

Revealing the (predictive) code of top-down signals in the brain

Christian David Martin (Márton in publications)

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

Imperial College London.

Department of Bioengineering & Centre for Neurotechnology

Imperial College London

February 27, 2020



To my parents, Lydia Martin (médecin sans frontières) & Andreas Martin (artiste sans frontières). Credit for drawing above goes to Andreas Martin / Márton Zsolt, artiste sans frontières [Martinez, 2009, Marton, 2019b, Marton, 2019a].

I, Christian David Martin (Márton in publications), confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

This research was carried out under the joint supervision of Dr. Simon R. Schultz (Imperial College London) and Dr. Bruno B. Averbeck (NIMH/NIH, USA). The data employed in this research was collected by collaborators in the Averbeck lab.

The first part of this research was carried out in the lab of Dr. Simon R. Schultz (Neural Coding Lab) in the Department of Bioengineering & the Centre for Neurotechnology at Imperial College London. The second part was carried out in the lab of Dr. Bruno B. Averbeck (Section on Learning & Decision Making) at the National Institute for Mental Health (NIMH/NIH) in Bethesda, MD (USA).

This work was jointly funded by the Wellcome Trust & the National Institutes of Health under a Wellcome Trust/NIH Four Year PhD studentship grant (106808/Z/15/Z).

Abstract

Reciprocal connections are common in the brain, yet little is known about their functional role. Top-down connections, in particular, remain functionally obscure in both neuroscience and in the nascent field of deep learning. On the theoretical side, predictive coding has been put forward as a framework that assigns specific roles to top-down and bottom-up connections in sensory information processing. It remains unclear, however, if and how the brain implements this predictive code. This work examined top-down signals in the auditory cortex and in the corticostriatal system in macaques in order to validate the claims put forward by predictive coding. This theory suggests there are imbalances in message passing up and down the cortical hierarchy; these imbalances imply cross-frequency couplings should predominate top-down. It is unknown whether these asymmetries are expressed in cross-frequency interactions in the brain. This work examined cross-frequency interactions across four sectors of the macaque auditory cortex. Predictive coding also applies in decision making, where it allows for action selection based on predicted reward (or value). This is more commonly known as reinforcement learning (RL) and is supported by the fronto-striatal systems in the brain. The computational mechanisms that drive learning in this system are unknown, however. This work drew on a recurrent neural network (RNN) model of the dlPFC-dSTR circuit in the brain together with recordings from macaques from the same regions to answer this question. Altogether, the findings are largely consistent with the predictive coding framework.

Acknowledgements

"I am not sure that I exist, actually. I am all the writers that I have read, all the people that I have met, all the women that I have loved; all the cities I have visited."

(Jorge Luis Borges)

I entered this PhD like Borges' protagonist wandering through the desert and chancing upon a strange structure emerging out of the sand. This structure looked like nothing I had seen before, strange dwellings and windy roads that would not let you see beyond a few steps. When you reached out naively to touch what was in front of you, to grab it, to get a hold of it, the shape shifted as if to tell you: the old ways won't do, you need to find new ways to interact with this world. Sometimes the road would lead you up a few steps and you continued in the assumption that up meant better, closer to the light. When you finally reached that door and overcame your Kafkaesque fear of pushing it open, it would yield in a resounding C-major like the gates in Bluebeard's castle. Only to open up a world stranger than the one before, seemingly one level deeper as if holding up what came before. Really, though, it dawned on you that you had been walking down and up the stairs at the same time like that little headless marionette in Escher's drawings.

But I kept walking, paving the road underneath me with the tiles that I picked up from behind. The pace kept accelerating, not because the clock was ticking as I wasn't sure the regular notion of time even held in this place. But because I was setting my own time and my expectations grew. The further I penetrated into this

land, the more I thought I knew about it. Until the music suddenly stopped and I was left recollecting my experiences in the search of lost time, as if after a dream when you are still not really here nor there.

When you finally manage to put it in words and you get a chance to reflect upon, re-tracing your path as if it was not yours but that of an unknown explorer, it dawns on you that you were given the chance to take part in a strangely beautiful, life-changing experience. The cities I dwelled in during my PhD travels, both real and theoretical, have reshaped who I am. It has given me confidence to explore unknown lands, to plunge into them and keep on learning new tools as I go.

It is to my advisors, Dr. Simon R. Schultz and Dr. Bruno B. Averbeck that I owe the chance to enter these lands - and the guidance to find my way out again. I learned a lot about self-discipline. I learned to learn from others, being a part of the Bioengineering department with its multitude of different voices and approaches.

But most importantly, both of my advisors taught me to believe in myself, to believe that I can learn new tools, that I am never limited by who I was, that I can push forward and learn from my mistakes, that I can bootstrap into a new, richer version of myself. They did this through their belief in me, pointing me to the right books, the right language to use. They did this by giving me plenty of opportunities to present my research at conferences throughout the world, to interact with others, to build confidence.

I also owe a lot to my companions in the Bioengineering department at Imperial and the students I was given the opportunity to teach. I learned a lot because others pushed back on things I said. I learned a lot by asking for help from others who knew better. I am also grateful for having become an honorary member of the Greek

gang (inducted by Georgios Zafeiropoulos & Panayotis Christodoulides).

I am also grateful to some of my oldest friends, William Scheinman and Moritz Mueller-Freitag. We shared many experiences along the way. They made life bearable, and even pleasurable at times. They encouraged me to push forward, to share my insights.

Lastly, but most importantly, I am grateful to my parents, Andreas & Lydia Márton. They encouraged me to keep my head high and march on forward.



Contents

1 General Introduction **36**

1.1 Predictive coding overview 36

1.2 Top-down vs. bottom-up information processing 41

1.3 Learning to predict 48

2 Signature patterns for top-down and bottom-up information processing via cross-frequency coupling in macaque auditory cortex **52**

2.1 Significance 54

2.2 Introduction 54

2.3 Methods and Materials 57

2.3.1 Subjects 57

2.3.2 Stimuli 57

2.3.3 Recordings 60

2.3.4 Stimulus presentation and recording parameters 61

2.3.5 Data Pre-Processing 61

2.3.6 Phase-amplitude (PAC) and amplitude-amplitude (AAC) cross-frequency coupling (CFC) Calculation 62

2.3.6.1 Frequency Transformation 62

2.3.6.2 Obtaining CFCs using canonical correlation . . . 63

2.3.6.3 Obtaining the difference between top-down (TD) and bottom-up (BU) coupling 66

2.3.6.4 Obtaining the difference in coupling strength between original (VOC) and synthetic stimuli (EPS/SPS) . . . 66
2.3.7 Statistical Analysis . . . 67
2.4 Results . . . 68
2.4.1 Distinct cross-frequency coupling signatures for bottom-up and top-down coupling, both involving low frequencies . . . 69
2.4.2 Cross-regional coupling mainly driven by interactions between A1 and higher order sectors . . . 71
2.4.3 Top-down and bottom-up cross-frequency signatures generalize to synthetic stimuli . . . 73
2.4.4 Synthetic stimuli show overall lower bi-directional coupling than natural stimuli . . . 75
2.5 Discussion . . . 77

3 Learning to select actions shapes recurrent dynamics in the corticostriatal system 85

3.1 Introduction . . . 87
3.2 Methods and Materials . . . 90
3.2.1 Neural Data . . . 90
3.2.1.1 Subjects . . . 90
3.2.1.2 Task and Stimuli . . . 90
3.2.1.3 Neural Data Analysis . . . 92
3.2.2 Corticostriatal model . . . 94
3.2.2.1 Architecture . . . 94
3.2.2.2 Network system training . . . 96
3.2.2.3 Task Coding . . . 96
3.2.2.4 Autonomous mode . . . 98

3.2.2.5 Model Analysis 99
3.3 Results 100
3.4 Discussion 127
3.5 Conclusion 130
3.6 Acknowledgments 131
3.7 Competing Interests 131

4 General Conclusions **132**

Appendices **135**

A Appendix Chapter 1 **135**

B Appendix Chapter 2 **149**

Bibliography **152**

List of Figures

1.1 **Predictive coding versus representational framework (adapted from [Keller and Mrsic-Flogel, 2018]). A** In the representational framework, internal representations are generated by bottom-up input, while top-down inputs act as modulatory signals. **B** In a hierarchical predictive processing framework, internal representations are updated based on comparison of a top-down prediction and bottom-up input. Prediction errors are sent forward in the hierarchy, while predictions are sent backward. The coordinate transformations between the different areas are the internal models (M). 38

1.2 **Schematic of the proposed hierarchical Bayesian inference framework in the cortex (adapted from [Lee and Mumford, 2003])** The different visual areas (boxes) are linked together as a Markov chain. The activity in V1, x_1 , is influenced by the bottom-up feedforward data x_0 and the probabilistic priors $P(x_1 \mid x_2)$ fed back from V2. The concept of a Markov chain is important computationally because each area is influenced mainly by its direct neighbors. 40

1.3 **Schematic of the canonical microcircuit for predictive coding (adapted from [Bastos et al., 2012a]) A** Here, prediction error populations are highlighted in pink. Inhibitory connections are shown in red, while excitatory connections are in black. (...) In this scheme, expectations (about causes and states) are assigned to (excitatory and inhibitory) interneurons in the supragranular layers, which are passed to infragranular layers. The corresponding prediction errors occupy granular layers, while superficial pyramidal cells encode prediction errors that are sent forward to the next hierarchical level. Conditional expectations and prediction errors on hidden causes are associated with excitatory cell types, while the corresponding quantities for hidden states are assigned to inhibitory cells. Dark circles indicate pyramidal cells. Finally, we have placed the precision of the feedforward prediction errors against the superficial pyramidal cells. This quantity controls the postsynaptic sensitivity or gain to (intrinsic and top-down) presynaptic inputs. **B** This schematic illustrates the functional asymmetry between the spectral activity of superficial and deep cells predicted theoretically. (...) The bottom panel shows the spectral density of deep pyramidal cell activity, given the spectral density of superficial pyramidal cell activity in the top panel. The equation expresses the spectral density of the deep cells as a function of the spectral density of the superficial cells (...). This schematic is meant to illustrate how the relative amounts of low (beta)- and high (gamma)-frequency activity in superficial and deep cells can be explained by the evidence accumulation implicit in predictive coding. 42

1.4 **Different types of cross-frequency coupling (adapted from [Jirsa and Mueller, 2013]).** **A** Signal X at a given constant frequency fluctuating in the amplitude over time (red line). **B** Power to power CFC: Signal Y1 at about 5 times higher frequency than in the signal X showing slow amplitude modulations over time like signal X (red line). **C** Phase to phase CFC: Signal Y2 showing 3:1 phase to phase coupling with signal X. One oscillation period of signal X corresponds to three periods of signal Y2. **D** Phase to power CFC: Signal Y3 with fast amplitude modulations, which are related or coupled with the phase of the signal X. **E** Phase to frequency CFC: Signal Y4 with frequency modulations, which are coupled with phase changes of signal X. **F** Power to frequency CFC: Signal Y5 with frequency modulations, which are coupled with the slow amplitude modulations of signal X (red line). **G** Frequency to frequency CFC: Signal Y6 with slower frequency modulations than in the signal Y5. The different types of CFC are not mutually exclusive [Jensen and Colgin, 2007]. It can be seen for instance, that slow amplitude modulations of Signal X are coupled not only with the amplitude changes of the signal Y1 but also with frequency changes of signals Y5 and Y6, which are at the same time coupled in their frequency modulations. 46

2.1 **Coupling Analysis Overview:** Recordings were made from four μECoG arrays spanning auditory cortex while monkeys listened to one of 20 natural vocalizations (VOC), 20 synthetic envelope-preserved sounds (EPSs) or 20 synthetic spectrum-preserved sounds (SPSs). Electrodes were partitioned into four sectors along the caudorostral axis (S1: A1/ML (primary auditory cortex/middle lateral belt), S2: R (rostral core region of the auditory cortex / AL (anterior lateral belt region of the auditory cortex), S3: RTL (lateral rostrotemporal belt region of the auditory cortex), S4: RTp (rostrotemporal pole area)). We decomposed the signal into its time-frequency representation, and obtained amplitude and phase components for each sector. We investigated two types of coupling across sectors: amplitude-amplitude and phase-amplitude coupling. We computed both types of coupling in the bottom-up and top-down direction. Top-down coupling was defined as coupling in which the source electrode came from a sector of higher order than the target electrode, and vice versa for bottom-up coupling. 59

2.2 **Top-down vs. bottom-up phase-amplitude and amplitude-amplitude coupling in natural vocalizations (VOC). (A-B)** Phase-amplitude coupling (PAC) strength in the top-down (A) and bottom-up direction (B). Depicted are canonical correlation-derived coupling coefficients (see Methods and Materials). **(C)** Difference in top-down and bottom-up in PAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). **(D-E)** Amplitude-amplitude coupling (AAC) strength in the top-down (D) and bottom-up direction (E). Depicted are canonical correlation-derived coupling coefficients. **(F)** Difference between top-down and bottom-up in AAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals. Supporting Figures: Fig. S1-1, S1-2, S1-3 70

2.3 **Top-down versus bottom-up phase-amplitude coupling strength (PAC) across the auditory hierarchy in natural stimuli** Difference in top-down and bottom-up *phase-amplitude* coupling (PAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals. Supporting Figures: Fig. S1-4. 72

2.4 **Top-down vs. bottom-up phase-amplitude and amplitude-amplitude coupling in synthetic envelope-preserved stimuli (EPS) and synthetic spectrum-preserved stimuli (SPS) (A)** Difference in top-down and bottom-up PAC strength in EPS stimuli. **(B)** Difference in top-down and bottom-up AAC strength in EPS stimuli. **(C)** Difference in top-down and bottom-up PAC strength in SPS stimuli. **(D)** Difference in top-down and bottom-up AAC strength in SPS stimuli. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals. Supporting Figures: Fig. S1-5, S1-6, S1-7, S1-8, S1-9, S1-10 74

2.5 **Phase-amplitude and amplitude-amplitude coupling strength in natural vocalizations (VOC) compared to synthetic envelope-preserved sounds (EPS) and spectrum-preserved sounds (SPS) (A-B)** Difference in PAC strength between natural (VOC) and synthetic envelope-preserved sounds (EPS), separately in the top-down (B) and bottom-up direction (C). **(C-D)** Difference in AAC strength between VOC and EPS, separately in the top-down (C) and bottom-up direction (D). Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals. Supporting Figures: Fig. S1-11 76

3.1 **Task and Model Overview.** (A) Sequence of events in the task. (B) Coronal section showing approximate location of recording chamber. (C) The eight possible movement-sequences used in the task. (D) Trial structure. Trials were arranged into blocks of 8 repeats of the correct movement-sequence plus a variable number of error trials (i.e. sequence S1 followed by sequence S5). (E) Corticostriatal model, consisting of a prefrontal network and a striatal network of recurrently connected units. The prefrontal network selects actions based on inputs from the striatal network. The striatal network outputs action values based on inputs specifying actions performed and rewards received. 102

3.2 **Task coding**. The correct movement sequence (S5) was signalled by the corresponding output units while units coding other movement directions stay flat. Actions (light blue) were indicated by pulses in the prefrontal network's output units corresponding to a particular direction. Rewards (dark blue, plotted on top) were delivered at the end of action pulses. Action values (red, plotted on top) were indicated by striatal output units corresponding to a particular direction; action values increased after reward delivery. 104

3.3 **Corticostriatal model: autonomous behavior.** Sample transition between two blocks in which the correct sequence switches (from S2, in magenta, to S4, in light brown; block transition indicated by vertical grey dotted line). Decoded Actions (light blue) are plotted together with rewards (dark blue) and action values (red) for all output units. Action values spike after correctly executed, rewarded actions and increase with successive correct actions. Action values for wrong, unrewarded moves decay quickly. Trials with erroneous movements remain unrewarded and are repeated, just like in the original experiment. Supporting Figures: Fig. S2-1 106

3.4 **Model Performance. A** Outputs of prefrontal model network (blue) versus targets (black) for a sample test block. **B** Outputs of striatal model network (red) versus targets (black) for the same sample test block as in A. **C-D** Mean squared error (MSE) between outputs and targets averaged over 25 test blocks. MSE is depicted as a function of trials after the sequence switched for all output units in the prefrontal (C) and striatal network (D). Supporting Figures: Fig. S2-2 108

3.5 **Comparison of model and recorded data A** CCA projections (canonical variables) for the *prefrontal* model network (left) and the neural data recorded from lPFC (right). These projections show the directions in state-space along which the data is maximally correlated with the model. Each row shows one canonical variable (CV 1-5). Traces are colored based on the mean value of the projection across the entire depicted duration of the trace. Traces show responses for the whole length one movement-sequence which is composed of three sequential movements (black dots mark movement onset). **B** CCA projections for the *striatal* model network (left) and the neural data recorded from dSTR (right). **C** Summary of canonical correlations between model and neural data. CCA analysis provides a spectrum of of correlation coefficients that can be used to assess model performance (see Methods and Materials). The canonical correlation coefficients are shown for the trained model (striatal network in dark red and prefrontal network in dark magenta), as well as for an untrained network with the same inputs which as a baseline (striatal network in light red, and prefrontal network in light magenta). **D** Fraction of correct decisions as a function of trials after the sequence switched, separately for the behavior of the two animals (solid line), the performance of the Q-learning algorithm (broken line; see Methods and Materials-Corticostriatal model-Coding), and the performance of the corticostriatal model network system (dotted line; see Methods and Materials-Corticostriatal model-Autonomous Mode). 109

3.6 **Evolution of latent task coding with learning for PFC recordings and model network. A-D** lPFC neural recordings. Latent task trajectories during the execution of the three-movement sequence depicted in three-dimensional latent space (obtained from dPCA using the sequence dimension; see Methods and Materials-Neural Data-Data Analysis). Latent trajectories for all eight possible movement sequences (S1-S8) are plotted together, for increasing certainty (fraction correct) levels (A-C). **D** Euclidean distance between all sequence trajectories within the two clusters (brightly and lightly colored trajectories) as a function of increasing certainty (see Methods and Materials-Neural Data-Data Analysis). **E-H** PFC model network. Latent task trajectories during the execution of the three-movement sequence depicted in three-dimensional latent space (obtained from dPCA using the sequence dimension, as in A-D). Latent trajectories for all eight possible movement sequences (S1-S8) are plotted together, for increasing certainty levels (E-G). **H** Euclidean distance between all sequence trajectories within the two clusters (brightly and lightly colored trajectories) as a function of increasing certainty level (calculated as in D). 112

3.7 **Evolution of potential surface with learning in PFC model network. A-C** Latent trajectories (obtained through dPCA) for two different movement sequences are depicted in two dimensional latent space (S1 in blue and S4 in red, on the bottom). Solid dots depict the beginning (green) and end (black) of the three-movement sequence. The surface depicts the potential energy of the network in the two dimensional space in which the two sequence trajectories sit. To obtain the potential surface, the network was iterated one step forward with inputs held fixed to a particular chosen timepoint (magenta dot; see Methods and Materials-Corticostriatal Model-Model Analysis). Latent movement-sequence trajectories and potential surface are depicted for increasing certainty (fraction correct) levels in A-C. **D** Minimum path length between the gradient minima of all sequence pairs in the test set. Path length was calculated along the joint gradient surface by using Dijkstra's algorithm (see Methods and Materials-Corticostriatal Model-Model Analysis). Results are depicted for increasing certainty levels. 114

3.8 **Evolution of latent task coding with learning for STR recordings and model network. A-D** dSTR recordings. Latent task trajectories during the execution of the three-movement sequence depicted in three-dimensional latent space (obtained from dPCA using the sequence dimension; see Methods and Materials-Neural Data-Data Analysis). Latent trajectories for all eight possible movement sequences (S1-S8) are plotted all together for increasing certainty (fraction correct) levels (A-D). **D** Euclidean distance between all sequence trajectories within the two clusters (brightly and lightly colored trajectories) as a function of increasing certainty level (see Methods and Materials-Neural Data-Data Analysis). **E-H** STR model network. Latent trajectories for all eight possible movement sequences (S1-S8) are plotted together in dPCA-derived latent space (obtained as in A-D) for increasing certainty (fraction correct) levels (E-G). **H** Euclidean distance between all sequence trajectories within the two clusters (brightly and lightly colored trajectories) as a function of increasing certainty level. 115

3.9 **Evolution of gradient landscape with learning in STR model network. A-C** Latent trajectories (obtained through dPCA) for two different movement sequences (S1 in blue, S4 in red) are depicted in two dimensional latent space (bottom). Solid dots depict the beginning (green) and end (black) of the three-movement sequence. The surface depicts the gradient of the network in the two dimensional space in which the two sequence trajectories sit. To obtain the gradient manifold, the network was iterated one step forward with inputs held fixed to a particular chosen timepoint (magenta dot; see Methods and Materials-Corticostriatal Model-Model Analysis). Latent movement-sequence trajectories and gradient surfaces are depicted for increasing certainty (fraction correct) levels (50%, 75% and 90% certainty in A-C). **D** Minimum path length between the gradient minima of all sequence pairs in the test set. Path length was calculated along the joint gradient surface by using Dijkstra's algorithm (see Methods and Materials-Corticostriatal Model-Model Analysis). Results are depicted for increasing certainty levels. . . . 116

3.10 **Shaping of latent movement sequence representations with learning. A-B** Latent trajectory for a single movement sequence (obtained through dPCA) depicted in two dimensional latent space for increasing certainty levels (from low, in light blue, to high in dark blue). Neural recordings from dSTR (A) and STR model network (B). **C-F** Euclidean distance between a sequence's centroid (multi-dimensional mean across time) and each point in time for all sequences in the test set (averaged across time points), as a function of increasing certainty (fraction correct). Neural recordings from dSTR (C) and STR model network (D). Neural recordings from lPFC (E) and PFC model networks (F). 118

3.11 **Relating the change in latent representations with learning to changes in mean firing rate**. **A** Data from neural recordings in dSTR. (a) The first principal component (dPCA) of neural population activity for the same sequence as in (Fig. 3.10A) plotted across time, for increasing certainty levels (from light to dark blue). (b) Mean firing rate (FR) in Hz over time for increasing certainty levels. FR is averaged over all units in the recording. The correlation between the change in latent trajectories and mean FR with learning is $\rho = 0.2950\,(p = 2.17e-56)$. **B** Data from neural recordings in PFC. (a),(b) as in A. The correlation between the change in latent trajectories and mean FR with learning is $\rho = 0.2400\,(p = 2.41e-37)$. Dots indicate whether the period immediately to the right of the dots was a hold period (red dot) or a move period (green dot). 119

3.12 **Latent representations of correct and wrong movements in the STR model network**. Latent trajectories of single movements (obtained through dPCA) from the first trial after a sequence switch imaged in three dimensional latent space. Representations for correctly executed movement representations (solid lines) are depicted alongside wrongly executed ones (dotted lines). 120

3.13 **Decoding sequence identity from latent trajectories**. Fraction of correctly decoded sequence identities (in %). Sequence identity was decoded by projecting the test sequence into latent space and computing the Euclidean distance to all sequence trajectories in latent space. The test sequence was then assigned the same label as the closest sequence trajectory. Results are averaged across all trials. Decoding performance was at 28% for dSTR (left) and at 59% for PFC latents (right; leave-one-out cross validation). Chance performance was at 12.5%. 121

3.14 **Decoding sequence identity from latent trajectories as a function of certainty level**. Fraction of correctly decoded sequence identities (in %). Sequence identity was decoded by projecting the test sequence into latent space and computing the Euclidean distance to all sequence trajectories in latent space. The test sequence was then assigned the same label as the closest sequence trajectory. Results are averaged across all trials. Decoding results per certainty level for PFC (**A**) and for STR (**B**) (leave-one-out cross validation). Chance performance was at 12.5%. 122

3.15 **Task structure of three-armed bandit task**. **A** Event structure of an individual trial in the three-armed bandit task where animals indicated their choice by saccading to a single choice option. **B** Each block of trials, n, began with the presentation of 3 novel images that served as choice options. A particular set, s, of choice options was repeatedly presented for a number, j, of trials. After $10-30$ trials, one of the three existing choice options was randomly selected and replaced with a novel image. Adapted from [Costa et al., 2019] . . . 122

3.16 **Shaping of latent image set representations in the three-armed bandit task with learning.** Neural responses to each set of images, s, (Fig. 3.15) were projected into latent space using the first three encoding vectors obtained through dPCA. The Euclidean distance was computed between a latent representation's centroid (multi-dimensional mean across time) and each point in time. This measure was computed for every trial $j \in 1...10$ after a novel image was introduced, for all trial blocks, n, and averaged across time. **A** Recordings from the amygdala. The linear fit showed a decreasing trend over trials 1-10 ($p = 0.001$). **B** Recordings from oPFC. The linear fit was not significant. **C** Recordings from PFC. The linear fit showed a decreasing trend ($p = 0.0005$. **D** Recordings from vSTR. The linear fit was not significant. 123

3.17 **Relating the change in latent representations in the three-armed band task to changes in mean firing rate**. **A** Recordings from PFC (a) The first principal component (dPCA) of neural population activity plotted across time, for trials $j \in 1...10$ (light blue to dark blue) after a novel image was introduced. Trajectories are averaged across all blocks, n. (b) Mean firing rate (FR) in Hz over time for trials $j \in 1...10$. FR is averaged over all units in the recording. The correlation between changes in latent trajectories and mean FR over trials was $\rho = 0.8525\,(p < 0.001)$. **B** Recordings from amygdala. (a),(b) computed as in A. The correlation between changes in latent trajectories and mean FR over trials was $\rho = 0.6258\,(p < 0.001)$. . 124

S1-1 **(A) Power vs. frequency of neural signal and (B) regression line fit.** The amplitude of the neural recording signal peaks in the α (7.5-12.5 Hz) frequency range and shows the typical 1/f decay pattern thereafter. The signal depicted here is from stimulus onset (time=0) for presentation of the intact stimulus (VOC), averaged across all variables. Standard deviation across all variables is depicted by the red shaded region. The regression line fit to the neural signal confirms a 1/f decay pattern (B). The δ frequency band (0-2.5 Hz) was excluded from analyses (see Methods and Materials). 136

S1-2 **Distribution of differences in PAC coupling strength for select CFC pairs in natural stimuli (VOC) (A)** θ (5Hz) to β (20Hz) PAC **(B)** θ (5Hz) to *high*γ (90Hz) PAC **(C)** β (20Hz) to *low*γ (55Hz) PAC 137

S1-3 **Distribution of differences in AAC coupling strength for select CFC pairs in natural stimuli (VOC) (A)** θ (5Hz) to β (20Hz) PAC **(B)** θ (5Hz) to *high*γ (90Hz) PAC **(C)** β (20Hz) to *low*γ (55Hz) PAC 138

S1-4 **Top-down versus bottom-up amplitude-amplitude coupling strength (AAC) across the auditory hierarchy in natural stimuli (VOC)** Difference in top-down and bottom-up *amplitude-amplitude* coupling (AAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals. 139

S1-5 **Top-down vs. bottom-up phase-amplitude and amplitude-amplitude coupling in synthetic envelope-preserved vocalizations (EPS). (A-B)** Phase-amplitude coupling (PAC) strength in the top-down (A) and bottom-up direction (B). Depicted are canonical correlation-derived coupling coefficients (see Methods and Materials). **(C)** Difference in top-down and bottom-up in PAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). **(D-E)** Amplitude-amplitude coupling (AAC) strength in the top-down (D) and bottom-up direction (E). Depicted are canonical correlation-derived coupling coefficients. **(F)** Difference between top-down and bottom-up in AAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals. 140

S1-6 **Top-down versus bottom-up phase-amplitude coupling strength (PAC) across the auditory hierarchy in synthetic envelop-preserved stimuli (EPS)** Difference in top-down and bottom-up *phase-amplitude* coupling (PAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals. 141

S1-7 **Top-down versus bottom-up amplitude-amplitude coupling strength (AAC) across the auditory hierarchy in synthetic envelope-preserved stimuli (EPS)** Difference in top-down and bottom-up *amplitude-amplitude* coupling (AAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals. 142

S1-8 **Top-down vs. bottom-up phase-amplitude and amplitude-amplitude coupling in synthetic spectrum-preserved vocalizations (SPS). (A-B)** Phase-amplitude coupling (PAC) strength in the top-down (A) and bottom-up direction (B). Depicted are canonical correlation-derived coupling coefficients (see Methods and Materials). **(C)** Difference in top-down and bottom-up in PAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). **(D-E)** Amplitude-amplitude coupling (AAC) strength in the top-down (D) and bottom-up direction (E). Depicted are canonical correlation-derived coupling coefficients. **(F)** Difference between top-down and bottom-up in AAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals. 143

S1-9 **Top-down versus bottom-up phase-amplitude coupling strength (PAC) across the auditory hierarchy in synthetic spectrum-preserved stimuli (SPS)** Difference in top-down and bottom-up *phase-amplitude* coupling (PAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals. 144

S1-10 **Top-down versus bottom-up amplitude-amplitude coupling strength (AAC) across the auditory hierarchy in synthetic spectrum-preserved stimuli (SPS)** Difference in top-down and bottom-up *amplitude-amplitude* coupling (AAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals. 145

S1-11 **Phase-amplitude and amplitude-amplitude coupling strength in synthetic envelope-preserved sounds (SPS) compared to envelope-preserved sounds (EPS) (A-B)** Difference in PAC strength between synthetic spectrum-preserved (SPS) and synthetic envelope-preserved sounds (EPS), separately in the top-down (B) and bottom-up direction (C). **(C-D)** Difference in AAC strength between SPS and EPS, separately in the top-down (C) and bottom-up direction (D). Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals. 146

S1-12 **Directional low frequency phase (2.5-12.5Hz) to low frequency amplitude (12.5-37.5Hz) coupling (PAC) over time in original stimuli (VOC). A** Bottom-up (blue) and top-down (red) coupling strength over the course of stimulus presentation. **B** Differences in directional coupling strength over the course of stimulus presentation. Significant differences are marked by thick black lines (FDR corrected, $q = 0.01$). Both, top-down and bottom-up coupling strength are elevated in the period immediately after stimulus presentation (0-200ms), whereas top-down coupling dominates in strength thereafter. Results were obtained using the short-time Fourier transform (FFT) and computing coupling strength per time window within the canonical correlation framework (see Methods and Materials). Results are presented averaged across all dimensions (stimuli, channels, cross-regional pairs and animals). 147

S1-13 **Directional low frequency phase (2.5-12.5Hz) to low frequency amplitude (12.5-37.5Hz) coupling (PAC) over time in original stimuli (VOC), broken out by regional pair.** Top-down and bottom-up coupling were found to alternate in strength in the period immediately after stimulus presentation (0-200ms), across all regional pairs. **A** S1 to S2 (blue) vs. S2 to S1 (red), **B** S1 to S3 (blue) vs. S3 to S1 (red), **C** S1 to S4 (blue) vs. S4 to S1 (red), **D** S2 to S3 (blue) vs. S3 to S2 (red), **E** S2 to S4 (blue) vs. S4 to S2 (red), **F** S3 to S4 (blue) vs. S4 to S3 (red). Results were obtained using the Hilbert transform and computing coupling strength per time window within the canonical correlation framework (see Methods and Materials). Results are presented averaged across all dimensions (stimuli, channels, and animals). 148

S2-1 **Corticostriatal model behavior on test set.** Sample transition between two blocks in which the correct sequence switches (from S6 to S4; block transition indicated by vertical grey dotted line). Actions (light blue; outputs of the prefrontal (PFC) network) are plotted together with rewards (dark blue; external inputs) and action values (red; outputs of the striatal (STR) network) for all output units. Action values spike after correctly executed, rewarded actions and increase with successive correct actions. Action values for wrong, unrewarded actions (i.e. on the second trial of the third move, bottom right plot) increase briefly but decay quickly, never reaching the height associated with correct moves. 150

S2-2 **Eigenvalue spectra after training of the network system**. **A** PFC model network, **B** STR model network. 151

Chapter 1

General Introduction

1.1 Predictive coding overview

Predictive coding is built around the idea that the brain uses an internal model to predict incoming sensory information. This internal model is built up from previous sensory experience and interactions with the world. Thus, before predictive coding, there was the idea of an internal model; and before the internal model, there was the question of how the brain could distinguish between self-generated and externally generated sensory input [Keller and Mrsic-Flogel, 2018].

This question arises during eye movements, for instance. Why do we not experience the very motion that drives the eyes during an eye movement? One (old) idea is that there is an internally generated signal that cancels the sense of movement generated from self-initiated actions, as Hermmann von Helmholtz argued [von Helmholtz, 1876]. What remains to be answered is what kind of signal this is. In the representational view (Fig. 1.1A), perception is driven primarily by bottom-up sensory input - a view in which neurons function as simple *feature detectors* as described by Hubel and Wiesel [Hubel and Wiesel, 1959] -, and modulatory feedback is relayed via a signal that provides direct information about the position of eye muscles. In this view, the modulatory feeedback signal is derived directly from

motor cortex, and not from a complex internal model, or a cascade (or hierarchy) of region-specific internal models. In the predictive processing view (Fig. 1.1B), in contrast, perception is more than just bottom-up inputs; it is also immediately affected by our predictions about the world and how our actions will affect new incoming sensory information. In this view, predictions from internal models are conveyed top-down, and prediction errors arising from the mismatch between predictions and new sensory evidence are conveyed bottom-up. Crucially, unlike in the predictive processing view, in the representational view there are no prediction errors because there are no predictions, and there are no predictions because there are no internal models.

The relevance of the predictive coding framework derives from its power to account for a wide range of phenomena [Griffiths et al., 2010, Koster-Hale and Saxe, 2013, Clark, 2013, Summerfield and de Lange, 2014, Keller and Mrsic-Flogel, 2018, Alexander et al., 2018, Singer et al., 2018]. Many visual phenomena can be explained within a predictive coding framework [Rao and Ballard, 1999, Lee and Mumford, 2003, Fiser et al., 2016, Singer et al., 2018]. The earliest evidence in favor came from explaining *endstopping* - the suppression of neural responses when a stimulus extends into the surround of a receptive field - as a prediction error [Rao and Ballard, 1999]. In this view, suppression results from top-down inhibition because of the prediction that a stimulus in a given location will extend into neighboring receptive fields. Visual receptive field properties such as sensory adaptation - the effect that neural activity decreases as a stimulus is repeated - can be explained by a positive prediction error: a top-down prediction that inhibits bottom-up input will result in the suppression of activity as a stimulus becomes more predictable [Summerfield and de Lange, 2014, Spratling, 2017]. Consistent with this view, contrast adaptation has been shown to be modulated by the behavioral relevance of the stimulus [Keller et al., 2017]. When a prediction is falsified, a

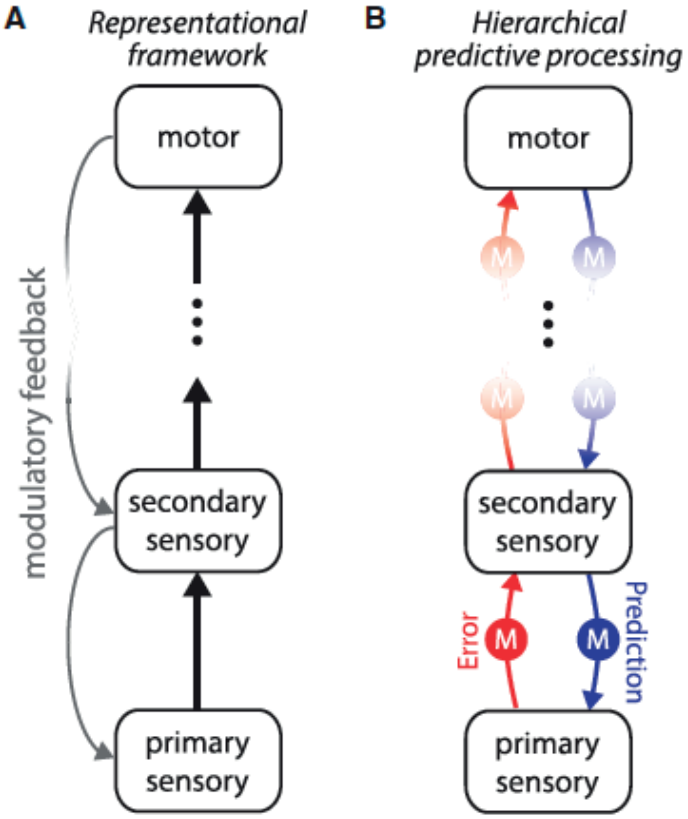


Fig 1.1. Predictive coding versus representational framework (adapted from [Keller and Mrsic-Flogel, 2018]). A In the representational framework, internal representations are generated by bottom-up input, while top-down inputs act as modulatory signals. **B** In a hierarchical predictive processing framework, internal representations are updated based on comparison of a top-down prediction and bottom-up input. Prediction errors are sent forward in the hierarchy, while predictions are sent backward. The coordinate transformations between the different areas are the internal models (M).

negative prediction error should result in a response increase. There were neurons in primary auditory cortex of marmosets that were selective to deviations from expected auditory stimuli [Eliades and Wang, 2008].

The predictive coding framework can potentially be relevant to brain function as a whole, not just to explaining the functioning of sensory cortices [Aitchison and Lengyel, 2017a]. Positive and negative prediction errors have been employed to describe the functioning of the dopaminergic system [Schultz et al., 1997, Langdon et al., 2018]. Cerebellar learning has been described with the help of internal models

and prediction errors [Wolpert et al., 1995, Wolpert et al., 1998]. Vision and action can be viewed as working together to minimize prediction errors [Friston et al., 2003, Friston, 2005, Friston, 2008, Friston and Kiebel, 2009, Friston et al., 2017]. Frontal cortex - key in orchestrating goal-directed behavior - with its wide-ranging functionality (drawing on working memory, attention, and reinforcement learning components, as further explained below) can be cast in a unifying framework based on predictive coding [Alexander et al., 2018]. Dysfunction of the brain such as the positive symptoms of schizophrenia and stereotyped activity as observed in autism could be regarded as a result of abnormally increased or decreased activity of top-down signals [Corlett et al., 2009, Sinha et al., 2014, Keller and Mrsic-Flogel, 2018]. Going further, one could view theory of mind - the way we think about others - as an instantiation of predictive coding [Koster-Hale and Saxe, 2013].

