



# Evidence accumulation under uncertainty - a neural marker of emerging choice and urgency

Elisabeth Parés-Pujolràs\*, Eoin Travers, Yoana Ahmetoglu, Patrick Haggard

*Institute of Cognitive Neuroscience, University College London, London WC1 3AR, UK*

## ARTICLE INFO

### Keywords:

Decision  
Action  
Evidence accumulation  
P3  
Urgency

## ABSTRACT

To interact meaningfully with its environment, an agent must integrate external information with its own internal states. However, information about the environment is often noisy. In this study, we identify a neural correlate that tracks how asymmetries between competing alternatives evolve over the course of a decision. In our task participants had to monitor a stream of discrete visual stimuli over time and decide whether or not to act, on the basis of either strong or ambiguous evidence. We found that the classic P3 event-related potential evoked by sequential evidence items tracked decision-making processes and predicted participants' categorical choices on a single trial level, both when evidence was strong and when it was ambiguous. The P3 amplitudes in response to evidence supporting the eventually selected option increased over trial time as decisions evolved, being maximally different from the P3 amplitudes evoked by competing evidence at the time of decision. Computational modelling showed that both the neural dynamics and behavioural primacy and recency effects can be explained by a combination of (a) competition between mutually inhibiting accumulators for the two categorical choice outcomes, and (b) a context-dependant urgency signal. In conditions where evidence was presented at a low rate, urgency increased faster than in conditions when evidence was very frequent. We also found that the readiness potential, a classic marker of endogenously initiated actions, was observed preceding movements in all conditions - even when those were strongly driven by external evidence.

## 1. Introduction

Human agents often face decisions where there are no compelling external or internal reasons to decide one way rather than another. However, successful behaviour requires the capacity to break the symmetry between alternatives and pick one of them. How do we decide what to do when the available information is ambiguous?

Decision-making is often described as a sequential sampling process, where evidence is continuously sampled and accumulated, triggering a decision or action once a given threshold is reached (Forstmann et al., 2016). Decisions in the absence of clear exogenous information may result from a combination of noisy accumulation processes either in sensory or motor areas (Bode et al., 2013; Schurger et al., 2012) and a time-varying urgency signal that modulates the accumulation process over time, adjusting the accuracy criterion as time elapses (Carland et al., 2019; Churchland et al., 2008; Thura et al., 2014). Here, we used a temporally-extended decision-making task involving sequential presentation of evidence items to investigate the neural correlates of how choices emerge when they are underdetermined (i.e. not strongly supported) by the available evidence. In particular, we investigated two neural signals that might reflect the evolution of a decision variable and

action initiation: the P3 event-related potential and the Readiness Potential (RP).

Research in monkeys has shown that the firing rate of single neurons in lateral and medial intraparietal areas (LIP/MIP) build to threshold tracking decision-making processes in macaque monkeys (de Lafuente et al., 2015; Kiani et al., 2008; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001; Shadlen et al., 1996). Because neural populations respond to specific decision-relevant stimuli, such as leftward- or rightward- drifting dots, these neural responses have been described as 'categorical' (Fitzgerald et al., 2011; Freedman and Assad, 2006; Swaminathan et al., 2013; Swaminathan and Freedman, 2012). While no EEG correlates of such categorical processing during evidence accumulation have been found in humans, the P3 event-related potential and the related Centro-Parietal positivity (CPP) are promising candidates. These signals have been shown to exhibit build-to-threshold dynamics, correlating with the task difficulty and reaction times (Kutas et al., 1977; Twomey et al., 2015), tracking the evolution of multimodal decisions (O'Connell et al., 2012; Steinemann et al., 2018) and random dot motion (RDM) stimuli (Kelly and O'Connell, 2013). However, to the best of our knowledge, it remains unknown whether the P3 or CPP can be used to track *which* decision people will make. Previous EEG studies

\* Corresponding author.

E-mail address: [eparesp@gmail.com](mailto:eparesp@gmail.com) (E. Parés-Pujolràs).

could not separately resolve the neural activity evoked by evidence for each of two alternative actions, due to the type of stimuli used. For example, many studies considered a single composite variable corresponding to the differential evidence for one choice over another rather than providing separate readouts of the evidence for each choice independently. The graded contrast measures of O'Connell et al. (2012), and the centro-parietal positivity during random dot motion (Kelly and O'Connell, 2013) are of this kind. Thus, previous EEG studies could not independently track the emerging neural representation of each choice alternative and could not identify human equivalents of the 'categorical' neural responses identified in studies of single unit activity during decision-making.

Finally, we aimed to investigate the potential role of motor noise accumulation as a symmetry-breaking mechanism. The Readiness Potential (RP) has been identified as a reliable neural precursor of spontaneous action (Deecke and Kornhuber, 1978; Keller and Heckhausen, 1990; Khalighinejad et al., 2018; Schultze-Kraft et al., 2016), and since it appears to be absent in cued-reaction tasks (Jahanshahi et al., 1995; Papa et al., 1991), it has traditionally been interpreted as a marker of intentional action initiation. However, RP shapes can be obtained by averaging of stochastic fluctuations in the motor system in humans (Schurger et al., 2012) and spontaneous behaviour in mice (Murakami et al., 2014, 2017). Thus, it has been suggested that the RP may reflect accumulation of internal noise and play a determinant role in breaking the symmetry to trigger action in arbitrary decisions, but not deliberate ones (Maoz et al., 2019). Here, we aimed to test this hypothesis by comparing neural precursors of actions either weakly or strongly informed by evidence.

In this experiment, participants had to decide whether to make an action on the basis of extended sequences of discrete stimuli. Each trial lasted for exactly 25 s and participants were not rewarded for acting fast, so the timing of their decisions and actions was largely self-paced. Each discrete target stimulus either favoured acting (positive evidence, +Ev), or withholding action (negative evidence, -Ev), and the balance of evidence types varied between an easy and a difficult condition. These stimuli crucially allowed us to track the progressive emergence of *choice-specific* decision signals with each successive, discrete item of evidence. To do so, we developed a novel 'snapshot measure' approach. Rather than measuring continuous parietal activity from trial onset to decision time, we looked at the P3 responses evoked by discrete, sequential evidence items. We assumed that parietal populations hold sustained representations of an internal decision variable (O'Connell et al., 2012), and that representations for alternative action possibilities are distinct (Fitzgerald et al., 2011; Freedman and Assad, 2006; Swaminathan et al., 2013; Swaminathan and Freedman, 2012). We hypothesised that the P3 response evoked by sequential task-relevant items, which should reflect activity of category-selective parietal neural populations, would track the accumulated evidence for that hypothesis (+Ev, evidence for Action; -Ev, evidence for No-Action). In that case, P3 responses should evolve differently for selected and unselected options and track participants' choices on each trial as they unfold. Further, comparing neural precursors of action in easy and difficult conditions allowed us to test whether the RP could reflect a symmetry-breaking mechanism at the motor level.

## 2. Methods

### 2.1. Participants

All participants were recruited from a departmental participant database (ICN Subject Database). All participants were healthy, right-handed, young adults with normal or corrected to normal vision, no known disabilities and no history of neurological or psychological disorder. The study was approved by the UCL Research Ethics Committee and written informed consent was obtained from all participants before beginning the experiment. Subjects were paid £7.50 per hour, plus a performance-dependant reward (see below).

Nineteen subjects were initially invited to a single EEG session. Three participants were excluded from EEG analysis due to excessive EEG noise in the electrodes of interest (Cz & Pz). Eventually, 16 participants (13 female) were included in the EEG analysis ( $M$  age = 22.66,  $SD$  = 3.19; range: 19–30 years). All participants ( $N$  = 19) were included for the computational modelling analysis.

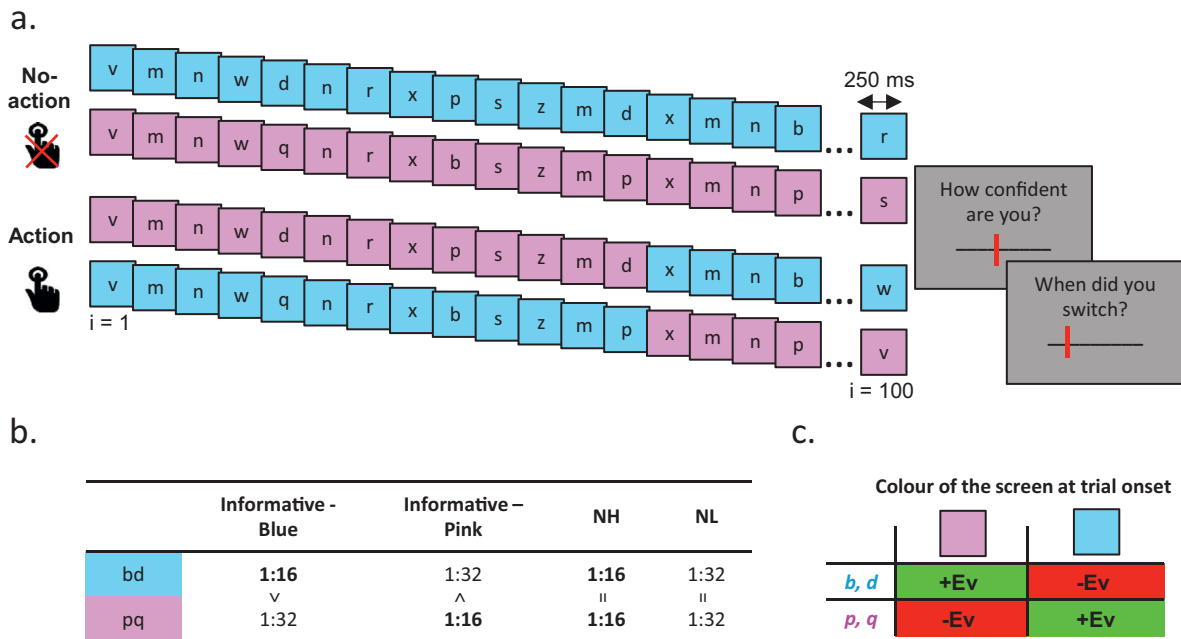
### 2.2. Procedure

Participants sat in a quiet room and were presented visual stimuli on a computer monitor. The instructions for the task were first displayed on the computer screen and then verbally repeated by the experimenter before the beginning of the experiment. Before the experiment, participants performed a practice version of the task (5 trials) to get familiar with the experiment.

