AN ACTIVE INFERENCE ACCOUNT OF P300-LIKE CORTICAL RESPONSES IN CARD SORTING: IMPLICATIONS FOR THEORIES OF PREFRONTAL EXECUTIVE FUNCTIONS

Francisco Barceló[1], Riccardo Proietti[2], Giovanni Pezzulo[2]

[1]Laboratory of Neuropsychology, University of the Balearic Islands, Palma de Mallorca, Spain

[2]Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy

[1] f.barcelo@uib.es ; [2] riccaproiet94@gmail.com ; [2] giovanni.pezzulo@istc.cnr.it

Background

- When one reconsiders the P300 literature from Active Inference views under the Bayesian brain hypothesis (Parr, Pezzulo & Friston 2022), then the two well-known classes of anterior and posterior P3-like responses can be recast in terms of precision-weighted prediction errors at hierarchically ordered levels across frontoparietal cortical networks (cf., Barceló, 2021).
- On this view, the anterior P300 (P3a, novelty P3) indexes perceptual inference for anticipatory action selection and/or inhibition. In turn, the posterior P300 (P3b, LPC) in published studies likely consists of a mixture of both inference and learning (Friston 2005), as these two processes are often mixed up in the grand-averaged event-related potentials (ERPs).
- Crucially, on this view there is NOT just ONE but TWO functionally distinct types of posterior "P3b-like" waves : (1) one always follows the anterior P300 during "context updating" in volatile task contexts and shows rapid "repetition suppression" (i.e., P3b'), whereas (2) another functionally distinct type of P3b is elicited during "context learning" in stable task contexts, and shows gradual "repetition enhancement". This hypothesis is a direct corollary of Friston's (2005) theory of cortical responses.
- Here, formal modeling of free energy minizimization via active inference is applied to revise the evidence about these two classes of P3b-like responses in rule-switching and card sorting (Fig. 1), as two distinct indexes of belief updating at posterior multimodal association cortices.

Computerized card sorting task

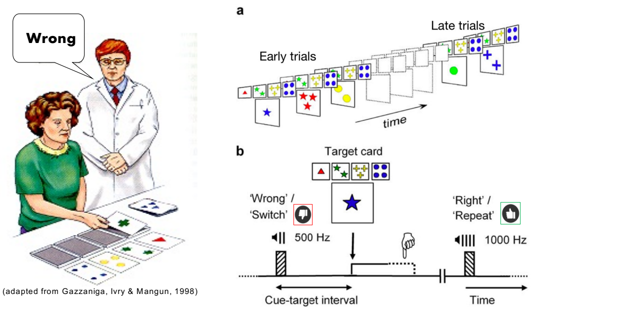

FIG. 1. (a) Schematic of one card sorting series where early and late trials broadly map onto the stages of inference and learning of perceptual categories. (b) Schematic of one card sorting trial where simple tonal sounds can be instructed either as negative and positive feedback or as "switch" and "repeat" cues informing about probabilistic updates in the policy for responding to the ensuing target card (cf., Barceló et al., 2006).

Formal modeling of card sorting in active inference

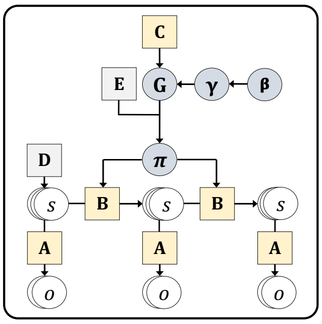

Symbols	Maths
π: Beliefs about policies	$\pi = \sigma(\gamma \mathbf{G} + \ln \mathbf{E})$
G: Expected free energy (scores policies during planning)	$G_\pi \propto -\ln P(\pi \mid \mathbf{C})$
A: Likelihood matrix (encodes the statistical relations between hidden states and outcomes)	$P(o_t \mid s_t)$
B: Transition matrix (encodes the statistical relations between states)	$P(s_{t+1} \mid s_t, \pi)$
C: Matrix encoding prior preferences over observations	$P(o_t \mid \mathbf{C})$
D: Vector encoding the prior belief about initial hidden states	$P(s_0 \mid \mathbf{D})$
E: Vector encoding the prior belief about initial policies	$P(\pi)$
β and **γ**: initial value of precision and precision parameter of expected free energy	$\boldsymbol{\beta} = 1/\gamma$

FIG. 2. Active inference graphical model that represents the Bayesian network used for state estimation (perception) and policy selection (action). The graph is a formal specification of the agent's generative model using the framework of Partially Observable Markov Decision Processes (POMDP). The nodes represent the agent's beliefs about states and task variables, encoded as discrete probability distributions. The edges illustrate the statistical relationships between these variables, defined by the A, B, C, D, and E matrices, while G represents expected free energy and γ and β its precision parameters.