Moving one level down Marr's hierarchy of analysis [Marr and Poggio, 1976], how can the computational mechanisms of predictive coding be implemented algorithmically? Traditionally, it has been conceived of as a hierarchical cascade of predictive processes (Fig. 1.2) in which activity in the higher-order area predicts and constrains activity in the area immediately below [Rao and Ballard, 1999, Lee and Mumford, 2003]. Predictive coding can be cast as *empirical Bayesian inference*. Predictions from the currently best model are fed downward via top-down connections and act as priors ($P(x_1 \mid x_2)$) to the inferential process taking place at a level below. As long as activity at the level below evolves as predicted, the model does not need to change. Only when there is a mismatch between predictions and new sensory evidence, a prediction error (X_1) is generated and relayed to the level above via bottom-up connections. This way of signalling the *difference* between actual and expected activity makes this coding scheme more efficient in terms of the bandwidth required for information transmission (as opposed to the naive approach of relaying all incoming information to the higher-order area). This

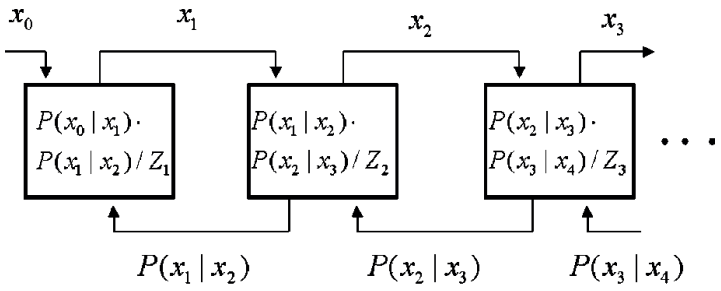


Fig 1.2. Schematic of the proposed hierarchical Bayesian inference framework in the cortex (adapted from [Lee and Mumford, 2003]) The different visual areas (boxes) are linked together as a Markov chain. The activity in V1, x_1 , is influenced by the bottom-up feedforward data x_0 and the probabilistic priors $P(x_1 \mid x_2)$ fed back from V2. The concept of a Markov chain is important computationally because each area is influenced mainly by its direct neighbors.

feature of predictive coding stems from the theory's origins in signal processing where it had been developed as a data compression technique [She and Sun, 1999]. Predictive coding can be equivalently cast as a form of particle filtering and belief propagation [Lee and Mumford, 2003], or as free-energy minimization [Friston et al., 2003,Friston, 2005,Friston, 2008,Friston and Kiebel, 2009,Friston et al., 2017].

If this (Bayesian) view of the brain is correct, then the brain should encode probability distributions and integrate prior expectations with new information in a Bayesian way. Behavioral evidence that humans act as rational Bayesian estimators [Kording and Wolpert, 2004, Kording et al., 2007] can thus be taken as indirect evidence in favor of the predictive coding framework. Indirect only, because one cannot conclude that a system actually implements Bayesian inference just because its outputs look Bayesian. Thus, there are still many details left to be worked out in how predictive coding is actually implemented in the brain [Bastos et al., 2012a, Keller and Mrsic-Flogel, 2018].

Moving another level down Marr's hierarchy [Marr and Poggio, 1976], there are three components to implementing the predictive coding algorithm in neural

circuits. (1) One needs a way to compute the prediction error between bottom-up inputs and predictions, (2) a way to maintain internal models from which to derive predictions, and (3) a way to modulate the weight or precision of the prediction error (potentially modulated by attention). These three functions could potentially be mapped onto different parts of the cortical column (Fig. 1.3A), with superficial pyramidal cells computing prediction errors and adjusting the precision of the update and deep pyramidal cells encoding predictions [Bastos et al., 2012a].

The implications of conveying prediction errors and predictions in a manner prescribed by predictive coding will be the focus of the next section of this introduction (Ch. 1.2) and testing these implications is the goal of the first part of this work (Ch. 2). The section thereafter (Ch. 1.3) will be an introduction to component (2), the maintaining and updating of predictions during learning, which is the focus of the rest of this work (Ch. 3).

1.2 Top-down vs. bottom-up information processing

Predictive coding theory says predictions are conveyed downwards (via top-down connections) while prediction errors resulting from the mismatch between predictions and actual experience are conveyed back upwards (via bottom-up connections) so that predictions can be updated. The differences in informational content carried by top-down and bottom-up connections (predictions and prediction errors, respectively) imply differences in the expected spectral profiles (Fig. 1.3B).

In the generative model that underlies predictive coding, non-linear functions appear only in the top-down direction. Predictions, communicated via top-down connections, are a *non-linear* function of conditional expectations. Conversely, bottom-up connections convey a linear mixture of prediction errors. In sum, in-

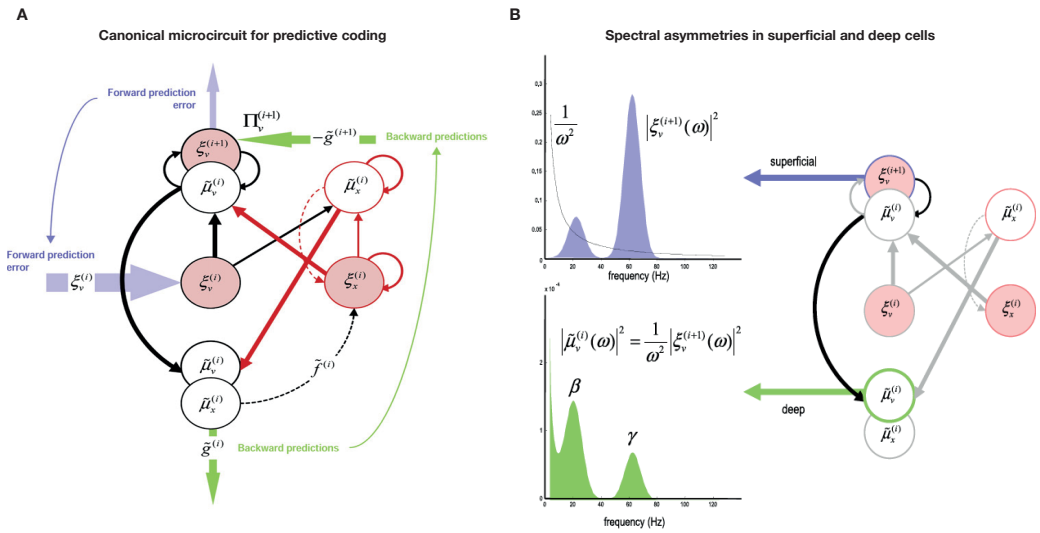


Fig 1.3. Schematic of the canonical microcircuit for predictive coding (adapted from [Bastos et al., 2012a]) A Here, prediction error populations are highlighted in pink. Inhibitory connections are shown in red, while excitatory connections are in black. (...) In this scheme, expectations (about causes and states) are assigned to (excitatory and inhibitory) interneurons in the supragranular layers, which are passed to infragranular layers. The corresponding prediction errors occupy granular layers, while superficial pyramidal cells encode prediction errors that are sent forward to the next hierarchical level. Conditional expectations and prediction errors on hidden causes are associated with excitatory cell types, while the corresponding quantities for hidden states are assigned to inhibitory cells. Dark circles indicate pyramidal cells. Finally, we have placed the precision of the feedforward prediction errors against the superficial pyramidal cells. This quantity controls the postsynaptic sensitivity or gain to (intrinsic and top-down) presynaptic inputs. **B** This schematic illustrates the functional asymmetry between the spectral activity of superficial and deep cells predicted theoretically. (...) The bottom panel shows the spectral density of deep pyramidal cell activity, given the spectral density of superficial pyramidal cell activity in the top panel. The equation expresses the spectral density of the deep cells as a function of the spectral density of the superficial cells (...). This schematic is meant to illustrate how the relative amounts of low (beta)- and high (gamma)-frequency activity in superficial and deep cells can be explained by the evidence accumulation implicit in predictive coding.

formation conveyed via top-down connections (feedback) should exhibit relatively more non-linearities than information conveyed via bottom-up connections (feedforward) [Bastos et al., 2012a].

Another way of viewing this is that neural populations representing predictions should exhibit relatively more low frequencies than those in which prediction errors arise (Fig. 1.3B). This is because the linear accumulation of evidence in higher-order areas is expected to result in the attenuation of higher (gamma-range) frequencies (through Bayesian filtering). Conversely, neural populations representing prediction errors should exhibit relatively more high frequencies. This is because prediction errors are a non-linear function of predictions, and the non-linearity can induce higher (gamma-range) frequencies (think of how a low frequency rhythm is converted into a high frequency rhythm by a squaring operation). In sum, higher-order areas where predictions are formed should exhibit lower frequencies than lower-order areas where prediction errors arise [Bastos et al., 2012a, Friston et al., 2017]. A few studies support this view [Fontolan et al., 2014, Bastos et al., 2015, Michalareas et al., 2016, Chao et al., 2018], but not without exceptions.

It is usually not easy to validate the tenets of the predictive coding framework experimentally and distinguish them from those of the representational framework (Fig. 1.1A vs. B). This is because experimentalists have much better control over bottom-up inputs than over top-down predictions [Keller and Mrsic-Flogel, 2018]. But here we have found a way to tease the two apart by looking at bi-directional information flow. Note that in the representational view (Fig. 1.1A), there is no reason to expect particular asymmetries in the frequency spectra of information conveyed upwards and downwards along the cortical hierarchy.

Cross-frequency coupling (CFC) offers a way to study the spectral profile of

information flow in the brain. Measures such as *coherence* examine linear interactions at the same frequency only, whereas CFC considers (non-linear) interactions across frequencies. Of the many forms CFC could theoretically take on [Jirsa and Mueller, 2013], phase-amplitude (Fig. 1.4D) and amplitude-amplitude (Fig. 1.4B) interactions have shown the highest biological relevance. Phase-amplitude CFC (PAC) occurs when the phase of a target signal at a specific frequency band is modulated by the amplitude of a source signal at a different frequency band; in the case of amplitude-amplitude CFC (AAC), it is the amplitude in the source signal that modulates the target signal.

Based on the predictive coding view outlined above, we would expect to see relatively more non-linear *cross*-frequency interactions in the top-down direction. Conversely, linear *same*-frequency interactions should predominate in the bottom-up direction. Importantly, we do not expect to find stronger (non-linear) cross-frequency interactions in the bottom-up direction. We will test this explicitly in (Ch. 2) using a method that allows us to obtain directional estimates of cross-frequency coupling - and thus obtain directional estimates of the information conveyed across the cortical hierarchy.

CFC, in particular coupling between low-frequency phase and high-frequency power, has been reported across many brain regions and during a wide range of tasks [Schack et al., 2002, Canolty et al., 2006, Jensen and Colgin, 2007, Cohen, 2008, Tort et al., 2008, Tort et al., 2009, Canolty and Knight, 2010, Kendrick et al., 2011, Bosman et al., 2012, Doesburg et al., 2012, Jirsa and Mueller, 2013, Fontolan et al., 2014, Hyafil et al., 2015b, Colgin, 2015]. Mechanistically, CFC could allow for the control of spiking activity through the modulation of the amplitude of high-frequency oscillations by the phase of low frequency oscillations [Buzsáki et al., 2012, Ray and Maunsell, 2011].

Bottom-up CFC has been proposed to help with stimulus encoding [Giraud and Poeppel, 2012, Ghitza et al., 2013, Hyafil et al., 2015a]. In this view, cortical theta oscillations track the syllabic rhythm of speech and, in turn, reset the spiking of neurons at a gamma frequency level through θ-γ PAC. This ensures that neural excitability is optimally aligned with incoming speech.

It is conceivable that top-down (predictive) signals exert influence on lower-order areas through oscillatory mechanisms. [von Stein et al., 2000, Wang, 2010, Salazar et al., 2012, Arnal and Giraud, 2012, Bastos et al., 2012a, Bosman et al., 2012, Bressler and Richter, 2015]. Inter-areal coupling in the same frequency band (also referred to as coherence, or synchrony), in particular in the low frequency range ($< 30Hz$), has been observed during many tasks that draw on top-down processes such as working memory or attention [von Stein et al., 2000, Lakatos et al., 2005, Lakatos et al., 2008, Wang, 2010, Salazar et al., 2012, Bressler and Richter, 2015, Park et al., 2015, Lakatos et al., 2016]. Low-frequency oscillations in human auditory cortex, for instance, have been shown to be modulated by frontal signals in the delta and theta frequency range; oscillations in auditory cortex then, in turn, increased their coupling to continuous speech [Park et al., 2015]. Top-down oscillatory processes may reflect the modulatory activity of predictions on auditory processing [Poeppel et al., 2008, Wild et al., 2010, Davis et al., 2011, Gagnepain et al., 2012, Arnal and Giraud, 2012, Chao et al., 2018]. The effect of *non*-linear, *cross*-frequency coupling in the top-down direction has been largely left unexplored, however.

To a large part, this is because previous studies were not able to properly assess the differential effect of CFC in the two directions. Many previous studies that explored CFC did not compute directional estimates [Cohen, 2008, Canolty et al.,

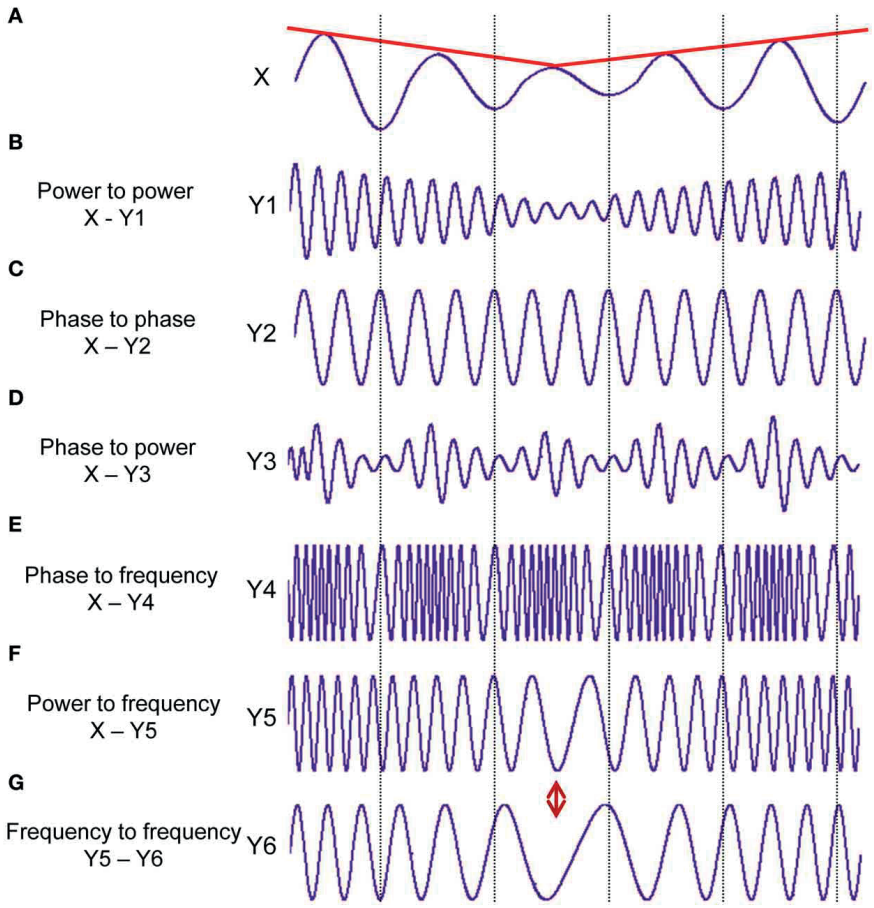


Fig 1.4. Different types of cross-frequency coupling (adapted from [Jirsa and Mueller, 2013]). A Signal X at a given constant frequency fluctuating in the amplitude over time (red line). **B** Power to power CFC: Signal Y1 at about 5 times higher frequency than in the signal X showing slow amplitude modulations over time like signal X (red line). **C** Phase to phase CFC: Signal Y2 showing 3:1 phase to phase coupling with signal X. One oscillation period of signal X corresponds to three periods of signal Y2. **D** Phase to power CFC: Signal Y3 with fast amplitude modulations, which are related or coupled with the phase of the signal X. **E** Phase to frequency CFC: Signal Y4 with frequency modulations, which are coupled with phase changes of signal X. **F** Power to frequency CFC: Signal Y5 with frequency modulations, which are coupled with the slow amplitude modulations of signal X (red line). **G** Frequency to frequency CFC: Signal Y6 with slower frequency modulations than in the signal Y5. The different types of CFC are not mutually exclusive [Jensen and Colgin, 2007]. It can be seen for instance, that slow amplitude modulations of Signal X are coupled not only with the amplitude changes of the signal Y1 but also with frequency changes of signals Y5 and Y6, which are at the same time coupled in their frequency modulations.

2006, Bruns and Eckhorn, 2004, Tort et al., 2010, Voytek et al., 2013, van Wijk et al., 2015] and were thus not able to distinguish between bottom-up (feedforward) or top-down (feedback) components. Other studies that considered directionality employed different measures to look at direction of information flow and coupling strength, and do not consider *cross*-frequency interactions [Bastos et al., 2012b, Michalareas et al., 2016, Bastos et al., 2015, Fontolan et al., 2014]. The dynamic causal modelling (DCM) framework [Friston et al., 2017, Friston et al., 2003, Chen et al., 2008] can be employed to compute directional estimates of interactions across frequencies [Furl et al., 2014]. Current implementations of DCM, however, apply a dimensionality reduction step in order to enhance computational efficiency by applying a singular value decomposition to the data from the source region, and then projecting the data from the target region into this reduced space to compute coupling estimates [Chen et al., 2008]. The variance preserving dimensions in the source and target space are not necessarily the dimensions that maximize correlations between source and target space, so interactions may be missed.

This work (Ch. 2) proposes a new method to obtain directional estimates of cross-frequency coupling (CFC) which remedies shortcomings of previous approaches. This method integrates Granger-causality within a canonical correlation framework to allow for computationally efficient calculation of CFC. This method is subsequently employed to study differences in bi-directional information processing through CFC in macaque auditory cortex. In this way, the work presented here examined whether the tenets of predictive coding hold using cross-frequency coupling.

The results of this work revealed signature patterns of directional information processing among cross-regional CFCs in macaque auditory cortex. These patterns were highly preserved across coupling types and stimulus types.

1.3 Learning to predict

Another important question in predictive coding is how neural circuits maintain and update predictions as learning develops (Ch. 1.1; [Keller and Mrsic-Flogel, 2018]). One key area where predictions play a role is in allowing agents to pick actions based on their values, which is known as reinforcement learning (RL). Fundamentally, RL is a form of predictive learning. It drew inspiration from *Pavlovian conditioning* and *instrumental conditioning* in the psychological literature [Sutton and Barto, 2018]. In instrumental conditioning, and in RL, an agent seeks to select those actions that will increase the probability of reward, or viewed over the long-term, maximize some measure of overall reward or value. This implies building accurate predictions of the rewards or values associated with certain states or actions. There are different ways to implement RL as a learning algorithm [Sutton and Barto, 2018], but not all are equally likely to be implemented in the brain. A special form of RL known as temporal-difference (TD) learning has been the most successful so far in terms of accounting for neuroscientific findings [Niv, 2009, Averbeck and Costa, 2017]. In TD-learning, reward prediction occurs with the first cue that predicts future reward (Eq. 1.1; r represents the reward at a particular time step, γ the discount factor, and V the value of a particular state); if a predicted reward is omitted or a reward is given that was not predicted, a reward prediction error is issued (again mapping onto the prediction error terminology in predictive coding). This prediction error is used to update value estimates (Eq. 1.2; α represents the learning rate).

$$\delta_t = r_t + \gamma V_{old}(S_{t+1}) - V_{old}(S_t) \tag{1.1}$$

$$V_{new}(S_t) = V_{old}(S_t) + \alpha\left(r_t + \gamma V_{old}(S_{t+1}) - V_{old}(S_t)\right) \tag{1.2}$$

Previous experimental and modeling work has mapped aspects of RL onto

the fronto-striatal systems in the brain [Niv, 2009, Lee et al., 2012, Botvinick, 2012, Botvinick and Weinstein, 2014, Averbeck and Costa, 2017, Wang et al., 2018, Niv, 2019]. More specifically, the striatum has been found to encode action values [Houk, 1995, Suri and Schultz, 1998, Doya, 1910, Nakahara et al., 2001, O'Doherty et al., 2004, Frank et al., 2004, Frank, 2005, Samejima et al., 2005, Pasupathy and Miller, 2005, Histed et al., 2009, Amemori et al., 2011, Daw et al., 2011, Sarvestani et al., 2011, Li and Daw, 2011, Seo et al., 2012, Averbeck and Costa, 2017] and dopaminergic neurons that project to the striatum have been shown to represent prediction errors [Schultz et al., 1997, Niv, 2009]. Meanwhile, areas within the prefrontal cortex (PFC) have been implicated in action selection and decision making, as well as planning and prediction of future outcomes [Wood and Grafman, 2003, Friston, 2005, Averbeck et al., 2006, Cheong et al., 2006, Koechlin and Summerfield, 2007, Tsujimoto et al., 2008, Sakai, 2008, Histed et al., 2009, Collins and Koechlin, 2012, Botvinick, 2012, Stokes et al., 2013, Lim and Goldman, 2013, Verschure et al., 2014, Botvinick and Weinstein, 2014, Domenech and Koechlin, 2015, Rich and Wallis, 2016, Balaguer et al., 2016, Alexander et al., 2018, Aitchison and Lengyel, 2017b, Wallis et al., 2019, Radulescu et al., 2019]. Value signals were found to be stronger in the striatum than in the PFC, while choice signals were stronger in the PFC [Pasupathy and Miller, 2005, Samejima et al., 2005, Averbeck et al., 2006, Seo et al., 2012].

By treating neural circuits as a dynamical system and modeling their behavior during task execution one can gain insight into the underlying computational mechanisms [Rabinovich et al., 2008, Buonomano and Maass, 2009, Sussillo and Abbott, 2009, Sutskever, 2013, Shenoy et al., 2013, Mante et al., 2013, Sussillo and Barak, 2013, Barak et al., 2013, Hennequin et al., 2014, Carnevale et al., 2015, Rajan et al., 2016, Gallego et al., 2017, Wang et al., 2017, Chaisangmongkon et al., 2017, Remington et al., 2018, Wang et al., 2018, Yang et al., 2018, Gallego et al.,

2018, Kurikawa et al., 2018, Botvinick et al., 2019, Musall et al., 2019, Richards et al., 2019, Haesemeyer et al., 2019, Maheswaranathan et al., 2019, Hasson et al., 2020]. Recurrent neural networks (RNNs) trained on the same behavioral task as experimental animals can serve as a neural circuit model; once trained, these networks can be probed for any structure that has emerged [Sussillo and Abbott, 2009, Sussillo and Barak, 2013, Musall et al., 2019]. While artificial RNNs do not display a high degree of verisimilitude to real biological neural networks (i.e. in using backpropagation in adjusting weights, and adjusting connection weights only once during training instead of continously), they do include recurrent connections (just as real brains) and are able to recapitulate a lot of the features observed in real neural data. In contrast to simple latent variable models (which abstract away from real behavioral complexity), RNNs allow for studying tasks with rich, complex behavior. By studying increasingly complex behavior one can glean further insight into the representational structure present in real neural networks [Musall et al., 2019]. If the population dynamics in the artificial network can account for a large percentage of the variance observed in the neural recordings, the cost function used to train the artificial network is likely related to how the animal solves the task. One can then *reverse-engineer* the artificial network, determining how the dynamics account for task performance. This approach has revealed how tasks are represented in neural population dynamics in prefrontal cortex [Mante et al., 2013, Chaisangmongkon et al., 2017, Wang et al., 2018]. Going beyond single-unit dynamics, this approach has shown how task execution is regulated by fixed and slow point regions within population dynamics [Mante et al., 2013, Chaisangmongkon et al., 2017]. So far, these methods have only been used to study dynamics once the values of choices have already been learned.

Here (Ch. 3), this method was adapted to study how network representations develop and evolve as an animal learns to make choices based on rewards. A

recurrent neural network (RNN) model of the corticostriatal system was developed and trained on a sequence-learning task. This model consisted of a 'striatal' network generating action values and a 'prefrontal' network choosing actions based on those values. Recordings were obtained from the dorsal striatum (dSTR) and the lateral prefrontal cortex (lPFC) of two monkeys trained on the same learning task. Model and recordings were probed for the dynamics underlying task learning.

Chapter 2

Signature patterns for top-down and bottom-up information processing via cross-frequency coupling in macaque auditory cortex

The material in this chapter was presented at the Society for Neuroscience (SFN) conference 2016 and was published on bioarxiv and eNeuro: [Márton et al., 2019]

Abstract

Predictive coding is a theoretical framework that provides a functional interpretation of top-down and bottom up interactions in sensory processing. The theory suggests there are differences in message passing up vs. down the cortical hierarchy. These differences result from the linear feedforward of prediction errors, and the nonlinear feedback of predictions. This implies that cross-frequency interactions should predominate top-down. But it remains unknown whether these differences are expressed in cross-frequency interactions in the brain. Here we examined bi-directional cross-frequency coupling across four sectors of the auditory hierarchy in the macaque. We computed two measures of cross-frequency coupling, phase-amplitude coupling (PAC) and amplitude-amplitude coupling (AAC). Our findings revealed distinct patterns for bottom-up and top-down information processing among *cross*-frequency interactions. Both top-down and bottom-up made prominent use of low frequencies: low-to-low frequency (θ, α, β) and low frequency-to-*high* γ couplings were predominant *top-down*, while low frequency-to-*low* γ couplings were predominant *bottom-up*. These patterns were largely preserved across coupling types (PAC and AAC) and across stimulus types (natural and synthetic auditory stimuli), suggesting they are a general feature of information processing in auditory cortex. Our findings suggest the modulatory effect of low frequencies on γ-rhythms in distant regions is important for bi-directional information transfer. The finding of low frequency-to-*low* γ interactions in the bottom-up direction suggest non-linearities may also play a role in feedforward message passing. Altogether, the patterns of cross-frequency

interaction we observed across the auditory hierarchy are largely consistent with the predictive coding framework.

2.1 Significance

The brain consists of highly interconnected cortical areas, yet the patterns in directional cortical communication are not fully understood, in particular with regards to interactions between different signal components across frequencies. We employed a novel, convenient and computationally advantageous Granger-causal framework to examine bi-directional cross-frequency interactions across four sectors of the auditory cortical hierarchy in macaques. Our findings reveal cross-frequency interactions are predominant in the top-down direction, with important implications for theories of information processing in the brain such as predictive coding. Our findings thereby extend the view of cross-frequency interactions in auditory cortex, suggesting they also play a prominent role in top-down processing.

2.2 Introduction

One fundamental yet poorly understood component of cortical computation is the presence of widespread reciprocal cortical connections [Scott et al., 2015, Romanski and Averbeck, 2009, Kveraga et al., 2007, Salin and Bullier, 1995]. Moreover, most previous work has focused on the visual pathways. It remains unknown whether there are fundamental differences between top-down and bottom-up information processing along the auditory pathway. Cross-regional top-down effects in particular have remained unexplored, as most work has focused on bottom-up [Hyafil et al., 2015a, Giraud and Poeppel, 2012] or intra-regional effects [Lakatos et al., 2016]. In the present study, we examined bi-directional information processing across the macaque auditory hierarchy through an analysis of cross-frequency coupling (CFC). We computed two types of CFC measures, amplitude-amplitude (AAC) and phase-amplitude (PAC) coupling, for both natural and synthetic stimuli.

Cross-frequency coupling refers to correlation between the phase or amplitude component of one frequency and the amplitude component of other frequencies. This differs from linear analyses like coherence, which only examine coupling between signals at a single frequency. CFC, in contrast, considers interactions across the frequency spectrum. CFC may occur in several forms [Jirsa and Mueller, 2013] and has been observed during various tasks across several species including humans, macaques, and rodents. PAC, in particular, has received widespread attention. Coupling of low-frequency (e.g. θ, α) phase to high-frequency (e.g. γ) amplitude has been associated with various behavioral mechanisms including learning, memory and attention [Hyafil et al., 2015b, Colgin, 2015, Doesburg et al., 2012, Kendrick et al., 2011, Canolty and Knight, 2010, Tort et al., 2009, Tort et al., 2008, Cohen, 2008, Jensen and Colgin, 2007, Canolty et al., 2006, Schack et al., 2002]. The reason for interest in PAC is that it may offer a mechanistic account of information coordination across neural populations and timescales. The phase of low-frequency oscillations could modulate the amplitude of high-frequency oscillations, putatively allowing control of spiking activity [Buzsáki et al., 2012, Ray and Maunsell, 2011]. AAC has also been shown to correlate with behavior [Canolty and Knight, 2010], but has been less explored mechanistically. Here we obtained estimates of both PAC and AAC strength as measures of cross-regional information processing in the bottom-up and top-down direction, allowing the two measures to be contrasted and compared in terms of informational content.

One prevalent theory of information processing in the brain which separates bottom-up and top-down components is *predictive coding*. In this (Bayesian) view of the brain, expectations are formed about incoming sensory information. These top-down predictions (expectations) are compared with bottom-up sensory information (outcomes) and representations are then potentially updated based on a

prediction error (surprise), which is the difference between sensory inputs and their expectations. Neural implementations of predictive coding suggest that predictions are formed in higher-order areas through a (linear) accumulation of prediction errors. The accumulation of prediction errors in higher-order areas results in the attenuation of high frequencies, thus feedforward connections can be cast as a low-pass or Bayesian filter. Conversely, prediction errors arising in lower-order areas are a non-linear function of predictions. This means high frequencies can be created in lower-order areas due to the inherent non-linearity. Hence, predictive coding implies that lower-order areas (where prediction errors arise) should display relatively higher frequencies (e.g. γ-range) than higher-order areas (where predictions are formed) [Bastos et al., 2012b, Friston, 2008]. Several studies have supported this view [Michalareas et al., 2016, Bastos et al., 2015, Bastos et al., 2012b, Fontolan et al., 2014], though not without exceptions.

It remains unknown whether this prediction is expressed in cross-frequency interactions in the brain. Here we probe this prediction by examining differences in cross-frequency coupling strength in the bottom-up and top-down direction across the frequency spectrum. Based on the predictive coding framework [Bastos et al., 2012b, Friston, 2008], we would expect there to be an asymmetry between bottom-up and top-down interactions. This asymmetry arises due to the linear effect of prediction errors on predictions in the bottom-up direction, and the nonlinear effect of predictions on prediction errors in the top-down direction. Thus, we would expect nonlinear, *cross*-frequency interactions to predominate in the top-down direction.

Previous work has shown differences in how auditory cortical areas process natural and synthetic sounds [Fukushima et al., 2014]. It is possible that these differences translate into systematic differences in CFC patterns. If the revealed CFC pattern is a more general hallmark of inter-areal communication though, it

should not be specific to stimulus type: natural and synthetic stimuli should show similar overall coupling patterns.

Our work revealed signature patterns of top-down and bottom-up processing in cross-regional PACs and AACs. These patterns were conserved across coupling types and across natural and synthetic stimuli. We found that low-to-low frequency (θ, α, β) and low frequency-to-*high* γ interactions were predominant *top-down*, while low frequency-to-*low* γ couplings were predominant *bottom-up*. Our findings suggest modulation of γ frequency power by low frequency rhythms plays a role in bi-directional information transfer across the auditory hierarchy. The finding of predominant cross-frequency interactions in the top-down direction across the auditory hierarchy largely agrees with the predictive coding framework.

2.3 Methods and Materials

2.3.1 Subjects

Three adult male rhesus monkeys (Macaca mulatta) weighing 5.5–10 kg were used for recordings. All procedures and animal care were conducted in accordance with the Institute of Laboratory Animal Resources Guide for the Care and Use of Laboratory Animals. All experimental procedures were approved by the National Institute of Mental Health Animal Care and Use Committee.

2.3.2 Stimuli

The stimuli used for the main experiment included 20 conspecific monkey vocalizations (VOC) and two sets of 20 synthetic stimuli each (envelope-preserved sound, EPS; and spectrum preserved sound, SPS) derived from the original VOC stimuli. The VOC stimulus set consisted of 20 macaque vocalizations employed in previous studies [Kikuchi et al., 2010, Fukushima et al., 2014].

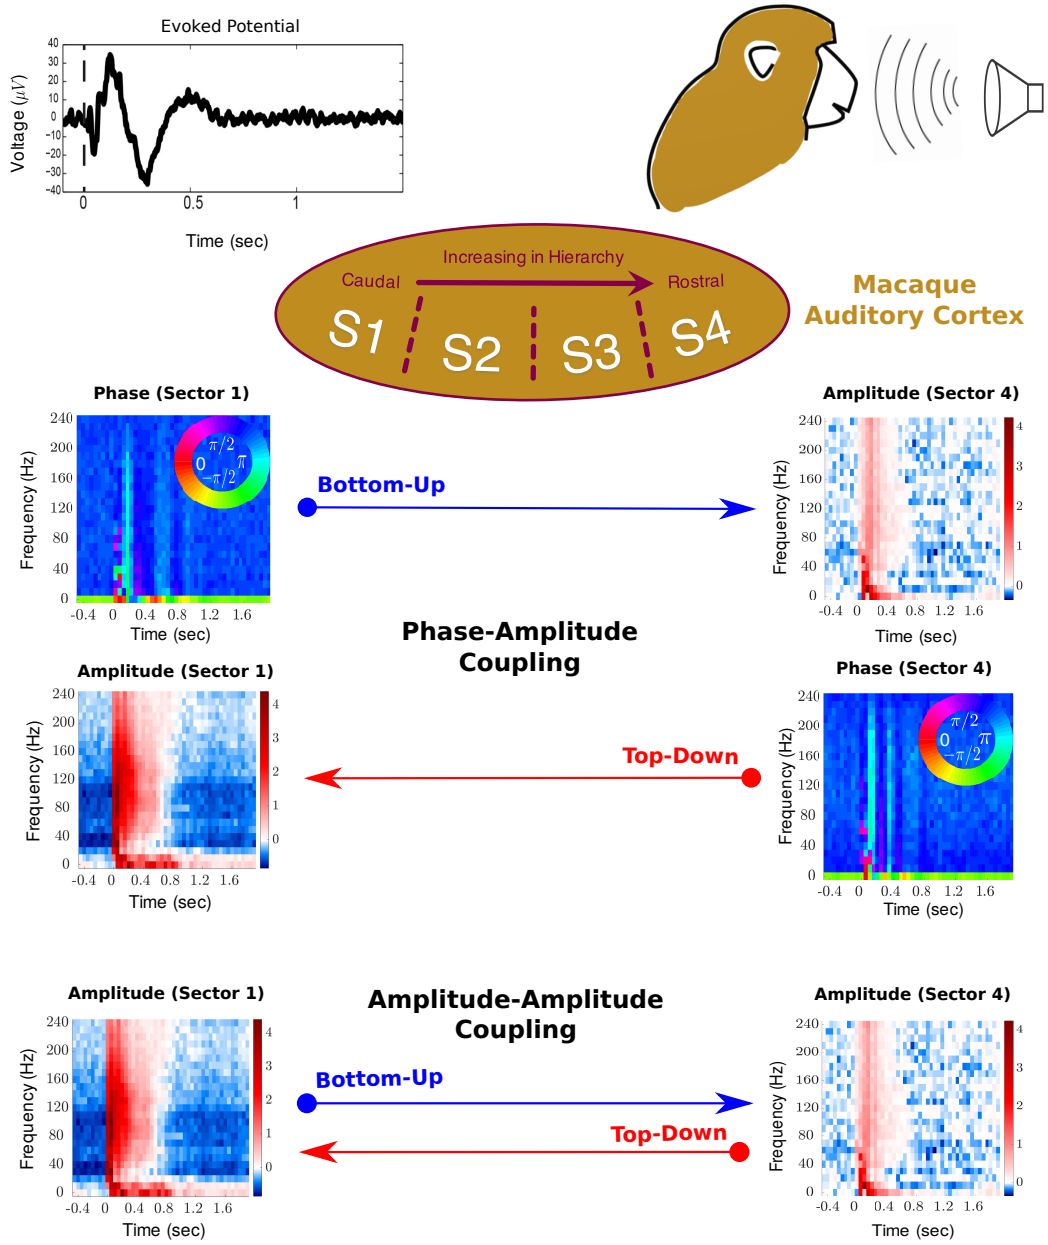

Evoked Potential
Voltage (μV)
Time (sec)
Increasing in Hierarchy
Caudal
Rostral
S1
S2
S3
S4
Macaque
Auditory Cortex
Phase (Sector 1)
Amplitude (Sector 4)
Bottom-Up
Frequency (Hz)
Time (sec)
Phase-Amplitude
Coupling
Amplitude (Sector 1)
Phase (Sector 4)
Top-Down
Amplitude-Amplitude
Coupling
Amplitude (Sector 1)
Amplitude (Sector 4)
Bottom-Up
Top-Down

Fig 2.1. Coupling Analysis Overview: Recordings were made from four μECoG arrays spanning auditory cortex while monkeys listened to one of 20 natural vocalizations (VOC), 20 synthetic envelope-preserved sounds (EPSs) or 20 synthetic spectrum-preserved sounds (SPSs). Electrodes were partitioned into four sectors along the caudorostral axis (S1: A1/ML (primary auditory cortex/middle lateral belt), S2: R (rostral core region of the auditory cortex / AL (anterior lateral belt region of the auditory cortex), S3: RTL (lateral rostrotemporal belt region of the auditory cortex), S4: RTp (rostrotemporal pole area)). We decomposed the signal into its time-frequency representation, and obtained amplitude and phase components for each sector. We investigated two types of coupling across sectors: amplitude-amplitude and phase-amplitude coupling. We computed both types of coupling in the bottom-up and top-down direction. Top-down coupling was defined as coupling in which the source electrode came from a sector of higher order than the target electrode, and vice versa for bottom-up coupling.

In order to obtain EPS from VOC stimuli, the envelope of a particular vocalization was estimated based on the amplitude component of the original stimulus' Hilbert transform. The amplitude envelope was then multiplied by broadband white noise to create the EPS stimulus. Thus, all 20 EPS stimuli exhibited flat spectral content; these stimuli could not be discriminated based on spectral features, while the temporal envelopes (and thus the durations) of the original vocalizations were preserved.

SPS stimuli were obtained by first generating broadband white noise with a duration of 500 ms and computing its Fourier transform. The SPS stimulus' amplitude in the Fourier domain was then replaced by the average amplitude of the corresponding VOC stimulus, before transforming back to the time domain by computing the inverse-Fourier transform. This resulted in a sound waveform that preserved the average spectrum of the original vocalization, while exhibiting a flat temporal envelope, random phase, with a duration of 500 ms. A 2 ms cosine rise/fall was then imposed on the stimulus to avoid abrupt onset/offset-effects. Hence, all 20 SPS stimuli exhibited nearly identical, flat temporal envelopes; these stimuli could not be discriminated using temporal features, while the average spectral power of

the original vocalizations was preserved.

A total of 60 different stimuli were presented in pseudorandom order, with an interstimulus interval of 3 s. Each stimulus was presented 60 times. The sound pressure levels of the stimuli measured by a sound level meter ranged from 65 to 72 dB at a location close to the animal's ear. Stimulus duration ranged from around .15 to 1 second [Fukushima et al., 2014].