### 2.3. Task & experimental design

The experiment was programmed in Matlab R2015a and Psychophysics Toolbox v3 (Brainard, 1997). Participants performed a decision-making task (Fig. 1a). Each trial started with either a blue or a pink background, and a 4 Hz, 25 s letter stream containing relevant and distractor stimuli was presented. All letters were lower-case black consonants presented without any blank interval between consecutive letters. There were two sets of task-relevant (target) letters, to which a task-relevant colour was assigned (in parenthesis): 'b' and 'd' (*bd*, blue), and 'p' and 'q' (*pq*, pink). Participants were instructed to monitor the letter stream and decide whether one set of target letters (*bd*) appeared more frequently than another set (*pq*). Participants' task was to make sure that the background colour matched the more frequent group of letters (i.e. if the most frequent set of targets was *bd*, the screen should be blue. If the most frequent set of targets was *pq*, the screen should be pink). If the initial colour of the screen on a given trial matched the set of targets participants perceived as most frequent, they were not required to do anything (*No-action* trials). If the initial colour of the screen on that trial did not match the most frequent set of targets, they had to press the key corresponding to the colour they wanted to switch it to (i.e. either pink or blue) with the left or right hand respectively (*Action trials*). They were allowed to switch the colour of the screen only once per trial. There was no incentive to respond fast, since trials were not terminated at the time of response and stimuli continued to appear on the screen until the end of the 25 s. At the end of each trial, participants were asked to report how confident they were that changing the colour of the screen (or not changing it) was the correct choice using a Visual analog Scale (VAS) by sliding a scrollbar with a standard computer mouse. Confidence ratings were obtained as an exploratory measure, as we had no a priori hypothesis regarding its relationship to the neural variables of interest. Given that the subjective time of commitment to an action has been linked to the neural precursors of movement (Libet et al., 1983), if participants had acted they were also asked to estimate at what point during the trial they acted, again using a VAS.

The frequencies of the two sets of targets, together with the colour of the screen at trial onset defined four experimental conditions (Fig. 1b). In the *Informative* conditions, the most frequent set of letters would appear twice as frequently as the other set. The stimuli were presented at a random time every 1 to 4 s (frequent stimuli) or every 1 to 8 s (infrequent stimuli). Thus, there was strong evidence either for or against action. In the *Informative anti-Action* condition, the colour of the screen at trial onset and the most frequent set of targets matched (i.e. the frequency of  $pq > bd$  and the screen was pink, or frequency of  $pq < bd$  and the screen was blue). Hence, participants should not press a key to change the colour of the screen. In the *Informative pro-Action* condition, the colour of the screen and the most frequent set of targets did not match (i.e. the frequency of  $pq > bd$  and the screen was blue, or the frequency of  $pq < bd$  and the screen was pink). Therefore, participants should press a key to switch the colour of the screen.



**Fig. 1. Task and experimental design.** **a.** Each trial started with either a blue or a pink background. A letter stream was presented for 25 s (100 letters, 250 ms each), and participants had to decide which of the two target letter sets was most frequent (blue targets: ‘b’ and ‘d’, or pink targets: ‘p’ and ‘q’). Their task was to make sure that the colour of the screen matched the most frequent set of targets (e.g. if ‘p’ and ‘q’ were most frequent, the screen should be pink). Sometimes, the starting colour of the screen would already match the set of targets participants perceived as most frequent. In those cases, they were instructed to not execute any action (No-Action). In other trials, the most frequent targets did not match the given colour of the screen (Action). In those cases, they could change the colour of the screen by pressing either the ‘f’ key with their left hand to turn the screen pink, or the ‘j’ key with their right hand to turn it blue. They were only allowed to change the colour of the screen once every trial. After 25 s, the letter stream was terminated and participants were asked confidence ratings and estimates of the time at which they made a decision in Action trials only. **b.** The relative frequency of relevant targets defined both conditions. In informative conditions, the most frequent set of targets appeared twice as often as the less frequent one (once every 16 letters, or once every 32). In the neutral conditions, both sets of targets appeared equally often. In the Neutral High (NH) condition they both appeared relatively often (once every 16 letters), while in the Neutral Low (NL) condition they appeared very rarely (once every 32 letters). **c.** Coding scheme. Relevant letters were coded with respect to action, given the initial colour of the screen. In trials where the screen was Blue at trial onset, pq were evidence for action (positive evidence, +Ev), and bd were evidence against action (negative evidence, -Ev).

In *Neutral* conditions, both sets of letters were presented at the same frequency. Two neutral conditions were created so that, on average, the overall number of relevant stimuli (i.e. rate) across the two neutral conditions would be equivalent to that in the informative ones. In the Neutral High (NH) input condition, evidence was abundant (i.e. one piece of +Ev and one piece of -Ev were presented at a random time every 1 to 4 s). In the Neutral Low (NL) input condition, evidence was scarce (i.e. one piece of +Ev and one piece of -Ev were presented at a random time every 1 to 8 s). In these two neutral conditions, there was no *net* evidence for or against switching in any given trial. Note, however, that participants were *not* told that in some conditions there was no net evidence. In all conditions, letters from both sets of targets were never presented within intervals shorter than 0.5 s. Possible distractors were the following letters: ‘a’, ‘c’, ‘e’, ‘i’, ‘m’, ‘n’, ‘o’, ‘r’, ‘s’, ‘u’, ‘v’, ‘w’, ‘x’, ‘z’. The precaution was taken not to include letters with ascending or descending strokes (‘f’, ‘g’, ‘h’, ‘j’, ‘k’, ‘l’, ‘t’, ‘y’) in the letter stream to ensure that morphological similarities between these and the target letters would not confound the results. The randomisation procedure ensured that participants could not build expectations regarding when target letters would appear (see below). Trials never started with a target stimulus, and the presentation times of the remaining target stimuli were randomly predetermined before trial start using an iterative procedure.

To randomise the position of target letters, we sequentially created a list of indices for each trial independently. If the frequency of the target letters on a given trial was 1/16, we picked a random number between 4 and 20 (e.g. 8) to determine the position of the first target letter on that trial. We then got another random number between 4 and 20 (e.g. 15), and added it to the first index of the vector to obtain the position of the

second letter on that trial ([8, 23]). We picked a random number starting from 4 rather than 1 to ensure that same-set letters would be separated by at least 3 distractors. We repeated this procedure until the obtained index was out of bounds (>100, end of trial). We generated a vector like this separately for each type of target letters and for each trial, and we repeated the procedure as many times as necessary to ensure that target letters were separated by at least 2 distractor letters.

In both Informative conditions, 1 penny was added to their bonus payment for correct decisions (i.e. acting in the pro-Action trials and not acting in the anti-Action trials) and 1 penny subtracted for incorrect decisions (i.e. acting in the anti-Action trials and not acting in the pro-Action trials). In the Neutral conditions, there were no correct or incorrect decisions. Hence, 1 penny reward or no reward was assigned at random on every trial. Participants were informed about the accumulated reward on the breaks between blocks.

Participants performed a total of 39 trials per condition. The experiment was divided in 3 blocks of 52 trials, and all conditions were randomised within blocks (13 trials per condition per block).

#### 2.4. EEG recording

EEG was recorded from 26 scalp sites (Fz, FCz, Cz, CPz, Pz, POz, FC1, FC2, C1, C2, CP1, CP2, F3, F4, F7, F8, C3, C4, CP5, CP6, FC5, FC6, P3, P4, O1, O2) using active electrodes (g.LADYbird) fixed to an EEG cap (g.GAMMAcap) according to the extended international 10/20 system (Jasper, 1958). EEG data were acquired using a g.GAMMAbox and g.USBamp with a sampling frequency of 256 Hz. Signal was recorded using g.Recorder (G.tec, medical engineering GmbH, Austria). All electrodes were online referenced to the right ear lobe. Vertical and hor-

horizontal electroocular activity was recorded from electrodes above and below the right eye and on the outer canthi of both eyes.

### 3. Behavioural data analysis

#### 3.1. Perceived time of switch decision

After trials where a decision to switch the colour of the screen was made and therefore participants acted (in Action trials), participants were required to provide an estimate of the time at which they decided to switch. In order to calculate the bias in their estimation with respect to the actual time of the decision to switch we first linearly converted the 100-point VAS on which participants provided the response dividing it by 4, thus providing a 25-point scale equivalent to the 25 s duration of the trial. We calculated the difference between the actual time and the perceived time of action for each individual trial (bias), averaged it across all Action trials for each participant and run a *t*-test against zero.

#### 3.2. Primacy effect

To investigate the extent to which decisions in highly ambiguous situations are driven by external evidence, we investigated whether the first piece of evidence participants saw in each trial biased their Action or No-Action choices in neutral conditions. We ran a multilevel logistic regression using the *lme4* package (Bates et al., 2015) to predict the probability of Action based on the first piece of evidence (+Ev/-Ev) on each single trial, the condition (NH/NL) and the interaction between both categorical variables. We included a random intercept to account for the between subject variability.

