S. Li, Q. Chevy, A. Siebels, M. Jing, Y. Li, B. Hangya, and A. Kepecs.
Basal forebrain-derived acetylcholine encodes valence-free reinforcement prediction error.
BioRxiv, pages 2020–02, 2020.

E. Sugisaki, Y. Fukushima, M. Tsukada, and T. Aihara.
Cholinergic modulation on spike timing-dependent plasticity in hippocampal ca1 network.
Neuroscience, 192:91–101, 2011.

R. S. Sutton.
Learning to predict by the methods of temporal differences.
Machine learning, 3:9–44, 1988.

J. D. Sweatt.
Hippocampal function in cognition.
Psychopharmacology, 174:99–110, 2004.

G. G. Szabó, N. Holderith, A. I. Gulyás, T. F. Freund, and N. Hájos.
Distinct synaptic properties of perisomatic inhibitory cell types and their different modulation by cholinergic receptor activation in the ca3 region of the mouse hippocampus.
European Journal of Neuroscience, 31(12):2234–2246, 2010.

H. Takahashi and J. C. Magee.
Pathway interactions and synaptic plasticity in the dendritic tuft regions of ca1 pyramidal neurons.
Neuron, 62(1):102–111, 2009.

L. M. Teles-Grilo Ruivo and J. R. Mellor.
Cholinergic modulation of hippocampal network function.
Frontiers in synaptic neuroscience, 5:2, 2013.

T. Tieleman, G. Hinton, et al.
Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude.
COURSERA: Neural networks for machine learning, 4(2):26–31, 2012.

R. Tremblay, S. Lee, and B. Rudy.
Gabaergic interneurons in the neocortex: from cellular properties to circuits.
Neuron, 91(2):260–292, 2016.

D. Tsay, J. T. Dudman, and S. A. Siegelbaum.
Hcn1 channels constrain synaptically evoked ca2+ spikes in distal dendrites of ca1 pyramidal neurons.
Neuron, 56(6):1076–1089, 2007.

I. Tsutsui-Kimura, H. Matsumoto, K. Akiti, M. M. Yamada, N. Uchida, and M. Watabe-Uchida.
Distinct temporal difference error signals in dopamine axons in three regions of the striatum in a decision-making task.
Elife, 9:e62390, 2020.

J. Turchi, R. C. Saunders, and M. Mishkin.
Effects of cholinergic deafferentation of the rhinal cortex on visual recognition memory in monkeys.
Proceedings of the National Academy of Sciences, 102(6):2158–2161, 2005.

L. Tyan, S. Chamberland, E. Magnin, O. Camiré, R. Francavilla, L. S. David, K. Deisseroth, and L. Topolnik.
Dendritic inhibition provided by interneuron-specific cells controls the firing rate and timing of the hippocampal feedback inhibitory circuitry.
Journal of Neuroscience, 34(13):4534–4547, 2014.

M. Udakis, V. Pedrosa, S. E. Chamberlain, C. Clopath, and J. R. Mellor.
Interneuron-specific plasticity at parvalbumin and somatostatin inhibitory synapses onto ca1 pyramidal neurons shapes hippocampal output.
Nature communications, 11(1):4395, 2020.

J. Urban-Ciecko and A. L. Barth.
Somatostatin-expressing neurons in cortical networks.
Nature Reviews Neuroscience, 17(7):401, 2016.

J. Urban-Ciecko, J.-S. Jouhanneau, S. E. Myal, J. F. Poulet, and A. L. Barth.
Precisely timed nicotinic activation drives sst inhibition in neocortical circuits.
Neuron, 97(3):611–625, 2018.

J. Viswanathan, F. Rémy, N. Bacon-Macé, and S. J. Thorpe.
Long term memory for noise: Evidence of robust encoding of very short temporal acoustic patterns.
Frontiers in Neuroscience, 10(NOV), 2016.
ISSN 1662453X.
doi: 10.3389/fnins.2016.00490.

T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner.
Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks.
Science, 334(6062):1569–1573, 2011.

J. Waters and S. J. Smith.
Vesicle pool partitioning influences presynaptic diversity and weighting in rat hippocampal synapses.
The Journal of physiology, 541(3):811–823, 2002.

J. Waters, T. Nevian, B. Sakmann, F. Helmchen, J. Byrne, and J. Sweatt.

Action potentials in dendrites and spike-timing-dependent plasticity.
Learning and memory, (4):803–828, 2008.

P. J. Werbos.
Generalization of backpropagation with application to a recurrent gas market model.
Neural networks, 1(4):339–356, 1988.

J. Werfel, X. Xie, and H. S. Seung.
Learning curves for stochastic gradient descent in linear feedforward networks.
In *Advances in neural information processing systems*, pages 1197–1204, 2004.