2.3.3 Recordings

Custom-designed μECoG arrays (NeuroNexus Technologies, MI, USA) were used to record field potentials from macaque auditory cortex. Arrays were machine fabricated on a very thin polyimide film (20 μm), with each array featuring 32 recording sites, 50 μm in diameter each, on a 4 x 8 grid with 1 mm spacing (i.e., 3 x 7 mm rectangular grid). Two animals, monkeys B and K, were implanted with arrays each in the left hemisphere, while one animal, monkey M, was implanted with arrays in the right hemisphere. Three of the arrays in each monkey were placed on top of STP, in a caudorostrally oriented row.

To implant the arrays, we removed a frontotemporal bone flap extending from the orbit ventrally toward the temporal pole and caudally behind the auditory meatus and then opened the dura to expose the lateral sulcus. The most caudal of the three ECoG arrays on the STP was placed first and aimed at area A1 by positioning it just caudal to an (imaginary) extension of the central sulcus and in close proximity to a small bump on the STP, both being markers of A1's approximate location. Each successively more rostral array was then placed immediately adjacent to the previous array to minimize interarray gaps. The arrays on the lateral surface of the STG were placed last. The probe connector attached to each array was temporarily attached with cyanoacrylate glue or Vetbond to the skull immediately above the cranial opening. Ceramic screws together with bone cement were used to fix the

connectors to the skull. The skin was closed in anatomical layers. Post- surgical analgesics were provided as necessary, in consultation with the National Institute of Mental Health veterinarian.

2.3.4 Stimulus presentation and recording parameters

The monkey was placed in a sound-attenuating booth for the experiment (Biocoustics Instruments). The sound stimuli were presented while the monkey sat in a primate chair and listened passively with its head fixed. Auditory evoked potentials from the 128 channels of the ECoG-array were bandpassed between 2-500 Hz, digitally sampled at 1500 Hz, and stored on hard-disk drives by a PZ2–128 preamplifier and the RZ2 base station (Tucker Davis Technology).

2.3.5 Data Pre-Processing

Data analysis was performed using MATLAB R2017a (MathWorks) software. Since there was little significant auditory evoked power above 250 Hz, recordings were low-pass filtered and resampled at 500 Hz to enhance calculation speed and reduce memory requirements. The signal was bandpassed at 60Hz with a narrow-band filter to eliminate the possible presence of line-noise; a 6th-order Butterworth filter was chosen to achieve a sharp dropoff in the stop-band, thus minimizing the effect of the filter on neighboring frequency bands.

The 96 sites on the supratemporal plane (STP) were grouped based on the characteristic frequency maps obtained from the high-gamma power of the evoked response to a set of pure-tone stimuli. The change in frequency tuning along the STP reverses across areal boundaries, and these reversals were therefore used to identify the areal boundaries. This resulted in a grouping of electrodes into four sectors (Fig 2.1), which were estimated to correspond to the following subdivisions at a caudorostral-level within macaque auditory cortex: Sec (Sector) 1, A1/ML (primary

auditory cortex/ middle lateral belt); Sec 2, R (rostral core region of the auditory cortex)/AL (anterior lateral belt region of the auditory cortex); Sec 3, RTL (lateral rostrotemporal belt region of the auditory cortex); Sec 4, RTp (rostro-temporal pole area). The recorded signal from each site partitioned into sectors 1-4 was re-referenced by subtracting the average of all sites within a particular sector [Fontolan et al., 2014, Kellis et al., 2010].

We defined the following terminology to refer to particular frequency spaces throughout this paper: *delta* (δ) range: 0 - 2.5 Hz; *theta* (θ) range: 2.5 - 7.5 Hz; *alpha* (α) range: 7.5 - 12.5 Hz; *low beta* ($low\,\beta$) range: 12.5 - 22.5 Hz; *high beta* ($high\,\beta$) range: 22.5 - 42.5 Hz; *low gamma* ($low\,\gamma$) range: 42.5 - 67.5 Hz; *high gamma* ($high\,\gamma$) range: $>$ 67.5 Hz. The delta frequency band (0 - 2.5 Hz) was not included in the analysis since the pre-amplifier employed during the recording sessions does not allow for recording frequency components below 2 Hz and the power in this band was expectedly low.

2.3.6 Phase-amplitude (PAC) and amplitude-amplitude (AAC) cross-frequency coupling (CFC) Calculation

2.3.6.1 Frequency Transformation

For Figures 1-5, the field potential data was first transformed into frequency space by computing the fast Fourier transform (FFT) of the signal in every channel using a 200 ms Hanning window, stepped by 200 ms. Thus, we obtained a frequency representation of the data at a 5Hz bandwidth ranging from 2.5Hz - 132.5Hz, which satisfies the Nyquist-criterium. We obtained the phase and amplitude information in a given source and target region for every channel and normalized the amplitude information in each frequency band by the corresponding power in that frequency band, averaged across the entire experiment, so as to account for decreasing power

with frequency (Fig. S1-1).

2.3.6.2 Obtaining CFCs using canonical correlation

We computed two types of cross-frequency coupling, phase-amplitude (PAC) and amplitude-amplitude (AAC) coupling. Couplings were calculated for every channel pair across the four sectors (S1-S4), resulting in 6 possible cross-sector pairs (S1-S2, S1-S3, S1-S4, S2-S3, S2-S4, S3-S4). We obtained estimates of coupling strength in two directions, bottom-up and top-down. In the former case, the source component was derived from a recording site in a sector lower in the auditory hierarchy than the target component (referred to as bottom-up coupling), while in the latter case the source component came from a sector higher in the auditory hierarchy than the target component (referred to as top-down coupling) (Fig. 2.1). Ultimately, we were interested in examining the difference in coupling strength in the top-down and bottom-up directions across the auditory hierarchy. Results are presented collapsed across all cross-sector pairs (Fig. 2.2), as well as separately for every cross-sector pair (Fig. 2.3), averaged across all channel pairs and across the three animals in both cases.

We used canonical correlation analysis (CCA), to compute both types of coupling. The advantage of CCA is that it computes coupling among multivariate observations in the source and target regions. This framework, therefore, allows CFC values to be computed across the entire frequency range simultaneously. Furthermore, some approaches to calculating AAC and PAC first carry out dimensionality reduction on the source and target region signals separately, and then calculate correlations in the low dimensional space. If the variance preserving dimensions in the source and target space are not the dimensions that maximize correlations between the source and target space, interactions may be missed. CCA on the other

hand finds the low dimensional representations that maximize coupling between the source (X) and target (Y) space. In CCA, one is seeking projections (linear combinations) $U = XA$ and $V = YB$ such that the correlation $corr(U,V)$ among these derived directions is maximal [Durstewitz, 2017].

We carried out our analyses using a Granger Causal framework. Therefore, we first carried out CCA between amplitude in the target region and time lagged versions of the amplitude also within the target region, obtaining an estimate of the the coupling strength $\widehat{Y}$ that could be inferred purely from coupling of the signal at the target electrode onto itself. For computational reasons, we used two time lags to estimate the target signal (we found that using one time lag did not affect our results). More explicitly, we transformed the current and lagged signal into the frequency domain as described above and computed the canonical correlation between the lagged components and the current signal in the target. We used the first canonical value only to obtain $\widehat{Y}$ (results did not differ from those obtained from including more entries). We then obtained the residual $\widetilde{Y}$ from this analysis, and used this residual when we examined the interactions between regions. Subsequent analysis was performed using the residuals $\widetilde{Y}$ from this regression as the signal for the target region.

We then computed phase-amplitude (PAC) and amplitude-amplitude (AAC) cross-frequency couplings across the four sectors (Fig. 2.1), by regressing the phase or amplitude component of the signal from the source sector, respectively, on the amplitude component of the signal in the target sector for every channel. More specifically, we calculated instantaneous coupling in 200ms windows between source and target region for each electrode pair. PAC and AAC was calculated in the CCA framework across all time windows and trials altogether, simultaneously across all frequency bands. In the case of PAC, X_{nxd1} contained the sine and cosine transformed phase component of the signal, while Y_{nxd2} contained the squared and

log-transformed amplitude component of the signal, with $n = 10800$ (1200 trials * 9 windows) observations for VOC and SPS stimuli and $n = 10773$ observations (1197 trials * 9 windows) for EPS stimuli, $d2 = 25$, and $d1 = 25$ or $d1 = 25 * 2$ for PAC and AAC, respectively (Eq 2.1).

After obtaining the canonical coefficient matrices B_{d2xd} and A_{d1xd} within the CCA framework (Eq 2.1), one can obtain the coupling matrix P_{d2xd1} by multiplying the B and A matrices with the canonical values S of the correlation matrix (Eq 2.2). We used the first 10 directions ($d = 10$) that captured most of the interaction; results did not differ from those obtained from including all entries. For AAC, the final cross-frequency coupling matrix is P. For PAC, one combines the sine and cosine components of P by computing their Euclidian norm, yielding the final coupling matrix P^*_{d2xd1} with $d2 = d1$.

$$U_{nxd} = X_{nxd1} A_{d1xd}, \quad V_{nxd} = \widetilde{Y}_{nxd2} B_{d2xd}, \quad \underset{A,B}{\operatorname{argmax}} \{corr(U,V)\} \tag{2.1}$$

$$P_{d2xd1} = BSA^T, \quad P^*_{d2xd1} = \begin{pmatrix} \sqrt{P_1^2 + P_{m/2}^2} \\ \dots \\ \sqrt{P_{m/2}^2 + P_m^2} \end{pmatrix} \tag{2.2}$$

2.3.6.3 Obtaining the difference between top-down (TD) and bottom-up (BU) coupling

We obtained AAC and PAC coupling matrices in this manner for both the bottom-up and the top-down direction. In bottom-up coupling, the source component was derived from a recording site in a sector lower in hierarchy than the target component (referred to as bottom-up coupling), while in top-down coupling, the source component came from a sector higher in hierarchy than the target component (Fig. 2.1). We then obtained the difference $\Delta P_{same\,stim}$ between the top-down and bottom-up coupling matrices $P_{Top-Down}$ and $P_{Bottom-Up}$ (Eq 2.3). This was done separately for each of the cross-sector pairs. Results are presented collapsed across all cross-sector pairs (Fig. 2.2), as well as separately for every cross-sector pair (Fig. 3.3).

$$\Delta P = P_{Top-Down} - P_{Bottom-Up} \tag{2.3}$$

2.3.6.4 Obtaining the difference in coupling strength between original (VOC) and synthetic stimuli (EPS/SPS)

In order to examine the difference in coupling strength between natural and synthetic stimulus types (Fig. 2.4,2.5), we computed the difference between VOC and EPS stimuli (Fig. 2.4) and between VOC and SPS stimuli (Fig. 2.5) separately in the top-down (TD, Eq 2.4) and bottom-up direction (BU, Eq 2.5).

$$\Delta P^{TD}_{VOC\,vs\,Synthetic} = P^{TD}_{VOC} - P^{TD}_{EPS/SPS} \tag{2.4}$$

$$\Delta P^{BU}_{VOC\,vs\,Synthetic} = P^{BU}_{VOC} - P^{BU}_{EPS/SPS} \tag{2.5}$$

2.3.7 Statistical Analysis

Paired t-tests were performed to establish statistical significance of the difference between top-down and bottom-up coupling (ΔP; Fig. 2.2, Fig. 2.3, Fig. 2.4), as well as the differences between original and synthetic stimuli ($\Delta P^{TD}_{VOC\,vs\,Synthetic}$, $\Delta P^{BU}_{VOC\,vs\,Synthetic}$; Fig. 2.5). This was done separately for each frequency pair, by testing whether the mean of the difference between top-down and bottom-up coupling (or the difference between original and synthetic stimuli in each direction) was significantly different from zero (null hypothesis). Lilliefors' test for normality was conducted for each frequency pair in order to ensure the normality assumption was satisfied, and histogram plots were inspected for select frequency pairs (Fig. S1-2, S1-3).

More specifically, in order to assess overall differences in coupling strength between the two directions (Fig. 3.2, Fig. 2.4) and to assess differences in coupling strength between natural and synthetic stimuli (Fig. 2.5), paired t-tests were conducted across all channel pairs, cross-regional pairs and animals altogether ($d.f. = 9344$). In order to assess coupling differences for each regional pair separately (Fig. 2.3), paired t-tests were conducted separately for each regional pair, across all channel pairs and animals altogether ($d.f. = 1056$).

Significant phase-amplitude and amplitude-amplitude coupling (highlighted by black outlines throughout the manuscript) was based on corrected p-values using non-parametric permutation tests that generated null distributions of the maximum cluster size [Maris and Oostenveld, 2007, Nichols and Holmes, 2001]. This procedure implicitly adjusts for searching over multiple frequencies. The null distribution was obtained by randomly reassigning electrodes to source or target region and recalculating differences in coupling strength for 100 repeats. Clusters of the 99th

percentile (corresponding to $p \leq 0.01$) were considered significant and are reported throughout the manuscript (marked by black outlines).

2.4 Results

We analyzed local field potentials recorded simultaneously from multiple cortical areas in the auditory cortex of three macaques while the animals listened to auditory stimuli (Fig. 2.1). These auditory stimuli consisted of 20 natural conspecific vocalizations, 20 envelope-preserved stimuli (EPS) and 20 spectrum-preserved stimuli (SPS), which were derived from the original vocalizations [Fukushima et al., 2014]. A previous study established the approximate rostrocaudal location of implanted micro-electrocorticography (μECoG) arrays along the auditory hierarchy by the characteristic frequency of each electrode contact [Fukushima et al., 2012, Fukushima et al., 2014]. This partitioned the recording sites into four sectors (S1-S4), putatively spanning the caudorostral levels from the core A1 (S1) to RTp (S4), including some of the surrounding belt areas (Fig. 2.1). We operationally defined feedforward to occur from earlier to later sectors (bottom-up direction), and feedback to occur from later to earlier sectors (top-down direction).

We wanted to examine cross frequency interactions across the auditory hierarchy. To characterise phase-amplitude and amplitude-amplitude coupling, we first decomposed the signal into its spectral component using the short-time fast Fourier transform and obtained phase and amplitude components (Fig. 2.1). We then used canonical correlation analysis (CCA) to parameterise the cross-frequency coupling under a general linear model (Methods and Materials). The fluctuations in amplitude in the target sector were predicted based on instantaneous amplitude fluctuations in the source sector within 200ms windows. Furthermore, the target and source time series were adjusted to remove the effect of predicted fluctuations from the target

region. In other words, by appealing to the notion of Granger causality, a significant mapping between the source and the target - that cannot be explained by the history of the target - is evidence of (cross-frequency) coupling.

We computed coupling strength across the frequency spectrum for each cross-regional pair, separately in the top-down and bottom-up direction (Fig. 2.1). Top-down coupling was defined as coupling in which the source electrode came from a sector of higher order than the target electrode, and vice versa for bottom-up coupling. For example, coupling between the phase (or amplitude) in S4 and the amplitude in S1 is defined as top-down, and coupling between the phase (or amplitude) in S1 and the amplitude in S4 is defined as bottom-up in the case of PAC (or AAC). This definition is in accordance with a previous approach [Fontolan et al., 2014].

2.4.1 Distinct cross-frequency coupling signatures for bottom-up and top-down coupling, both involving low frequencies

We first examined PAC and AAC across the four sectors, in auditory evoked potentials to 20 conspecific vocalizations (VOC). Both CFC measures peaked in the θ frequency range and showed relatively strong coupling across all the target amplitude frequencies (Fig. 2.2A,B,D,E). The field potential signal showed the highest power in the same frequency band decaying with the typical $1/f$ pattern thereafter (Fig S1-1), thus satisfying a pre-requisite for the presence of physiologically meaningful CFC [Aru et al., 2015].

Among PACs, both top-down and bottom-up coupling was dominant between θ phase and broadband amplitude (Fig. 2.2A,B). Examining differences in coupling strength in the two directions more closely (Fig. 2.2C), PACs showed dominant

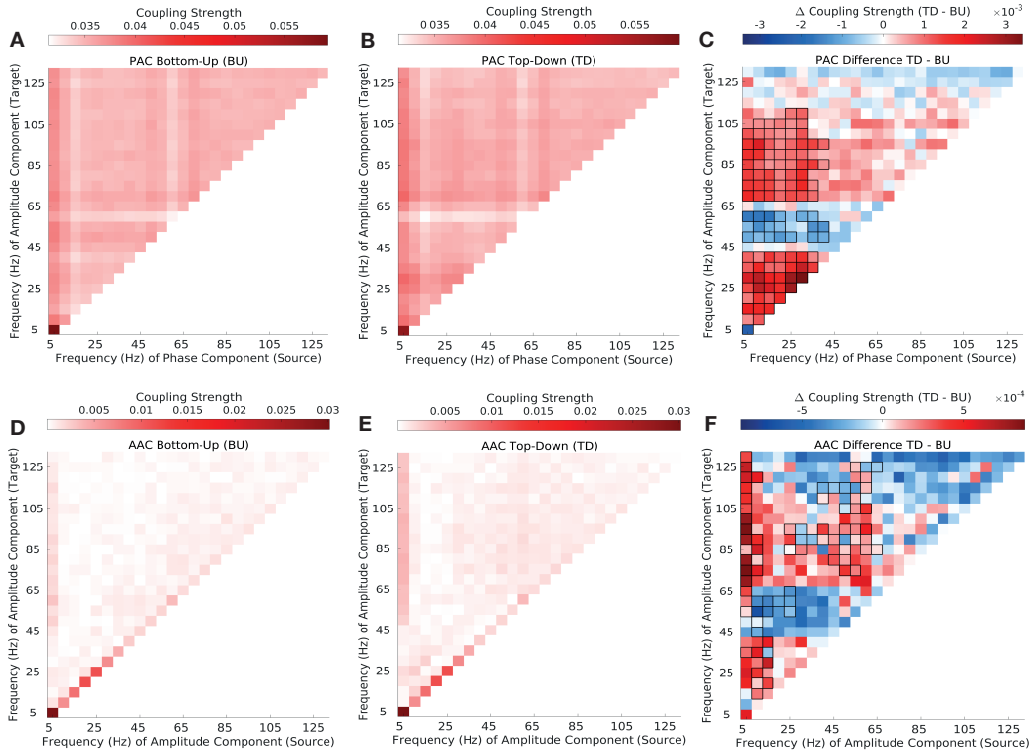


Fig 2.2. Top-down vs. bottom-up phase-amplitude and amplitude-amplitude coupling in natural vocalizations (VOC). (A-B) Phase-amplitude coupling (PAC) strength in the top-down (A) and bottom-up direction (B). Depicted are canonical correlation-derived coupling coefficients (see Methods and Materials). **(C)** Difference in top-down and bottom-up in PAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). **(D-E)** Amplitude-amplitude coupling (AAC) strength in the top-down (D) and bottom-up direction (E). Depicted are canonical correlation-derived coupling coefficients. **(F)** Difference between top-down and bottom-up in AAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals. Supporting Figures: Fig. S1-1, S1-2, S1-3

top-down coupling in the low-frequency space (θ, α, β phase to θ, α, β amplitude), as well as between low frequency phase and high-γ (>67.5Hz) amplitude ($p \leq 0.01$, cluster-corrected). Bottom-up PAC was dominant between low frequency phase (θ, α, β) and low-γ (47.5-62.5Hz) amplitude ($p \leq 0.01$, cluster-corrected).

AACs showed a similar pattern to PACs in both directions, but coupling strength and differences in coupling strength were overall by an order of magnitude lower than for PACs (Fig. 2.2D,E). AACs replicated the pattern observed in PACs, with dominant θ to broadband coupling in both directions. Additionally, and unlike in PACs, coupling strength among same-frequency couplings was pronounced in both directions with strongest coupling in the low-frequency range (Fig. 2.2D,E). Examining differences in coupling strength in the two directions more closely (Fig. 2.2F), AACs displayed a similar pattern to PACs but showed smaller significant clusters than PACs across the entire frequency spectrum.

2.4.2 Cross-regional coupling mainly driven by interactions between A1 and higher order sectors

To examine cross-regional patterns in coupling across the auditory hierarchy (Fig. 2.1) more closely, we analyzed differences between top-down and bottom-up coupling strength separately across the six cross-regional pairs (Fig. 2.3). This revealed that the overall coupling pattern (Fig. 2.2) was mostly driven by interactions between S1 (A1/ML) and S2-S4 (higher-order sectors) (Fig. 2.3 A,B,C), rather than by interactions among higher-order sectors (Fig. 2.3 D,E,F).

More specifically, interactions between S1 (A1/ML) and S2 (R/AL)/ S3 (RTL) / S4 (RTp) show dominant top-down coupling of low frequency phase to low and high frequency amplitude (Fig. 2.3A-C), whereas these asymmetries are largely

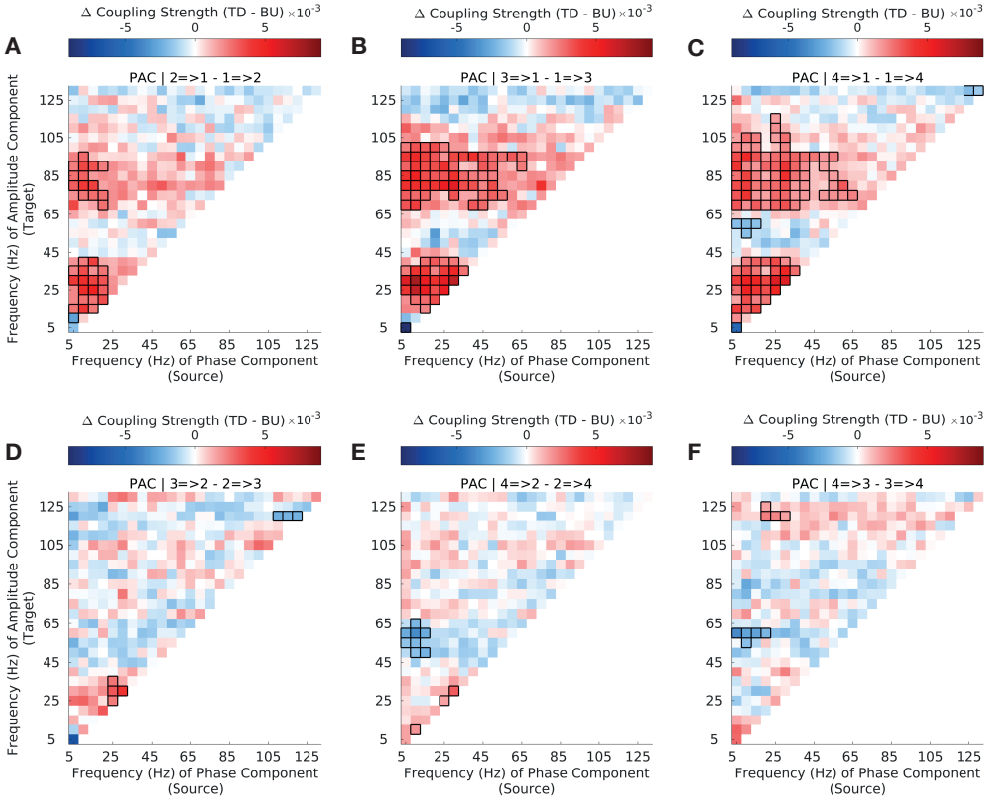


Fig 2.3. Top-down versus bottom-up phase-amplitude coupling strength (PAC) across the auditory hierarchy in natural stimuli Difference in top-down and bottom-up *phase-amplitude* coupling (PAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals. Supporting Figures: Fig. S1-4.

absent among interactions between S2-S4 (Fig. 2.3D-F). The interaction between S1 and S2 (Fig. 2.3B) shows overall less widespread top-down coupling strength compared to interactions between S1 and S3/S4. The full pattern (as manifested in Fig. 2.2) only gains full prominence among cross-regional pairs S1-S4 (Fig. 2.3C), which also includes predominant bottom-up coupling between low frequency phase and low-γ amplitude. Interactions between higher-order sectors (Fig. 2.3 D,E,F) show fewer differences between top-down and bottom-up coupling, but they do show predominant bottom-up coupling between low frequency phase and low-γ amplitude

(Fig. 2.3 E,F). A similar trend was observed among AACs, though differences were overall smaller and less widespread compared to PACs (Fig. S1-4).

2.4.3 Top-down and bottom-up cross-frequency signatures generalize to synthetic stimuli

To examine whether the observed coupling patterns are a general feature of inter-areal communication, we also analyzed coupling strength in two types of synthetic stimuli derived from the natural vocalizations. Envelope-preserved sounds (EPS) were obtained by leaving the original temporal envelope of the vocalization intact but flattening its spectral content, while spectrum-preserved sounds (SPS) were obtained by leaving the spectral information of the original vocalization intact but flattening its temporal envelope [Fukushima et al., 2014] (see Methods and Materials 3.2).

EPS and SPS stimuli showed similar bottom-up and top-down coupling patterns to natural stimuli, with pronounced θ-to-broadband coupling strength in both directions (Fig. S1-5,S1-8). Examining differences between top-down and bottom-up coupling strength in both types of synthetic stimuli (Fig. 2.4,S1-5, S1-8), we found that they reflected the same coupling asymmetries across the frequency spectrum as natural stimuli did (Fig. 2.2). This was true for PAC (Fig. 2.4A,C) and AAC (Fig. 2.4B,D) alike. More specifically, both types of synthetic stimuli showed dominant top-down coupling among low-to-low frequency and low-to-high γ interactions, and dominant bottom-up coupling among low-to-low γ interactions, like natural stimuli did. Both types of synthetic stimuli also showed similar coupling asymmetries to natural stimuli when examining cross-regional pairs separately (Fig. S1-6, S1-7, S1-9, S1-10), with most of the pattern driven by interactions between A1 and higher-order sectors rather than interactions among higher-order sectors, as in natural vocalizations.

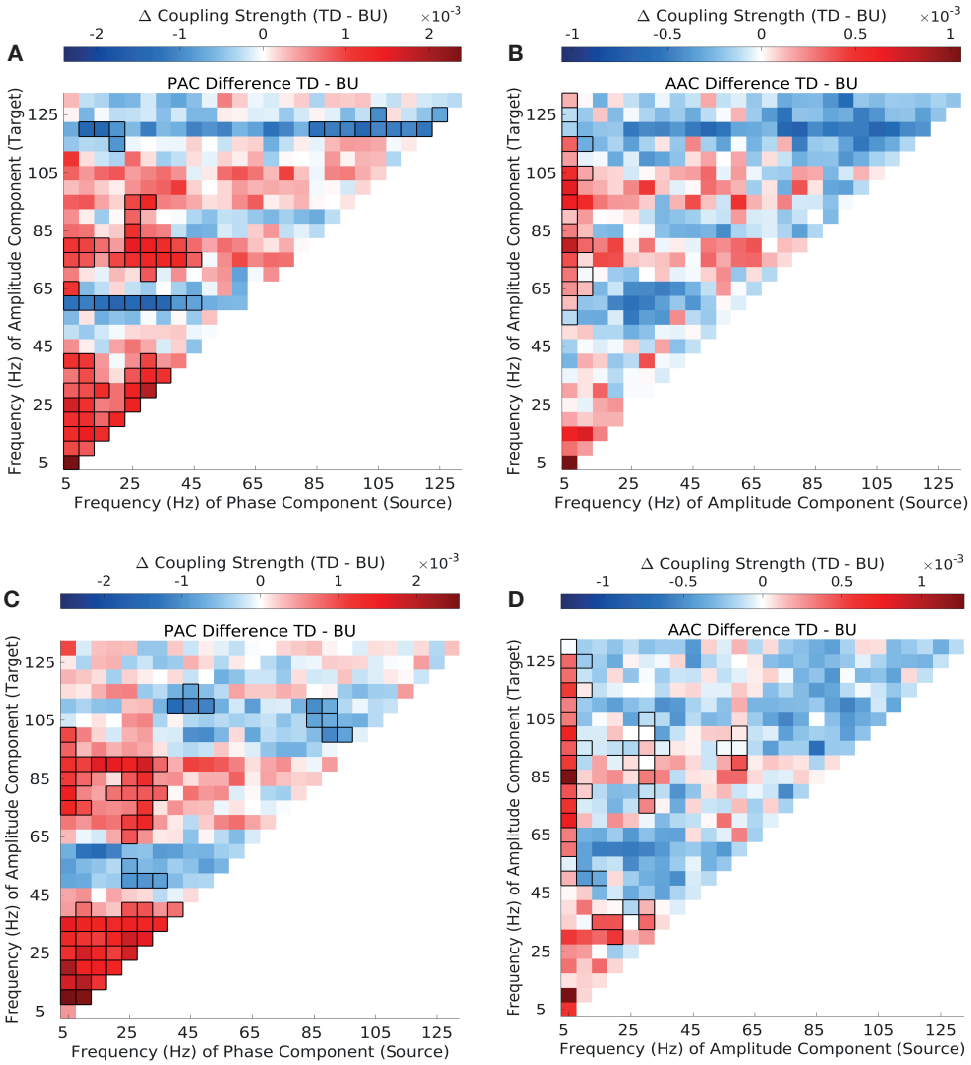


Fig 2.4. Top-down vs. bottom-up phase-amplitude and amplitude-amplitude coupling in synthetic envelope-preserved stimuli (EPS) and synthetic spectrum-preserved stimuli (SPS) (A) Difference in top-down and bottom-up PAC strength in EPS stimuli. **(B)** Difference in top-down and bottom-up AAC strength in EPS stimuli. **(C)** Difference in top-down and bottom-up PAC strength in SPS stimuli. **(D)** Difference in top-down and bottom-up AAC strength in SPS stimuli. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals. Supporting Figures: Fig. S1-5, S1-6, S1-7, S1-8, S1-9, S1-10

2.4.4 Synthetic stimuli show overall lower bi-directional coupling than natural stimuli

To examine differences in coupling strength between natural and synthetic stimuli more closely, we compared coupling strength across stimuli separately in the top-down and bottom-up direction (Fig. 2.5).

EPS stimuli showed overall lower coupling strength than natural stimuli in both directions for both PACs and AACs (Fig. 2.5). Among PACs more specifically (Fig. 2.5A-B), natural stimuli showed higher coupling strength in the top-down direction for select frequency pairs in the low frequency space as well as among low frequency to high-γ couplings (Fig. 2.5A). In the bottom-up direction, coupling strength for natural stimuli was enhanced among low frequency to low and high γ couplings (Fig. 2.5B). In the bottom up direction, there were also select coupling values in the high γ frequency target amplitude space for which coupling was stronger in EPS stimuli.

Among AACs, coupling strength was increased for natural compared to EPS stimuli across the entire frequency spectrum in both directions (Fig. 2.5 E-F). In particular, θ-γ amplitude-amplitude interactions were significantly lower for EPS stimuli in both directions. Altogether, coupling strength is enhanced for natural stimuli over envelope-preserved but spectrally flat synthetic stimuli (EPS) in both directions and across the entire frequency spectrum.

SPS stimuli also showed overall lower coupling strength compared to natural stimuli. Among PACs, SPS stimuli showed lower coupling strength for select low-to-low and low-to-high frequency couplings in both directions (Fig. 2.5C,D). In the top-down direction, SPS stimuli showed higher coupling strength among γ-γ interactions (Fig. 2.5D). SPS and natural stimuli displayed similar differences

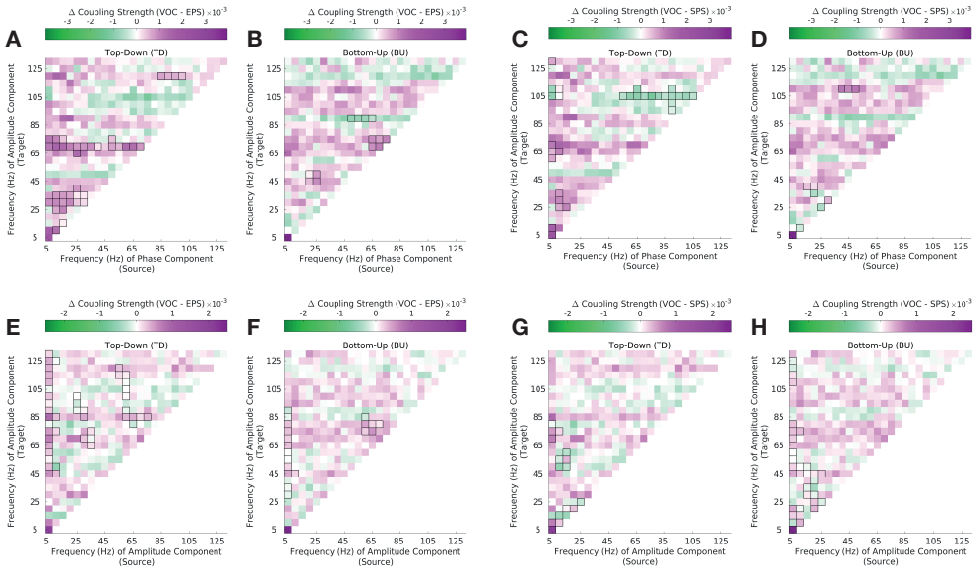


Fig 2.5. Phase-amplitude and amplitude-amplitude coupling strength in natural vocalizations (VOC) compared to synthetic envelope-preserved sounds (EPS) and spectrum-preserved sounds (SPS) (A-B) Difference in PAC strength between natural (VOC) and synthetic envelope-preserved sounds (EPS), separately in the top-down (B) and bottom-up direction (C). **(C-D)** Difference in AAC strength between VOC and EPS, separately in the top-down (C) and bottom-up direction (D). Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals. Supporting Figures: Fig. S1-11

among AACs (Fig. 2.5G,H). Altogether, coupling strength was mostly increased for natural compared to to SPS stimuli across the entire frequency spectrum in both directions.

Comparing SPS and EPS stimuli directly (Fig. S1-11), coupling strength was largely reduced for EPS stimuli across the frequency spectrum in both directions among PACs (Fig. S1-11A,B). This was also true among AACs (Fig. S1-11C,D), though there were fewer differences here than among PACs, especially in the high frequency space.

Overall, top-down and bottom-up cross-frequency coupling signatures were preserved among synthetic stimuli, albeit at overall reduced coupling strength.

We also examined coupling over time (Figs. S1-12,S1-13). We focused on coupling in the phase band that displayed the highest power (2.5Hz), and examined how strongly it influenced low frequency amplitude (12.5-37.5Hz) over time, in both directions (Fig. S1-12). Coupling strength was computed on the FFT-ed signal within the canonical correlation framework, separately for each time bin. We found that coupling strength in both directions was enhanced in the period immediately after stimulus onset (0-200ms). There was a period after stimulus onset in which bottom-up coupling strength was significantly enhanced, whereas later (¿200ms) in the course of stimulus presentation top-down coupling strength was relatively more enhanced.

Additionally, we computed coupling strength over time in the same frequency bands separately for all cross-regional pairs (Fig. S1-13). Coupling strength within the canonical correlation framework using the Hilbert-transformed signal. Coupling strength in the two directions was found to alternate in the period following stimulus onset (0-200ms). Across pais involving sector 1 (A1), bottom-up coupling strength was predominant early on and top-down coupling strength predominated somewhat later (Fig. S1-13A-C). Across higher order sectors, it was the other way round, with top-down coupling strength predominant early on (Fig. S1-13D-F). Altogether, this suggests that directional coupling strength across frequencies varies over time, and that bottom-up and top-down coupling strength alternates in time.

2.5 Discussion

We found distinct cross-frequency signatures for top-down and bottom-up information processing in auditory cortex. These patterns were similar across phase-amplitude and amplitude-amplitude couplings, and across stimulus types (albeit

at lower coupling strength in synthetic spectrum-flat stimuli). More specifically, we observed dominant top-down coupling among low-to-low and low-to-high γ frequency interactions, and dominant bottom-up coupling among low-to-low γ frequency interactions. These patterns were largely preserved across coupling types (PAC and AAC) and across stimulus types (natural and synthetic stimuli), and mostly driven by interactions between A1 and higher-order sectors. Differences in coupling strength across the auditory hierarchy are overall by an order of magnitude smaller among AACs than among PACs. Although the observed effects are subtle, they are significant and highly preserved. Altogether, our results suggest these cross-frequency signatures are a general hallmark of bottom-up and top-down information processing in auditory cortex.

The top-down coupling profile we observed is in accordance with previous results, but we also observed predominant low frequency to $low\,\gamma$ coupling in the bottom-up direction. Previous studies found that feedforward (or bottom-up) influences were mediated by gamma frequency range rhythms [Michalareas et al., 2016, Bastos et al., 2015, Fontolan et al., 2014, Bosman et al., 2012]. However, across each of these studies, there were examples of a broader frequency range being employed for bottom-up information transmission. [Fontolan et al., 2014] reported higher bottom-up coupling of delta (1-3Hz) frequency phase in left A1 to high gamma (80-100Hz) amplitude in left association auditory cortex using human depth recordings. [Bastos et al., 2015] found that theta-channels were being used in addition to gamma-frequency channels in feedforward information transmission among primate visual areas. Finally, [Michalareas et al., 2016] did not observe stronger alpha-beta frequency range feedback (top-down) modulation in 8 out of 21 primate visual cortex area pairs (with 2 pairs showing the opposite pattern), and no stronger gamma frequency range feedforward (bottom-up) modulation in 5 out of 21 pairs. We observed predominant bottom-up coupling between low frequency (θ,α,β)

phase (or amplitude) and low-γ amplitude among phase-amplitude couplings (and amplitude-amplitude couplings,respectively). Differences in exact frequency ranges reported may be due to regional (visual versus auditory cortex) variation [Tort et al., 2009, Canolty et al., 2006] and species-related variation. We also found that directional coupling strength within a particular cross-frequency pair varied over time, with periods of predominant bottom-up coupling strength. Overall, our findings agree with accounts of predominant (nonlinear) low-to-high cross-frequency interactions top-down. But together with the exceptions above, our findings suggest non-linearities may also play a significant role in bottom-up processing, pointing to a picture of top-down and bottom-up processing that is more complex than what was hypothesized by predictive coding theory ('β-down vs. γ-up') [Bastos et al., 2012a, Bastos et al., 2015].

Bottom-up PAC can be relevant to stimulus encoding. For example, an oscillation-based theory of speech encoding [Hyafil et al., 2015a, Giraud and Poeppel, 2012] predicts that cortical theta oscillations track the syllabic rhythm of speech and, in turn, reset the spiking of neurons at a gamma frequency level through θ-γ PAC. This ensures that neural excitability is optimally aligned with incoming speech. Our finding of dominant low frequency phase to γ amplitude coupling bottom-up fits within a framework that requires theta-gamma coupling for accurate speech encoding [Hyafil et al., 2015a]. Cross-regional θ-γ PAC as we observed here across A1 and higher order sectors could be responsible for coordinating information processing along the auditory cortical hierarchy. Our observation that these patterns are consistent across stimulus classes implies that this may be a generalized mechanism of auditory information processing which is co-opted during speech processing. However, our finding of enhanced low-to-high frequency phase-amplitude coupling in the top-down direction shows these couplings also mediate top-down influences on information processing.