#### 3.3. Recency effect

We developed a measure to estimate the extent to which participant's actions were dependant on immediate external evidence. We assumed that, if participants were making a decision completely independently of the environment, the distribution of +Ev and -Ev preceding an action should be uniform, as expected at random. However, if their actions are driven to some extent by immediate external evidence, the observed distribution just prior to actions should deviate from the one expected at random. To measure the magnitude of this deviation, we compared the observed distribution of -Ev (i.e. evidence *against* making a key press) and +Ev (i.e. *for* making a key press) to the expected uniform distribution during the 2.5 s preceding an action, divided in 250 ms time bins. For this comparison, we calculated a *Deviation Score* (DS) for +Ev and -Ev separately ( $DS_{(+Ev)}$  and  $DS_{(-Ev)}$ ). For each participant and condition, we subtracted the expected number of  $\pm Ev$  ( $\pm Ev_e = \text{probability of any given letter being } \pm Ev \times \text{number of Action trials}$ ) from the observed number of ( $\pm Ev_o$ ) in each time bin. We then divided the result by the expected number of targets to normalise the score (Generic formula:  $DS = (\pm Ev_o - \pm Ev_e) / \pm Ev_e$ ). Positive values in the deviation score indicate that there were more targets than expected, while negative values indicate that there were fewer targets than expected. Finally, we combined these two deviation scores in a single measure: the *Recency Index* (RI). The Recency Index was calculated by subtracting the  $DS_{(-Ev)}$  from the  $DS_{(+Ev)}$  (Generic formula:  $RI = DS_{(+Ev)} - DS_{(-Ev)}$ ). The RI is thus a measure of how much the distribution of both +Ev and -Ev combined deviates from the one expected at random. The greater the RI, the greater the dependency of actions on immediate evidence. For statistical analysis, we averaged the RI in each time bin for the whole 2.5 s epoch preceding actions and a repeated measures ANOVA and post-hoc pairwise comparisons for each condition pair.

#### 3.4. Statistical analysis

Statistical analysis of the recency effect described above was run on Matlab R2014b. All other analysis of behavioural data were performed

in R using mixed-models regression with the *lme4* (Bates et al., 2015) package. We fit models to compare the percentage of actions and the time of action, and all models included a random intercept to account for the between subject variability.

### 4. EEG analysis

#### 4.1. Preprocessing

EEG data were processed using Matlab R2014b (MathWorks), EEGLab (Delorme and Makeig, 2004) and the Berlin Brain Computer Interface toolbox (Blankertz et al., 2016). First, scalp and eye electrodes were re-referenced to the average of two mastoid electrodes. Continuous EEG and EOG data were low-pass filtered under 30 Hz using a zero-phase 8th order Butterworth filter. No high-pass filter was applied to avoid possible distortions of low-frequency signal components contributing to the RP. Second, EEG signals were epoched. For P3 analysis, EEG signal was locked from -0.5 before to 1 s after the presentation of relevant letters ('p', 'q', 'b' and 'd'). For RP analysis, EEG signal was locked from -2.5 s to 0.5 s after the presentation of the letter following a keypress. Next, baseline correction was performed using the 500 ms at the beginning of the epoch [-0.5 to 0 s in P3 epochs, or -2.5 s to -2 s in RP epochs]. Ocular movements were identified using the *runica* algorithm in Matlab and removed from the epoched signal. Finally, artefact rejection was performed by removing all epochs with  $>200 \mu V$  fluctuations from baseline in the preselected channels of interest (Cz and Pz).

#### 4.2. Sequential P3 analysis

To study the dynamics of evidence accumulation as encoded in the P3, we extracted the P3 at electrode Pz in response to every instance of relevant evidence ('b', 'd', 'p' and 'q'), and we obtained the average amplitude of the whole duration of the component [0.3 to 0.8 s post stimulus], as determined by visual inspection of the grand-averaged traces. We used these raw values to investigate three research questions.

First, we investigated how the P3 amplitude evoked by single stimulus changed over sequential stimuli presentations as a decision unfolded, throughout the whole trial duration. For action trials only, we fitted a model to test whether P3 amplitudes could be predicted based on the time to action (in s), the type of evidence (+Ev/-Ev) and the condition (Informative/NH/NL).

Second, we investigated whether differences in the P3 components evoked by +Ev and -Ev at different times within a trial would reflect the evolving states of an internal decision variable. For this purpose, we ran a model to test whether the P3 amplitudes evoked by the *first* and *last* pieces of evidence presented from trial onset and before action execution in Action trials varied between different kinds of evidence (+Ev/-Ev), conditions (Informative/NH/NL).

Finally, we directly compared Action and No-Action trials. For this analysis, we averaged all the P3 components in response to +Ev and -Ev separately for each single trial, and we used these values for statistical inference. We fit a mixed model to test whether the single-trial average amplitude of the P3 significantly varied between different kinds of evidence (+Ev/-Ev), conditions (Informative/NH/NL) and Action decisions (Action/No-Action).

All models were fit using the *lme4* package (Bates et al., 2015) and included nested random intercepts for each subject and individual trial to control for trial-by-trial and individual variability. All reported *p*-values were Benjamini-Hochberg corrected to control for the number of comparisons performed within each model (Benjamini et al., 2009; Benjamini and Hochberg, 1995).

Post-hoc grouping of trials and P3 response averaging cannot, however, provide evidence for a link between P3 evolution and participants' choice formation. To test whether P3 amplitudes had predictive power at a single trial level, we used a simple classification approach. For each single trial, we used the following heuristic: if the P3 evoked by +Ev

was greater than that of  $-Ev$ , an Action response was predicted for that trial. Conversely, if the P3 evoked by  $-Ev$  was greater than that evoked by  $+Ev$  on average for a given trial, then a No-Action response was predicted. We then fitted a logistic regression to predict whether participants would decide to act or not act on any given trial, based on the P3-based prediction (Action/NoAction), the condition (Informative/NH/NL) and their interaction. Further, we repeated this analysis using the continuous difference value of the P3 magnitudes rather than its sign to test whether the predictive power was linearly related to the strength of the signal.

We also investigated how the P3-based classification accuracy evolved over trial time. To do so, we calculated the difference between the P3 amplitudes evoked by consecutive and competing evidence items. For example, if the P3 amplitude evoked by  $+Ev$  presented at time  $t$  was greater than the P3 amplitude evoked by  $-Ev$  presented at time  $t - 1$ , an “Action” prediction was made. A prediction of this kind was made for each evidence item up to the end of the trial. If the P3 is tracking a decisional process, the accuracy of the predictions based on it should improve as the decision is formed. If it is tracking a metacognitive monitoring process, it might continue increasing even after the decision time. We statistically tested whether the predictive accuracy of the P3 improved towards the time of the decision in Action trials. We could not do this for no-Action trials, for which the decision time is unknown. We ran a logistic regression predicting the probability that a prediction would be accurate based on the time with respect to action (in s), and whether the prediction was made before or after the action was executed for each condition (Informative/NH/NL).

#### 4.3. RP analysis

Statistical tests on averaged EEG data were run using FieldTrip toolbox (Oostenveld et al., 2011) cluster-based permutation analysis (Maris and Oostenveld, 2007). The main contrast of interest was the comparison between neural activity preceding actions between conditions (pro-Action, NH, NL). One electrode over the medial frontal areas (Cz) was preselected for analysis. The cluster-based tests were performed on the individual participant averages using the following parameters: two-tailed dependent samples  $t$ -test, time interval =  $[-2\ 0$  s relative to action], number of draws from the permutation distribution = 10,000.

### 5. Computational modelling

#### 5.1. Model description

We modelled the integration of evidence over time using two competing accumulators (Usher and McClelland, 2001). Such a model is a compromise between single-accumulator drift diffusion models (Ratcliff et al., 2016), which track the balance of evidence and therefore cannot account for differences between the Neutral-High and Neutral-Low conditions, and race models (Brown and Heathcote, 2008) which track the evidence in favour of each response but disregard the balance of evidence (Bogacz et al., 2006).

We modelled the state of two accumulators,  $x^1$  and  $x^2$  which integrated positive (pro-Action) and negative (anti-Action) evidence respectively. The state of the accumulators was updated at every time step,  $dt = 0.25$ , using the equations:

$$\begin{aligned} dx_t^1 &= g \times (v^1 + \varepsilon_t^1) - b \times x_t^2 \\ dx_t^2 &= g \times (v^2 + \varepsilon_t^2) - b \times x_t^1 \\ x_{t+1} &= x_t + dx_t \times dt \\ x_0 &= z \end{aligned}$$

where  $g$  is a gain parameter,  $v^1$  and  $v^2$  are the inputs to the positive and negative accumulators respectively,  $\varepsilon$  is independent Gaussian noise with variance  $\sigma^2$ ,  $b$  is a lateral inhibition parameter, and  $z$  is a starting point parameter. Each condition was simulated separately, with values

of  $v^1$  and  $v^2$  set to match the probability of positive and negative evidence in each condition. These values are therefore not free parameters. We simulated 100 time steps (25 s) for each trial. Simulated responses (Action or No-Action) were produced by whichever accumulator was first to reach a criterion of 1, and simulated response times corresponded to the time at which the criterion was reached. Parameters were estimated separately for each participant.

In the *constant gain* model, all free parameters were kept constant across conditions (number of free parameters  $k = 4$ , parameter vector  $\theta = [g, b, \sigma^2, z]$ ). In the *rising urgency* model, gain was set to  $g_\alpha$  at the start of each trial and increased by  $g_\beta$  with each time step ( $k = 5$ ,  $\theta = [g_\alpha, g_\beta, b, \sigma^2, z]$ ). In the *context-dependant urgency* model,  $g_\beta$  was estimated separately for each condition, allowing urgency to increase at different rates ( $k = 8$ ,  $\theta = [g_\alpha, g_\beta^{\text{Pro-Action}}, g_\beta^{\text{Neutral-High}}, g_\beta^{\text{Neutral-Low}}, g_\beta^{\text{Anti-Action}}, b, \sigma^2, z]$ ).