J. R. Whitlock, A. J. Heynen, M. G. Shuler, and M. F. Bear.
Learning induces long-term potentiation in the hippocampus.
science, 313(5790):1093–1097, 2006.

L. E. Williams and A. Holtmaat.
Higher-order thalamocortical inputs gate synaptic long-term potentiation via disinhibition.
Neuron, 101(1):91–102, 2019.

S. R. Williams and L. N. Fletcher.
A dendritic substrate for the cholinergic control of neocortical output neurons.
Neuron, 101(3):486–499, 2019.

S. R. Williams and G. J. Stuart.
Dependence of epsp efficacy on synapse location in neocortical pyramidal neurons.
Science, 295(5561):1907–1910, 2002.

A. C. Wilson, R. Roelofs, M. Stern, N. Srebro, and B. Recht.
The Marginal Value of Adaptive Gradient Methods in Machine Learning.
In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4148–4158. Curran Associates, Inc., 2017.
URL `http://papers.nips.cc/paper/7003-the-marginal-value-of-adaptive-gradient-methods-in`
`pdf`.

J. Winson.
Loss of hippocampal theta rhythm results in spatial memory deficit in the rat.
Science, 201(4351):160–163, 1978.

R. Wong, D. Prince, and A. Basbaum.
Intradendritic recordings from hippocampal neurons.
Proceedings of the National Academy of Sciences, 76(2):986–990, 1979.

H. Xu, H.-Y. Jeong, R. Tremblay, and B. Rudy.
Neocortical somatostatin-expressing gabaergic interneurons disinhibit the thalamorecipient layer 4.
Neuron, 77(1):155–167, 2013.

N.-l. Xu, M. T. Harnett, S. R. Williams, D. Huber, D. H. O'Connor, K. Svoboda, and J. C. Magee.
Nonlinear dendritic integration of sensory and motor input during an active sensing task.
Nature, 492(7428):247–251, 2012.

D. Yi, S. Ji, and J. Park.
An adaptive optimization method based on learning rate schedule for neural networks.
Applied Sciences, 11(2):850, 2021.

O. Yizhar, L. E. Fenno, M. Prigge, F. Schneider, T. J. Davidson, D. J. O'shea, V. S. Sohal, I. Goshen, J. Finkelstein, J. T. Paz, et al.
Neocortical excitation/inhibition balance in information processing and social dysfunction.
Nature, 477(7363):171–178, 2011.

A. J. Yu and P. Dayan.
Uncertainty, neuromodulation, and attention.
Neuron, 46(4):681–692, 2005.

R. Yuste and D. W. Tank.
Dendritic integration in mammalian neurons, a century after cajal.
Neuron, 16(4):701–716, 1996.

L. Záborszky, P. Gombkoto, P. Varsanyi, M. R. Gielow, G. Poe, L. W. Role, M. Ananth, P. Rajebhosale, D. A. Talmage, M. E. Hasselmo, et al.
Specific basal forebrain–cortical cholinergic circuits coordinate cognitive operations.
Journal of Neuroscience, 38(44):9446–9458, 2018.

S. Zannone, Z. Brzosko, O. Paulsen, and C. Clopath.
Acetylcholine-modulated plasticity in reward-driven navigation: a computational study.
Scientific reports, 8(1):9486, 2018.

M. D. Zeiler.
Adadelta: an adaptive learning rate method.
arXiv preprint arXiv:1212.5701, 2012.

F. Zenke and W. Gerstner.
Limits to high-speed simulations of spiking neural networks using general-purpose computers.
Frontiers in neuroinformatics, 8:76, 2014.

F. Zenke, B. Poole, and S. Ganguli.
Continual learning through synaptic intelligence, 2017.

F. Zhang, M. Prigge, F. Beyrière, S. P. Tsunoda, J. Mattis, O. Yizhar, P. Hegemann, and K. Deisseroth.
Red-shifted optogenetic excitation: a tool for fast neural control derived from volvox carteri.
Nature neuroscience, 11(6):631–633, 2008.

H. Zhang, S.-C. Lin, and M. A. Nicolelis.
Spatiotemporal coupling between hippocampal acetylcholine release and theta oscillations in vivo.
Journal of Neuroscience, 30(40):13431–13440, 2010.

F. Zheng, T. Seeger, B. E. Nixdorf-Bergweiler, and C. Alzheimer.
Layer-specific processing of excitatory signals in ca1 interneurons depends on postsynaptic m2 muscarinic receptors.
Neuroscience letters, 494(3):217–221, 2011.

T. A. Zolnik, J. Ledderose, M. Toumazou, T. Trimbuch, T. Oram, C. Rosenmund, B. J. Eickholt, R. N. Sachdev, and M. E. Larkum.
Layer 6b is driven by intracortical long-range projection neurons.
Cell reports, 30(10):3492–3505, 2020.