Coupling in the top-down direction has been of particular interest in theories of the effect of attention on auditory information processing. Attention is thought to sample stimuli rhythmically, fluctuating at a 4-10 Hz rhythm [Landau and Fries, 2012], and was shown to enhance the processing of degraded speech in the anterior and posterior STS [Wild et al., 2012]. Frontal top-down signals in the delta and theta frequency range have been shown to modulate low-frequency oscillations in human auditory cortex, increasing their coupling to continuous speech [Park et al., 2015]. The same delta and theta frequency bands have been shown to be more strongly coupled to gamma-range burst rate within monkey A1 during periods of task entrainment [Lakatos et al., 2016]. In visual cortex, delta frequency oscillations have been shown to entrain with the rhythm of stimuli in the attended stream [Schroeder et al., 2010, Lakatos et al., 2008], and top-down beta-band modulation was found to be enhanced with attention [Bastos et al., 2015]. Increased power in low-frequency range oscillations may also, in turn, enhance bottom-up rhythms through top-down (cross-frequency) coupling [Bressler and Richter, 2015, Lee et al., 2013]; indeed, top-down beta-band activity was found to be maximally correlated with bottom-up gamma-band activity when top-down preceded bottom-up activity [Richter et al., 2017]. In addition, top-down couplings may also reflect the modulatory activity of predictive processes on auditory/speech perception & processing [Chao et al., 2018, Arnal and Giraud, 2012, Gagnepain et al., 2012, Davis et al., 2011, Wild et al., 2010, Poeppel et al., 2008]. Our findings of dominant low-to-low and low-to-high frequency coupling top-down are consistent with these accounts, and could well serve to enhance the animal's engagement in the task [Lakatos et al., 2016, Park et al., 2015]. In addition, the top-down coupling signal we observed may also be carrying predictive information, being consistent with frequency ranges which have been shown to carry prediction-update signals [Chao et al., 2018]. We cannot dissociate between attentive or predictive signals here, however.

We found PAC and AAC to offer similar accounts of information processing. This could be due to a high degree of shared informational content between the two measures; it remains unclear how these measures differ, and how AAC could function mechanistically in the brain. To the degree that PAC and AAC offer complementary, non-overlapping information, it is possible that top-down influences are exerted more broadly through various CFC types including both PAC and AAC, and others such as phase synchronization [Fries, 2005].

We observed the same pattern of results in synthetic stimuli. The overall lower coupling strength we observed in synthetic stimuli could be a consequence of the coding differences that were previously observed between synthetic and intact stimuli [Fukushima et al., 2014], or they could be driving these coding differences. Lower coupling strength in synthetic stimuli could also be a result of the modulatory effect of attention being lower for synthetic compared to natural stimuli, perhaps because animals pay less attention to unnatural, ecologically irrelevant stimuli. The observation of higher coupling strength in spectrally rich (SPS) over spectrally poor (EPS) stimuli, envelope-intact stimuli suggests enhanced encoding and processing of information-rich stimuli. This is also consistent witho our previous results which showed that information about spectrally rich stimuli was better maintained than information about spectrally flat stimuli, at least across the first two sectors (S1,S2) of the auditory hierarchy. In the highest-order auditory sector (S4), decoding performance showed differences across a broader frequency range among VOC and EPS than among VOC and SPS.

The fact that we find the same coupling patterns across natural and synthetic stimuli points to the existence of conserved cross-frequency signatures of bottom-up and top-down information processing in auditory cortex. These findings have

important implications for theories of neural processing such as predictive coding. In theoretical implementations of predictive coding, the presence of nonlinearities in going from predictions to prediction errors may induce cross-frequency interactions in the top-down direction. These should be less prevalent if not absent in the bottom-up direction as prediction errors are thought to accumulate linearly. Our finding of predominant low frequency to $low\gamma$ cross-frequency coupling bottom-up suggests nonlinearities may also play a role in bottom-up processing, allowing the modulation of γ power in higher-order areas by the phase of low frequency rhythms in lower order areas, for example. Overall though, our finding of widespread predominant cross-frequency interactions in the top-down direction is in agreement with the predictive coding framework.

Many of the theoretical propositions for spectral asymmetries in forward and backwards message passing in predictive coding do not consider the role of attention, however. In predictive coding, attention is mediated by optimising the precision of ascending prediction errors in a context-sensitive fashion. Physiologically, this corresponds to changing the gain of lower levels (usually considered to involve fast spiking inhibitory interneurons and some form of coherence or synchronous gain). Crucially, the effect of attention adds an extra level of nonlinearity that could be mediated by backward connections, and, implicitly, cross frequency coupling.

We employed a method based on the canonical correlation framework that allowed us to estimate directed cross-frequency coupling strength in a computationally efficient framework. Many previous studies and methods to estimate cross-frequency coupling do not obtain directional estimates [Cohen, 2008,Canolty et al., 2006,Bruns and Eckhorn, 2004, Tort et al., 2010, Voytek et al., 2013, van Wijk et al., 2015], and thus do not allow for feedfoward and feedback components to be teased apart. Studies that do take direction into account [Bastos et al., 2012b, Michalareas et al.,

2016, Bastos et al., 2015, Fontolan et al., 2014] employ separate measures to look at coupling strength (e.g. through coherence) and directional information flow (e.g. through Granger causality), and do not take *cross*-frequency interactions into account. The dynamic causal modelling (DCM) framework [Friston et al., 2017, Friston et al., 2003, Chen et al., 2008] does allow for the computation of directed coupling across frequencies [Furl et al., 2014]. Practically however, to enhance computational efficiency, current implementations of DCM reduce the dimensionality of the data by applying a singular value decomposition to the data from the source region X, and then projecting Y into this reduced space to compute coupling estimates [Chen et al., 2008]. The variance preserving dimensions in the source and target space are not necessarily the dimensions that maximize correlations between source and target space, in which case interactions may be missed. In contrast, the canonical correlation framework employed here has the advantage of finding low dimensional representations which maximize coupling between the source (X) and target (Y) data. A similar approach has been used to obtain estimates of cross-frequency interactions in MEG data [Sato et al., 2010], however without the directional Granger-causal component included here.

In summary, we revealed distinct cross-frequency signatures of top-down and bottom-up information processing. These signatures are largely preserved across coupling types and stimulus types, suggesting they are a general hallmark of information processing in auditory cortex. Our finding of prominent low-to-high frequency coupling top-down extends the current view of low-to-high frequency interactions in auditory cortex, which primarily emphasized their involvement in bottom-up processes. We employed a method that allows for the computationally efficient estimation of directed cross-frequency interactions. Altogether, this extends our understanding of how information is propagated up and down the cortical hierarchy, with significant implications for theories of neural information processing such

predictive coding.

Chapter 3

Learning to select actions shapes recurrent dynamics in the corticostriatal system

The material in this chapter at various stages of evolution was presented previously at the FENS 2018 conference, the Bernstein Computational Neuroscience Conference 2018, the Society for Neuroscience (SFN) 2018 conference, and the FENS Dynamics of the Brain 2019 conference. It was published on bioarxiv and is under review at *Neural Networks*.

Abstract

Learning to select appropriate actions based on their values is fundamental to adaptive behavior. This form of learning is supported by fronto-striatal systems. The dorsal-lateral prefrontal cortex (dlPFC) and the dorsal striatum (dSTR), which are strongly interconnected, are key nodes in this circuitry. Substantial experimental evidence, including neurophysiological recordings, have shown that neurons in these structures represent key aspects of learning. The computational mechanisms that shape the neurophysiological responses, however, are not clear. To examine this, we developed a recurrent neural network (RNN) model of the dlPFC-dSTR circuit and trained it on an oculomotor sequence learning task. We compared the activity generated by the model to activity recorded from monkey dlPFC and dSTR in the same task. This network consisted of a striatal component which encoded action values, and a prefrontal component which selected appropriate actions. After training, this system was able to autonomously represent and update action values and select actions, thus being able to closely approximate the representational structure in corticostriatal recordings. We found that learning to select the correct actions drove action-sequence representations further apart in activity space, both in the model and in the neural data. The model revealed that learning proceeds by increasing the distance between sequence-specific representations. This makes it more likely that the model will select the appropriate action sequence as learning develops. Our model thus supports the hypothesis that learning in networks drives the neural representations of actions further apart, increasing the probability that

the network generates correct actions as learning proceeds. Altogether, this study advances our understanding of how neural circuit dynamics are involved in neural computation, showing how dynamics in the corticostriatal system support task learning.

3.1 Introduction

Human and nonhuman primates are capable of complex adaptive behavior. Adaptive behavior requires predicting the values of choices, executing actions on the basis of those predictions, and updating predictions following the rewarding or punishing outcomes of choices. Reinforcement learning (RL) is a formal, algorithmic framework useful for characterizing these behavioral processes. Experimental work suggests that RL maps onto fronto-striatal systems, dopaminergic interactions with those systems, and other structures including the amygdala and thalamus [Averbeck and Costa, 2017]. Little is known, however, about how the RL formalism and the associated behaviors map onto mechanisms at the neural population level across these systems. How do neural population codes evolve with learning across these systems, and what are the underlying network mechanisms that give rise to these population codes?

Experimental work and modeling has implicated fronto-striatal systems in aspects of RL [Niv, 2009, Lee et al., 2012, Botvinick, 2012, Botvinick and Weinstein, 2014, Wang et al., 2018, Langdon et al., 2018, Niv, 2019]. Several studies have suggested that the striatum codes action values [Houk, 1995, Suri and Schultz, 1998, Doya, 1910, Nakahara et al., 2001, O'Doherty et al., 2004, Frank et al., 2004, Frank, 2005, Samejima et al., 2005, Pasupathy and Miller, 2005, Histed et al., 2009, Amemori et al., 2011, Daw et al., 2011, Sarvestani et al., 2011, Li and Daw, 2011, Seo et al., 2012, Averbeck and Costa, 2017]. These studies have further suggested that the phasic activity of dopamine, which codes reward prediction errors,

drives updates of the striatal action value representations following reward feedback. Several areas in the PFC have also been implicated in action selection and decision making [Wood and Grafman, 2003, Friston, 2005, Averbeck et al., 2006, Cheong et al., 2006, Koechlin and Summerfield, 2007, Tsujimoto et al., 2008, Sakai, 2008, Histed et al., 2009, Collins and Koechlin, 2012, Botvinick, 2012, Stokes et al., 2013, Lim and Goldman, 2013, Verschure et al., 2014, Botvinick and Weinstein, 2014, Domenech and Koechlin, 2015, Rich and Wallis, 2016, Balaguer et al., 2016, Alexander et al., 2018, Aitchison and Lengyel, 2017b, Wallis et al., 2019, Radulescu et al., 2019]. These studies further suggest that PFC plans future actions and predicts future outcomes. While both striatum and PFC have been found to represent action value and choice signals, value signals were found to be stronger in dSTR than lPFC, while action related signals were stronger in lPFC [Pasupathy and Miller, 2005, Samejima et al., 2005, Averbeck et al., 2006, Seo et al., 2012].

Previous work in the motor system and in prefrontal cortex has shown that insight into the computational mechanisms that underlie complex tasks can be gained by treating neural populations as a dynamical system and studying how their trajectories evolve with time [Rabinovich et al., 2008, Buonomano and Maass, 2009, Sussillo and Abbott, 2009, Sutskever, 2013, Shenoy et al., 2013, Mante et al., 2013, Sussillo and Barak, 2013, Barak et al., 2013, Hennequin et al., 2014, Carnevale et al., 2015, Rajan et al., 2016, Gallego et al., 2017, Wang et al., 2017, Chaisangmongkon et al., 2017, Remington et al., 2018, Wang et al., 2018, Yang et al., 2018, Gallego et al., 2018, Sussillo et al., 2015, Botvinick et al., 2019, Musall et al., 2019, Richards et al., 2019]. In prefrontal cortex this work has helped shed light on how task execution is driven by dynamics around fixed and slow points in neural population space [Mante et al., 2013, Chaisangmongkon et al., 2017]. These studies have examined representations and computational mechanisms in decision making tasks, where the values of choices have already been learned previously. In the present

study, we have used a similar approach to study how representations *develop* as animals learn to make choices that deliver rewards.

In accordance with these findings, we built a joint recurrent network model of the corticostriatal system in which the striatal network represents RL-derived action values and the prefrontal cortex, via recurrent basal ganglia loops, selects appropriate actions based on this signal. We trained this system on a complex decision making task. We also obtained neural recordings from these two regions in two macaques trained on the same task.

We hypothesized that tasks are encoded in fronto-striatal network dynamics and that learning would shape these dynamics at the network level. If so, learning would drive specific structure in state space dynamics. Alternatively, learning could be driven mainly by single-unit dynamics (by modulating persistant activity of single units, for example), without any significant changes to overall state space dynamics. We further hypothesized that a system designed to learn RL-derived action values and select appropriate actions based on them would be computationally similar to the fronto-striatal system in the brain. If so, the representational structure during learning in the network should be similar to that found in neural recordings. Moreover, the differing roles assigned to the striatal and prefrontal networks should suffice to induce a difference in representational structure across the two regions in a way that matches the asymmetries in action value and choice representation observed previously [Seo et al., 2012].

Investigating the change in representational structure with learning, we found that dynamic movement-sequence representations moved apart from each other in latent space with learning, in both the model and the neural data. We found that this process was driven by the evolution of potential surfaces in the networks such that

movement-sequence specific gradient minima moved farther apart in activity space. This increase in distance, or in other words, the increase in the height of the gradient hill between sequence representations, makes it less likely that the wrong action is selected as learning proceeds.

3.2 Methods and Materials

3.2.1 Neural Data

The neural data employed here has been previously published in [Seo et al., 2012], though not with the analysis that has been carried out here.

3.2.1.1 Subjects

Two adult male rhesus monkeys (Macaca mulatta) weighing 5.5–10 kg were used for recordings. All procedures and animal care were conducted in accordance with the Institute of Laboratory Animal Resources Guide for the Care and Use of Laboratory Animals. Experimental procedures for the first animal were in accordance with the United Kingdom Animals (Scientific Procedures) Act of 1986. Procedures for the second animal were in accordance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals and were approved by the Animal Care and Use Committee of the National Institute of Mental Health (NIMH).

3.2.1.2 Task and Stimuli

The two animals performed an oculomotor sequential decision-making task (Fig. 3.1A). A particular trial began when the animals acquired fixation on a green circle (Fixate). If the animal maintained fixation for 500 ms, the green target was replaced by a dynamic pixelating stimulus with a varied proportion of red and blue pixels and the target stimuli were presented (Stim On). The fixation circle stimulus was generated by randomly choosing the color of each pixel in the stimulus

(n = 518 pixels) to be blue (or red) with a probability q. The color of a subset (10%) of the pixels was updated on each video refresh (60 Hz). Whenever a pixel was updated its color was always selected with the same probability q. The set of pixels that was updated was selected randomly on each refresh. In the original experiment we focused on differences between choices driven by reinforcement learning and choices driven by immediately available information, in alternating blocks of trials. In the present manuscript we will only consider the learning blocks. The pixelating stimulus was relevant for the blocks in which choices were driven by immediately available information. Therefore, we will not consider this stimulus further.

The animal's task was to saccade to the correct target (Fig. 3.1A). The animal could make their decision at any time after the target stimuli appeared. After the animal made a saccade to the peripheral target, it had to maintain fixation for 300 ms to signal its decision (first Move + Hold). If the saccade was to the correct target, the target then turned green and the animal had to maintain fixation for an additional 250 ms (Fixate). After this fixation period, the green target was again replaced by a fixation stimulus and two new peripheral targets were presented (Stim On). In the case that the animal made a saccade to the wrong target, the target was extinguished and the animal was forced back to repeat the previous decision step. This was repeated until the animal made the correct choice. For every trial the animal's task was to correctly execute a sequence of three correct decisions for which the animal received either a juice reward (0.1 ml) or a food pellet reward (TestDiet 5TUL 45 mg). After that, a 2000ms inter-trial interval began. The animals always received a reward if they reached the end of the sequence of three correct decisions, even if errors were made along the way. If the animal made a mistake, it only had to repeat the previous decision, it was not forced back to the beginning of the sequence. The full task included both fixed and random conditions, as explained in detail in [Seo et al., 2012]. In the present study, however, only data from the fixed condition was

used.

There were eight possible sequences in this task as every trial was composed of three binary decisions (Fig. 3.1C). The eight sequences were composed of ten different possible individual movements. Every movement occurred in at least two sequences. We also used several levels of color bias, q as defined above. On most recording days in the fixed sets we used $q \in (0.50, 0.55, 0.60, 0.65)$. The color bias was selected randomly for each movement and was not held constant within a trial. Choices on the 50% color bias condition were rewarded randomly. The sequences were highly overlearned. One animal had 103 total days of training, and the other 92 days before chambers were implanted. The first 5–10 days of this training were devoted to basic fixation and saccade training.

In the fixed condition employed here, the correct spatial sequence of eye movements remained fixed for blocks of eight correct trials (Fig. 3.1D). After eight trials were executed without any mistakes, the sequence switched pseudorandomly to a new one. Thus, the animal could draw on its memory to execute a particular sequence, except following a sequence switch.

3.2.1.3 Neural Data Analysis

The neural data analyzed comprised 365 units from dSTR and 479 units from lPFC (Fig. 3.1B, and as explained in more detail in [Seo et al., 2012]). Data from individual movements was not analyzed if animal failed to maintain fixation or did not saccade to one of the choice targets. We fitted an ANOVA model with a 200ms sliding window applied in 25ms steps aligned to movement onset, as done previously [Seo et al., 2012], and identified units that did not show a significant effect for the *sequence* factor at any point across the entire recording session. These

units were excluded from the analysis, as were units that showed average firing rates below $1Hz$. For subsequent analysis, data was pooled across animals and recording sessions and averaged across runs.

To analyze neural population responses, we applied demixed principal component analysis (dPCA) [Brendel et al., 2011, Kobak et al., 2016] to the firing rate traces. As a dimensionality reduction technique, dPCA strives to find an encoding (or latent representation) which captures most of the variance in the data but also expresses the dependence of the representation on different task parameters such as stimuli or decisions. More specifically, it decomposes neural activity into different task parameters, in our case time ($\mathbf{X}_t$), sequence ($\mathbf{X}_s$), and certainty ($\mathbf{X}_t$), and any combination of those ($\mathbf{X}_{sc}$, $\mathbf{X}_{st}$, $\mathbf{X}_{ct}$):

$$\begin{aligned} \mathbf{X} = \mathbf{X}_t + \mathbf{X}_s + \mathbf{X}_c + \mathbf{X}_{sc} + \mathbf{X}_{st} + \mathbf{X}_{ct} = \sum_{\phi} \mathbf{X}_{\phi} + \mathbf{X}_{noise} \\ L_{dPCA} = \sum_{\phi} \|\mathbf{X}_{\phi} - \mathbf{F}_{\phi}\mathbf{D}_{\phi}\mathbf{X}\|^2 \end{aligned} \tag{3.1}$$

In order to obtain task trajectories in the reduced dPCA space, the data ($\mathbf{X}$) was smoothed with a Gaussian kernel and then projected into 3-dimensional latent space spanned by the first three vectors of the *sequence*-decoder matrix ($\mathbf{D}_s$).

Distance measures were obtained on the full datasets in full-dimensional neural space, not in the reduced subspace. Euclidean distance between all sequences was computed across all time points for each of the 8 trial repeats and averaged across all possible sequence combinations. For PFC, distances were computed between sequences *within* the two clusters (defined by whether a sequence ended in the upper or lower visual hemisphere (sequences S1, S2, S5, S6 and S3, S4, S7, S8,

respectively; see Fig.3.1C). As a measure of how compact the trajectories were, the Euclidian distance of every sequence to its centroid was computed across all time points for each of the 8 trial repeats and averaged across sequences. A sequence's centroid was defined as the mean across time of a trajectory in N-dimensional space with $N = 365$ for dSTR and $N = 479$ for lPFC.

3.2.2 Corticostriatal model

3.2.2.1 Architecture

We jointly trained a connected system of two recurrent neural networks to perform the movement-sequence task (Fig. 3.1A,E). Single-unit dynamics in these networks are governed by:

$$\begin{aligned}
\tau \dot{\mathbf{x}}^s(t) &= -\mathbf{x}^s(t) + \mathbf{W}^s_{rr}\mathbf{r}^s(t) + \mathbf{W}^s_{ir}[\mathbf{u}_a(t), \mathbf{u}_r(t)] + \eta(t) \\
\mathbf{r}^s(t) &= \tanh(\mathbf{x}^s) \\
\mathbf{y}^s &= \mathbf{W}^s_{ro}\mathbf{x}^s \\
\tau \dot{\mathbf{x}}^p(t) &= -\mathbf{x}^p(t) + \mathbf{W}^p_{rr}\mathbf{r}^p(t) + \mathbf{W}^p_{ir}[\mathbf{y}^s(t), \mathbf{u}_{ins}(t)] + \eta(t) \\
\mathbf{r}^p(t) &= \tanh(\mathbf{x}^p) \\
\mathbf{y}^p &= \mathbf{W}^p_{ro}\mathbf{x}^p
\end{aligned} \tag{3.2}$$

where $x^s_i(t)$ and $x^p_i(t)$ are synaptic current variables of unit i at time t in the striatal and prefrontal network, respectively, and activity (firing rate variables r^s_i and r^p_i) is a nonlinear function of x ($r^s_i = \tanh(x^s_i)$ and $r^p_i = \tanh(x^p_i)$), $\mathbf{W}^s_{rr}$ and $\mathbf{W}^p_{rr}$ are the recurrent weight matrices, $\mathbf{W}^s_{ir}$ and $\mathbf{W}^p_{ir}$ are the input weight matrices, and $[\mathbf{u}_a, \mathbf{u}_r]$ and $[\mathbf{r}^s, \mathbf{u}_{ins}]$ are the inputs of the striatal and prefrontal networks, respectively, and added noise $\eta_i(t)$.

The striatal and the prefrontal network had $N^s_{rr} = 1000$ and $N^p_{rr} = 800$ units,

respectively. The connectivity weight matrices $\mathbf{W}_{rr}^{s}$ and $\mathbf{W}_{rr}^{p}$ were initially drawn from the standard normal distribution and multiplied by a scaling factor of $\sqrt{N_{rr}^{s}}$ and $\sqrt{N_{rr}^{p}}$, respectively. The neural time constant was $\tau = 10\,ms$. Each unit received an independent white noise input, η_i, with zero mean and $SD = 0.01$. Inputs were fed into the networks through the input weight matrices $\mathbf{W}_{ir}^{s}$ and $\mathbf{W}_{ir}^{p}$ which were initially drawn from the standard normal distribution and multiplied by a scaling factor of $\sqrt{N_{ir}^{s}}$ with $N_{ir}^{s} = 20$ and $\sqrt{N_{ir}^{p}}$ with $N_{ir}^{p} = 510$ for the prefrontal network. Otherwise, the model parameters were set in the same range as [Mante et al., 2013].

The striatal network (Fig. 3.1E) received a 15-dimensional input vector, $\mathbf{u}_s = [\mathbf{u}_a, \mathbf{u}_r]$, composed of a 10-dimensional vector $\mathbf{u}_a$ specifying actions taken and a 5-dimensional vector $\mathbf{u}_r$ specifying the rewards received for those actions. Outputs for each of the networks were linearly read out from the synaptic currents of the recurrent circuit (Eq 3.2). The striatal network read-out was a 10-dimensional vector, $\mathbf{y}^s$, of action values derived from TD learning (as described further below).

The prefrontal network (Fig. 3.1E) received a 20-dimensional input, $\mathbf{u}_p = [\mathbf{y}^s, \mathbf{u}_{ins}]$, composed of a 10-dimensional vector $\mathbf{r}^s$ of action value outputs from the striatal network, together with a 10-dimensional instruction vector $\mathbf{u}_{ins}$ specifying when to initiate fixation (Fixate) and when to move towards or hold the target (Move + Hold). The prefrontal network read-out was a 11-dimensional output vector, $\mathbf{y}_p = [\mathbf{y}_a, y_v]$ composed of a 10-dimensional vector $\mathbf{y}_a$ of actions and an additional unit coding for the visual hemifield (upper or lower) in which a particular sequence terminated.

3.2.2.2 Network system training

All synaptic weight matrices ($\mathbf{W}_{rr}$, $\mathbf{W}_{ir}$, and $\mathbf{W}_{ro}$) were updated with the gradient of the loss function (Eq 3.3), which was designed to minimize the square of the difference between network and target output:

$$l = 2\sum_{k=1}^{K}\sum_{t=0}^{T}\sum_{i=1}^{N_{rrs}}(\hat{y}_i - y_i)^2 + \sum_{k=1}^{K}\sum_{t=0}^{T}\sum_{i=1}^{N_{rrp}}(\hat{y}_i - y_i)^2 \tag{3.3}$$

The error was thus obtained by taking the difference between network output y and target output $\hat{y}$ and summing over all trials K in a batch (with $K = 10$), time points T and recurrent units N_{rr}. The total loss function (Eq 3.3) was obtained by combining the loss terms from the striatal and the prefrontal network while assigning double weight to the striatal loss term. The striatal and the prefrontal networks were jointly trained by obtaining the gradient of the combined loss function (Eq 3.3) through automatic differentiation with *autograd* [Maclaurin et al., 2017] and custom implementations of specific functions with GPU-based acceleration using *JAX* [Johnson et al., 2018]. The network was trained with an initial learning rate of $\alpha = 0.001$ for 10 steps with 1000 iterations each, while the learning rate was decayed by $\frac{2}{3}$ at every step. After this training phase, all the synaptic weight matrices remained fixed.

3.2.2.3 Task Coding

Actions ($\mathbf{u}_a$/$\mathbf{y}_p$) and action values ($\mathbf{y}_s$) were coded as 10-D vectors in which every unit coded for one of the movement choice options (see Methods and Materials - Task and Stimuli). So, for instance, for the first binary choice option (Fig 3.1C, S1-center) one unit coded the right movement choice and the other the left movement choice. Rewards ($\mathbf{u}_r$) and visual inputs which indicated target locations and fixation

points ($\mathbf{u}_{ins}$) were coded as 5-D vectors with every unit coding for reward delivered at one of the 5 decision points (center, upwards, downwards, upper, lower; Fig 3.1C). Actions and rewards were coded as brief transients. The reward signal interval was $\frac{2}{10}$th as long as the action signal interval. Rewards were delivered after the end of the action signal interval.

Rewards drove the update of action values according to a temporal-difference reinforcement learning algorithm (Q-Learning) [Sutton and Barto, 2018]:

$$\begin{aligned} Q(s_t,a_t) \longleftarrow Q(s_t,a_t) + \alpha\,[r_{t+1} + \gamma max(Q(s_{t+1},:)) - Q(s_t,a_t)] \\ \mathbf{p}_t = softmax(\beta\,\mathbf{Q}_t) \\ \mathbf{Q}_{t+1} = \tau\,\mathbf{Q}_t \end{aligned} \tag{3.4}$$

The learning rate parameter α, the discount factor γ and an additional inverse temperature parameter β were fitted to one of the training sessions of monkey 1 using *fminsearch* in Matlab, with the decay parameter set to $\tau = 0.95$. The values obtained for the parameters were $\alpha = 0.8100$, $\gamma = 0.2010$, and $\beta = 3.050$.

The training data set for the corticostriatal model system was composed of the behavioral data from the recording sessions of the two monkeys. Action values ($\mathbf{y}_s$) were derived by feeding the actions chosen and rewards received through the Q-learning algorithm (Eq 3.4). A subset of 25 blocks from the recordings was left out as a test set. Additionally, the training data set was augmented with artificial data generated by randomly choosing one of the eight sequences for the current block and drawing action movement outcomes with the same error probabilities that the animals displayed in the real task (see Fig. 3.4E - Behavior). During training, a batch

composed of 10 blocks was randomly chosen from across the entire training set at each step.

3.2.2.4 Autonomous mode

After training, we also tested the model's ability to autonomously produce movement-sequence blocks. This meant obtaining trial-by-trial value and action estimates. More specifically, we obtained an initial set of outputs ($\mathbf{y}^s$; Eq. 3.2) from the striatal network by feeding in a vector of white noise inputs for $\mathbf{u}_s$ (with zero mean and $SD = 0.01$). Striatal outputs were then fed into the prefrontal network together with instructions $\mathbf{u}_{ins}$ specifying the particular movement stage as part of the prefrontal input vector $\mathbf{u}_p$. We obtained a vector of actions $\mathbf{y}_p$ as outputs from the prefrontal network. Actions were decoded from among the choice options that were possible at a particular movement stage (after the first movement stage, certain choice options become unavailable, i.e. for sequences S1 and S5 in Fig. 3.1C the choice options in the lower visual hemifield are no longer possible at the second movement stage). Actions were decoded probabilitiscally from the result of passing action outputs through a softmax function ($\mathbf{a}_{choice} = softmax(\mathbf{y}_p)$).

The chosen action was then fed back into the dSTR network. If the action was correct, a reinforcement signal was delivered to the dSTR. If the action was incorrect, no reinforcement signal was delivered. Then, the signal for the next movement stage was fed into the dPFC network and a new set of dSTR outputs, for the next movement stage, was fed into the prefrontal network, etc.

If the decoded action was wrong (in that it was not the correct action for the current movement-sequence block), the network system was forced to repeat this particular movement stage before proceeding (just like the animals in the original

experiment). This was done by concatenating the set of wrong actions with the striatal action input vector $\mathbf{u}_a$ and feeding it into the striatal network together with the reward vector $\mathbf{u}_r$ set to white noise. The striatal outputs obtained as such were then fed into the prefrontal network together with the instruction vector $\mathbf{u}_{ins}$ set to the same movement stage as before.

3.2.2.5 Model Analysis

The neural population activity from the striatal and prefrontal networks of the corticostriatal model was imaged in the same way as the real neural recordings, by projecting activity into a 3-dimensional latent space spanned by the first three vectors of the *sequence*-decoding matrix ($\mathbf{D}_s$) obtained through dPCA (Eq. 3.1) [Brendel et al., 2011, Kobak et al., 2016].

Canonical correlation analysis (CCA) was used to compare model and neural population responses, as was done previously [Sussillo et al., 2015]. First, both the model and the recorded neural data were averaged across trials, smoothed with a Gaussian kernel and reduced to 15 dimensions using dPCA. This ensured CCA was not performed along dimensions of high correlation but low variance. Then the first 10 canonical correlation coefficients were obtained using the entire duration of a chosen movement-sequence as a comparison window.

In order to image the potential surface, we first projected neural population activity into 3-dimensional latent space and obtained a mesh of points ($\mathbf{X}^*$) around the sample trajectories. Then we projected this mesh of points back out into neural population space using the *stimulus*-encoder matrix ($\mathbf{F}_s$; Eq. 3.1). We started the networks off at each of these points by providing them as the initial vector of firing rates (**r(t)**; Eq. 3.2), and let the network system run through a whole block of

movement-sequence trials. For illustration purposes, we imaged the potential surface at the same point for all trials in a block. We obtained the potential surface at the chosen moment in time by calculating the magnitude of the gradient for the mesh of points. In order to ensure that the trough of the potential surface at the chosen moments in time really pointed to fixed point regions, we kept the input fixed and continued running the network system through 10 iterations (which corresponded to the length of the fixation period) to make sure the location of the trough remained fixed on the timescale of the network. If the location of the minimum remained fixed for the length of this period, we labelled this location as a fixed point. We then obtained the joint potential surface of two chosen sequences by taking the point-wise minimum across the two sequences' manifolds.

In order to calculate the distance between the minima of the potential surfaces of two different sequences, we first confirmed the location of the fixed points in 3-dimensional latent space as described above. Then, we used Dijkstra's algorithm to obtain the minimal path length along the joint potential surface between two particular sequences' fixed points. We repeated this for all possible sequence combinations in the test set and averaged the result. In order to calculate the distance between sequences in latent space as well as the distance of a particular sequence to its centroid, we first projected the data into 10-dimensional dPCA-space (for computational reasons) and then proceeded the same way as in the neural data (see Methods and Materials - Neural Data - Data Analysis). We found no difference in these metrics when projecting into 10- or 20-dimensional dPCA- space.

3.3 Results

We investigated dynamics during learning in the fronto-striatal system (Fig. 3.1). The task consisted of a sequence of three movements which the animal had to execute

by saccading to the correct targets (Fig. 3.1A). While the animal exectued these movement trajectories, recordings were obtained from dPCA and lPFC (Fig. 3.1B). There were a total of 8 possible arrangements (or sequences) of 3 movements (Fig. 3.1C). These sequences, in turn, were arranged into blocks in which one particular sequence remained fixed for the duration of the block (Fig. 3.1D). Thus the animal was able to learn which sequence was correct in the current block using feedback about chosen movements. At the start of a new block, a new sequence was chosen randomly (see Methods and Materials for further details).

Two animals performed this task while we obtained neural recordings from lateral prefrontal cortex (lPFC) and dorsal striatum (dSTR) (Fig. 3.1B, see Methods and Materials and [Seo et al., 2012] for further details). To investigate learning dynamics, we built a model of the fronto-striatal system. In this model, consistent with the neural data, the striatum represents action values, and the prefrontal cortex selects actions (Fig. 3.1E). Whether actions are selected via return loops through the thalamus, or in other structures downstream from the striatum is currently unclear. However, at least in some conditions in the in-vivo study, actions were selected in cortex before they were selected in the striatum [Seo et al., 2012]. Therefore, we organized our model consistent with this. The prefrontal network received action value information from the striatum and selected actions based on this information. The prefrontal network received additional inputs signifying fixation and move/hold periods, which were presented to the animal as visual cues during the task. The two networks in this system were trained using actions and rewards from the animals' real recorded behavior. Corresponding action values were generated by feeding actions and rewards through a temporal-difference reinforcement learning algorithm (Q-learning; see Methods and Materials).

After training, the fronto-striatal model learned to produce correct movement

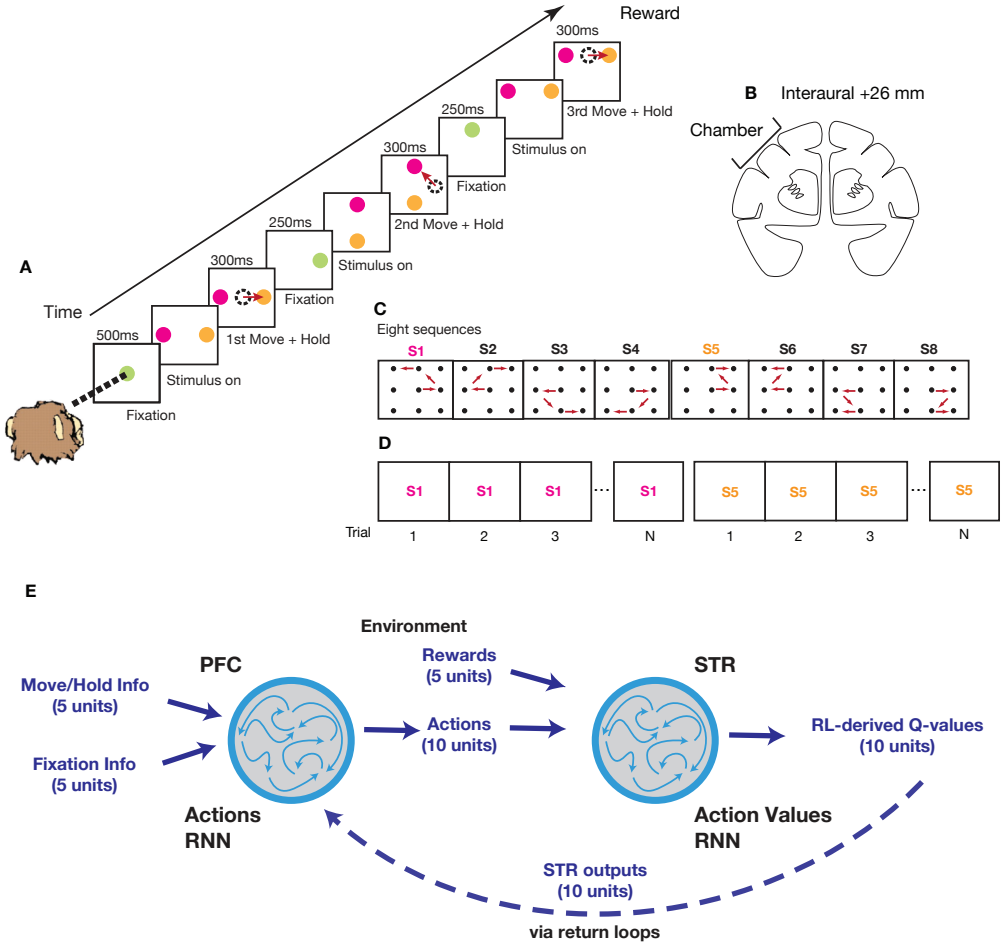


Fig 3.1. Task and Model Overview. (A) Sequence of events in the task. (B) Coronal section showing approximate location of recording chamber. (C) The eight possible movement-sequences used in the task. (D) Trial structure. Trials were arranged into blocks of 8 repeats of the correct movement-sequence plus a variable number of error trials (i.e. sequence S1 followed by sequence S5). (E) Corticostriatal model, consisting of a prefrontal network and a striatal network of recurrently connected units. The prefrontal network selects actions based on inputs from the striatal network. The striatal network outputs action values based on inputs specifying actions performed and rewards received.

sequences (Fig. 3.2). In this example, sequence S5 (consisting of a rightwards movement, followed by upwards, and another rightwards movement) was the correct sequence for the current block. The prefrontal network units coding for these movements selected the correct action (light blue). Following selection of the correct action the network received a reward input (dark blue). The reward, in turn, made the value signal in the striatal network (red) increase for the rewarded movement. The output of the prefrontal and striatal network units coding for other movement directions remained flat. Altogether, the system learned to produce movement sequences with the correct action and action value output.