#### 5.2. Model fitting

We fit models to each participant’s behavioural data by maximum-likelihood. For a given vector of parameters  $\theta$ , we simulated 500 trials in each condition using these parameters and the appropriate values of  $v^1$  and  $v^2$ . We used the proportion of simulated trials where an action occurred as an estimate of the likelihood of action given these parameters,  $P(\text{Act}|\theta)$ . We estimated the response time probability density function,  $P(\text{RT} = t|\text{Act}, \theta)$ , by fitting an Epikanhov kernel density estimate (Turner and Sederberg, 2014) to the distribution of simulated response times on simulated trials where actions did occur. The log-likelihood on each trial was therefore  $\log\text{lik} = \log(P(\text{Act}|\theta)) + \log(P(\text{RT} = t|\text{Act}, \theta))$  on trials where the participant acted, and  $\log(1 - P(\text{Act}|\theta))$  on trials where they did not. The total log-likelihood for each participant’s data,  $\text{LL} = \text{Log}(P(\text{Data}|\theta))$  was calculated by summing the log-likelihood of each trial, across conditions. We used differential-evolution optimisation, implemented in the scipy package for python (Virtanen et al., 2020), to maximise total log-likelihood for each participant.

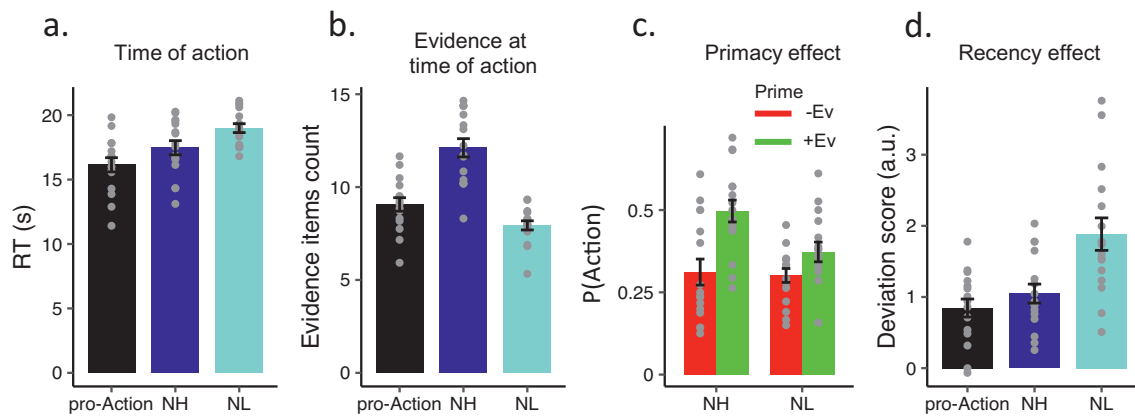
#### 5.3. Model comparison

To compare models, we calculate Akaike Information Criterion (AIC) for each participant,  $\text{AIC} = -2 \times \text{LL} + 2 \times k$ , where  $k$  is the number of parameters per model. Lower AIC values indicate greater support for the model. The differences in the AIC,  $\delta(\text{AIC})$ , were calculated as the difference between the AIC of the model being assessed and the AIC of the best model. These were used to calculate AIC weights,  $w_i(\text{AIC}) = \exp\{\delta_i(\text{AIC})\} / \sum_i^3 \exp\{\delta_i(\text{AIC})\}$  (Burnham and Anderson, 2001), indicating the probability that each model is the best, in terms of expected out-of-sample predictive accuracy, of the three considered.

### 6. Results

#### 6.1. Behavioural results

Participants performed the task well. On average, they duly switched on most pro-Action trials ( $M = 89.58\%$ ,  $SD = 6.03\%$ ) and switched rarely in anti-Action trials ( $M = 3.04\%$ ,  $SD = 2.14\%$ ). As they were very few in number, trials with errors of omission ( $M = 4.06$ ,  $SD = 2.35$ ) and commission ( $M = 1.46$ ,  $SD = 0.66$ ) were excluded from EEG analysis. Thus, all Informative trials reported correspond to Action trials in the pro-Action condition, and No-Action trials in the anti-Action condition. Participants decided to switch the colour of the screen more often ( $t_{(15)} = 2.57$ ,  $p = 0.021$ ) and sooner ( $t_{(15)} = 3.31$ ,  $p = 0.004$ ) in the NH (Percentage of Action trials:  $M = 40.54\%$ ,  $SD = 10.89\%$ , Time of action:  $M = 17.50$  s,  $SD = 2.20$  s), than in the NL condition (Percentage of Action trials:  $M = 34.13\%$ ,  $SD = 8.25\%$ ; Time of Action:  $M = 18.89$  s,  $SD = 1.56$  s; Fig. 2a). However, the average number of evidence items seen before action in the NH condition ( $M = 12.11$ ,  $SD = 1.97$ ) was



**Fig. 2. Behavioural results.** **a.** Grand-averaged time of action ( $\pm$ SEM) in all conditions. **b.** Grand-averaged mean ( $\pm$ SEM) total number of +Ev and -Ev items presented at the time of action for all conditions. Although the net balance of +Ev and -Ev was equivalent in the NH and NL conditions, a smaller number of evidence items sufficed to trigger Action in the NL condition. **c.** Grand-averaged probability ( $\pm$ SEM) of Action decisions given the first letter seen on a trial, for NH and NL conditions. In both conditions, participants were more likely to act if the first target letter they saw on a trial was +Ev (e.g. 'b'/'d' in a pink screen, or 'p'/'q' on a blue screen). **d.** Grand-averaged ( $\pm$ SEM) effect of recent evidence on action triggering (see Fig. S1 for full details). Higher values of the Recency Index (RI) indicate a greater dependence of action on recent evidence. In the NL condition, single pieces of evidence items were more likely to be followed by an action than in the other two conditions. Overlaid grey dots correspond to single-participant data in all plots.

significantly higher ( $t_{(15)} = 11.79$ ,  $p < 0.001$ ) than in the NL condition ( $M = 7.94$ ,  $SD = 0.98$ ), suggesting that participants adjusted their decision criterion to the amount of available evidence (see Fig. 2b). Since the net balance of +Ev and -Ev was the same in those two conditions, this strategic adjustment must have depended on an internal signal evolving throughout the trial and sensitive to differences between conditions. We explore this hypothesis in the computational section below.

Participants' subjective reports of their action times significantly preceded the actual movement time by an average of 2.63 s ( $SD = 1.45$ ) across all conditions ( $t$ -test against zero:  $t_{(15)} = 7.24$ ,  $p < 0.001$ ). Unsurprisingly, participants were more confident about their decisions in the informative ( $M = 67.55$ ,  $SD = 3.99$ ) than both in the NH ( $M = 47.56$ ,  $SD = 4.35$ ;  $t_{(15)} = 11.90$ ,  $p < 0.001$ ) and NL ( $M = 48.79$ ,  $SD = 4.11$ ;  $t_{(15)} = 10.54$ ,  $p < 0.001$ ) conditions, which did not significantly differ between them ( $t_{(15)} = 1.11$ ,  $p = 0.281$ ).

We then investigated how evidence presented early and late in the trial influenced participants' decisions in our neutral conditions. We found that the order in which evidence was presented early in the trial strongly influenced participants' decision to act or not act (*primacy effect*, see Fig. 2c). In particular, the first evidence item participants saw in each trial (+Ev/-Ev) biased their Action or No-Action choices in neutral conditions (NH/NL). A main effect of the prime on the probability of deciding to act ( $X^2_{(1)} = 20.48$ ,  $p < 0.001$ ) showed that participants decided to switch the colour of the screen more often when the first letter they saw on that trial was +Ev (NH:  $M = 50.84\%$ ,  $SEM = 2.90\%$ ; NL:  $M = 37.10\%$ ,  $SEM = 2.75\%$ ) than when it was -Ev (NH:  $M = 31.08\%$ ,  $SEM = 2.57\%$ ; NL:  $M = 31.21\%$ ,  $SEM = 2.62\%$ ) in both the NH and NL conditions. Importantly, the first target stimulus seen on any given trial was independent of the condition and the colour of the screen. A significant interaction between the first letter seen on the trial and the condition indicated that the primacy effect was stronger in the NH condition ( $X^2_{(1)} = 5.48$ ,  $p = 0.019$ ).