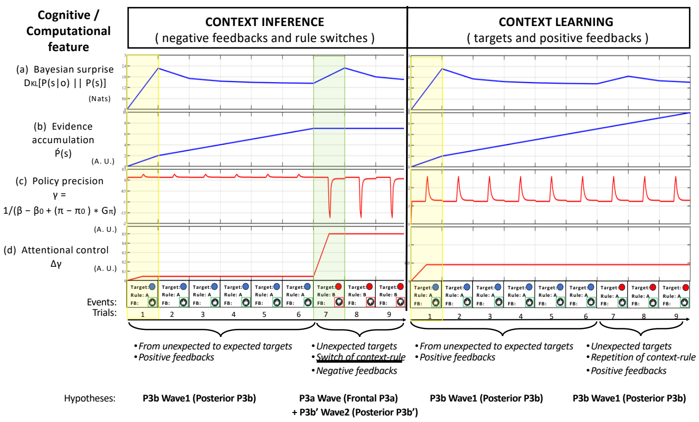

FIG. 3. Active inference simulations of behaviour during CONTEXT INFERENCE and CONTEXT LEARNING. Plotted here are the four cognitive-computational features that emerge naturally from the process of free-energy minimization given a discrete POMDP generative model of a rule-switching task. These are: (a) *Bayesian surprise*, which scores the difference (KL divergence) between the expected stimulus (i.e., target or task rule) and the observed stimulus. (b) *Evidence accumulation*, which represents the evidence about the current expected context, in terms of (unnormalized) Bayesian belief about hidden states. (c) *Policy precision*, which represents the dynamics of optimizing the γ parameter, the precision term of the expected free energy. Essentially, it entails a positive peak when the observed outcomes are in line with the expected outcomes given the agent's action plan, while a negative peak occurs when outcomes are not those predicted by the action plan. (d) *Attentional control*, scored as the rate of change of the precision. When the dynamics of policy precision change, it indicates that action plans are being revised, suggesting a form of cognitive control. In blue are highlighted the dynamics of processes associated with perception. In red, dynamics of processes associated with action/cognitive control.

These modeling results can be interpreted to support the hypothesis of two functionally distinct types of posterior P3b: (1) ***P3b Wave1*** (regular Posterior P3b) that builds-up over several trials when there is both a gradual *decrease* in Bayesian surprise and an *increase* in Evidence accummulation under stable contexts (i.e., to repetitive expected targets & positive feedbacks; cf., FIG. 5a & FIG. 8), and (2) ***P3b' Wave2*** (Posterior P3b') that fires concurrently with the frontal P3a when there is both an increase in Bayesian surprise and in Attentional control during context updating and perceptual inference (i.e., to unexpected targets & negative feedbacks; green shaded area in the figure; cf., FIG. 5a & FIG. 6a), and posibly also during first target trials (yellow shaded areas).

For simplicity, this model discretizes time so that the timing of some observations made by the simulated agent could not reflect precisely the timing of the observations made by real participants. In the future this could be addressed by implementing active inference deep temporal models.

Acknowledgements. Funded by the Spanish Ministry of Science and Innovation (PID2019-106045GB-I00 / AEI / 10.13039/501100011033).

Trial-by-trial P300 dynamics during context updating and learning

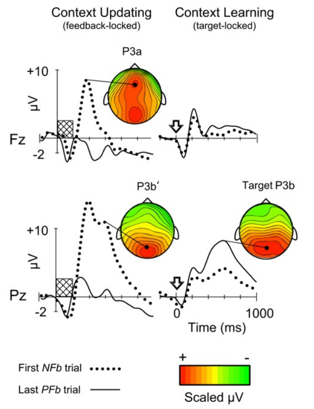

FIG. 4. Cortical responses to feedback cues and WCST target cards. Grand ERPs time-locked to feedback cues (shaded rectangle) and target cards (arrow) are displayed for first negative feedback (NFb) trials and last positive feedback (PFb) trials in a card sorting series, at midfrontal (Fz) and midparietal (Pz) regions. Scalp maps show mean P3a and P3b' amplitudes to first NFb cues and mean target P3b amplitudes to last correct target cards in the series. Early NFb trials foster context updating and inference, whereas late PFb trials foster context learning (Barceló, 2021).

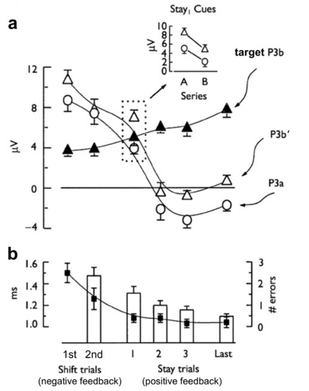

FIG. 5. Cortical and behavioral responses to negative and positive feedback trials in a card sorting series. (a) Group-averaged mean ± SEM amplitudes of feedback-locked P3a and P3b' as well as target P3b responses plotted across negative feedback ('switch') and positive feedback ('stay') trials from midfrontal (Fz) and midparietal (Pz) scalp regions. (b) Mean ± SEM RTs from efficiently completed WCST series without errors (solid squares) and mean ± SEM number of set-loss errors from failed series (bars) are shown during the inference and learning of the sorting category (Barceló, 2021).