After training, the network system could be run autonomously. That is dlPFC selected actions based on the dSTR action value inputs, and external inputs to dlPFC indicating the ordinal position of the current movement in the sequence (see Methods and Materials - Corticostriatal Model - Autonomous mode). In the first step, a vector of white noise inputs and a vector of external inputs specifying the ordinal position of the movement were used to generate value signals in the striatal network. These value signals were input to the prefrontal network which selected the first movement. This movement was then fed back into the striatal network together with the corresponding reward outcome. The signal representing the new ordinal position in the sequence was used to generate the next value signal, and so on. In this manner, we generated several trials of the same movement sequence, which were strung together into blocks, as in the original experiment (Fig. 3.1D). We imaged the outputs of the network system as the correct sequence switched from one block to the next (Fig. 3.3). There were a total of 10 output units, one for each of the available movement directions. Decoded actions (in light blue) from the outputs of the prefrontal network, action value outputs (in red) from the striatal network and rewards (in dark blue) delivered externally were all imaged together for these two sequential blocks. The correct sequence switched from one block to the next

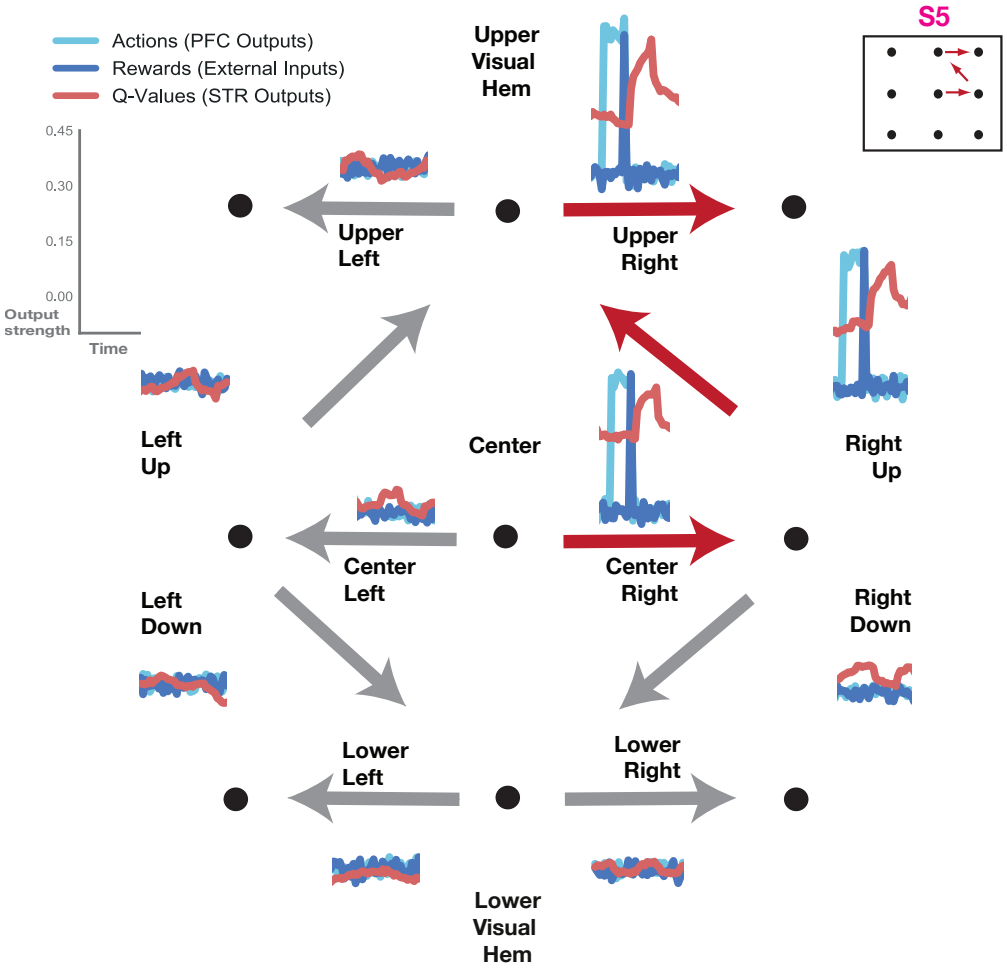


Fig 3.2. Task coding. The correct movement sequence (S5) was signalled by the corresponding output units while units coding other movement directions stay flat. Actions (light blue) were indicated by pulses in the prefrontal network's output units corresponding to a particular direction. Rewards (dark blue, plotted on top) were delivered at the end of action pulses. Action values (red, plotted on top) were indicated by striatal output units corresponding to a particular direction; action values increased after reward delivery.

(sequence S2, magenta, for the first block, and S4, light brown, for the second; Fig. 3.3 inlays).

The correct sequence for the first block (S2, magenta) is composed of a left move in the center, followed by an upwards move on the left side and a rightwards move at the top (see Fig. 3.3 inlay). Accordingly, the three output units at center left, left up and upper right were active. The correct actions could be decoded probabilistically from the outputs of the prefrontal network; actions were signalled by sequential pulses (light blue) in the three output units that made up the correct sequence. Rewards were presented as short pulses at the end of an action (dark blue). The striatal network subsequently produced the corresponding action value signal (red). The value signal decayed over time, but recovered when correctly executed trials followed upon each other (as in the upper right movement panel, Fig. 3.3).

As the new block began (grey dotted line), the correct sequence changed (from S2 in magenta to S4 in light brown). The new sequence was composed of a center right move, followed by a downwards move on the right and a move to the left at the bottom (see Fig. 3.3 inlay), ending up in the lower visual hemifield. Accordingly, the three output units at center right, right down and lower left were now active. The output units representing the movements that had been correct during the previous sequence (center left, left up, upper right) were now inactive; the striatal network's value signal in these units decayed back towards zero.

Occasionally, the prefrontal network commits errors and the wrong movement was decoded. Errors occur with the same frequency as in the behavioral data from the two monkeys carrying out this task (as the training set for the networks was derived from real behavior, see Methods and Materials). Upon a block switch, errors can occur at all three movement stages. At the first movement stage, the PFC

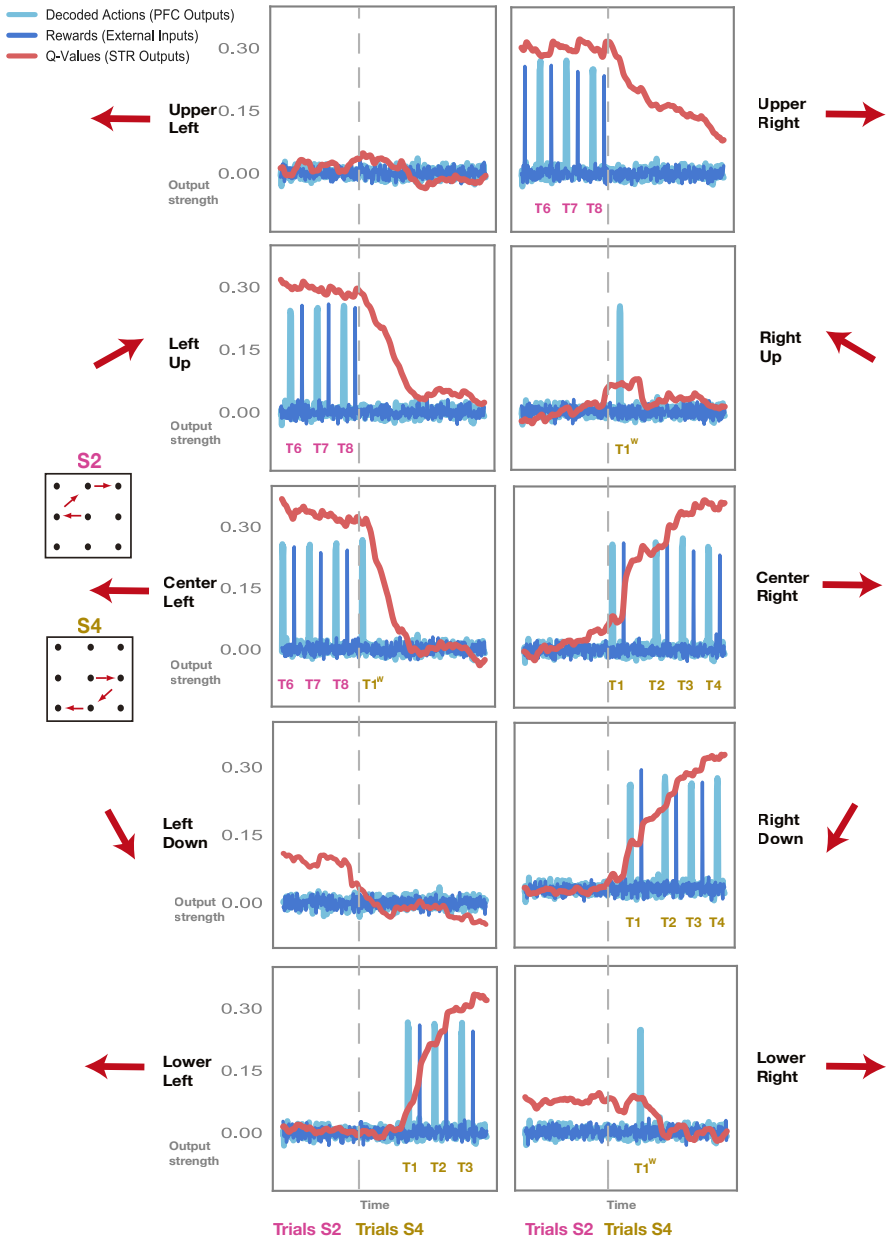


Fig 3.3. Corticostriatal model: autonomous behavior. Sample transition between two blocks in which the correct sequence switches (from S2, in magenta, to S4, in light brown; block transition indicated by vertical grey dotted line). Decoded Actions (light blue) are plotted together with rewards (dark blue) and action values (red) for all output units. Action values spike after correctly executed, rewarded actions and increase with successive correct actions. Action values for wrong, unrewarded moves decay quickly. Trials with erroneous movements remain unrewarded and are repeated, just like in the original experiment. Supporting Figures: Fig. S2-1

network produced the same movement which had been correct in the previous block (center left, $T1^w$ in light brown; Fig. 3.3). This move remained unrewarded. The striatal value signal decayed quickly after unrewarded moves, instead of spiking as it did after rewarded moves (i.e. center left, T8 in magenta). The network was forced to repeat the move before progressing further (just like the animals were in the original experiment; see Methods and Materials). The striatal value signal for the opposite movement direction now increased, prompting the prefrontal network to produce the correct move at the second attempt (T1 in light brown, center right). The correct action was now rewarded and the striatal value signal for this move increased further, while the value signal for the unrewarded action decayed further (center left). Further erroneous moves occurred at the second and third movement stages (right up and lower right, respectively). These outputs remained unrewarded, and the network was again forced to repeat these moves. The correct moves were generated at the second attempt, and the network was allowed to progress to the next trial.

We also plotted the behavior of the network system on the test set (Fig. S2-1). In this case, actions from the test set (a portion of the original dataset which was not used in network training) were fed into the striatal network to obtain value signals which, in turn, were fed into the prefrontal network to obtain action outputs. Actions were not fed back into the striatal network this time; rather, the next action from the test set was picked as the next input to the striatal network. The network behaved similarly to the autonomous mode, with value signals increasing after successive correctly executed moves and decaying quickly upon the omission of reward (Fig. S2-1). To analyze the system's performance in more detail, we plotted striatal and prefrontal outputs together with their target outputs for a sample block transition from the test set (Fig. 3.4A,B; see Methods and Materials). Output traces generally follow target traces. To quantify this, we computed the mean squared error between targets and outputs over the duration of a block averaged over all blocks in the test

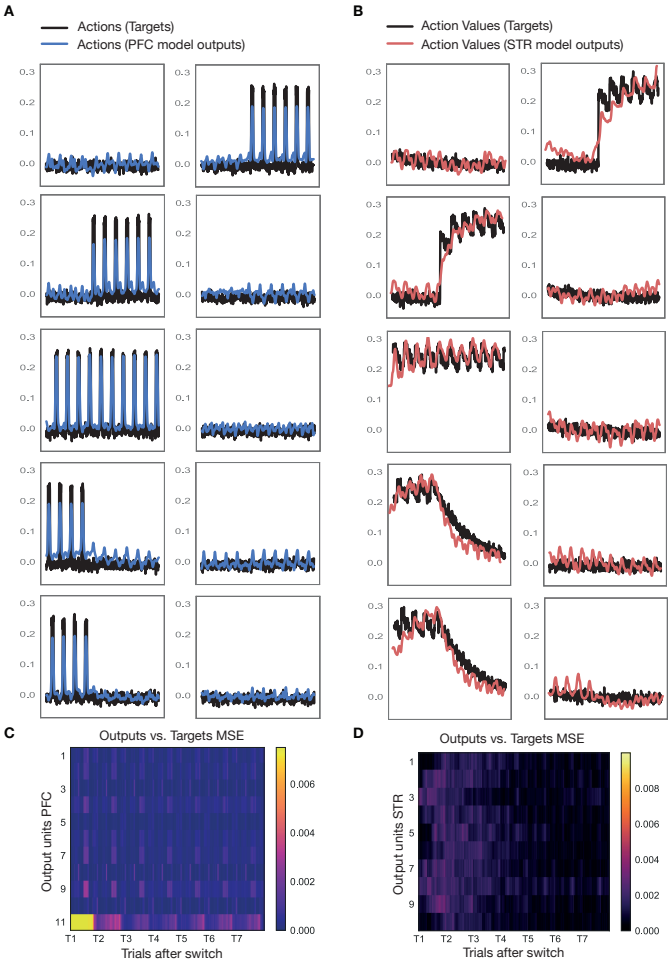


Fig 3.4. Model Performance. **A** Outputs of prefrontal model network (blue) versus targets (black) for a sample test block. **B** Outputs of striatal model network (red) versus targets (black) for the same sample test block as in A. **C-D** Mean squared error (MSE) between outputs and targets averaged over 25 test blocks. MSE is depicted as a function of trials after the sequence switched for all output units in the prefrontal (C) and striatal network (D). Supporting Figures: Fig. S2-2

set (Fig. 3.4C-D). The largest difference between outputs and targets occurred for the first trial when the correct sequence was unknown. Over the course of the block, as certainty increased, the error decreased for both prefrontal (Fig. 3.4C) and striatal networks (Fig. 3.4D). We also imaged the eigenvalue spectra of the two networks' weight matrices after training (Fig. S2-2). We found that variance was spread across many dimensions in the prefrontal network (Fig. S2-2A), while most of the variance was concentrated around 5 large eigenvalues in the striatal network (Fig. S2-2B).

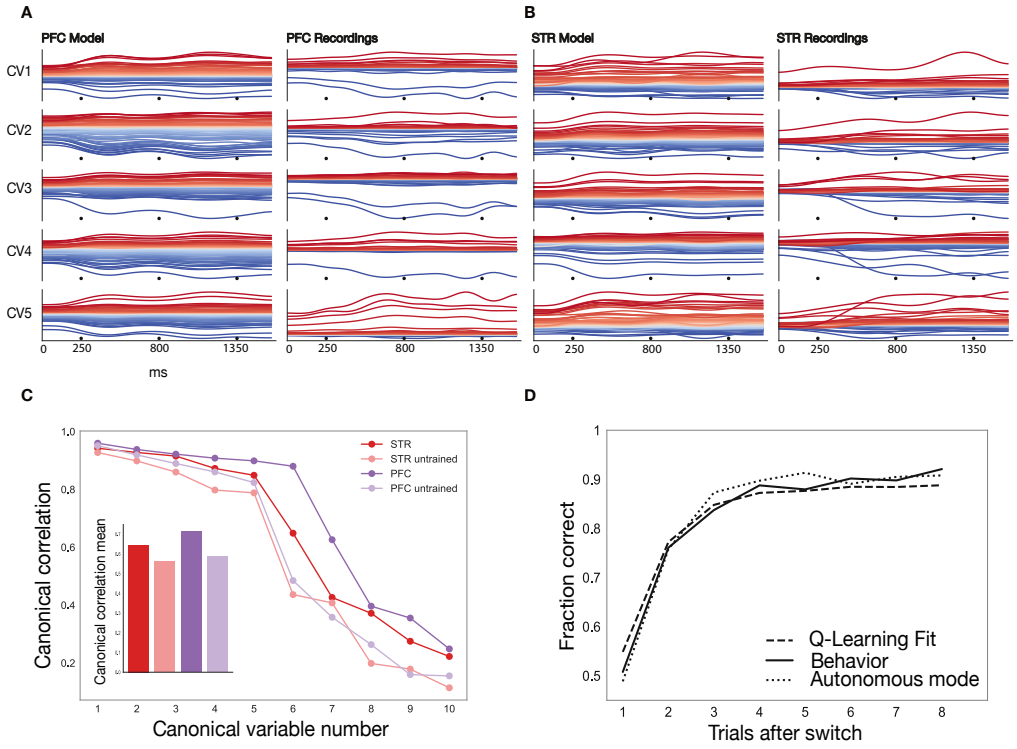


Fig 3.5. Comparison of model and recorded data A CCA projections (canonical variables) for the *prefrontal* model network (left) and the neural data recorded from lPFC (right). These projections show the directions in state-space along which the data is maximally correlated with the model. Each row shows one canonical variable (CV 1-5). Traces are colored based on the mean value of the projection across the entire depicted duration of the trace. Traces show responses for the whole length one movement-sequence which is composed of three sequential movements (black dots mark movement onset). **B** CCA projections for the *striatal* model network (left) and the neural data recorded from dSTR (right). **C** Summary of canonical correlations between model and neural data. CCA analysis provides a spectrum of of correlation coefficients that can be used to assess model performance (see Methods and Materials). The canonical correlation coefficients are shown for the trained model (striatal network in dark red and prefrontal network in dark magenta), as well as for an untrained network with the same inputs which as a baseline (striatal network in light red, and prefrontal network in light magenta). **D** Fraction of correct decisions as a function of trials after the sequence switched, separately for the behavior of the two animals (solid line), the performance of the Q-learning algorithm (broken line; see Methods and Materials-Corticostriatal model-Coding), and the performance of the corticostriatal model network system (dotted line; see Methods and Materials-Corticostriatal model-Autonomous Mode).

We compared neural activity from the model with the recorded neural data using canonical correlation analysis (Fig. 3.5A-C; see Methods and Materials). We plotted the first five canonical variables for the neural responses from the model and the recorded data (Fig. 3.5A,B). The canonical variables are the directions in state-space along which model and recorded data are maximally correlated, and offer a way to assess the similarity between model and recorded population responses. The model (left) and the neural data (right) shared many population-level response patterns. The model was able to pick up the slow oscillatory patterns observed in the neural data [Seo et al., 2012]. We also plotted the canoncial correlation coefficient for the striatal and the prefrontal network across the first ten canonical variables (Fig. 3.5C). We also obtained the correlation coefficients for an untrained model with the same inputs as a baseline. The canonical correlations were consistently higher than the baseline for both the striatal and the prefrontal network. The first seven correlation coefficients were significant for the trained striatal and prefrontal networks ($p < 4.823e-11$). The average correlation coefficients were higher for the trained networks (0.65 striatal and 0.71 prefrontal) than for the untrained networks (0.56 striatal and 0.59 prefrontal). We retrained the model 10 times with the same configuration and parameter settings, but starting with different randomly initiated weights (see Methods and Materials), and obtained a similar fit to recordings.

We also determined the behavior of the model system by measuring the fraction of correct decisions over the course of the movement-sequence block and compared it to the recorded behavior from the animals (Fig. 3.5D). The recorded behavioral data from the animals (solid line, as in [Seo et al., 2012]) shows chance performance at the beginning of the block, and rapid, steady improvement over the course of the block. In order to determine the fraction of correct decisions of the model system, we let the networks produce movement-sequences autonomously (see Methods and

Materials-Corticostriatal model). That is, we used the action outputs of the prefrontal network as inputs to the striatal network. In this way we obtained a distribution of activations for the next movement step in the prefrontal output units, from which we decoded the predicted action. The fraction of correct decisions of the autonomous model system averaged over 25 blocks (dotted line) approximated that of the animals' behavior (solid line). We also plotted the fraction of correct decisions obtained from the fit of the Q-learning algorithm (dashed line), and it agreed well with behavior. This shows the model system was able to capture the animals' behavior in this task.

To study how neural representations evolved with learning, we imaged PFC neural population activity in 3-dimensional latent space using demixed principal component analysis (see Methods and Materials). Trajectories from neural recordings in lPFC (Fig. 3.6A-D) are plotted alongside trajectories from the prefrontal network of the corticostriatal model (Fig. 3.6E-H). The plots capture representations for progressive trials over the course of a block (Fig. 3.1D), as the animals advance from 50% certainty at the start of the block (Fig. 3.6A&B) to close to 90% certainty by the fifth trial into the block (Fig. 3.6C&G).

Sequence representations from neural recordings in lPFC (Fig. 3.6A-D) showed a separation by visual hemifield: sequences S1, S2, S5 and S6 which progressed along the upper visual hemifield (Fig. 3.1C) were clustered to the left while the remaining sequences which progressed along the lower visual hemifield were clustered to the right. This separation appeared to be present from the very start of a block. We also examined the evolution of representations across learning. As the animals learned in each block, they more frequently selected the correct option (Fig. 3.4E), which suggests they become more certain of their choices. We found that sequence trajectories within each particular cluster separated more from each other with increasing certainty (Fig. 3.6A-C). To capture this effect, we computed

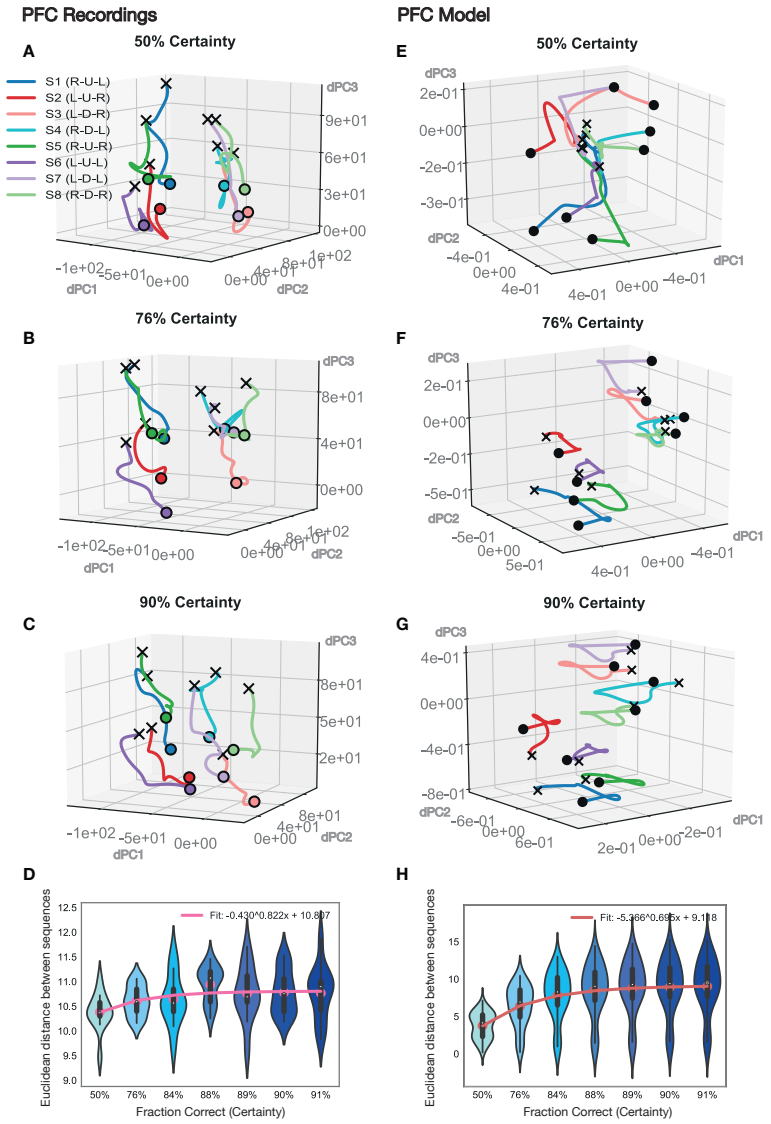


Fig 3.6. Evolution of latent task coding with learning for PFC recordings and model network. A-D lPFC neural recordings. Latent task trajectories during the execution of the three-movement sequence depicted in three-dimensional latent space (obtained from dPCA using the sequence dimension; see Methods and Materials-Neural Data-Data Analysis). Latent trajectories for all eight possible movement sequences (S1-S8) are plotted together, for increasing certainty (fraction correct) levels (A-C). **D** Euclidean distance between all sequence trajectories within the two clusters (brightly and lightly colored trajectories) as a function of increasing certainty (see Methods and Materials-Neural Data-Data Analysis). **E-H** PFC model network. Latent task trajectories during the execution of the three-movement sequence depicted in three-dimensional latent space (obtained from dPCA using the sequence dimension, as in A-D). Latent trajectories for all eight possible movement sequences (S1-S8) are plotted together, for increasing certainty levels (E-G). **H** Euclidean distance between all sequence trajectories within the two clusters (brightly and lightly colored trajectories) as a function of increasing certainty level (calculated as in D).

the Euclidean distance between sequences within clusters in neural population space (in the full-dimensional space of recordings, not in the reduced latent space; see Methods and Materials). We found this measure increased with certainty as learning progressed over the block (Fig. 3.6D). Sequence representations from the prefrontal network model (Fig. 3.6E-H) showed the same clustering by visual hemifield. Task trajectories were somewhat more distinctly clustered by hemifield at the start in the neural data (Fig. 3.6A) than in the model (Fig. 3.6E). Overall, though, we also found Euclidian distance between sequences within clusters increased for the model trajectories (Fig. 3.6H).

To study what underlies increasing separation of movement-sequence representations with learning, we probed the prefrontal model network further (Fig. 3.7). We obtained the potential surface in the region around the latent sequence trajectories as learning progressed across the block (see Methods and Materials - Model Analysis). We obtained the joint surface across two particular movement-sequence trajectories in latent space by taking the point-wise minimum across the potential surfaces of the individual sequences, and imaged this common surface for increasing levels of certainty across a block (Fig. 3.7A-C). The potential surface shows how activity evolves at different points in neural latent space in the vicinity of sequence trajectories. Neural activity has a high propensity to be pushed away from locations where the magnitude of the gradient is high (yellow), and remain in locations where the magnitude is low (dark blue). We ascertained that the troughs of the potential surface pointed to energy minima (or fixed points) in neural activity (see Methods and Materials - Model Analysis).

Observing the evolution of joint potential surfaces with learning (Fig. 3.7A-C) for a particular point along the movement trajectory (magenta dot), one notices how energy minima for different regions become increasingly well separated. Along with

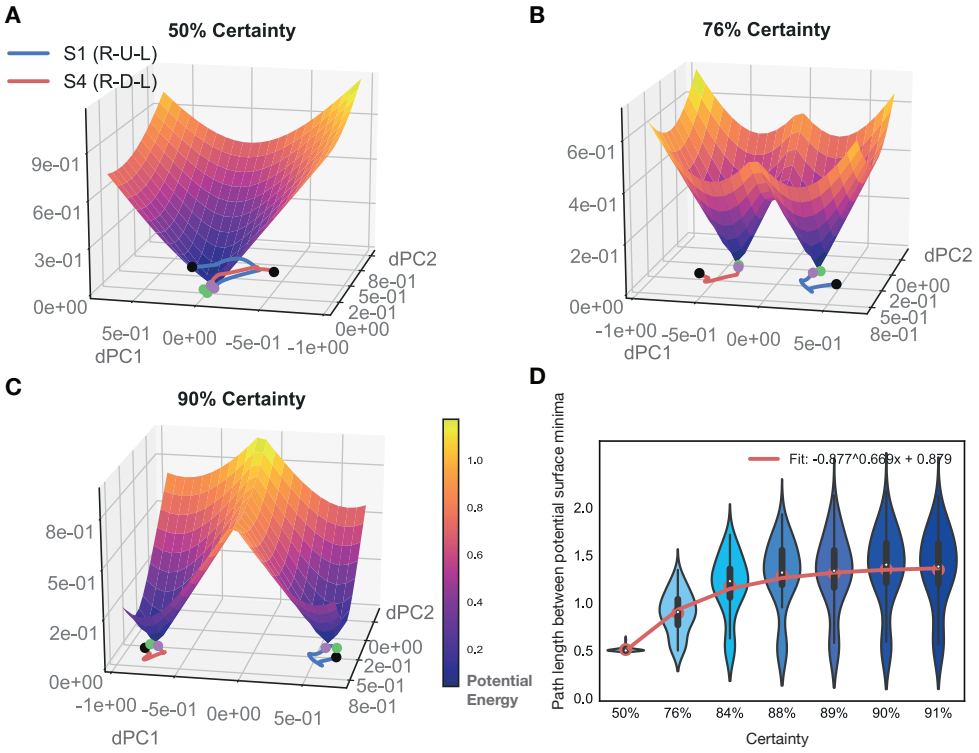


Fig 3.7. Evolution of potential surface with learning in PFC model network. A-C Latent trajectories (obtained through dPCA) for two different movement sequences are depicted in two dimensional latent space (S1 in blue and S4 in red, on the bottom). Solid dots depict the beginning (green) and end (black) of the three-movement sequence. The surface depicts the potential energy of the network in the two dimensional space in which the two sequence trajectories sit. To obtain the potential surface, the network was iterated one step forward with inputs held fixed to a particular chosen timepoint (magenta dot; see Methods and Materials-Corticostriatal Model-Model Analysis). Latent movement-sequence trajectories and potential surface are depicted for increasing certainty (fraction correct) levels in A-C. **D** Minimum path length between the gradient minima of all sequence pairs in the test set. Path length was calculated along the joint gradient surface by using Dijkstra's algorithm (see Methods and Materials-Corticostriatal Model-Model Analysis). Results are depicted for increasing certainty levels.

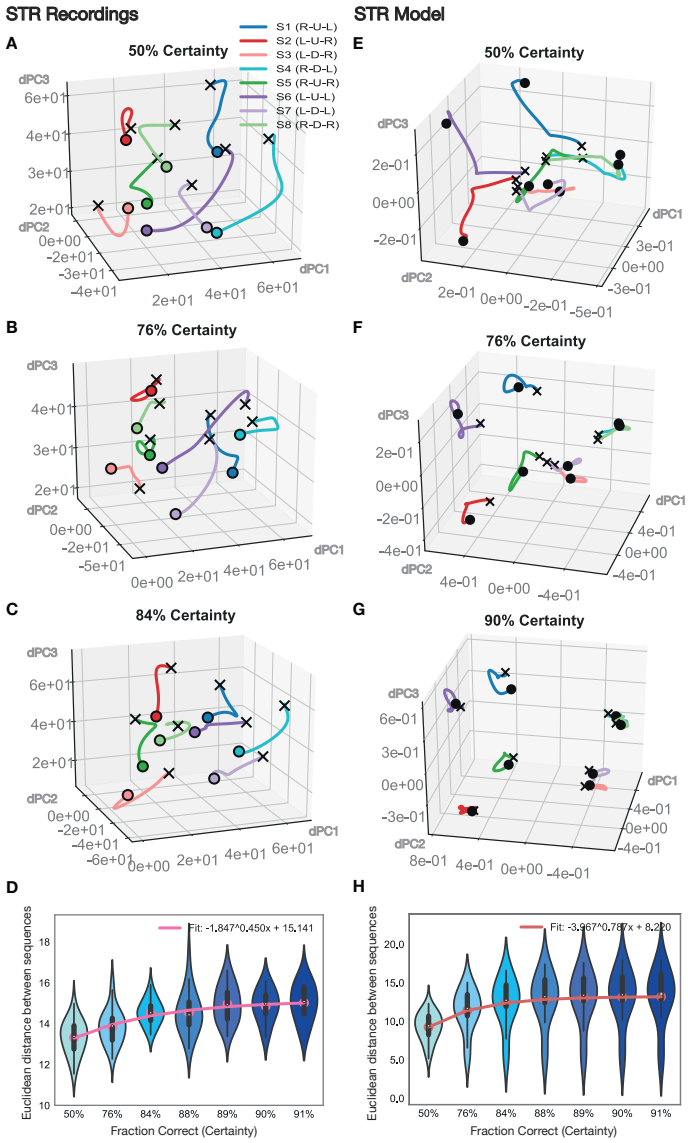


Fig 3.8. Evolution of latent task coding with learning for STR recordings and model network. A-D dSTR recordings. Latent task trajectories during the execution of the three-movement sequence depicted in three-dimensional latent space (obtained from dPCA using the sequence dimension; see Methods and Materials-Neural Data-Data Analysis). Latent trajectories for all eight possible movement sequences (S1-S8) are plotted all together for increasing certainty (fraction correct) levels (A-D). **D** Euclidean distance between all sequence trajectories within the two clusters (brightly and lightly colored trajectories) as a function of increasing certainty level (see Methods and Materials-Neural Data-Data Analysis). **E-H** STR model network. Latent trajectories for all eight possible movement sequences (S1-S8) are plotted together in dPCA-derived latent space (obtained as in A-D) for increasing certainty (fraction correct) levels (E-G). **H** Euclidean distance between all sequence trajectories within the two clusters (brightly and lightly colored trajectories) as a function of increasing certainty level.

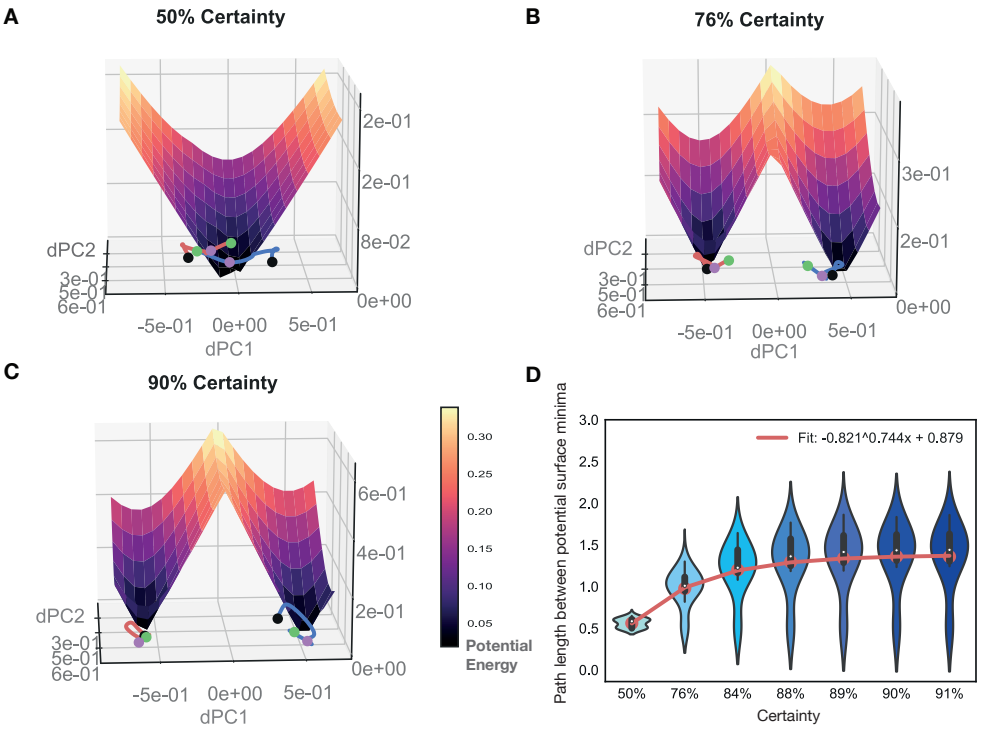


Fig 3.9. Evolution of gradient landscape with learning in STR model network. A-C Latent trajectories (obtained through dPCA) for two different movement sequences (S1 in blue, S4 in red) are depicted in two dimensional latent space (bottom). Solid dots depict the beginning (green) and end (black) of the three-movement sequence. The surface depicts the gradient of the network in the two dimensional space in which the two sequence trajectories sit. To obtain the gradient manifold, the network was iterated one step forward with inputs held fixed to a particular chosen timepoint (magenta dot; see Methods and Materials-Corticostriatal Model-Model Analysis). Latent movement-sequence trajectories and gradient surfaces are depicted for increasing certainty (fraction correct) levels (50%, 75% and 90% certainty in A-C). **D** Minimum path length between the gradient minima of all sequence pairs in the test set. Path length was calculated along the joint gradient surface by using Dijkstra's algorithm (see Methods and Materials-Corticostriatal Model-Model Analysis). Results are depicted for increasing certainty levels.

this separation, the ridge in the joint potential surface between the two sequence minima heightened with increasing certainty during learning. As this happens, it becomes increasingly less likely to commit errors: the gradient along a particular side of the ridge drives activity more strongly towards a particular sequence's fixed point region, so that the chance to end up in the other minima decreases with learning. To quantify this effect, we determined the minimal path length between the various sequence's fixed point regions as learning progressed (see Methods and Materials - Model Analysis). We observed that minimal path length between fixed point regions increased with certainty during learning (Fig. 3.7D). Altogether, we established that the energy landscape in the network changes with learning such that fixed point regions for different movement-sequences are pushed farther apart from each other, underlying the increase in behavioral accuracy.

We also examined how neural representations evolved with learning in the striatum (Fig. 3.8). Trajectories from neural recordings in dSTR (Fig. 3.8A-D) are plotted alongside trajectories from the striatal model network (Fig. 3.8E-H). Sequence representations in the dSTR (Fig. 3.8A-D) did not display any particular clustering by visual hemifield, unlike representations in lPFC (Fig. 3.6A-D). Sequence representations were instead scattered around in latent space. As learning progressed, sequence representations spread further apart from each other (Fig. 3.8A-C). We computed the Euclidean distance between all sequences in neural population space and found it to be increasing with learning (Fig. 3.8D). Trajectories from the striatal model network (Fig. 3.8E-H) were also scattered around latent space, similar the neural recordings (Fig. 3.8A-D). Like in the recordings, sequence representations in the model spread further apart from each other as learning progressed. To quantify this effect, we again computed the Euclidean distance between sequences (see Methods and Materials) and found it to be increasing with learning (Fig. 3.8H), as in the neural data (Fig. 3.8D).

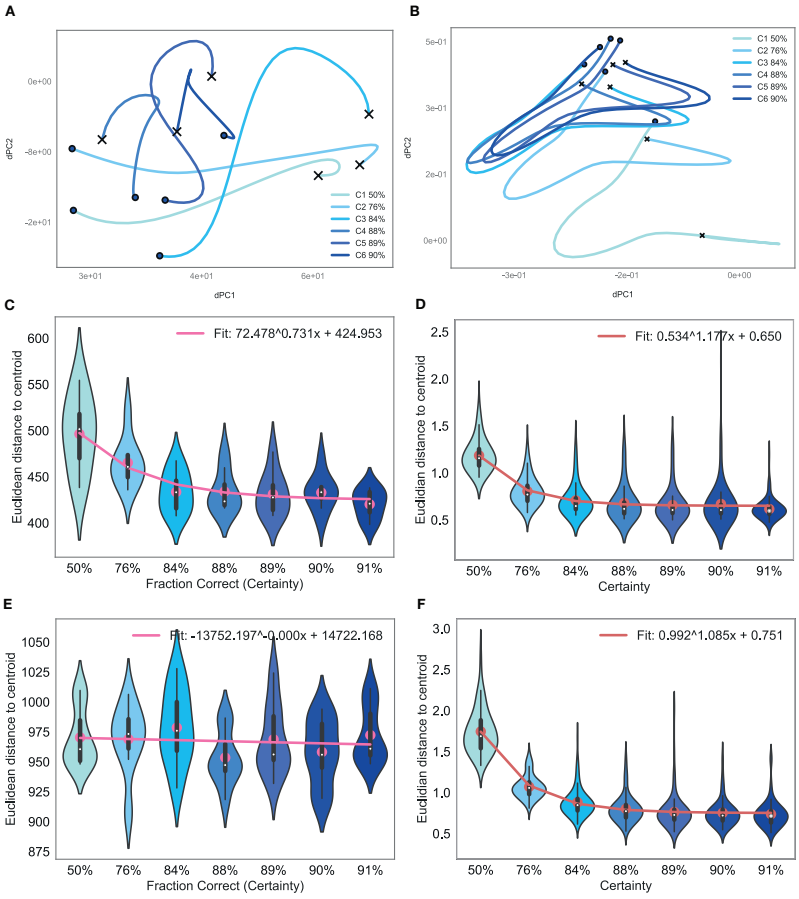


Fig 3.10. Shaping of latent movement sequence representations with learning. A-B Latent trajectory for a single movement sequence (obtained through dPCA) depicted in two dimensional latent space for increasing certainty levels (from low, in light blue, to high in dark blue). Neural recordings from dSTR (A) and STR model network (B). **C-F** Euclidean distance between a sequence's centroid (multi-dimensional mean across time) and each point in time for all sequences in the test set (averaged across time points), as a function of increasing certainty (fraction correct). Neural recordings from dSTR (C) and STR model network (D). Neural recordings from lPFC (E) and PFC model networks (F).