We further quantified the extent to which the time of actions was influenced by immediately-preceding evidence (*recency effect*). We assumed that, if participants were making a decision completely independently of the environment, the temporal distribution of +Ev and -Ev events prior to action should be uniform distributed, as expected at random. However, if their actions depend on immediate external evidence, the observed distribution just prior to actions should deviate from random. We therefore compared the observed distribution of -Ev (i.e. evidence *against* making a key press) and +Ev (i.e. *for* making a key press) to the expected uniform distribution during the 2.5 s preceding an action

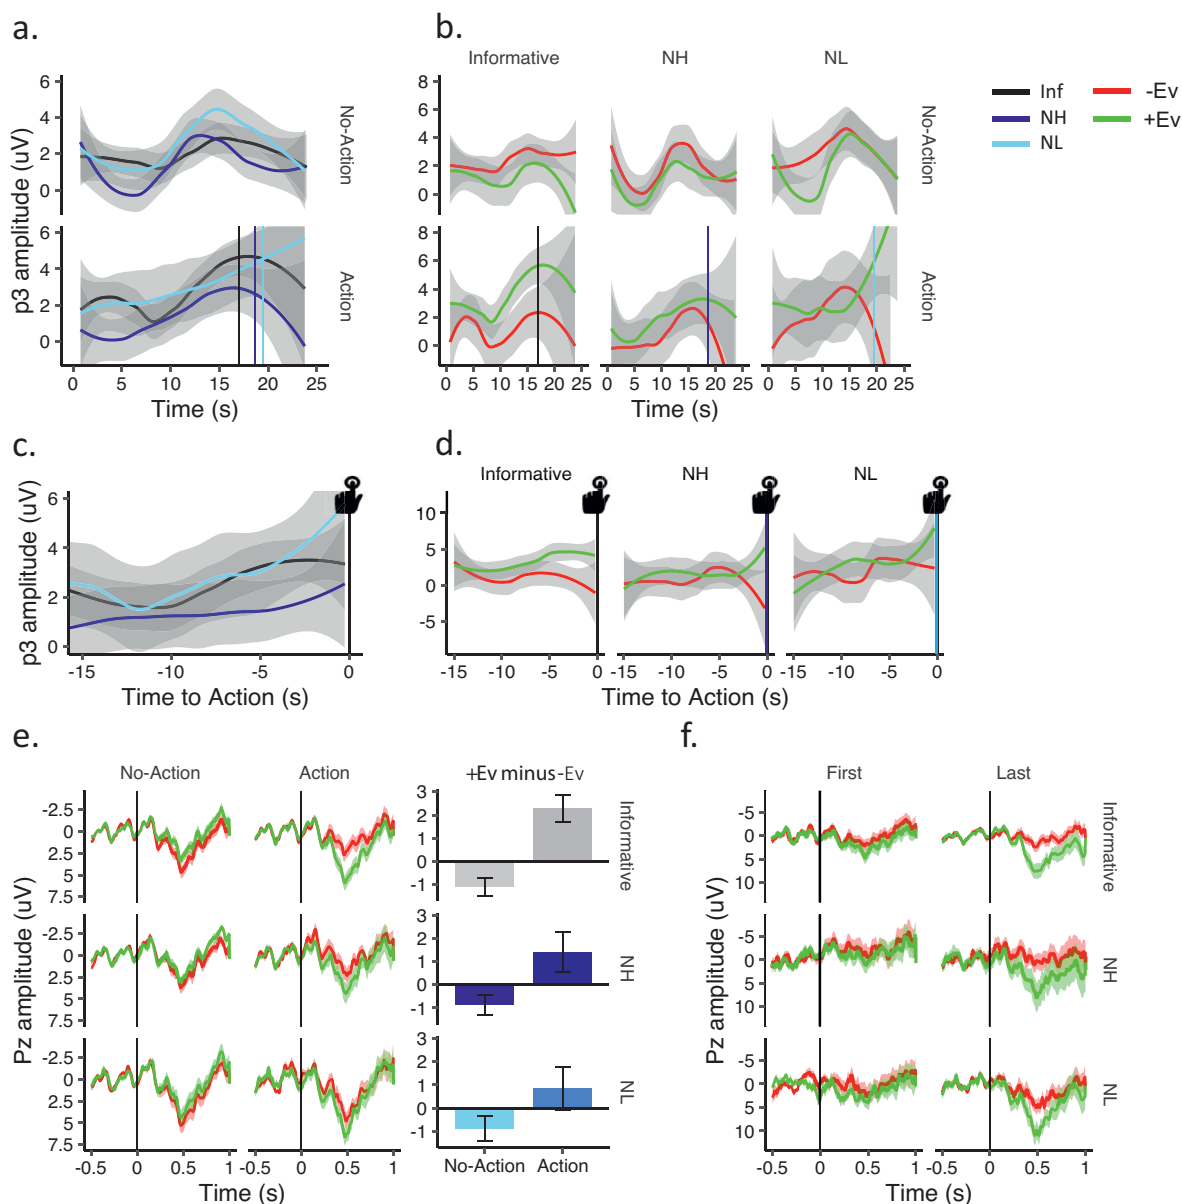
(see *Methods* for full details) and obtained a summary measures (*Recency Index, RI*). The greater the RI, the more likely participants were to move shortly after a piece of +Ev evidence, and *not* to move shortly after a -Ev item. In all conditions, participants tended to move shortly after a piece of +Ev (Figs. 2d; and S1 for full details). However, the deviation from the expected distribution was greatest in the NL condition ( $F_{(2,15)} = 11.03.03$ ,  $p < 0.001$ ; see *Supplementary note 1* for full statistical details).

These behavioural results therefore suggested two primarily different mechanisms of symmetry-breaking in the neutral conditions: a primacy effect driven by early evidence when evidence was abundant (NH), and reaction-like responses to late evidence when evidence was scarce (NL).

## 6.2. The P3 encodes a categorical decision variable

Based on previous literature suggesting that the P3 encodes a decision variable (Twomey et al., 2015), we first investigated the evolution of P3 amplitude in response to sequential, discrete presentations of target letters (i.e. 'b', 'd', 'p' and 'q'). We hypothesized that the evolution of sequential P3 components throughout each trial would exhibit a pattern of build-up associated with participants' decision-making process. In particular, we hypothesized that the P3 evoked by each stimulus would reflect the accumulated evidence for that hypothesis (+Ev, evidence for Action; -Ev, evidence for No-Action), and thus encode a decision variable driving participants' choices on a trial-by-trial basis. Since the time of decision in the No-action trials is unknown, all stimuli up to the end of the trial were included for the sequential P3 analysis. In Action trials, only stimuli up to the time of decision were included.

To study the dynamics of evidence accumulation as encoded in the P3, we extracted the P3 at Pz in response to every instance of relevant evidence ('b', 'd', 'p' and 'q'), and we obtained the average amplitude of the whole P3 component (0.3 to 0.8 s post stimulus). Fig. 3a shows the evolution of sequential stimulus-locked P3 amplitudes, as a function of the time within the trial that each target letter appeared. In Action trials, an overall increase in the P3 amplitude is observed from trial onset, peaking at the time of the action. This increase is driven by the increase in the P3 amplitude in response to evidence in favour of the eventual decision to act (+Ev, Fig. 3b). In No-Action trials, the time of any decision to withhold action is unknown, since participants were not required to produce any response. However, the evolution of the P3 amplitude throughout those trials exhibited a pattern similar to that observed in Action trials. That is, the overall P3 amplitude increased from trial on-



**Fig. 3. The p3 encodes a categorical decision variable.** Trials in which participants pressed a key to change the colour of the screen are labelled “Action” trials, and include p3 data only up to the time of movement. Trials where no movement was executed are labelled “No-Action” trials, and include p3 data up to the end of the trial. **a.** Evolution of p3 amplitude at Pz evoked by sequential evidence items from trial onset, pooled across participants and types of evidence in No-action (*top*) and Action (*bottom*) trials (loess smoothed,  $\pm$ SE). The amplitude of the p3 increased over time in all conditions. Single-participant data for this panel can be found in Fig. S2. **b.** Evolution of the p3 amplitude at Pz from trial onset, pooled across participants for each type of evidence separately in No-Action (*top*) and Action (*bottom*) trials (loess smoothed,  $\pm$ SE). Vertical bars indicate mean time of Action in each condition. **c.** Evolution of p3 amplitude at Pz aligned to the time of action in Action trials, pooled across participants and types of evidence ( $\pm$ SE). **d.** Evolution of the p3 amplitude at Pz aligned to the time of action in Action trials, pooled across participants for each type of evidence separately ( $\pm$ SE). The amplitude of the p3 only increased significantly in response to +Ev. **e.** ERP traces (*left*) show the grand-average ( $\pm$ SEM) p3 in No-Action and Action trials, for +Ev and -Ev (positive values are plotted down). Bar graphs (*right*) show the trial-by-trial averaged difference ( $\pm$ SE) between +Ev and -Ev (+Ev minus -Ev), showing that evidence in favour of the selected option (i.e. +Ev in Action, -Ev in No-Action) evoked significantly higher p3 amplitudes, on average, within any given trial. **f.** Grand-averaged ( $\pm$ SEM) p3 evoked in response to the first (*left*) and last (*right*) piece of +Ev and -Ev presented before action in Action trials only (positive values are plotted down). Single-participant data for this panel can be found in Fig. S3.

set, reaching a peak between 10 and 20 s, and was characterised by a greater increase in the P3 amplitudes evoked by evidence supporting the eventual decision (i.e. in this case evidence *against* action, -Ev).

To investigate the build-up dynamics of decision-relevant information prior to action, we aligned the P3 data to the time of action on Action trials only (Fig. 3c, d). We fitted a linear multilevel regression to predict the amplitude of the P3 evoked by each target letter (*bd/pq*) using the time to action (in seconds), the type of letter that evoked it

(+Ev/-Ev), participants’ decision (Action/No-Action), the condition (Informative/NH/NL) and all possible interactions as predictors. An interaction between the time to action and type of evidence ( $X^2_{(1)} = 4.52$ ,  $p = 0.033$ , uncorrected ( $p = 0.077$ , BH-corrected)) indicated that the amplitude of the P3 tended to increase in response to +Ev towards the time of action ( $X^2_{(1)} = 18.32$ ,  $p < < 0.001$ ), but not in response to -Ev ( $X^2_{(1)} = 0.42$ ,  $p = 0.512$ ). To visualise this gradual increase of the P3 in response to +Ev towards the time of action, we plotted the ERPs for the

3 target letters immediately preceding action (Fig. S4). See Table S1 for full statistical details, and Fig. S5 for the unsmoothed data used in the analysis.