Inference of sorting categories and novel percepts

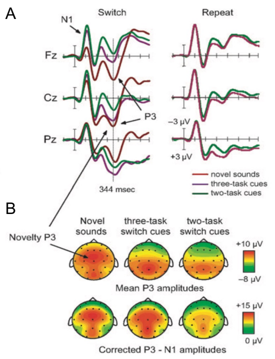

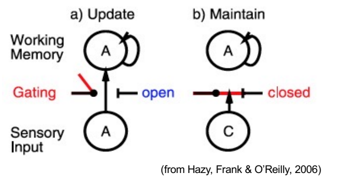

FIG. 7. Working memory gating during context updating and learning. (a) Both anterior P3a and posterior P3b' to predictive cues are elicited when the gate is open during perceptual inference (Context updating). (b) Posterior P3b to target cards are elicited when the gate is closed, mainly during context learning (cf., Barceló, 2021).

FIG. 6. (A) Cortical P3-like responses to predictive switch and repeat cues in the three-task (purple line) and two-task (green line) conditions from three midline electrodes. Overlaid for comparison are the P3-like responses to the highly surprising novel sounds (red line). Positive voltage values are displayed downwards. (B) Scalp potential maps for mean P3-like responses to novel sounds and switch cues in the three-task and two-task conditions (upper row), and for N1-corrected P3 voltages (lower row). Adapted from Barceló et al. (2006).

Evidence accumulation during context learning

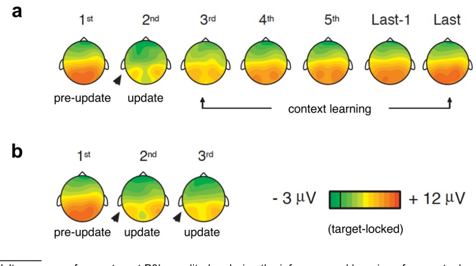

FIG. 8. Voltage maps of mean target P3b amplitudes during the inference and learning of perceptual categories. Arrowheads mark target trials preceded by a negative feedback (i.e., update trials). (a) Series with only one negative feedback trial. (b) Series with two negative feedback trials (postupdate trials not shown). Card sorting series (a) and (b) evoked similar repetition enhancement of target P3b amplitudes after the first positive feedback (Barceló, 2021).

Conclusions

- Active inference and the free-energy principle can explain many paradoxes of the frontal lobe riddle (cf., Barceló, 2021).
- On this view, the engagement of prefrontal cortices depends on the magnitude of precision-weighted prediction errors.
- The two broad classes of P300-like responses would index high- and low-level belief updating at frontal and posterior multimodal association cortices, during inference and learning of perceptual categories, respectively (cf., Friston 2005).
- Frontal-central P3a would index surprise minimization over unknown perceptual categories (e.g., expected policies), whereas parietal P3bs would index surprise minimization over task parameters (e.g., stimulus–response mappings).
- Crucially, on this view, perception and action are closely intertwined into perception–action cycles, thus reinstating old ideas about reafference in the neuropsychology of the frontal lobes (Luria 1966).
- Finally, our modeling of expected free energy supports the presence of two functionally distinct types of posterior "P3b-like" waveforms : (1) *P3b wave1* gradually builds-up during perceptual learning in stable contexts and is not compromised by frontal damage (cf., Knight 1997); and (2) *P3b' wave2* fires concurrently with the frontal P3a during context updating, and thus, it engages frontal cortical resources for the attentional control and the inference of novel response policities.

References

- Barceló, F. (2021). *J. Cog. Neurosci.* https://doi.org/10.1162/jocn_a_01662
- Barceló, F., Escera, C., et al., (2006). *J Cog Neurosci*, 18, 1734–1748. https://doi.org/10.1162/jocn.2006.18.10.173
- Friston, K. (2005). *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, 360, 815–836. https://doi.org/10.1098/rstb.2005.1622
- Knight, R.T. (1997). *J. Cog. Neurosci.*, 9(1), 75-91 https://doi.org/10.1162/jocn.1997.9.1.75
- Parr, T., Pezzulo, G., & Friston, K.J. (2022). Active Inference. Cambridge, MA: The MIT Press.
- Hazy, T.E., Frank, M.J., O'Reilly, R.C. (2006) *Neurosci.*, 139(1):105-18.
- Luria, A. R. (1966). *Higher cortical functions in man*. London: Tavistock Publications.