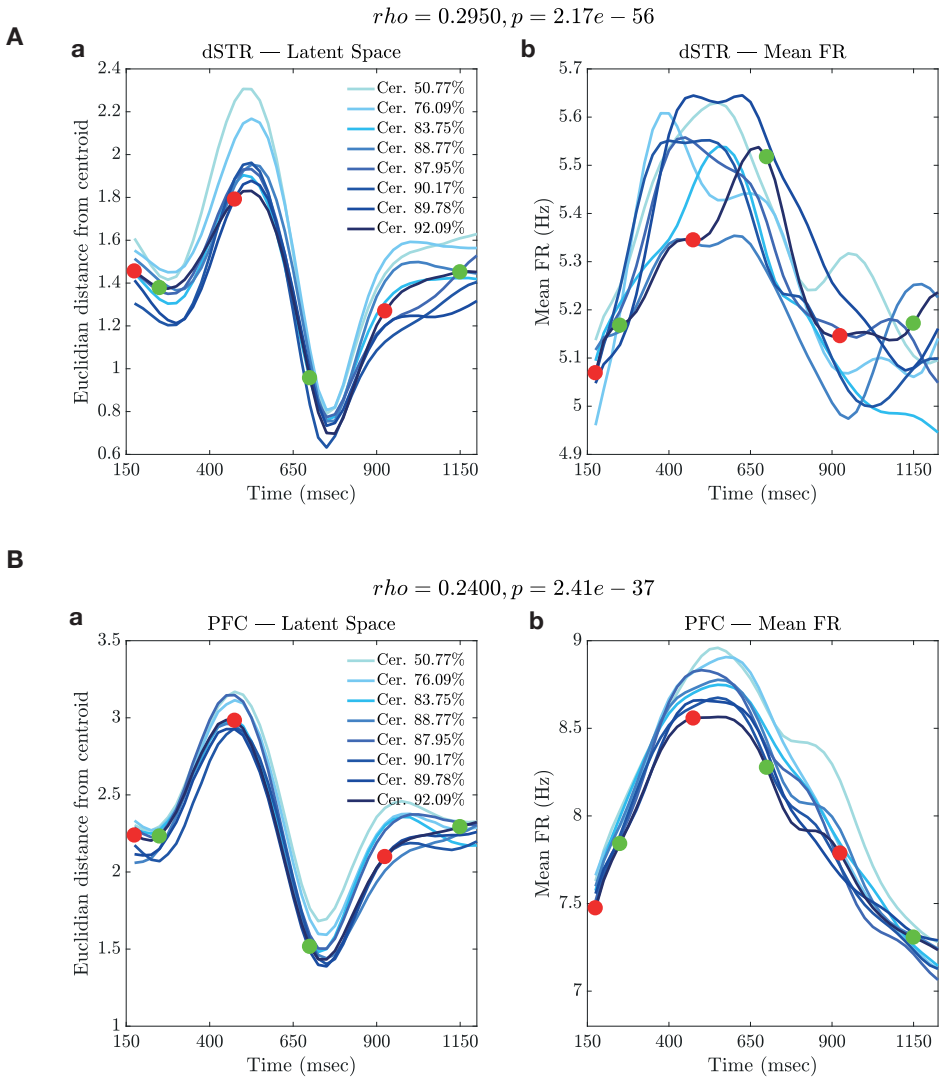


Fig 3.11. Relating the change in latent representations with learning to changes in mean firing rate. **A** Data from neural recordings in dSTR. (a) The first principal component (dPCA) of neural population activity for the same sequence as in (Fig. 3.10A) plotted across time, for increasing certainty levels (from light to dark blue). (b) Mean firing rate (FR) in Hz over time for increasing certainty levels. FR is averaged over all units in the recording. The correlation between the change in latent trajectories and mean FR with learning is $\rho = 0.2950\,(p = 2.17e-56)$. **B** Data from neural recordings in PFC. (a),(b) as in A. The correlation between the change in latent trajectories and mean FR with learning is $\rho = 0.2400\,(p = 2.41e-37)$. Dots indicate whether the period immediately to the right of the dots was a hold period (red dot) or a move period (green dot).

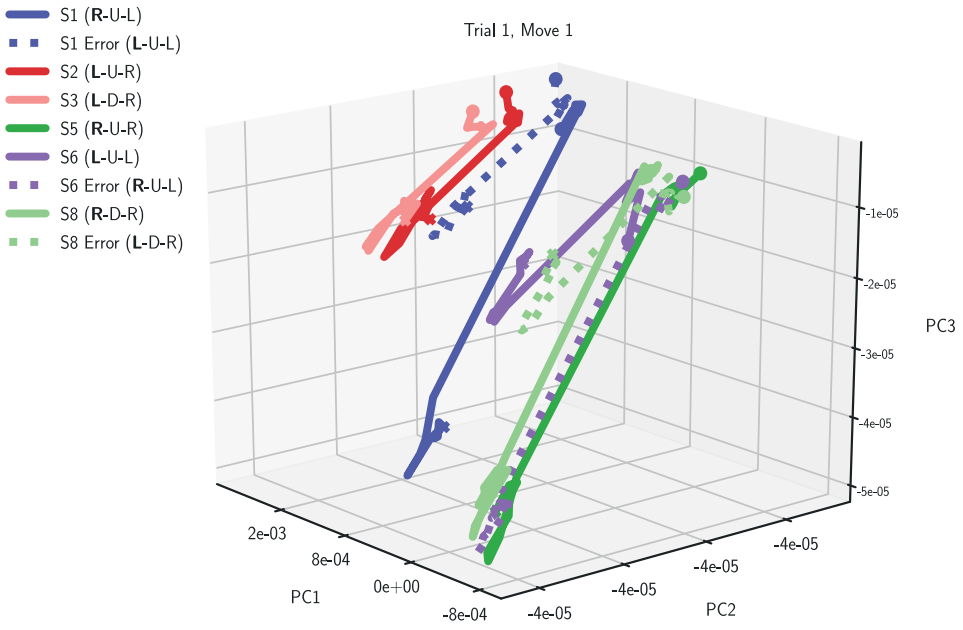


Fig 3.12. Latent representations of correct and wrong movements in the STR model network. Latent trajectories of single movements (obtained through dPCA) from the first trial after a sequence switch imaged in three dimensional latent space. Representations for correctly executed movement representations (solid lines) are depicted alongside wrongly executed ones (dotted lines).

To examine this effect in further detail, We computed the potential surface for various points as learning progressed (Fig. 3.9), as done previously for the PFC model. In the STR model network we also observed energy minima for different sequences becoming increasingly well separated as learning progressed (Fig. 3.9A-C). Along with this separation, the ridge in the joint potential surface between different sequence minima heightened with increasing certainty during learning. To quantify this effect, we determined the minimal path length between the various sequence's fixed point regions as learning progressed, as before. We observed that minimal path length between fixed point regions increased with certainty during learning in the STR model network (Fig. 3.9D). Altogether, we established that the energy landscape in the network changes with learning such that fixed point regions for different movement-sequences are pushed farther apart from each other.

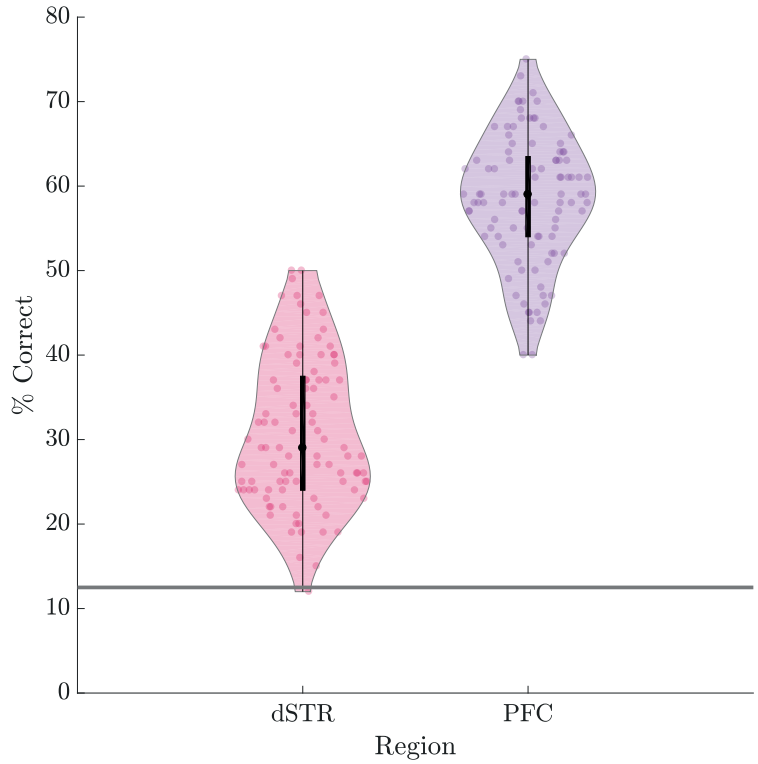


Fig 3.13. Decoding sequence identity from latent trajectories. Fraction of correctly decoded sequence identities (in %). Sequence identity was decoded by projecting the test sequence into latent space and computing the Euclidean distance to all sequence trajectories in latent space. The test sequence was then assigned the same label as the closest sequence trajectory. Results are averaged across all trials. Decoding performance was at 28% for dSTR (left) and at 59% for PFC latents (right; leave-one-out cross validation). Chance performance was at 12.5%.

We also examined how the shape of a particular movement-sequence representation in latent space changes with learning in the two regions (Fig. 3.10). We observed that trajectories became more compact with learning (from lighter to darker blue) as could be seen in neural recordings from dSTR (Fig. 3.10A) and in the striatal model network (Fig. 3.10B). To quantify this effect, we computed the Euclidean distance of a particular sequence to its centroid for increasing certainty levels (see Methods and Materials) and found this measure to be decreasing with learning in both striatal recordings (Fig. 3.10C) and in the striatal model (Fig. 3.10D). The

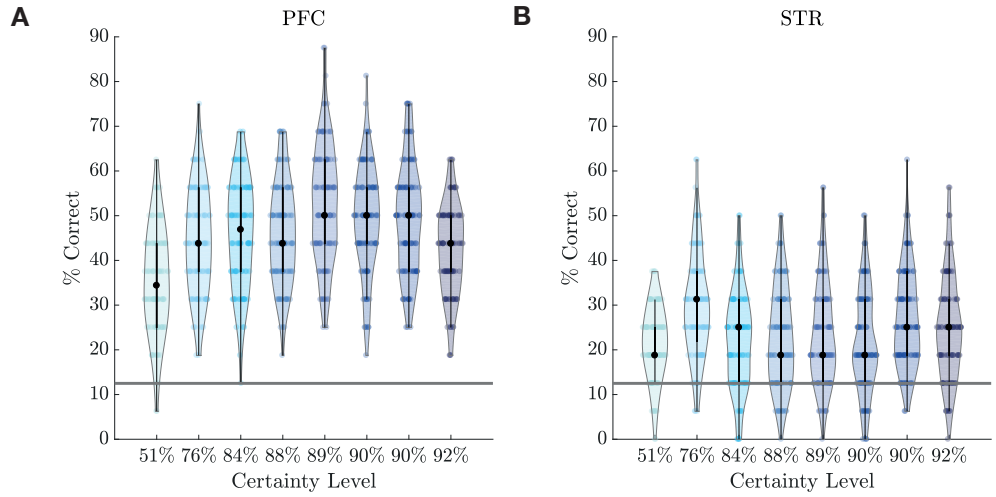


Fig 3.14. Decoding sequence identity from latent trajectories as a function of certainty level. Fraction of correctly decoded sequence identities (in %). Sequence identity was decoded by projecting the test sequence into latent space and computing the Euclidean distance to all sequence trajectories in latent space. The test sequence was then assigned the same label as the closest sequence trajectory. Results are averaged across all trials. Decoding results per certainty level for PFC (**A**) and for STR (**B**) (leave-one-out cross validation). Chance performance was at 12.5%.

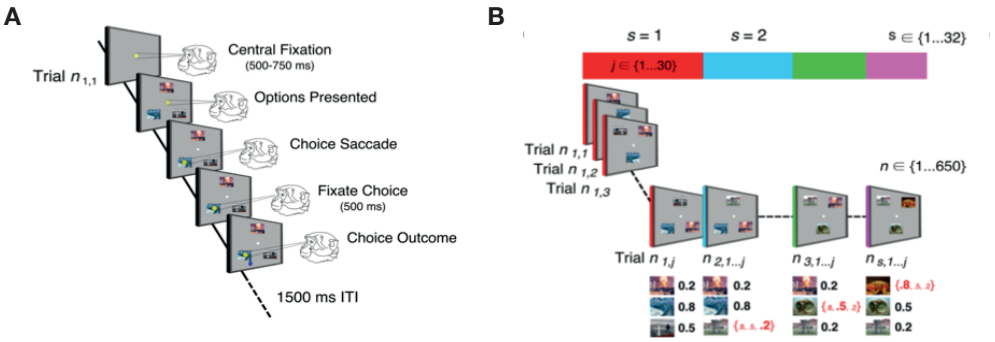


Fig 3.15. Task structure of three-armed bandit task. **A** Event structure of an individual trial in the three-armed bandit task where animals indicated their choice by saccading to a single choice option. **B** Each block of trials, n, began with the presentation of 3 novel images that served as choice options. A particular set, s, of choice options was repeatedly presented for a number, j, of trials. After $10-30$ trials, one of the three existing choice options was randomly selected and replaced with a novel image. Adapted from [Costa et al., 2019]

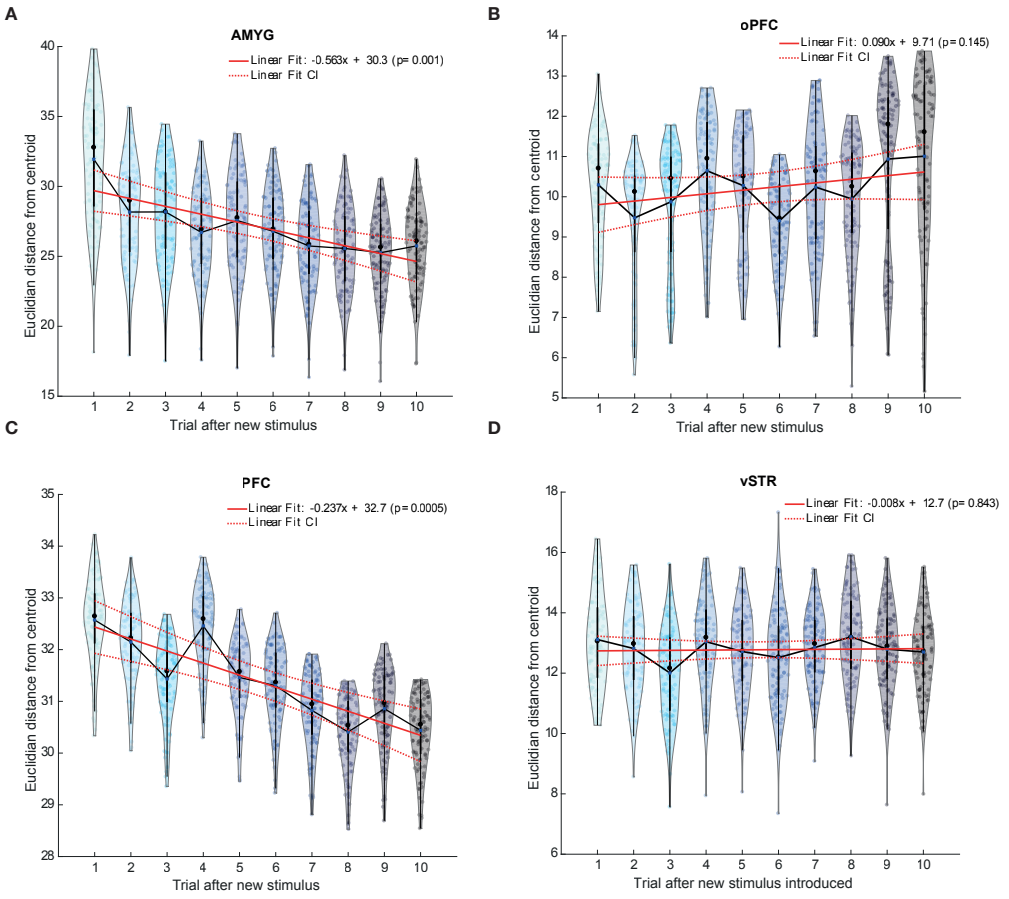


Fig 3.16. Shaping of latent image set representations in the three-armed bandit task with learning. Neural responses to each set of images, s, (Fig. 3.15) were projected into latent space using the first three encoding vectors obtained through dPCA. The Euclidean distance was computed between a latent representation's centroid (multi-dimensional mean across time) and each point in time. This measure was computed for every trial $j \in 1...10$ after a novel image was introduced, for all trial blocks, n, and averaged across time. **A** Recordings from the amygdala. The linear fit showed a decreasing trend over trials 1-10 ($p = 0.001$). **B** Recordings from oPFC. The linear fit was not significant. **C** Recordings from PFC. The linear fit showed a decreasing trend ($p = 0.0005$. **D** Recordings from vSTR. The linear fit was not significant.

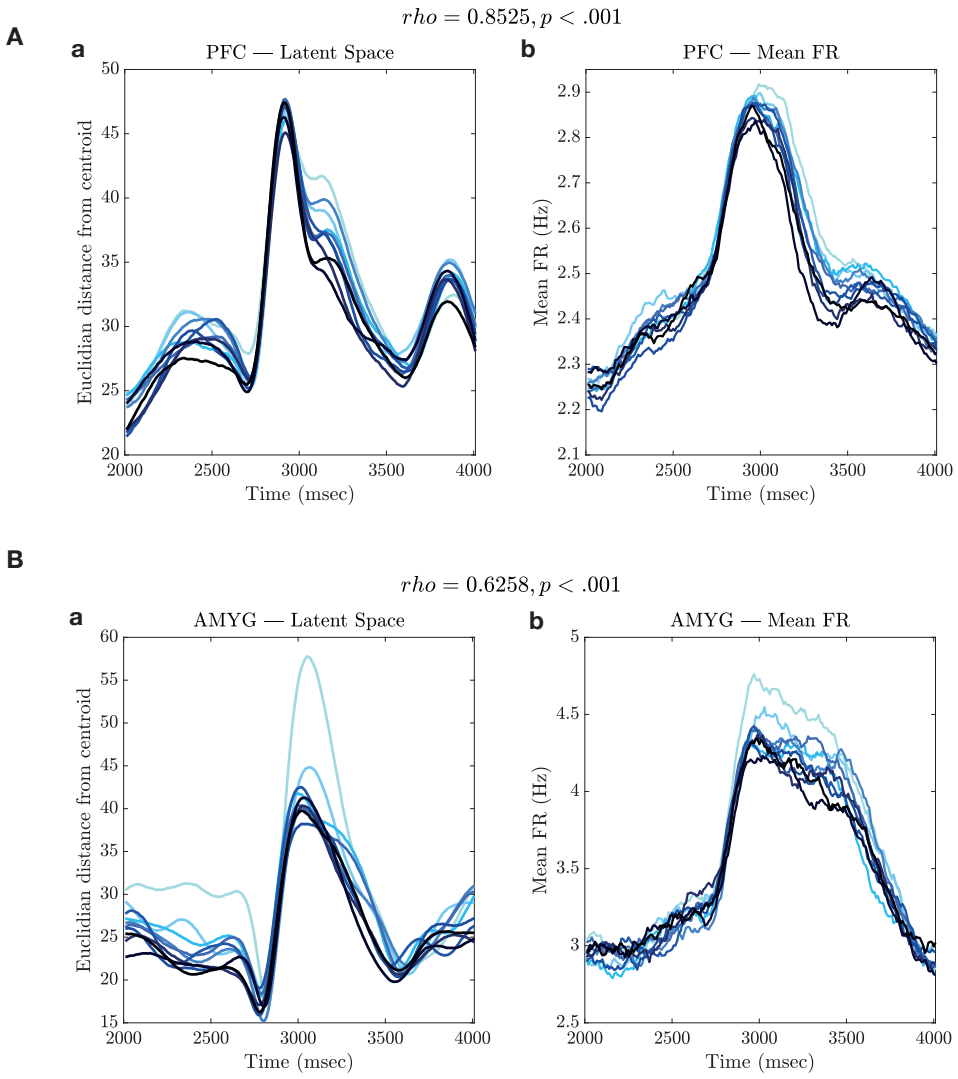


Fig 3.17. Relating the change in latent representations in the three-armed band task to changes in mean firing rate. **A** Recordings from PFC (a) The first principal component (dPCA) of neural population activity plotted across time, for trials $j \in 1...10$ (light blue to dark blue) after a novel image was introduced. Trajectories are averaged across all blocks, n. (b) Mean firing rate (FR) in Hz over time for trials $j \in 1...10$. FR is averaged over all units in the recording. The correlation between changes in latent trajectories and mean FR over trials was $\rho = 0.8525\,(p < 0.001)$. **B** Recordings from amygdala. (a),(b) computed as in A. The correlation between changes in latent trajectories and mean FR over trials was $\rho = 0.6258\,(p < 0.001)$.

decreasing trend in prefrontal recordings was not significant (Fig. 3.10E), while our PFC model network showed a more clear decreasing trend (Fig. 3.10F).

In order to examine this effect in further depth, We plotted the evolution of the first principal component across time for increasing certainty levels and compared it to changes in mean firing rate across the recorded population (Fig. 3.11). In data from dSTR (Fig. 3.11A) we observed a correlation of $r = 0.295$ between changes in the first principal component and changes in the mean firing rate of the recorded population of neurons ($p = 2.17e-56$). In PFC (Fig. 3.11B) we observed a correlation of $r = 0.240$ between the two ($p = 2.41e-37$). Thus, changes in the state space trajectories with learning are partially related to changes in mean firing rate across the population, but cannot be fully accounted for by them.

We further compared latent representations for correct and wrong movements in the prefrontal model network (Fig. 3.12). We imaged representations of the first movement during the first trial in the block (when error rate is highest) in 3-dimensional dPCA-derived latent space for a few different sample sequences. We found that when the wrong movement was executed (i.e. S1 Error, dotted blue line), the representation moved away from that of the correct movement (right move, R in bold, for sequence S1, solid blue line) and closer to that of sequences which shared the same executed movement (left move, L in bold, for sequence S2, solid red line, and sequence S3, solid rose line). Similarly, the movement representation for the wrong move in sequence 8 (S8, dotted light green line) moved away from the correct move trajectory (solid light green line) and closer to the movement representation of sequence 6 (solid magenta line) which shared the same movement direction.

In order to quantify the amount of information present in latent representations in the two regions, we decoded sequence identity from latent trajectories (Fig. 3.13).

This was done by projecting a particular test sequence into latent dPC-space and computing the Euclidean distance to all of the eight sequence's trajectories in latent space. The test sequence was then assigned the same label as the trajectory to which it was closest. We obtained a decoding performance of 28% correct for dSTR and 59% for PFC (leave-one-out cross validated), with chance performance at 12.5% (1/8 sequences). We also broke this down further by certainty level (Fig. 3.14). Sequence identity was decoded the same way, but separately for every certainty level. Decoding performance was again better in PFC than in dSTR across all certainty levels. In PFC (Fig. 3.14A), decoding performance increased with certainty.

Finally, we examined changes in latent representation with learning in another dataset (Fig. 3.14, [Costa et al., 2019]). Animals chose one of three choice options in the three-armed bandit task by saccading to the target (Fig. 3.14A). All three choice options were updated on each block of trials and a particular set of choice options was repeated for a certain number, j, of trials (Fig. 3.14B); after $j \in 10...30$ trials, one of the three choice options was randomly selected and replaced with a novel option. Neural recordings were obtained from four regions including amygdala, orbitofrontal prefrontal cortex (oPFC), PFC and ventral striatum (vSTR) (see [Costa et al., 2019] for further details).

Here, we examined neural population responses after a novel option was introduced (Fig. 3.15). We projected neural population responses onto the first three encoding vectors obtained through dPCA (as done previously in (Fig. 3.10)). We then calculated the Euclidean distance between a latent trajectories' centroid and each point in time. We obtained this measure for all trial blocks and plotted it averaged across time for trials $j \in 1...10$ after a novel option was introduced. We observed a decreasing trend in this measure in recordings from amygdala (Fig. 3.10A; $p = 0.001$) and PFC (Fig. 3.10C;$p = 0.0005$), while no significant trend was ob-

served in oPFC (Fig. 3.10B) or vSTR (Fig. 3.10D). Thus, latent representations became more compact with learning (of the value of the new choice option) in amygdala and PFC, reproducing the same effect we observed earlier (Fig. 3.10) in another dataset.

We again related changes in these latent representations with learning to changes in mean firing rate during the three-armed band task (Fig. 3.16). We observed a high correlation between changes in the first principal component of population activity and mean firing rate with learning across the first ten trials. We observed a correlation of $r = 0.853$ ($p < 0.001$) in PFC (Fig. 3.16A), and $r = 0.626$ ($p < 0.001$) in amygdala (Fig. 3.16B). This suggests compactification in latent trajectories may be driven by changes in mean firing rate.

3.4 Discussion

The results reveal how learning shapes neurophysiological responses in the corticostriatal system in the brain. Examining data from recordings in the dlPFC-dSTR circuit during an oculomotor sequence learning task, we found that learning increased the distance between action-sequence representations in activity space. To examine this further, we built a model of the corticostriatal system composed of a striatal network encoding action values and a prefrontal network selecting actions. This model was able to autonomously perform the task, matching the animals' behavioral accuracy and closely approximating the representational structure present in neural recordings. Our model revealed that the gradient landscape in the network is shaped during learning such that fixed point regions corresponding to different action sequences are pushed farther apart from each other. This makes it more likely the network generates the correct action. Altogether, this work shows how learning is expressed in network-level dynamics.

Our findings lend support to the view that task-learning is best understood as a network-level phenomenon [Stokes, 2015, Masse et al., 2020]. Our findings suggest that tasks are encoded in network dynamics (as opposed to single unit dynamics), and that learning shapes the dynamical landscape in the network (as opposed to selectively modulating single-unit responses, for example through persistent activity [Masse et al., 2020]). We would thus predict that task learning should be more robust to single unit perturbations, whereas interfering with network-level dynamics could significantly interfere with task learning. This could be tested, for instance, by identifying patterns of stimulation that most effectively push neural activity along the principle axes of the task space [Sadtler et al., 2014, Bashivan et al., 2019], and then interfering (e.g. optogenetically) during task learning.

Our model and results agree with previous findings from the corticostriatal system [Pasupathy and Miller, 2005, Samejima et al., 2005, Seo et al., 2012]. We found that task representations in neural recordings from lPFC were organized by visual hemifield, unlike dSTR representations, which showed no particular organizational structure. lPFC representations were also found to be less malleable with learning; the increase in distance between fixed point regions was smaller than in the dSTR. This agrees with previous findings that show less units in lPFC coding for a learning signal than in dSTR ([Seo et al., 2012]). Seo et al found less units in lPFC that coded for a reinforcement learning (RL) variable and more that encoded sequence information. Conversely, more units in dSTR showed a significant effect for RL and less units showed an effect for sequence information.

We also found that striatal representations became more compact with learning, both in recordings and in our model. This happened at the same time as sequence-specific representations spread further apart from each other in activity space. This

effect could be partially driven by changes in the mean and variance of the neural population firing rate with learning, as well as by changes in higher order statistics. A response decrease as rewards become more predictable fits well within a predictive coding framework [Keller and Mrsic-Flogel, 2018]. Previously it was found that the Fano factor, a measure of variability (variance of spike count divided by its mean), decreased in prefrontal neurons with learning [Qi and Constantinidis, 2012]. Also, changes in firing patterns within the neural population with learning - such as changes in synchronous firing [Baeg et al., 2007] - might be responsible for the changes in latent representation observed here.

Representations in recordings from lPFC were not found to become significantly more compact with learning, unlike in dSTR. This could be a consequence of less units in lPFC encoding a learning signal [Seo et al., 2012]. In our model we did not see a difference in how compact the representations became with learning in the two networks. This could be amended by additional processing steps which transform the output of the striatal network before it reaches the prefrontal network, which were not included in our model.

Our corticostriatal model system respects neuroscientific evidence that implicates the striatum in action value representation [Pasupathy and Miller, 2005, Samejima et al., 2005, Averbeck and Costa, 2017] and the PFC in action selection [Averbeck et al., 2006]. The methods used to train this system, however, are not biologically validated, similar to previous approaches [Mante et al., 2013, Yamins et al., 2014, Chaisangmongkon et al., 2017, Yang et al., 2018]. Alternatively, one could implement this system with spiking networks [Nicola and Clopath, 2017]. Another approach may be to use a reinforcement learning paradigm, rather than the gradient of an error signal, to train the network system [Song et al., 2017]. However, the latent dynamics underlying task learning uncovered here evolve over a longer

timescale than the learning dynamics underlying network training, so findings are likely not impacted by different training protocols.

Our findings are in line with previous work showing that activity in a large population of neurons is confined to a lower dimensional manifold [Mante et al., 2013, Shenoy et al., 2013, Chaisangmongkon et al., 2017, Remington et al., 2018, Wang et al., 2018, Gallego et al., 2018]. We found striatal and prefrontal responses encoded the sequence task on a low-dimensional manifold. Previous work was confined to investigating dynamics on this manifold once the task was acquired (and the manifold fixed). We investigated how learning shaped the manifold during task acquisition, and found that learning is expressed in dynamics acting upon that manifold. These dynamics re-shaped the manifold in such a way that it became less likely for the network to commit errors.

3.5 Conclusion

We used a model of the corticostriatal system together with recordings from dlPFC-dSTR circuit in macaques during an oculomotor sequence learning task to investigate how learning shapes neural responses at the population level. The corticostriatal model was able to autonomously perform the task and to approximate neural task representations. Probing the model, we found that learning shapes latent representations such that it becomes less likely to commit errors as learning increases; the potential surface is reshaped with learning such that fixed point regions representing different action choices move farther apart from each other. All in all, this work shows how neural circuit dynamics in the corticostriatal system drive task learning and argues that learning is best understood as a network-level phenomenon. It predicts that task learning should be most affected by disruptions of dynamics at the network- rather than the single-unit level.

3.6 Acknowledgments

This work was supported by a Wellcome Trust NIH-fellowship (C.D.M., S.S. & B.B.A.) and by NIH grant ZIA MH002928-01 (B.B.A.).

3.7 Competing Interests

The authors report no conflicts of interest.

Chapter 4

General Conclusions

Predictive coding theory is a framework that considers the brain a predictive organ. It is unclear, however, if and how predictive coding is implemented in the brain. This work examined top-down signals in the brain in order to reveal if and how they implement this (predictive) code.

According to the theory, predictions are conveyed top-down to lower-order regions and prediction errors, arising from the mismatch between predictions and new sensory information, are conveyed bottom-up, to higher-order regions where predictions are generated. While prediction errors are accumulated linearly, predictions have a nonlinear effect on incoming information. These differences imply that nonlinear, cross-frequency interactions should predominate in the top-down direction. In order to test these predictions, cross-frequency interactions need to be computed in a *directed* manner. Previous approaches, however, assessed directional influences and cross-frequency interactions separately, and there is no straightforward approach to computing them in an integrated manner. In Ch. 2 of this work, a new method was developed to accurately and efficiently compute directed cross-frequency coupling. This method was used to assess differences between top-down and bottom-up cross-frequency interactions across the auditory hierarchy of macaques.

This revealed distinct, signature patterns for bottom-up and top-down information processing among cross-frequency interactions. Low-to-low frequency interactions (θ, α, β) and low frequency to high-γ interactions were more pronounced top-down, while low frequency to low-γ interactions were more pronounced bottom-up. These patterns were preserved across coupling types (phase-amplitude and amplitude-amplitude couplings) and across stimulus types (natural and synthetic stimuli). This suggest they are a general feature of auditory processing. The findings are largely consistent with the predictive coding framework, but suggest nonlinear interactions may also play a role in the bottom-up (feedforward) message passing. It specifically predicts an important role for the regulation of spiking activity in higher-order regions through modulation of low-γ rhythms.

The predictions that are computed according to predictive coding theory need to be maintained and updated as the agent learns. It is currently unclear, however, how this learning is expressed at the neural circuit level. In Ch. 3 of this work, a recurrent neural network model of the corticostriatal system together with recordings from the dlPFC-dSTR circuit in macaques were employed together to reveal how learning shapes population-level representations.

The corticostriatal model learned to perform the decision-making task autonomously. It learned to update value representations based on rewards and predict the next action to perform. In doing so, the model was able to approximate the neural data well. Task representations were found to be pushed farther apart in state space with increasing certainty during learning, both in the model and in the neural data.

The model further revealed that learning shapes dynamics in the network such that fixed-point regions representing different task outcomes move farther apart from each other. This makes it less likely to choose the wrong action as learning develops.

Task representations were also found to become more compact with learning. This agrees with predictive coding theory which suggests top-down responses should decrease as stimuli become more predictable. This work revealed that neural circuit dynamics serve a crucial computational function: enabling task learning. Altogether, this work argues that learning is best understood as a network-level phenomenon, and predicts that disruptions of dynamics at the network- (rather than the single-unit) level should significantly impede task learning.

Appendix A

Appendix Chapter 1

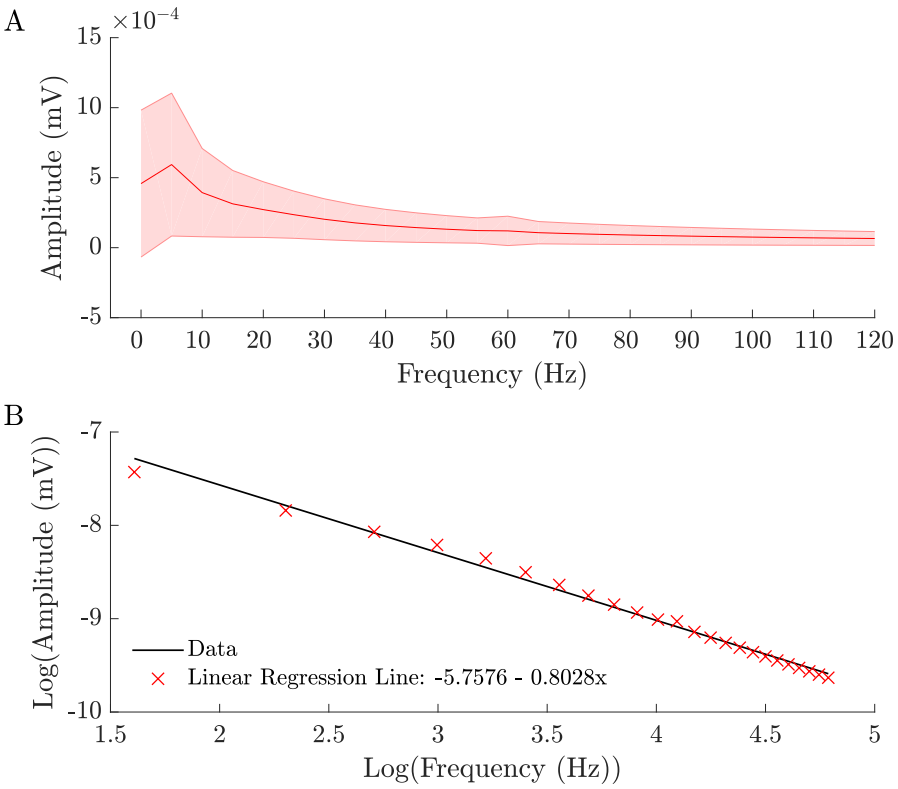


Fig S1-1. (A) Power vs. frequency of neural signal and (B) regression line fit. The amplitude of the neural recording signal peaks in the α (7.5-12.5 Hz) frequency range and shows the typical 1/f decay pattern thereafter. The signal depicted here is from stimulus onset (time=0) for presentation of the intact stimulus (VOC), averaged across all variables. Standard deviation across all variables is depicted by the red shaded region. The regression line fit to the neural signal confirms a 1/f decay pattern (B). The δ frequency band (0-2.5 Hz) was excluded from analyses (see Methods and Materials).