We further analysed the trial-by-trial P3 amplitudes to investigate whether the selective increase in P3 amplitude would be reflected in the averaged ERPs (Fig. 3e). We averaged all P3 amplitudes for each trial (i.e. up to the time of action in Action trials, and up to the end of the trial in No-action trials) and fitted a linear mixed model to predict the mean P3 amplitude based on the type of evidence (+Ev/-Ev), the decision (Action/No-Action) and the condition (Informative/NH/NL), ignoring the time at which evidence was presented. We found an interaction between the decision and the type of evidence ( $X^2_{(1)} = 31.45$ ,  $p < 0.001$ ). In trials in which participants decided to switch the colour of the screen (Action), the amplitude of the P3 was significantly higher ( $X^2_{(1)} = 18.10$ ,  $p < 0.001$ ) in response to +Ev ( $M = 2.73 \mu V$ ,  $SEM = 0.43 \mu V$ ) than to -Ev ( $M = 0.86 \mu V$ ,  $SEM = 0.47 \mu V$ ). In trials in which participants decided to not switch the colour of the screen (No-Action), the effect was in the opposite direction ( $X^2_{(1)} = 12.80$ ,  $p = 0.001$ ). The amplitude of the P3 in response to +Ev ( $M = 1.16 \mu V$ ,  $SEM = 0.38 \mu V$ ) was, on average, lower than that in response to -Ev ( $M = 1.96 \mu V$ ,  $SEM = 0.36 \mu V$ ). There was no three-way interaction between condition, type of evidence and decision ( $X^2_{(2)} = 3.22$ ,  $p = 0.465$ ), suggesting that the effect was present in both easy (Informative) and difficult (Neutral, NH & NL) trials (see Table S2 for full statistical details). This shows that the P3 encoded an internal decision variable reflecting the most likely decision, rather than simply the absolute amount of external decision-relevant evidence. While the momentary evidence at the time of action in the Neutral conditions did favour acting when participants moved (Fig. S7), the balance of evidence supporting that choice was significantly lower in Neutral trials than in the Informative condition. Note that given the trial-by-trial variability, the effect is more apparent in the trial-by-trial differences (bar graphs in Fig. 3e, right) than in the grand-averaged ERPs (Fig. 3e, left).

Finally, we compared the initial and final amplitudes reached by the P3 in Action trials (Fig. 3f). If the P3 genuinely encodes a signed or 'categorical' decision variable, accumulator models predict that differences between alternative options should be minimal at the beginning of the decision-making process, and maximal at the time a decision is made. That is, while the P3 responses to the first (initial) pieces of +Ev and -Ev presented in a trial should not differ, the P3 responses to the last (final) pieces of +Ev and -Ev should reflect the decision participants will make. We fitted a mixed regression to predict the P3 amplitude based on the time within the trial (first/last), the type of evidence (+Ev/-Ev) and the condition (Inf/NH/NL). Although an interaction between evidence type and time within trial did not survive within-model multiple comparisons correction ( $X^2_{(1)} = 3.95$ ,  $p = 0.046$  uncorrected,  $p = 0.109$ , BH-corrected), we further explored differences in P3 amplitude evoked by the first and last different types of evidence presented in a trial as originally planned. While at the beginning of the trial the P3 responses to +Ev and -Ev did not significantly differ ( $X^2_{(1)} = 2.73$ ,  $p = 0.149$ ), the P3 responses to the last piece of +Ev were significantly higher than responses to the last piece of -Ev ( $X^2_{(1)} = 22.43$ ,  $p < 0.001$ ). These results suggest that if any small differences in P3 amplitudes evoked by initial evidence were present, those differences were more robust by the end of a trial, as illustrated in Fig. 3f.

### 6.3. The P3 predicts choices on a single-trial level

The previous analyses were based on post-hoc groupings of trials according to participants' choice to act or not act. Thus, they remain inconclusive regarding the predictive power of the P3 with respect to participants' choices on a single trial level. Here, we used a simple classification heuristic to predict choices based on the average P3 amplitudes evoked by relevant letters. The sign of the difference between +Ev and -Ev averaged throughout the whole trial predicted participants' choices to act or not to act ( $X^2_{(1)} = 21.09$ ,  $p < 0.001$ ), across all conditions.

Classification accuracy did not differ across conditions: ( $X^2_{(2)} = 3.31$ ,  $p = 0.190$ ). Further, we found a significant linear relationship between the signed magnitude of the difference and the probability that participants would act or not ( $X^2_{(1)} = 18.77$ ,  $p < 0.001$ ; see Fig. 4b). Again, classification accuracy did not differ across conditions ( $X^2_{(2)} = 3.25$ ,  $p = 0.196$ ).

We further investigated how this predictive power progressively emerged over time by looking at the differences in the P3 amplitude evoked by consecutive items. Fig. 4c shows that classification accuracy first improved over time as the trial progressed, and then decreased again. This shows that the discriminative power of the P3 as a marker of choice increases over time, as the decision progresses, and then decreases, perhaps reflecting the fact that a decision has been made by this point. To get a closer insight into these temporal dynamics, we focused on the Action trials only, and tested whether predictive accuracy varied depending on whether predictions were made before or after action, how long before or after it (time in s) for each condition (Informative/NH/NL; see Fig. 4d). The predictive accuracy of the P3 amplitudes evoked by competing evidence items significantly improved prior to action ( $X^2_{(1)} = 24.05$ ,  $p < 0.001$ ), while there was no evidence of an effect of time on predictive accuracy post-action ( $X^2_{(1)} = 1.47$ ,  $p = 0.671$ ). Prediction accuracy based on P3s evoked after action was significantly worse overall than predictions based on P3 occurring before movement ( $X^2_{(1)} = 28.81$ ,  $p < 0.001$ ). These single-trial analyses support the idea that P3 modulations track an evidence accumulation process related to decision-making as it unfolds (see Table S3 for full statistical details).

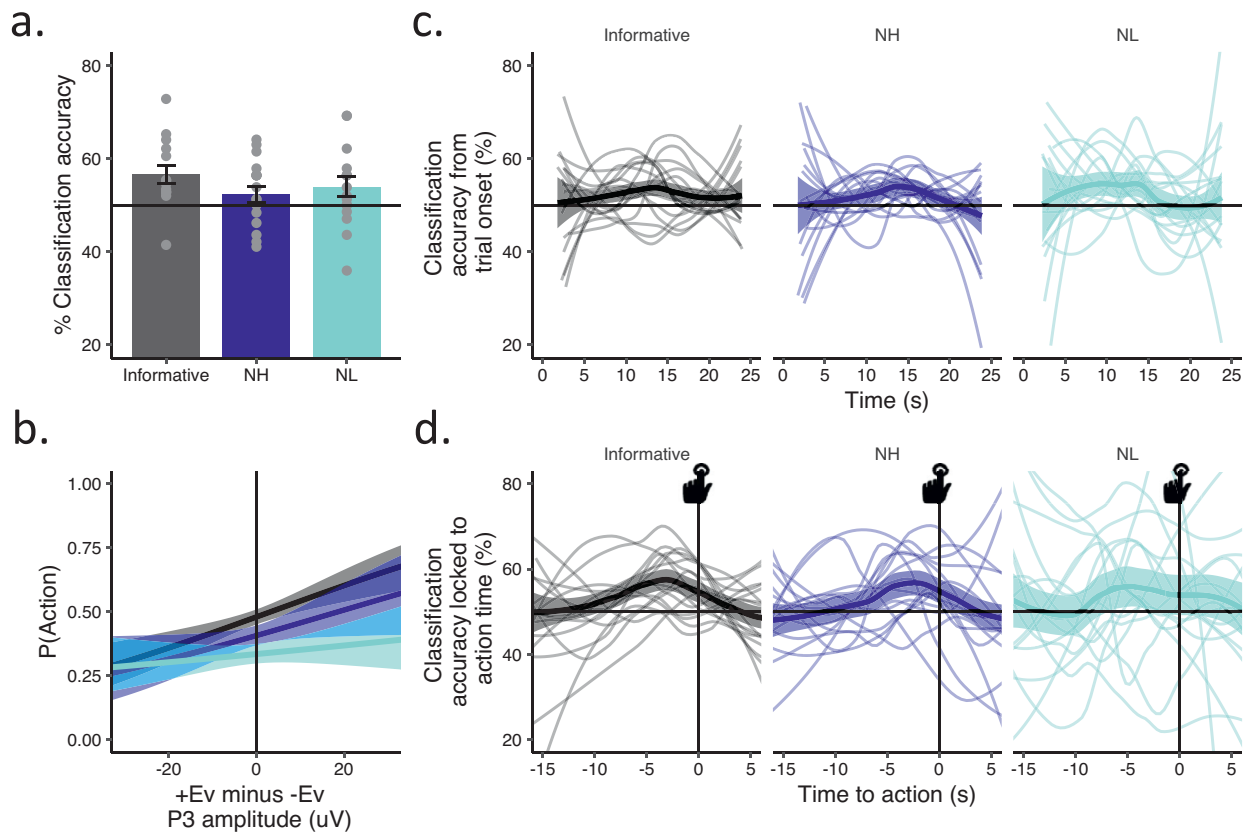
### 6.4. The evolution of the categorical decision variable encoded by the P3 is modulated by a context-dependant urgency signal

We next explored whether participants' behaviour and the dynamic changes in the P3 response could be captured by computational models of evidence accumulation. We modelled the integration of evidence over time using two competing accumulators (Usher and McClelland, 2001), and we compared the most simple form of the model to two additional models featuring additional time-dependant urgency signals that multiplicatively modulated the integration of incoming evidence.