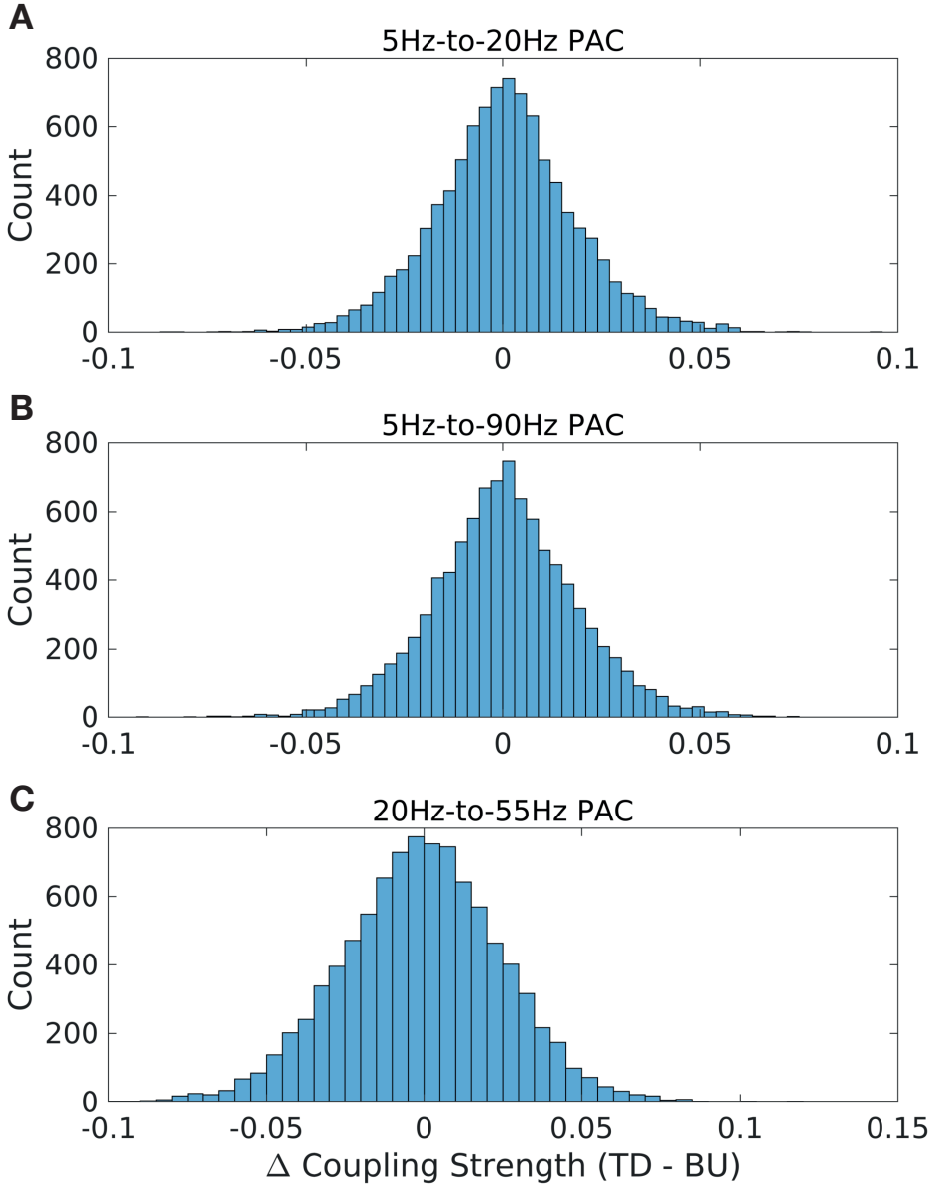


Fig S1-2. Distribution of differences in PAC coupling strength for select CFC pairs in natural stimuli (VOC) (A) θ (5Hz) to β (20Hz) PAC **(B)** θ (5Hz) to *high*γ (90Hz) PAC **(C)** β (20Hz) to *low*γ (55Hz) PAC

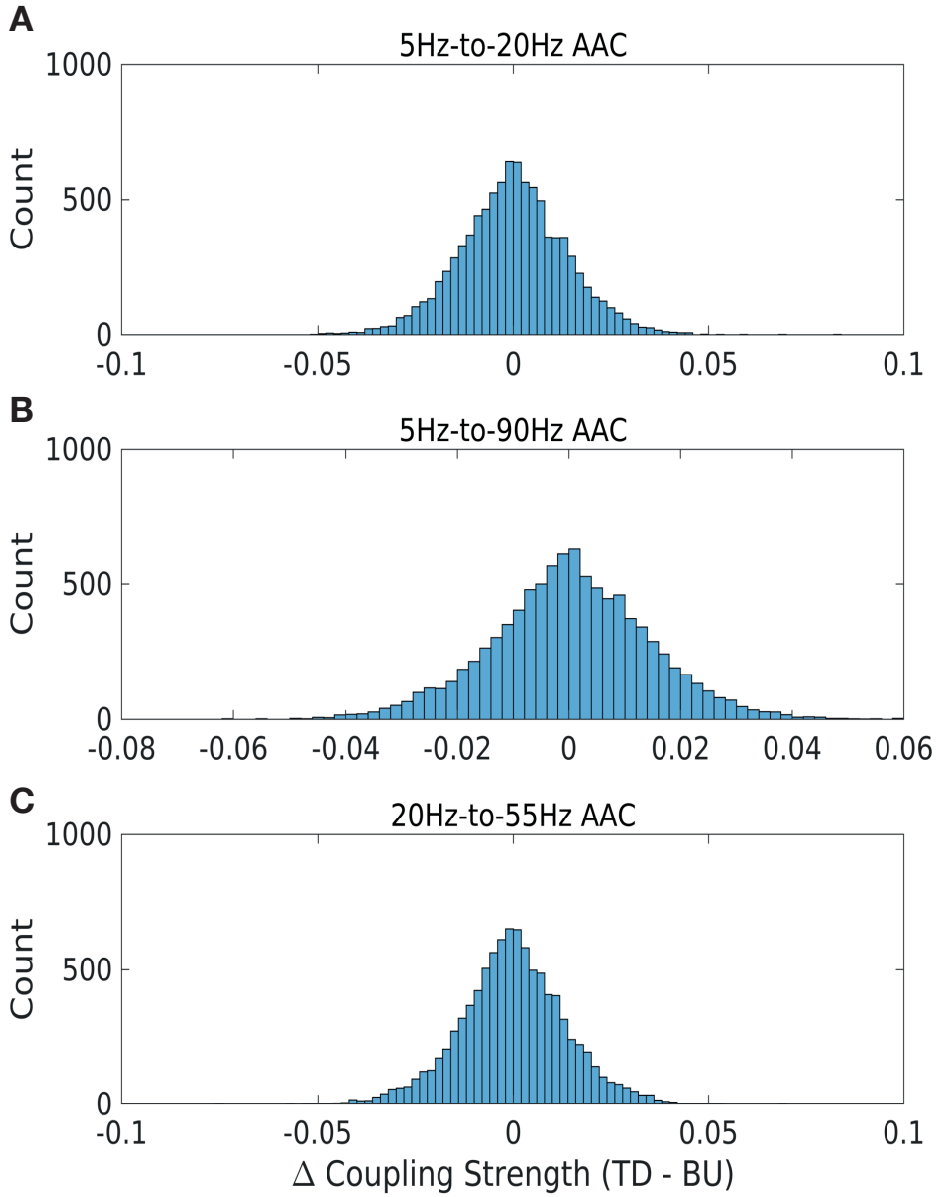


Fig S1-3. Distribution of differences in AAC coupling strength for select CFC pairs in natural stimuli (VOC) (A) θ (5Hz) to β (20Hz) PAC **(B)** θ (5Hz) to *high*γ (90Hz) PAC **(C)** β (20Hz) to *low*γ (55Hz) PAC

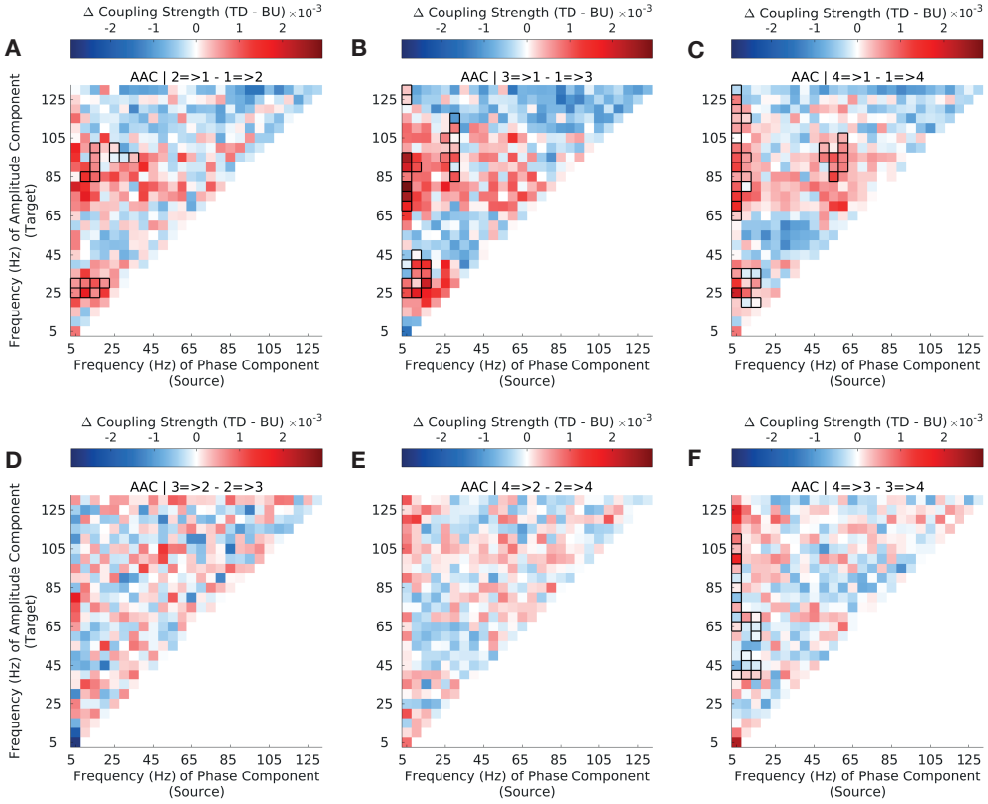


Fig S1-4. Top-down versus bottom-up amplitude-amplitude coupling strength (AAC) across the auditory hierarchy in natural stimuli (VOC) Difference in top-down and bottom-up *amplitude-amplitude* coupling (AAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals.

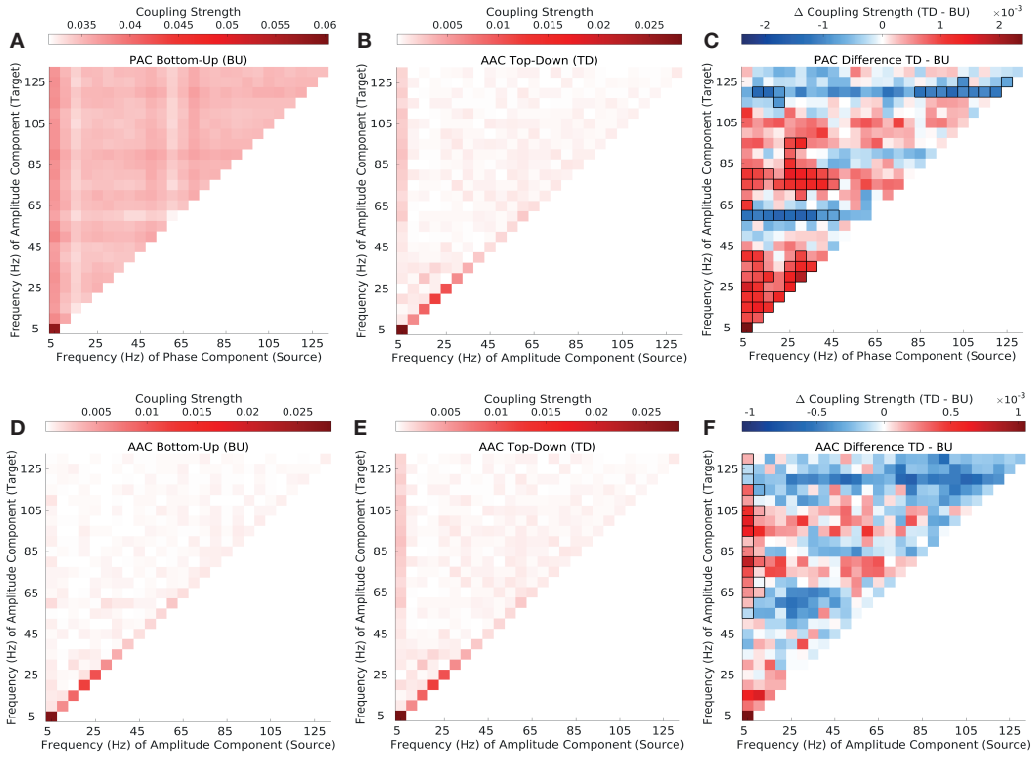


Fig S1-5. Top-down vs. bottom-up phase-amplitude and amplitude-amplitude coupling in synthetic envelope-preserved vocalizations (EPS). (A-B) Phase-amplitude coupling (PAC) strength in the top-down (A) and bottom-up direction (B). Depicted are canonical correlation-derived coupling coefficients (see Methods and Materials). **(C)** Difference in top-down and bottom-up in PAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). **(D-E)** Amplitude-amplitude coupling (AAC) strength in the top-down (D) and bottom-up direction (E). Depicted are canonical correlation-derived coupling coefficients. **(F)** Difference between top-down and bottom-up in AAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals.

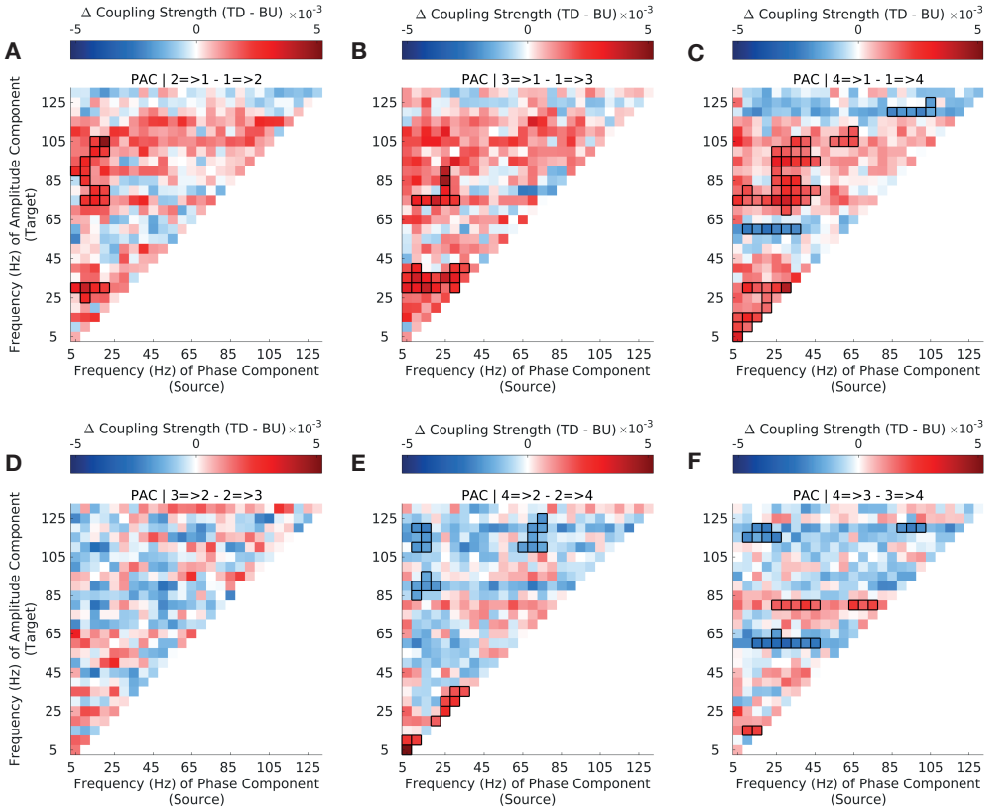


Fig S1-6. Top-down versus bottom-up phase-amplitude coupling strength (PAC) across the auditory hierarchy in synthetic envelop-preserved stimuli (EPS) Difference in top-down and bottom-up *phase-amplitude* coupling (PAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals.

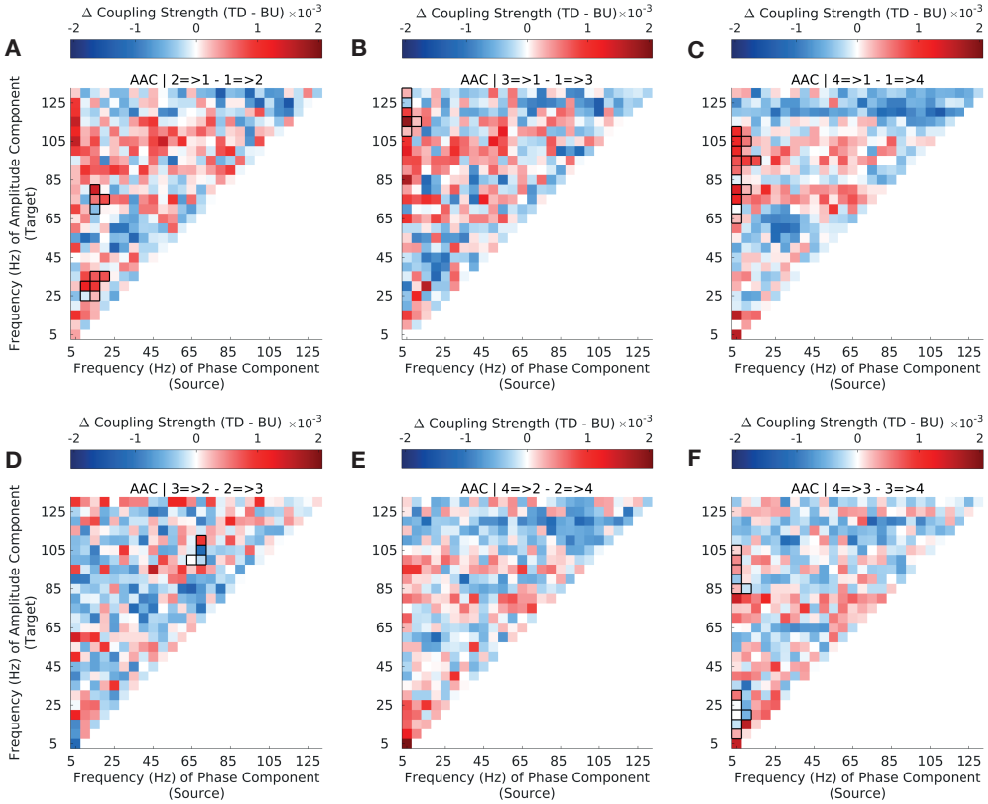


Fig S1-7. Top-down versus bottom-up amplitude-amplitude coupling strength (AAC) across the auditory hierarchy in synthetic envelope-preserved stimuli (EPS) Difference in top-down and bottom-up *amplitude-amplitude* coupling (AAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals.

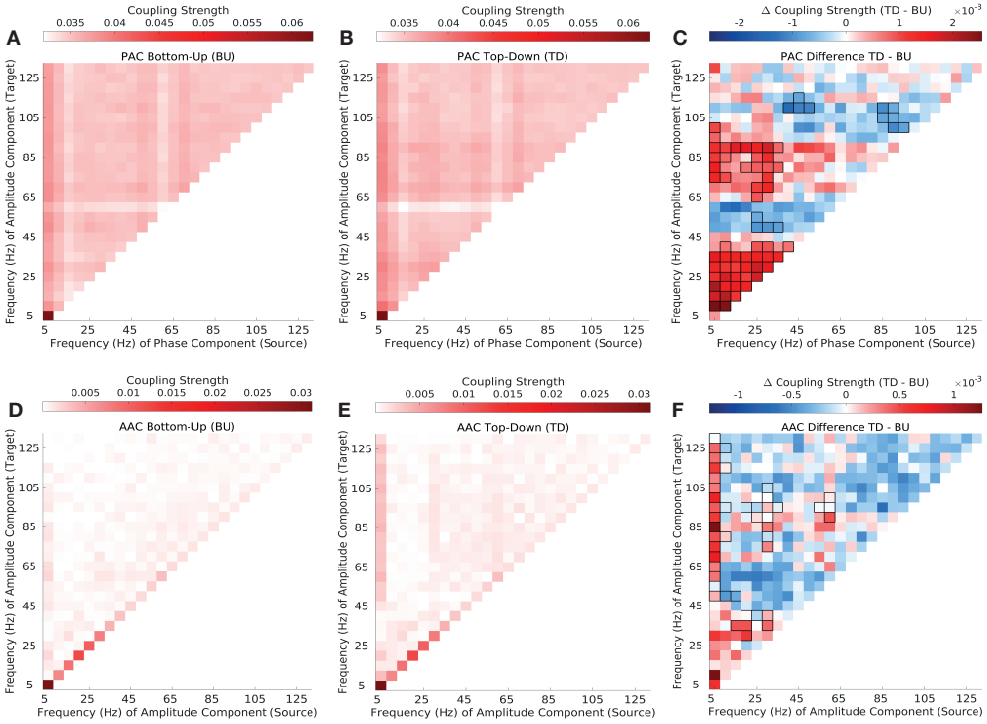


Fig S1-8. Top-down vs. bottom-up phase-amplitude and amplitude-amplitude coupling in synthetic spectrum-preserved vocalizations (SPS). (A-B) Phase-amplitude coupling (PAC) strength in the top-down (A) and bottom-up direction (B). Depicted are canonical correlation-derived coupling coefficients (see Methods and Materials). **(C)** Difference in top-down and bottom-up in PAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). **(D-E)** Amplitude-amplitude coupling (AAC) strength in the top-down (D) and bottom-up direction (E). Depicted are canonical correlation-derived coupling coefficients. **(F)** Difference between top-down and bottom-up in AAC strength. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals.

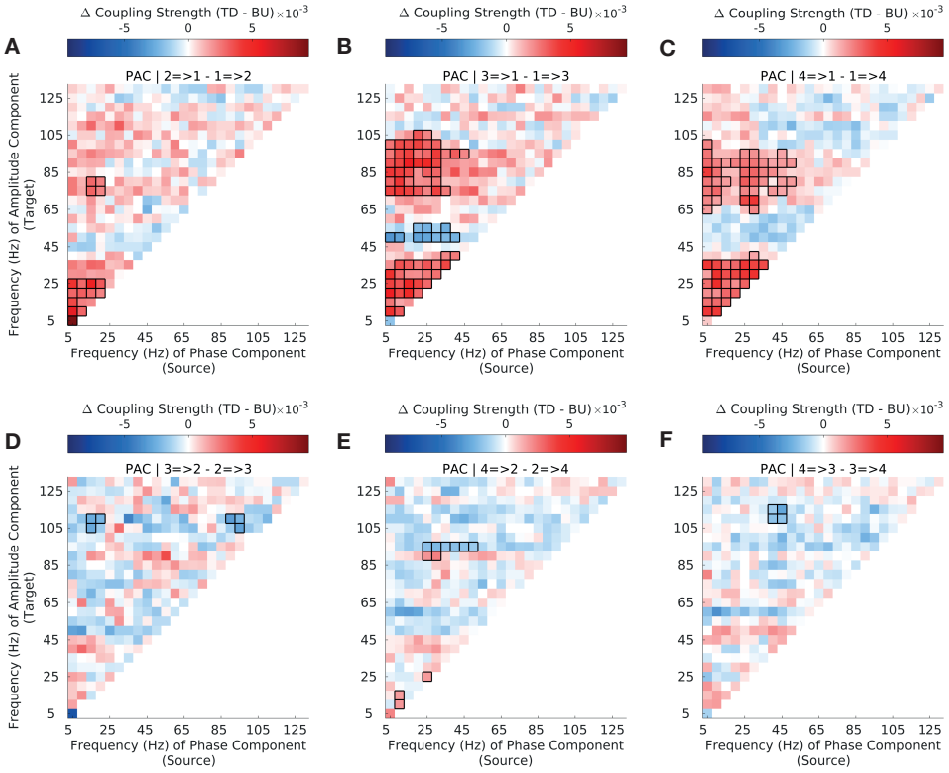


Fig S1-9. Top-down versus bottom-up phase-amplitude coupling strength (PAC) across the auditory hierarchy in synthetic spectrum-preserved stimuli (SPS) Difference in top-down and bottom-up *phase-amplitude* coupling (PAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals.

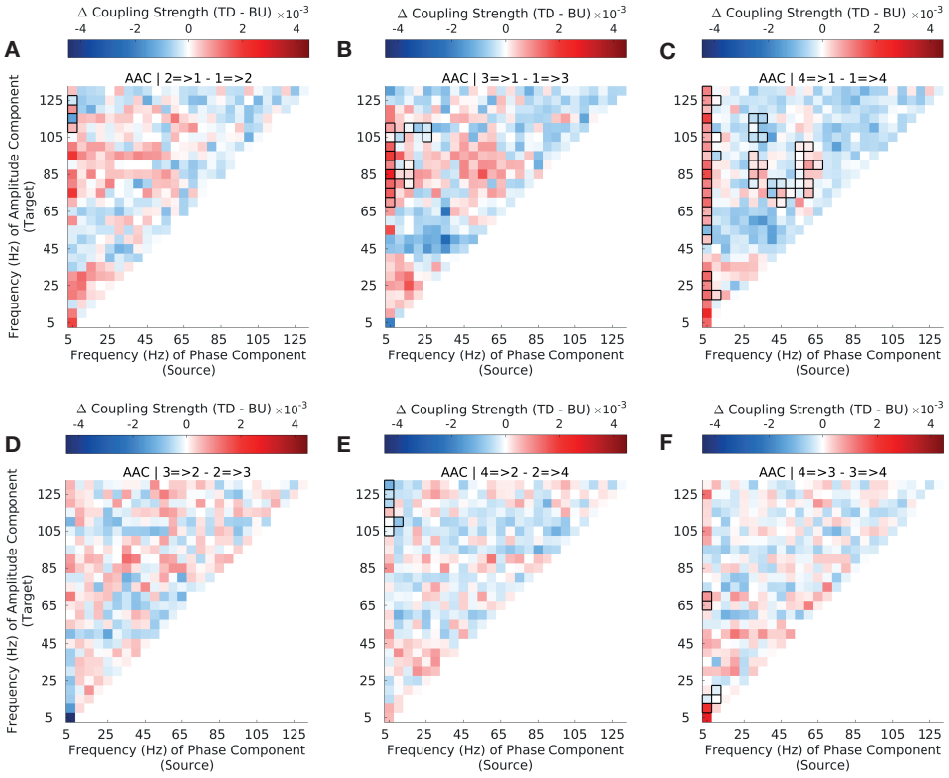


Fig S1-10. Top-down versus bottom-up amplitude-amplitude coupling strength (AAC) across the auditory hierarchy in synthetic spectrum-preserved stimuli (SPS) Difference in top-down and bottom-up *amplitude-amplitude* coupling (AAC) strength for CFC between sectors 1 (A1/ML) and 2 (R/AL) **(A)**, sectors 1 (A1/ML) and 3 (RTL) **(B)**, sectors 1 (A1/ML) and 4 (RTp) **(C)**, sectors 2 and 3 **(D)**, sectors 2 and 4 **(E)** and sectors 3 and 4 **(F)** (see Methods and Materials 3.5 for sector definitions). Interactions between sector 1 (A1/ML) and higher order sectors show strong asymmetries in bottom-up and top-down coupling strength across the frequency spectrum; interactions among higher-order sectors (S2-S4) show less widespread asymmetries. Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels and animals.

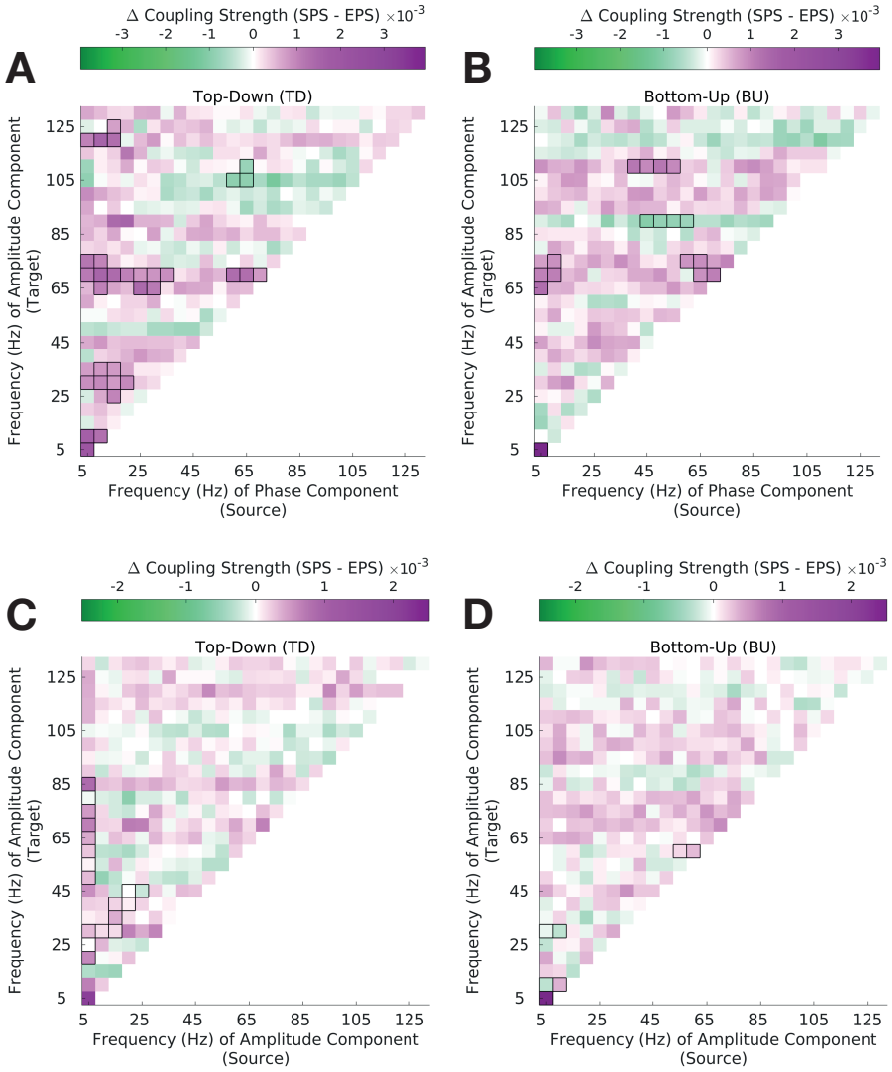


Fig S1-11. Phase-amplitude and amplitude-amplitude coupling strength in synthetic envelope-preserved sounds (SPS) compared to envelope-preserved sounds (EPS) (A-B) Difference in PAC strength between synthetic spectrum-preserved (SPS) and synthetic envelope-preserved sounds (EPS), separately in the top-down (B) and bottom-up direction (C). **(C-D)** Difference in AAC strength between SPS and EPS, separately in the top-down (C) and bottom-up direction (D). Significant differences are enclosed by black rectangles ($p \leq 0.01$, cluster-corrected). Results are depicted averaged across all channels, cross-regional pairs, and animals.

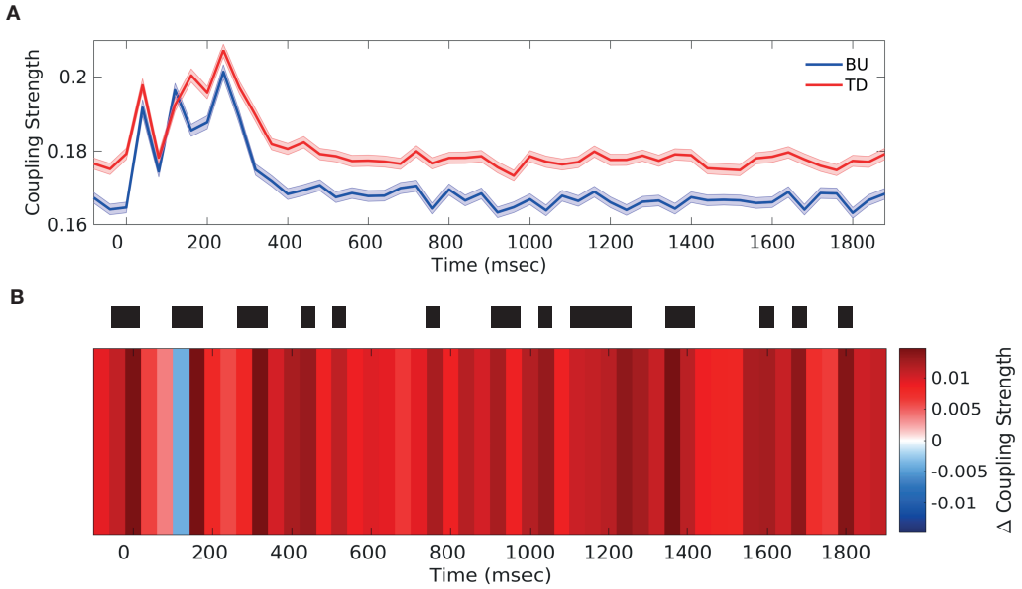


Fig S1-12. Directional low frequency phase (2.5-12.5Hz) to low frequency amplitude (12.5-37.5Hz) coupling (PAC) over time in original stimuli (VOC). A Bottom-up (blue) and top-down (red) coupling strength over the course of stimulus presentation. **B** Differences in directional coupling strength over the course of stimulus presentation. Significant differences are marked by thick black lines (FDR corrected, $q = 0.01$). Both, top-down and bottom-up coupling strength are elevated in the period immediately after stimulus presentation (0-200ms), whereas top-down coupling dominates in strength thereafter. Results were obtained using the short-time Fourier transform (FFT) and computing coupling strength per time window within the canonical correlation framework (see Methods and Materials). Results are presented averaged across all dimensions (stimuli, channels, cross-regional pairs and animals).

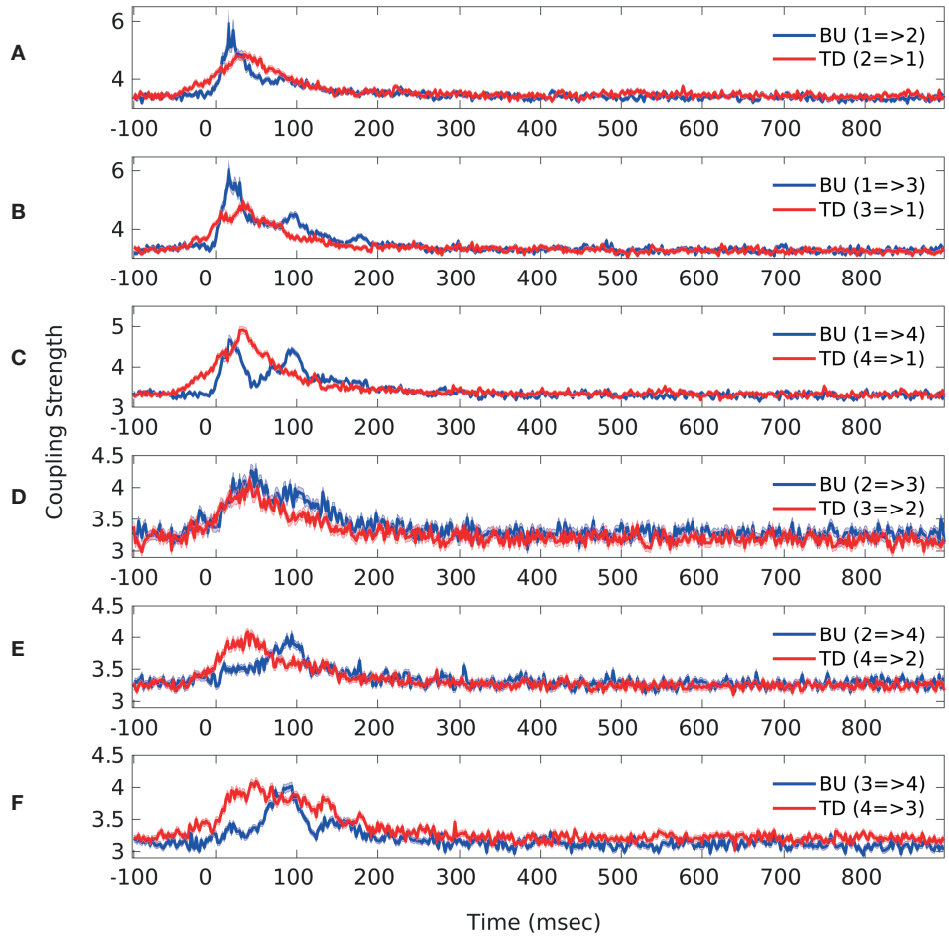


Fig S1-13. Directional low frequency phase (2.5-12.5Hz) to low frequency amplitude (12.5-37.5Hz) coupling (PAC) over time in original stimuli (VOC), broken out by regional pair. Top-down and bottom-up coupling were found to alternate in strength in the period immediately after stimulus presentation (0-200ms), across all regional pairs. **A** S1 to S2 (blue) vs. S2 to S1 (red), **B** S1 to S3 (blue) vs. S3 to S1 (red), **C** S1 to S4 (blue) vs. S4 to S1 (red), **D** S2 to S3 (blue) vs. S3 to S2 (red), **E** S2 to S4 (blue) vs. S4 to S2 (red), **F** S3 to S4 (blue) vs. S4 to S3 (red). Results were obtained using the Hilbert transform and computing coupling strength per time window within the canonical correlation framework (see Methods and Materials). Results are presented averaged across all dimensions (stimuli, channels, and animals).

Appendix B

Appendix Chapter 2

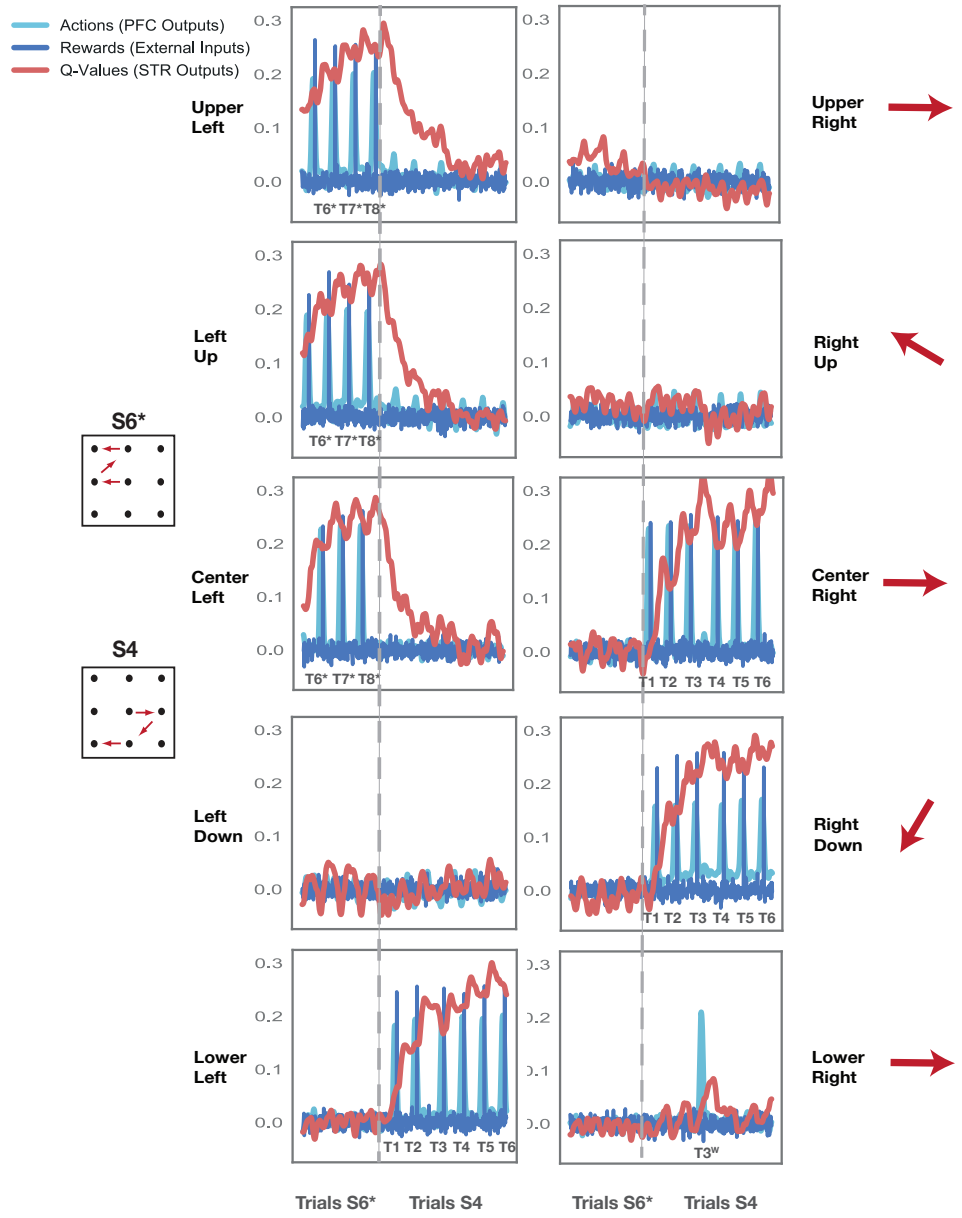


Fig S2-1. Corticostriatal model behavior on test set. Sample transition between two blocks in which the correct sequence switches (from S6 to S4; block transition indicated by vertical grey dotted line). Actions (light blue; outputs of the prefrontal (PFC) network) are plotted together with rewards (dark blue; external inputs) and action values (red; outputs of the striatal (STR) network) for all output units. Action values spike after correctly executed, rewarded actions and increase with successive correct actions. Action values for wrong, unrewarded actions (i.e. on the second trial of the third move, bottom right plot) increase briefly but decay quickly, never reaching the height associated with correct moves.