In our first model (Constant gain), the decision parameters were not allowed to vary across conditions and the weight given to incoming evidence (gain) was constant over time. Based on previous literature (Braunlich and Seger, 2016; Cisek et al., 2009; Murphy et al., 2016; Thura and Cisek, 2014), in the second model (Rising urgency) we incorporated a growing urgency parameter that dynamically modulated gain. Due to the multiplicative interaction postulated by this model, as urgency monotonically increased over time, the gain for both evidence accumulators also increased. This time-varying gain means that, as time elapses, less evidence is required to reach a decision. It has been suggested that such a modulation of evidence integration by a growing urgency signal might facilitate decision-making and provide a better account of behavioural and neural data in decision-making tasks (Thura et al., 2012). Finally, in the third model (Context-dependant urgency), we further allowed the urgency parameter, and thus the increase in gain over time, to vary separately for each condition. Condition-specific modulations in an internal urgency signal could potentially account for the different decision thresholds observed in the NL and NH conditions in our task (Fig. 2b).

Both urgency models performed significantly better than the first model, but did not differ significantly from each other. In aggregate, the Context-dependant Urgency (total AIC = 6638.40) was numerically better than the Rising Urgency model (AIC = 6664.76). Further, the context-dependant urgency model was the best performing one for the majority of participants ( $n = 12$  of 19). This suggests that condition-specific modulations of evidence accumulation occurred, and that they may be explained by a context-dependant urgency signal. Full details of the model specification, comparison and parameter estimation can be found in the





**Fig. 4. Decoding choice based on P3 amplitudes.** **a.** Mean ( $\pm$ SEM) percentage of trials that were accurately classified as Action or No-Action trials based on the whole-trial averaged P3 amplitudes. Overlaid dots correspond to single-subject average data. **b.** Probability of action as a function of the difference in P3 amplitudes evoked by +Ev and -Ev items on each trial (mean  $\pm$ SE). Positive values in the x axis indicate that the average amplitude of the P3 evoked by +Ev was greater than that evoked by -Ev. Increasing values lead to increasing probability of action. Negative values are associated with greater P3s evoked by -Ev compared to +Ev, and stronger negative differences increase the probability that participants did not move. **c.** Group mean ( $\pm$ SEM, *thick lines*) and individual subjects (*thin lines*) classification accuracy from trial onset. Predictions were made every time a stimulus was presented by comparing the P3 amplitude evoked by it to the P3 amplitude by the immediately preceding item of competing evidence. For example, an “Action” prediction was made if the P3 evoked by a +Ev item at time  $t = 3$  was greater than the P3 evoked by a -Ev item presented at time  $t = 2$ . **d.** Group mean ( $\pm$ SEM, *thick lines*) and individual subjects (*thin lines*) classification accuracy aligned to the time of action (vertical line). The accuracy of P3-based predictions significantly improved towards the time of action. After action execution, the predictive value of the P3 was significantly smaller than that preceding it. Note: data in A, B and C include both Action and No-Action trials; data illustrated in D and the corresponding analysis include Action trials only. Line graphs in C and D are loess smoothed.

Supplementary Material. Below we describe the *Context-dependent Urgency* model in more detail.

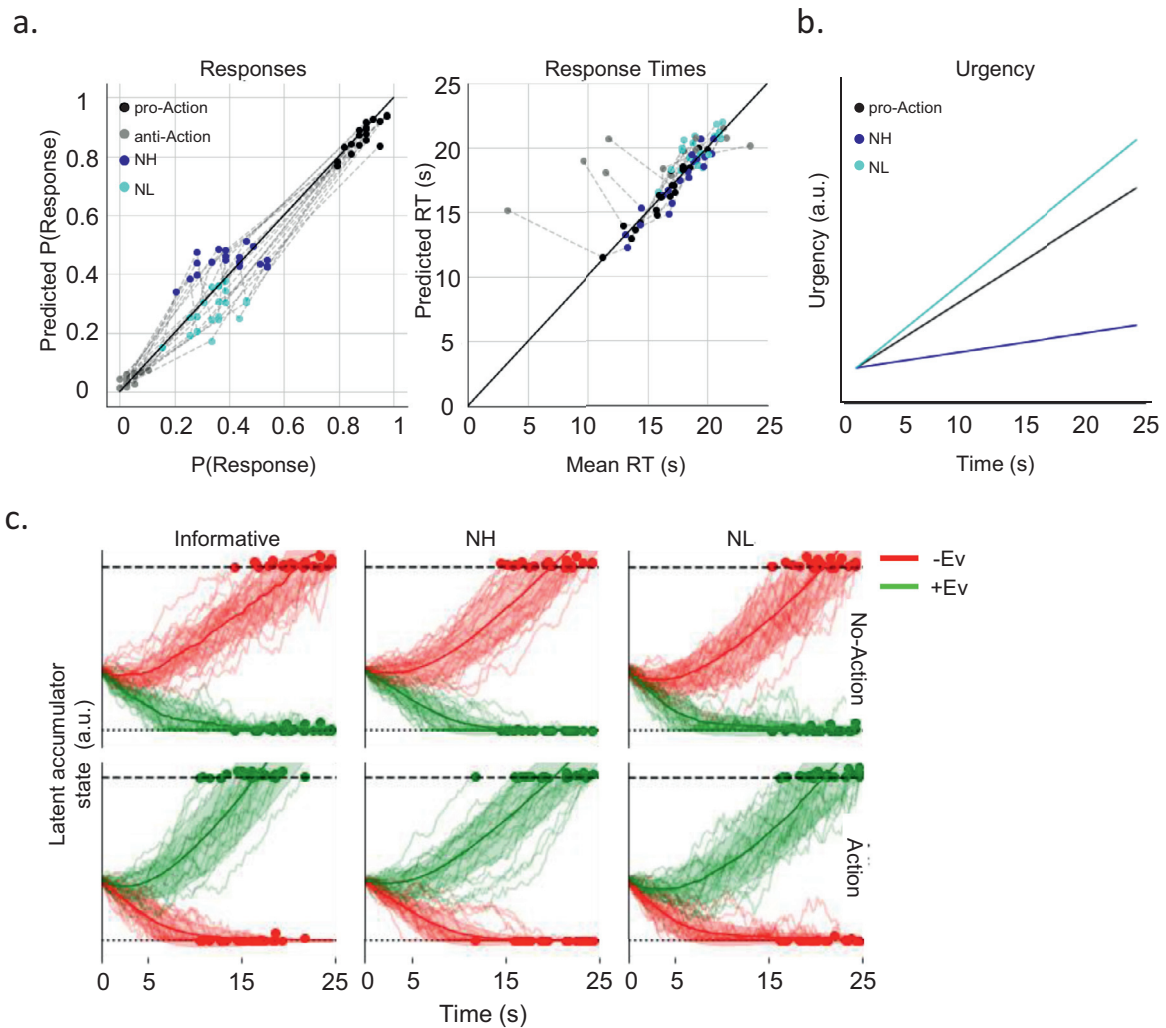
The context-dependant urgency model accurately predicted the probabilities of action and the reaction times for all 19 participants (Fig. 5a). The estimated urgency parameters in this model indicated that the weight given to individual pieces of evidence was greater in conditions where evidence was less frequent (Fig. 5b). When the rate of evidence was high, the estimated urgency was low (Urgency<sub>NH</sub> = 0.009). Instead, when the rate of evidence was low, the estimated urgency was higher (Urgency<sub>NL</sub> = 0.045). This means that while in the NH condition a decision variable evolved towards a threshold mostly driven by the accumulation of abundant external evidence, in the NL condition the decision variable approached the threshold thanks to a strong contribution from an internal urgency signal, which boosted the impact of the scarce available evidence in the accumulation process. Such a modulation can account for the fact that participants acted after having seen significantly fewer evidence items in the NL condition as compared to the NH one (see Fig. 2a). The inhibition parameter was on average greater than 0 ( $M = 0.19$ ,  $SD = 0.12$ ,  $t_{(18)} = 6.58$ ,  $p < 0.001$ ), and was greater than 0 for 16 out of 19 participants, indicating lateral inhibition between accumulators.

To explore the dynamics of the model, we set each parameter to its average across participants and we simulated 1000 trials per condition

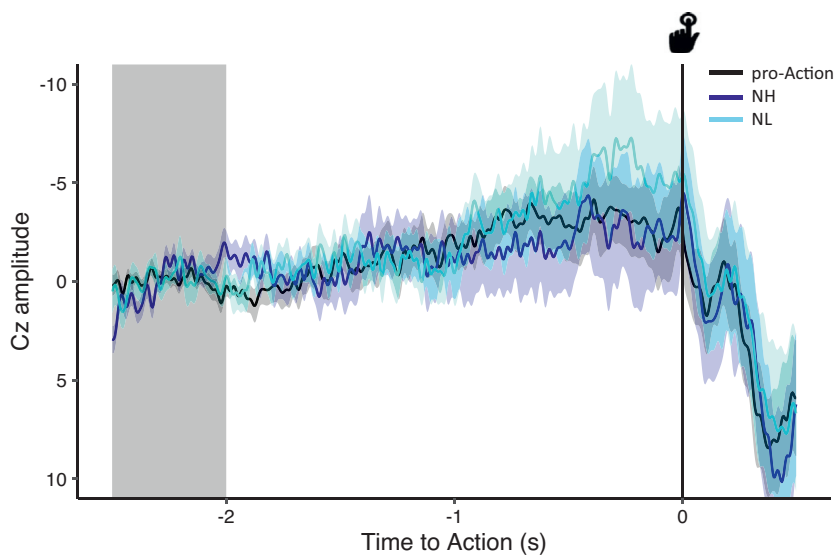
(Fig. 5c). These simulations broadly recapitulate the P3 evolution during trials (Fig. 2b). On trials where participants act, the activity of the +Ev accumulator (P3 responses to +Ev stimuli) is stronger than that of the -Ev accumulator (responses to -Ev stimuli). On trials where participants do not act, this effect is reversed. Furthermore, the difference between the two accumulators is weak at the start of each trial and increases over time.