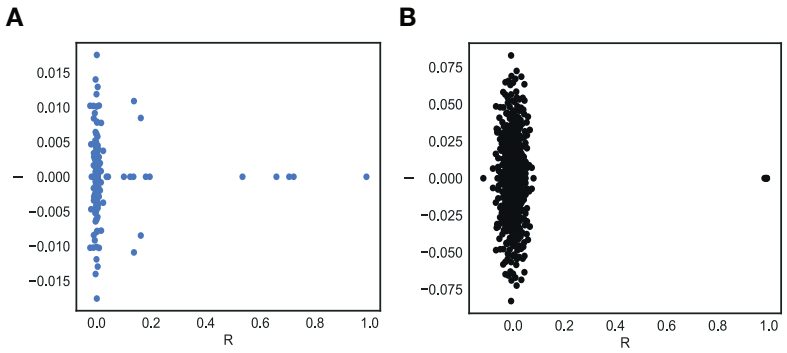


Fig S2-2. Eigenvalue spectra after training of the network system. **A** PFC model network, **B** STR model network.

Bibliography

Aitchison and Lengyel, 2017a. Aitchison, L. and Lengyel, M. (2017a). ScienceDirect With or without you: predictive coding and Bayesian inference in the brain. *Current Opinion in Neurobiology*, 46:219–227.

Aitchison and Lengyel, 2017b. Aitchison, L. and Lengyel, M. (2017b). With or without you: predictive coding and Bayesian inference in the brain. *Current Opinion in Neurobiology*, 46:219–227.

Alexander et al., 2018. Alexander, W. H., , , and Brown, J. W. (2018). Frontal cortex function as derived from hierarchical predictive coding. *Scientific Reports*, 8(1):1–11.

Amemori et al., 2011. Amemori, K.-i., Graybiel, A. M., and Gibb, L. G. (2011). Shifting responsibly: the importance of striatal modularity to reinforcement learning in uncertain environments. *Frontiers in Human Neuroscience*, 5(47):1–20.

Arnal and Giraud, 2012. Arnal, L. H. and Giraud, A.-L. (2012). Cortical oscillations and sensory predictions. *Trends in Cognitive Sciences*, 16(7):390–398.

Aru et al., 2015. Aru, J., Aru, J., Priesemann, V., Wibral, M., Lana, L., Pipa, G., Singer, W., and Vicente, R. (2015). Untangling cross-frequency coupling in neuroscience. *Current Opinion in Neurobiology*, 31:51–61.

Averbeck and Costa, 2017. Averbeck, B. B. and Costa, V. D. (2017). Motivational neural circuits underlying reinforcement learning. *Nature Neuroscience*, 20(4):505–512.

Averbeck et al., 2006. Averbeck, B. B., Sohn, J.-W., and Lee, D. (2006). Activity in prefrontal cortex during dynamic selection of action sequences. *Nature Neuroscience*, 9(2):276–282.

Baeg et al., 2007. Baeg, E. H., Kim, Y. B., Kim, J., Ghim, J. W., Kim, J. J., and Jung, M. W. (2007). Learning-Induced Enduring Changes in Functional Connectivity among Prefrontal Cortical Neurons. *Journal of Neuroscience*, 27(4):909–918.

Balaguer et al., 2016. Balaguer, J., Spiers, H., Hassabis, D., and Summerfield, C. (2016). Neural Mechanisms of Hierarchical Planning in a Virtual Subway Network. *Neuron*, 90(4):893–903.

Barak et al., 2013. Barak, O., Sussillo, D., Romo, R., Tsodyks, M., and Abbott, L. F. (2013). From fixed points to chaos: Three models of delayed discrimination. *Progress in Neurobiology*, 103:214–222.

Bashivan et al., 2019. Bashivan, P., Kar, K., and DiCarlo, J. J. (2019). Neural population control via deep image synthesis. *Science*, 364(6439):eaav9436–13.

Bastos et al., 2012a. Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012a). Canonical Microcircuits for Predictive Coding. *Neuron*, 76(4):695–711.

Bastos et al., 2012b. Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012b). Canonical Microcircuits for Predictive Coding. *Neuron*, 76(4):695–711.

Bastos et al., 2015. Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., De Weerd, P., Kennedy, H., and Fries, P. (2015). Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. *Neuron*, 85(2):390–401.

Bosman et al., 2012. Bosman, C. A., Schoffelen, J.-M., Brunet, N., Oostenveld, R., Bastos, A. M., Womelsdorf, T., Rubehn, B., Stieglitz, T., De Weerd, P., and Fries, P. (2012). Attentional Stimulus Selection through Selective Synchronization between Monkey Visual Areas. *Neuron*, 75(5):875–888.

Botvinick et al., 2019. Botvinick, M., Ritter, S., Wang, J. X., Kurth-Nelson, Z., Blundell, C., and Hassabis, D. (2019). Reinforcement Learning, Fast and Slow. *Trends in Cognitive Sciences*, 23(5):408–422.

Botvinick and Weinstein, 2014. Botvinick, M. and Weinstein, A. (2014). Model-based hierarchical reinforcement learning and human action control. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1655):20130480–20130480.

Botvinick, 2012. Botvinick, M. M. (2012). Hierarchical reinforcement learning and decision making. *Current Opinion in Neurobiology*, 22(6):956–962.

Brendel et al., 2011. Brendel, W., Romo, R., and Machens, C. K. (2011). Demixed Principal Component Analysis. *Advances in Neural Information Processing Systems*, pages 2654–2662.

Bressler and Richter, 2015. Bressler, S. L. and Richter, C. G. (2015). ScienceDirect Interareal oscillatory synchronization in top-down neocortical processing. *Current Opinion in Neurobiology*, 31:62–66.

Bruns and Eckhorn, 2004. Bruns, A. and Eckhorn, R. (2004). Task-related coupling from high- to low-frequency signals among visual cortical areas in human subdural recordings. *International Journal of Psychophysiology*, 51(2):97–116.

Buonomano and Maass, 2009. Buonomano, D. V. and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10(2):113–125.

Buzsáki et al., 2012. Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6):1–14.

Canolty et al., 2006. Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., Berger, M. S., Barbaro, N. M., and Knight, R. T. (2006). High Gamma Power Is Phase-Locked to Theta Oscillations in Human Neocortex. *Science*, 313(5793):1626–1628.

Canolty and Knight, 2010. Canolty, R. T. and Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, 14(11):506–515.

Carnevale et al., 2015. Carnevale, F., de Lafuente, V., Romo, R., Barak, O., and Parga, N. (2015). Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty. *Neuron*, 86(4):1067–1077.

Chaisangmongkon et al., 2017. Chaisangmongkon, W., Swaminathan, S. K., Freedman, D. J., and Wang, X.-J. (2017). Computing by Robust Transience: How the Fronto- Parietal Network Performs Sequential, Category- Based Decisions. *Neuron*, 93(6):1504–1517.e4.

Chao et al., 2018. Chao, Z. C., Takaura, K., Wang, L., Fujii, N., and Dehaene, S. (2018). Large-Scale Cortical Networks for Hierarchical Prediction and Prediction Error in the Primate Brain. *Neuron*, pages 1–30.

Chen et al., 2008. Chen, C. C., Kiebel, S. J., and Friston, K. J. (2008). Dynamic causal modelling of induced responses. *NeuroImage*, 41(4):1293–1312.

Cheong et al., 2006. Cheong, I., Huang, X., Bettegowda, C., Diaz, L. A., Kinzler, K. W., Zhou, S., and Vogelstein, B. (2006). A Bacterial Protein Enhances the Release and Efficacy of Liposomal Cancer Drugs. *Science*, 314(5803):1308–1311.

Clark, 2013. Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. In *Behavioral and Brain Sciences*, pages 181–204. Behavioral and Brain Sciences.

Cohen, 2008. Cohen, M. X. (2008). Assessing transient cross-frequency coupling in EEG data. *Journal of Neuroscience Methods*, 168(2):494–499.

Colgin, 2015. Colgin, L. L. (2015). Theta–gamma coupling in the entorhinal–hippocampal system. *Current Opinion in Neurobiology*, 31:45–50.

Collins and Koechlin, 2012. Collins, A. and Koechlin, E. (2012). Reasoning, Learning, and Creativity: Frontal Lobe Function and Human Decision-Making. *PLoS Biology*, 10(3):e1001293.

Corlett et al., 2009. Corlett, P. R., Frith, C. D., and Fletcher, P. C. (2009). From drugs to deprivation: a Bayesian framework for understanding models of psychosis. *Psychopharmacology*, 206(4):515–530.

Costa et al., 2019. Costa, V. D., Mitz, A. R., and Averbeck, B. B. (2019). Subcortical Substrates of Explore-Exploit Decisions in Primates. *Neuron*, 103(3):533–545.e5.

Davis et al., 2011. Davis, M. H., Ford, M. A., Kherif, F., and Johnsrude, I. S. (2011). Does Semantic Context Benefit Speech Understanding through. *Journal of Cognitive Neuroscience*, pages 1–21.

Daw et al., 2011. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., and Dolan, R. J. (2011). Model-Based Influences on Humans' Choices and Striatal Prediction Errors. *Neuron*, 69(6):1204–1215.

Doesburg et al., 2012. Doesburg, S. M., Green, J. J., McDonald, J. J., and Ward, L. M. (2012). Theta modulation of inter-regional gamma synchronization during auditory attention control. *Brain Research*, 1431(C):77–85.

Domenech and Koechlin, 2015. Domenech, P. and Koechlin, E. (2015). Executive control and decision-making in the prefrontal cortex. *Current Opinion in Behavioral Sciences*, 1:101–106.

Doya, 1910. Doya, K. (1910). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology*, 10:732–739.

Durstewitz, 2017. Durstewitz, D. (2017). Advanced Data Analysis in Neuroscience. *Springer International Publishing*.

Eliades and Wang, 2008. Eliades, S. J. and Wang, X. (2008). Neural substrates of vocalization feedback monitoring in primate auditory cortex. *Nature*, 453(7198):1102–1106.

Fiser et al., 2016. Fiser, A., Mahringer, D., Oyibo, H. K., Petersen, A. V., Leinweber, M., and Keller, G. B. (2016). Experience-dependent spatial expectations in mouse visual cortex. *Nature Neuroscience*, 19(12):1658–1664.

Fontolan et al., 2014. Fontolan, L., Morillon, B., Liegeois-Chauvel, C., and Giraud, A.-L. (2014). The contribution of frequency-specific activity to hierarchical

information processing in the human auditory cortex. *Nature Communications*, 5(1):1–10.

Frank, 2005. Frank, M. J. (2005). Dynamic Dopamine Modulation in the Basal Ganglia: A Neurocomputational Account of Cognitive Deficits in Medicated and Nonmedicated Parkinsonism. *Journal of Cognitive Neuroscience*, 17(1):51–72.

Frank et al., 2004. Frank, M. J., Seeberger, L. C., and OReilly, R. C. (2004). By Carrot or by Stick: Cognitive Reinforcement Learning in Parkinsonism. *Science*, 306(1):71–100.

Fries, 2005. Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10):474–480.

Friston, 2005. Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456):815–836.

Friston, 2008. Friston, K. (2008). Hierarchical Models in the Brain. *PLOS Computational Biology*, 4(11):e1000211–24.

Friston and Kiebel, 2009. Friston, K. and Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521):1211–1221.

Friston et al., 2003. Friston, K. J., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, 19(4):1273–1302.

Friston et al., 2017. Friston, K. J., Preller, K. H., Mathys, C., Cagnan, H., Heinzle, J., Razi, A., and Zeidman, P. (2017). Dynamic causal modelling revisited. *NeuroImage*, pages 1–0.

Fukushima et al., 2012. Fukushima, M., Saunders, R. C., Leopold, D. A., Mishkin, M., and Averbeck, B. B. (2012). Spontaneous high-gamma band activity reflects functional organization of auditory cortex in the awake macaque. *Neuron*, 74(5):899–910.

Fukushima et al., 2014. Fukushima, M., Saunders, R. C., Leopold, D. A., Mishkin, M., and Averbeck, B. B. (2014). Differential Coding of Conspecific Vocalizations in the Ventral Auditory Cortical Stream. *Journal of Neuroscience*, 34(13):4665–4676.

Furl et al., 2014. Furl, N., Coppola, R., Averbeck, B. B., and Weinberger, D. R. (2014). Cross-Frequency Power Coupling Between Hierarchically Organized Face-Selective Areas. *Cerebral cortex (New York, N.Y. : 1991)*, 24(9):2409–2420.

Gagnepain et al., 2012. Gagnepain, P., Henson, R. N., and Davis, M. H. (2012). Temporal Predictive Codes for Spoken Words in Auditory Cortex. *Current Biology*, 22(7):615–621.

Gallego et al., 2017. Gallego, J. A., Perich, M. G., Miller, L. E., and Solla, S. A. (2017). Neural Manifolds for the Control of Movement. *Neuron*, 94(5):978–984.

Gallego et al., 2018. Gallego, J. A., Perich, M. G., Naufel, S. N., Ethier, C., Solla, S. A., and Miller, L. E. (2018). Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications*, 9(1):1–13.

Ghitza et al., 2013. Ghitza, O., Giraud, A.-L., and Poeppel, D. (2013). Neuronal oscillations and speech perception: critical-band temporal envelopes are the essence. *Frontiers in Human Neuroscience*, 6:1–4.

Giraud and Poeppel, 2012. Giraud, A.-L. and Poeppel, D. (2012). Cortical oscillations and speech processing: emerging computational principles and operations. *Nature Neuroscience*, 15(4):511–517.

Griffiths et al., 2010. Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., and Tenenbaum, J. B. (2010). Probabilistic models of cognition: exploring representations and inductive biases. *Trends in Cognitive Sciences*, 14(8):357–364.

Haesemeyer et al., 2019. Haesemeyer, M., Schier, A. F., and Engert, F. (2019). Convergent Temperature Representations in Artificial and Biological Neural Networks. *Neuron*, 103(6):1123–1134.e6.

Hasson et al., 2020. Hasson, U., Nastase, S. A., and Goldstein, A. (2020). Direct Fit to Nature: An Evolutionary Perspective on Biological and Artificial Neural Networks. *Neuron*, 105(3):416–434.

Hennequin et al., 2014. Hennequin, G., Vogels, T. P., and Gerstner, W. (2014). Optimal Control of Transient Dynamics in Balanced Networks Supports Generation of Complex Movements. *Neuron*, 82(6):1394–1406.

Histed et al., 2009. Histed, M. H., Pasupathy, A., and Miller, E. K. (2009). Learning Substrates in the Primate Prefrontal Cortex and Striatum: Sustained Activity Related to Successful Actions. *Neuron*, 63(2):244–253.

Houk, 1995. Houk, J. C. (1995). A Model of How the Basal Ganglia Generate and Use Neural Signals that Predict Reinforcement. *Models of Information Processing in the Basal Ganglia MIT Press*, pages 249–274.

Hubel and Wiesel, 1959. Hubel, D. H. and Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *J. Physiol.*, pages 574–591.

Hyafil et al., 2015a. Hyafil, A., Fontolan, L., Kabdebon, C., Gutkin, B., and Giraud, A.-L. (2015a). Speech encoding by coupled cortical theta and gamma oscillations. *eLife*, 4:3958.

Hyafil et al., 2015b. Hyafil, A., Giraud, A.-L., Fontolan, L., and Gutkin, B. (2015b). Neural Cross-Frequency Coupling: Connecting Architectures, Mechanisms, and Functions. *Trends in Neurosciences*, 38(11):725–740.

Jensen and Colgin, 2007. Jensen, O. and Colgin, L. L. (2007). Cross-frequency coupling between neuronal oscillations. *Trends in Cognitive Sciences*, 11(7):267–269.

Jirsa and Mueller, 2013. Jirsa, V. and Mueller, V. (2013). Cross-frequency coupling in real and virtual brain networks. *Frontiers in Computational Neuroscience*, pages 1–25.

Johnson et al., 2018. Johnson, M., Townsend, J., Hawkins, P., and Maclaurin, D. (2018). JAX: Autograd and XLA. *Retrieved from https://github.com/google/jax*.

Keller et al., 2017. Keller, A. J., Houlton, R., Kampa, B. M., Lesica, N. A., Mrsic-Flogel, T. D., Keller, G. B., and Helmchen, F. (2017). Stimulus relevance modulates contrast adaptation in visual cortex. *eLife*, pages 1–12.

Keller and Mrsic-Flogel, 2018. Keller, G. B. and Mrsic-Flogel, T. D. (2018). Predictive Processing: A Canonical Cortical Computation. *Neuron*, 100(2):424–435.

Kellis et al., 2010. Kellis, S., Miller, K., Thomson, K., Brown, R., House, P., and Greger, B. (2010). Decoding spoken words using local field potentials recorded from the cortical surface. *Journal of Neural Engineering*, 7(5):056007–20.

Kendrick et al., 2011. Kendrick, K. M., Zhan, Y., Fischer, H., Nicol, A. U., Zhang, X., and Feng, J. (2011). Learning alters theta amplitude, theta-gamma coupling

and neuronal synchronization in inferotemporal cortex. *BMC Neuroscience*, 12(1):55.

Kikuchi et al., 2010. Kikuchi, Y., Horwitz, B., and Mishkin, M. (2010). Hierarchical auditory processing directed rostrally along the monkey's supratemporal plane. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(39):13021–13030.

Kobak et al., 2016. Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C. E., Kepecs, A., Mainen, Z. F., Qi, X.-L., Romo, R., Uchida, N., and Machens, C. K. (2016). Demixed principal component analysis of neural population data. *eLife*, e10989:1–36.

Koechlin and Summerfield, 2007. Koechlin, E. and Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. *Trends in Cognitive Sciences*, 11(6):229–235.

Kording et al., 2007. Kording, K. P., Tenenbaum, J. B., and Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience*, 10(6):779–786.

Kording and Wolpert, 2004. Kording, K. P. and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427:1–4.

Koster-Hale and Saxe, 2013. Koster-Hale, J. and Saxe, R. (2013). Theory of Mind: A Neural Prediction Problem. *Neuron*, 79(5):836–848.

Kurikawa et al., 2018. Kurikawa, T., Haga, T., Handa, T., Harukuni, R., and Fukai, T. (2018). Neuronal stability in medial frontal cortex sets individual variability in decision-making. *Nature Neuroscience*, 21(12):1–16.

Kveraga et al., 2007. Kveraga, K., Ghuman, A. S., and Bar, M. (2007). Top-down predictions in the cognitive brain. *Brain and Cognition*, 65(2):145–168.

Lakatos et al., 2016. Lakatos, P., Barczak, A., Neymotin, S. A., McGinnis, T., Ross, D., Javitt, D. C., and O'Connell, M. N. (2016). Global dynamics of selective attention and its lapses in primary auditory cortex. *Nature Neuroscience*, 19(12):1707–1717.

Lakatos et al., 2008. Lakatos, P., Karmos, G., Mehta, A. D., Ulbert, I., and Schroeder, C. E. (2008). Entrainment of Neuronal Oscillations as a Mechanism of Attentional Selection. *Science*, 320(5872):110–113.

Lakatos et al., 2005. Lakatos, P., Shah, A. S., Knuth, K. H., Ulbert, I., Karmos, G., and Schroeder, C. E. (2005). An Oscillatory Hierarchy Controlling Neuronal Excitability and Stimulus Processing in the Auditory Cortex. *Journal of Neurophysiology*, 94(3):1904–1911.

Landau and Fries, 2012. Landau, A. N. and Fries, P. (2012). Attention Samples Stimuli Rhythmically. *Current Biology*, 22(11):1000–1004.

Langdon et al., 2018. Langdon, A. J., Sharpe, M. J., Schoenbaum, G., and Niv, Y. (2018). ScienceDirect Model-based predictions for dopamine. *Current Opinion in Neurobiology*, 49:1–7.

Lee et al., 2012. Lee, D., Seo, H., and Jung, M. W. (2012). Neural Basis of Reinforcement Learning and Decision Making. *Annual Review of Neuroscience*, 35(1):287–308.

Lee et al., 2013. Lee, J. H., Whittington, M. A., and Kopell, N. J. (2013). Top-Down Beta Rhythms Support Selective Attention via Interlaminar Interaction: A Model. *PLOS Computational Biology*, 9(8):e1003164–23.

Lee and Mumford, 2003. Lee, T. S. and Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am.*, 20(7):1–15.

Li and Daw, 2011. Li, J. and Daw, N. D. (2011). Signals in Human Striatum Are Appropriate for Policy Update Rather than Value Prediction. *Journal of Neuroscience*, 31(14):5504–5511.

Lim and Goldman, 2013. Lim, S. and Goldman, M. S. (2013). Balanced cortical microcircuitry for maintaining information in working memory. *Nature Neuroscience*, 16(9):1306–1314.

Maclaurin et al., 2017. Maclaurin, D., Duvenaud, D., Johnson, M., and Townsend, J. (2017). Autograd. *Retrieved from https://github.com/HIPS/autograd*.

Maheswaranathan et al., 2019. Maheswaranathan, N., Williams, A. H., Golub, M. D., Ganguli, S., and Sussillo, D. (2019). Universality and individuality in neural dynamics across large populations of recurrent networks . *arXiv*, pages 1–16.

Mante et al., 2013. Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84.

Maris and Oostenveld, 2007. Maris, E. and Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164(1):177–190.

Marr and Poggio, 1976. Marr, D. and Poggio, T. (1976). From Understanding Computation to Understanding Neural Circuitry. *Artificial Intelligence Laboratory. A.I. Memo. Massachusetts Institute of Technology.*

Martinez, 2009. Martinez, J. A. (2009). BESOS oder J'aime. *BoD. Book: https://www.amazon.com/BESOS-Jaime-German-Andrade-Martinez/dp/3833489278; BoD Award (Part I):*

https://www.youtube.com/watch?v=zjkXAcvdjh8; BoD Award (Part II): https://www.youtube.com/watch?v=clxR8A9wC8k; BoD Award (Part III): https://www.youtube.com/watch?v=F4fTGZzl-D4; Poem Sample: https://www.youtube.com/watch?v=KPx7MEkO$_S$A.

Márton et al., 2019. Márton, C. D., Fukushima, M., Camalier, C. R., Schultz, S. R., and Averbeck, B. B. (2019). Signature Patterns for Top-Down and Bottom-Up Information Processing via Cross-Frequency Coupling in Macaque Auditory Cortex. *eneuro*, 6(2):ENEURO.0467–18.2019–14.

Marton, 2019a. Marton, Z. (2019a). Tordenelem. *BoD. https://www.bod.ch/buchshop/toerdenelem-zsolt-marton-9783748133612*.

Marton, 2019b. Marton, Z. (2019b). Zetelaki Halasto: Pitch utazasai II. *BoD. https://www.bod.de/buchshop/zetelaki-halasto-zsolt-marton-9783746017785*.

Masse et al., 2020. Masse, N. Y., Rosen, M. C., and Freedman, D. J. (2020). Reevaluating the Role of Persistent Neural Activity in Short-Term Memory. *Trends in Cognitive Sciences*, pages 1–17.

Michalareas et al., 2016. Michalareas, G., Vezoli, J., van Pelt, S., Schoffelen, J.-M., Kennedy, H., and Fries, P. (2016). Alpha-Beta and Gamma Rhythms Subserve Feedback and Feedforward Influences among Human Visual Cortical Areas. *Neuron*, 89(2):384–397.

Musall et al., 2019. Musall, S., Urai, A. E., Sussillo, D., and Churchland, A. K. (2019). ScienceDirect Harnessing behavioral diversity to understand neural computations for cognition. *Current Opinion in Neurobiology*, 58:229–238.

Nakahara et al., 2001. Nakahara, H., Doya, K., and Hikosaka, O. (2001). Parallel Cortico-Basal Ganglia Mechanisms for Acquisition and Execution of Visuomo-

tor Sequences - A Computational Approach. *Journal of Cognitive Neuroscience*, 13(5):626–647.

Nichols and Holmes, 2001. Nichols, T. E. and Holmes, A. P. (2001). Nonparametric Permutation Tests For Functional Neuroimaging: A Primer with Examples. *Human Brain Mapping*, 15:1–25.

Nicola and Clopath, 2017. Nicola, W. and Clopath, C. (2017). Supervised learning in spiking neural networks with FORCE training. *Nature Communications*, 8(1):1–15.

Niv, 2009. Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53(3):139–154.

Niv, 2019. Niv, Y. (2019). Learning task-state representations. *Nature Neuroscience*, 22(10):1–10.

O'Doherty et al., 2004. O'Doherty, J. P., Dayan, P., Schultz, J., Deichmann, R., Friston, K., and Dolan, R. J. (2004). Dissociable Roles of Ventral and Dorsal Striatum in Instrumental Conditioning. *Science*, 304:452–454.

Park et al., 2015. Park, H., Ince, R. A. A., Schyns, P. G., Thut, G., and Gross, J. (2015). Frontal Top-Down Signals Increase Coupling of Auditory Low-Frequency Oscillations to Continuous Speech in Human Listeners. *CURBIO*, 25(12):1649–1653.

Pasupathy and Miller, 2005. Pasupathy, A. and Miller, E. K. (2005). Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature*, 433:873–876.

Poeppel et al., 2008. Poeppel, D., Idsardi, W. J., and van Wassenhove, V. (2008). Speech perception at the interface of neurobiology and linguistics. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1493):1071–1086.

Qi and Constantinidis, 2012. Qi, X.-L. and Constantinidis, C. (2012). Variability of Prefrontal Neuronal Discharges before and after Training in a Working Memory Task. *PLoS ONE*, 7(7):e41053–8.

Rabinovich et al., 2008. Rabinovich, M., Huerta, R., and Laurent, G. (2008). Transient Dynamics for Neural Processing. *Science*, 321(5885):47–48.

Radulescu et al., 2019. Radulescu, A., , , Niv, Y., , and Ballard, I. (2019). Holistic Reinforcement Learning: The Role of Structure and Attention. *Trends in Cognitive Sciences*, 23(4):1–15.

Rajan et al., 2016. Rajan, K., Harvey, C. D., and Tank, D. W. (2016). Recurrent Network Models of Sequence Generation and Memory. *Neuron*, 90(1):128–142.

Rao and Ballard, 1999. Rao, R. P. N. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87.

Ray and Maunsell, 2011. Ray, S. and Maunsell, J. H. R. (2011). Different Origins of Gamma Rhythm and High-Gamma Activity in Macaque Visual Cortex. *PLoS Biology*, 9(4):e1000610–15.

Remington et al., 2018. Remington, E. D., Narain, D., Hosseini, E. A., and Jazayeri, M. (2018). Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics. *Neuron*, 98(5):1005–1019.e5.

Rich and Wallis, 2016. Rich, E. L. and Wallis, J. D. (2016). Decoding subjective decisions from orbitofrontal cortex. *Nature Neuroscience*, 19(7):973–980.

Richards et al., 2019. Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R. P., Berker, A., Ganguli, S., Gillon, C. J., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay,

G. W., Miller, K. D., Naud, R., Pack, C. C., Poirazi, P., Roelfsema, P., Sacramento, J., Saxe, A., Scellier, B., Schapiro, A. C., Senn, W., Wayne, G., Yamins, D., Zenke, F., Zylberberg, J., Therien, D., and Kording, K. P. (2019). A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11):1–10.

Richter et al., 2017. Richter, C. G., Thompson, W. H., Bosman, C. A., and Fries, P. (2017). Top-Down Beta Enhances Bottom-Up Gamma. *Journal of Neuroscience*, 37(28):6698–6711.

Romanski and Averbeck, 2009. Romanski, L. M. and Averbeck, B. B. (2009). The Primate Cortical Auditory System and Neural Representation of Conspecific Vocalizations. *Annual Review of Neuroscience*, 32(1):315–346.

Sadtler et al., 2014. Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., and Batista, A. P. (2014). Neural constraints on learning. *Nature*, 512(7515):423–426.

Sakai, 2008. Sakai, K. (2008). Task Set and Prefrontal Cortex. *Annual Review of Neuroscience*, 31(1):219–245.

Salazar et al., 2012. Salazar, R. F., Dotson, N. M., Bressler, S. L., and Gray, C. M. (2012). Content-Specific Fronto-Parietal Synchronization During Visual Working Memory. *Science*, 338(6110):1097–1100.

Salin and Bullier, 1995. Salin, P.-A. and Bullier, J. (1995). Corticocortical Connections in the Visual System: Structure and Function. *Physiological Review*, 75(1):1–48.

Samejima et al., 2005. Samejima, K., Ueda, Y., Doya, K., and Kimura, M. (2005). Representation of Action-Specific Reward Values in the Striatum. *Science*, 310(5752):1333–1337.

Sarvestani et al., 2011. Sarvestani, I. K., Lindahl, M., Hellgren-Kotaleski, J., and Ekeberg, O. (2011). The arbitration–extension hypothesis: a hierarchical interpretation of the functional organization of the basal ganglia. *Frontiers in Systems Neuroscience*, 5(13):1–12.

Sato et al., 2010. Sato, J. R., Fujita, A., Cardoso, E. F., Thomaz, C. E., Brammer, M. J., and Amaro Jr, E. (2010). Analyzing the connectivity between regions of interest: An approach based on cluster Granger causality for fMRI data analysis. *NeuroImage*, 52(4):1444–1455.

Schack et al., 2002. Schack, B., Vath, N., Petsche, H., Geissler, H. G., and Moeller, E. (2002). Phase-coupling of theta-gamma EEG rhythms during short-term memory processing. *International Journal of Psychophysiology*, 44:143–163.

Schroeder et al., 2010. Schroeder, C. E., Wilson, D. A., Radman, T., Scharfman, H., and Lakatos, P. (2010). Dynamics of Active Sensing and perceptual selection. *Current Opinion in Neurobiology*, 20(2):172–176.

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

Scott et al., 2015. Scott, B. H., Leccese, P. A., Saleem, K. S., Kikuchi, Y., Mullarkey, M. P., Fukushima, M., Mishkin, M., and Saunders, R. C. (2015). Intrinsic Connections of the Core Auditory Cortical Regions and Rostral Supratemporal Plane in the Macaque Monkey. *Cerebral cortex (New York, N.Y. : 1991)*, 7:bhv277–32.

Seo et al., 2012. Seo, M., Lee, E., and Averbeck, B. B. (2012). Action Selection and Action Value in Frontal-Striatal Circuits. *Neuron*, 74(5):947–960.

She and Sun, 1999. She, Y. Q. and Sun, H. (1999). Image and video compression for multimedia engineering: Fundamentals, algorithms and standards. *CRC Press*.

Shenoy et al., 2013. Shenoy, K. V., Sahani, M., and Churchland, M. M. (2013). Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience*, 36(1):337–359.

Singer et al., 2018. Singer, Y., Teramoto, Y., Willmore, B. D., Schnupp, J. W., King, A. J., and Harper, N. S. (2018). Sensory cortex is optimized for prediction of future input. *eLife*, pages 1–31.

Sinha et al., 2014. Sinha, P., Kjelgaard, M. M., Gandhi, T. K., Tsourides, K., Cardinaux, A. L., Pantazis, D., Diamond, S. P., and Held, R. M. (2014). Autism as a disorder of prediction. *Proceedings of the National Academy of Sciences of the United States of America*, 111(42):15220–15225.

Song et al., 2017. Song, H. F., Yang, G. R., and Wang, X.-J. (2017). Reward-based training of recurrent neural networks for cognitive and value-based tasks. *eLife*, 6(e21492):1–24.

Spratling, 2017. Spratling, M. W. (2017). A review of predictive coding algorithms. *Brain and Cognition*, 112:92–97.

Stokes, 2015. Stokes, M. G. (2015). 'Activity-silent' working memory in prefrontal cortex: a dynamic coding framework. *Trends in Cognitive Sciences*, 19(7):394–405.

Stokes et al., 2013. Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., and Duncan, J. (2013). Dynamic Coding for Cognitive Control in Prefrontal Cortex. *Neuron*, 78(2):364–375.

Summerfield and de Lange, 2014. Summerfield, C. and de Lange, F. P. (2014). Expectation in perceptual decision making: neural and computational mechanisms. *Nature Publishing Group*, 15(11):1–12.

Suri and Schultz, 1998. Suri, R. E. and Schultz, W. (1998). Learning of sequential movements by neural network model with dopamine-like reinforcement signal. *Exp Brain Res*, 121:350–354.

Sussillo and Abbott, 2009. Sussillo, D. and Abbott, L. F. (2009). Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron*, 63(4):544–557.

Sussillo and Barak, 2013. Sussillo, D. and Barak, O. (2013). Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks . *Neural Computation*, pages 1–24.

Sussillo et al., 2015. Sussillo, D., Churchland, M. M., Kaufman, M. T., and Shenoy, K. V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience*, 18(7):1025–1033.

Sutskever, 2013. Sutskever, I. (2013). Training Recurrent Neural Networks. *PhD Thesis*, pages 1–101.

Sutton and Barto, 2018. Sutton, R. S. and Barto, A. G. (2018). Reinforcement Learning: An Introduction. *MIT Press, Cambridge, MA*.

Tort et al., 2010. Tort, A. B. L., Komorowski, R., Eichenbaum, H., and Kopell, N. (2010). Measuring Phase-Amplitude Coupling Between Neuronal Oscillations of Different Frequencies. *Journal of Neurophysiology*, 104(2):1195–1210.

Tort et al., 2009. Tort, A. B. L., Komorowski, R. W., Manns, J. R., Kopell, N. J., and Eichenbaum, H. (2009). Theta–gamma coupling increases during the learning of item– context associations. *PNAS*, 106(49):20942–20947.

Tort et al., 2008. Tort, A. B. L., Kramer, M. A., Thorn, C., Gibson, D. J., Kubota, Y., Graybriel, A. M., and Kopell, N. J. (2008). Dynamic cross-frequency couplings of local field potential oscillations in rat striatum and hippocampus during performance of a T-maze task. *PNAS*, 105(51):20517–20522.

Tsujimoto et al., 2008. Tsujimoto, S., Genovesio, A., and Wise, S. P. (2008). Transient Neuronal Correlations Underlying Goal Selection and Maintenance in Prefrontal Cortex. *Cerebral cortex (New York, N.Y. : 1991)*, 18(12):2748–2761.

van Wijk et al., 2015. van Wijk, B. C. M., Jha, A., Penny, W., and Litvak, V. (2015). Parametric estimation of cross-frequency coupling. *Journal of Neuroscience Methods*, 243:94–102.

Verschure et al., 2014. Verschure, P. F. M. J., Pennartz, C. M. A., and Pezzulo, G. (2014). The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1655):20130483–20130483.

von Helmholtz, 1876. von Helmholtz, H. (1876). Handbuch der physiologischen Optik. *Stuttg.*

von Stein et al., 2000. von Stein, A., Chiang, C., and König, P. (2000). Top-down processing mediated by interareal synchronization. *PNAS*, 97(26):1–6.

Voytek et al., 2013. Voytek, B., D'Esposito, M., Crone, N., and Knight, R. T. (2013). A method for event-related phase/amplitude coupling. *NeuroImage*, 64(C):416–424.

Wallis et al., 2019. Wallis, J. D., Anderson, K. C., and Miller, E. K. (2019). Single neurons in prefrontal cortex encode abstract rules. *Nature*, 411:953–956.

Wang et al., 2017. Wang, J., Narain, D., Hosseini, E. A., and Jazayeri, M. (2017). Flexible timing by temporal scaling of cortical responses. *Nature Neuroscience*, 21(1):1–13.

Wang et al., 2018. Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., Hassabis, D., and Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6):1–14.

Wang, 2010. Wang, X.-J. (2010). Neurophysiological and Computational Principles of Cortical Rhythms in Cognition. *Physiological Reviews*, 90(3):1195–1268.

Wild et al., 2010. Wild, C. J., Davis, M. H., and Johnsrude, I. S. (2010). Human auditory cortex is sensitive to the perceived clarity of speech. *NeuroImage*, 60(2):1–13.

Wild et al., 2012. Wild, C. J., Yusuf, A., Wilson, D. E., Peelle, J. E., Davis, M. H., and Johnsrude, I. S. (2012). Effortful Listening: The Processing of Degraded Speech Depends Critically on Attention. *Journal of Neuroscience*, 32(40):14010–14021.

Wolpert et al., 1995. Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). An Internal Model for Sensorimotor Integration. *Science*, 269:1–4.

Wolpert et al., 1998. Wolpert, D. M., Miall, C. R., and Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, 2(9):1–10.

Wood and Grafman, 2003. Wood, J. N. and Grafman, J. (2003). Human prefrontal cortex: processing and representational perspectives. *Nature Reviews Neuroscience*, 4(2):139–147.

Yamins et al., 2014. Yamins, D. L. K., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. (2014). Performance-optimized hierarchical

models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 111(23):8619–8624.

Yang et al., 2018. Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T., and Wang, X.-J. (2018). Task representations in neural networks trained to perform many cognitive tasks. *Nature Neuroscience*, 22(2):1–16.