#### 6.5. Markers of self-paced action in evidence-informed decisions

Finally, we investigated the neural activity preceding action execution to test whether the readiness potential would appear in all, some or no conditions in our task. We found an RP preceding actions in all conditions. We performed pairwise cluster comparisons between all conditions including the two seconds preceding the time of action (-2 to 0). An RP-like shape was visible in all conditions (Fig. 6), and pairwise cluster comparisons between all conditions found no significant clusters in any of the comparisons (all  $p > 0.05$ ), suggesting that the RPs were not significantly different in amplitude in the different conditions. Given the lack of differences between conditions in the neural precursors of action, which have been linked to subjective experience of intention (Libet et al., 1983), behavioural reports of subjective time of decision were not further analysed.



**Fig. 5. Behavioural predictions and dynamic simulations of the context-dependant urgency model.** a. Observed and predicted response probabilities (*left*) and response times (*right*) for each condition separately. b. Estimated evolution of the urgency signal in arbitrary units (a.u.) for each condition separately. The urgency was estimated to be highest for the NL condition, where evidence was scarce, and lowest for the NH condition, where evidence was abundant. c. Simulated accumulator states over time recapitulate our analyses of P3 amplitudes in response to +Ev and -Ev evidence (compare Fig. 2b). The state of the +Ev accumulator is greater than that of the -Ev accumulator on trials when action occurs (*bottom*, Action), and vice versa when no action takes place (*top*, No-Action). The magnitude of this effect increases as the time advances.



**Fig. 6. The RP preceding evidence-informed voluntary actions.** Grand-averaged EEG activity ( $\pm$ SEM) at electrode Cz preceding actions in all conditions. Shaded grey area indicates baseline period (-2.5 to -2 s prior to action).

## 7. Discussion

The ability to decide when faced with ambiguous information is crucial for functional voluntary behaviour. In this experiment, we developed a paradigm to investigate the neural mechanisms enabling decision and action under uncertainty. We focussed on the P3 as a marker of evidence accumulation and investigated the RP as a neural precursor of self-paced action. We show that the P3 components evoked by sequential evidence items encode a decision variable driving participants' choices on a trial-by-trial basis. Further, we show that a combination of competitive evidence accumulation and context-sensitive urgency can account for two seemingly different behavioural symmetry-breaking strategies: a primacy effect, where choices are strongly driven by early evidence, and reaction-like actions, mostly influenced by immediate evidence at late stages in the trial.

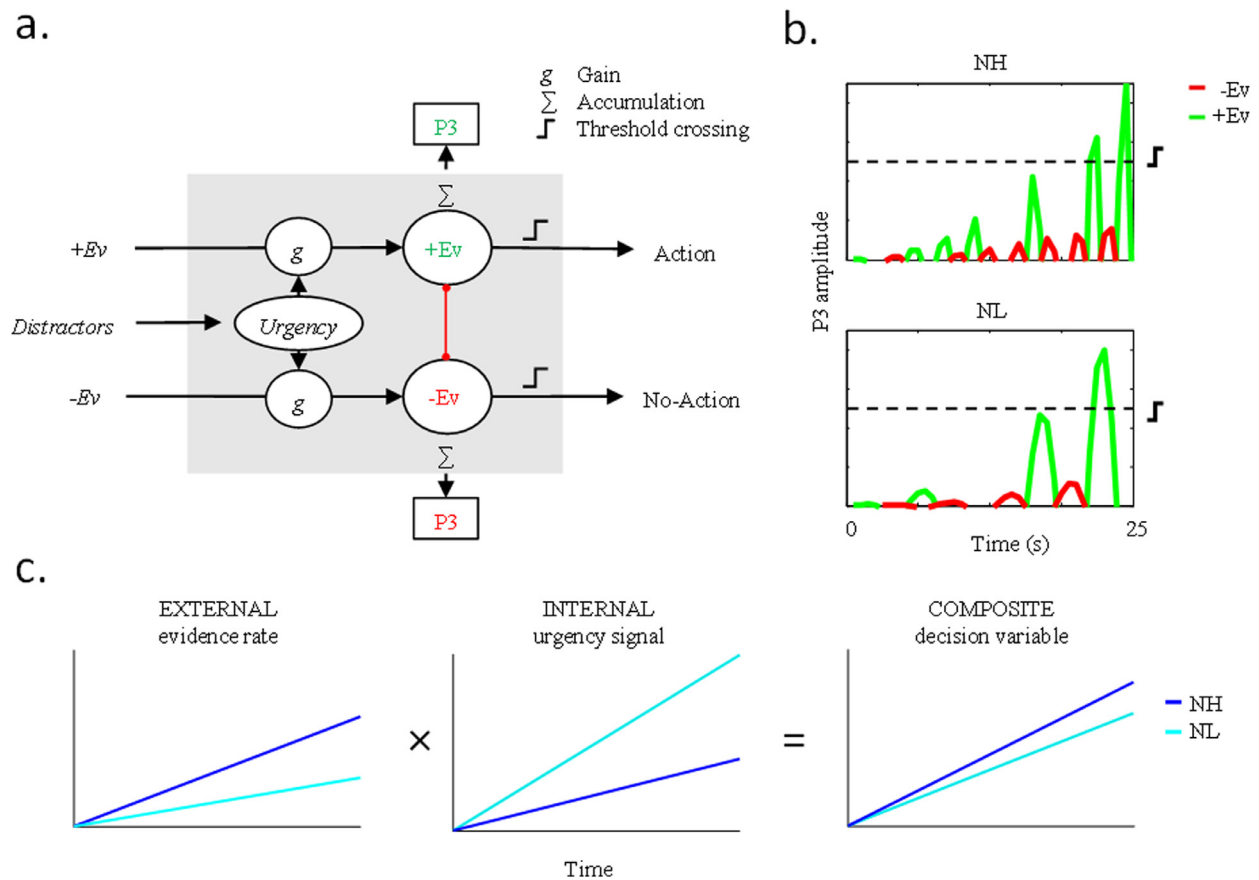
To investigate whether we could track how competing options evolve during a decision-making process, we used P3 responses evoked by sequential pieces of evidence. Previous research has shown that the morphology of a single P3 ERP tracked reaction times (Twomey et al., 2015). However, that study required responses to individual evidence items, and did not involve any evidence accumulation. In contrast, in our task, participants did not have to respond to individual evidence items – rather, they had to perform a sequential sampling task involving multiple successive stimuli and evaluate the evidence across the entire sequence. Our new ‘snapshot’ measure approach allowed us to use the individual response evoked by discrete stimulus items to probe the cumulative state of a signed internal accumulator at different time points during the trial as a decision unfolded. We found that the P3 responses evoked by each individual target stimuli encoded information about participants' impending choices (Fig. 3e). In trials where participants decided to act, the amplitude of the P3 in response to evidence favouring Action (+Ev) was significantly greater than the P3 amplitude in response to evidence favouring No-Action (-Ev). This was due to a selective increase in the P3 response to selected evidence as a decision was formed (Fig. 3). This modulation was visible both in the Informative and Neutral conditions, suggesting that the P3 captures the state of an internal decision rather than the objective evidence available to participants. Interestingly, in Informative conditions, we found that the P3 evoked by the most frequent targets grew progressively larger in amplitude as the stimulus sequence progressed. This goes against theories that link the P3 to surprise. Surprise theories would predict decreasing P3 amplitude as a stimulus becomes less unexpected (Neville et al., 1982), and would also predict an inverse relation between P3 amplitude and stimulus probability (Duncan-Johnson and Donchin, 1977). Further, these dynamics recall the activity of neural populations in monkey's lateral intra-parietal (LIP) area during decision-making tasks (Churchland et al., 2008; Huk and Shadlen, 2005; Roitman and Shadlen, 2002), in which the activity of neural populations selective for alternative choices diverges over time, and builds up to threshold only for the population encoding the eventual choice. Further, the fact that the P3 amplitude reached a peak well before the end No-Action trials suggests that participants did not continue to accumulate evidence until the end of the letter stream, but rather made a covert decision to *not* move before hand and then stopped accumulating evidence. This is in agreement with previous literature showing that once a decision bound is reached, further evidence is not further processed for perceptual decisions (Balsdon et al., 2020; Kiani et al., 2008).

The P3 could further be used to predict participants' choices on a single trial basis (Fig. 4). Predictive power was at chance level at the beginning of the trial and progressively converged to a maximum at the time of action (see Fig. 4c). Further, predictive accuracy decreased sharply after movement in Action trials. This suggests that the P3 does indeed track the evolution of an action-related decision, rather than merely post-hoc correlate with the eventual choice. While the P3-based classification accuracy was modest (the whole-trial P3 average accuracy across conditions was  $M = 54.76\%$ ,  $SEM = 1.15\%$ ), it was significantly better

than chance, and comparable to classification performance in other neuroimaging datasets (e.g. Hubbard et al., 2019). Similar P3 dynamics and predictive power were visible in all conditions, including neutral trials which contained no net evidence. Since the P3 evolution could nevertheless be used to predict responses on a single-trial level, it seems to track the state of an internal decision variable rather than objective evidence. However, the P3 cannot track *only* the mere amount of accumulated evidence items, since the amplitude of the P3 evoked by, say, the third evidence item in the NH condition is substantially lower than that evoked by the same item in the NL condition (see Fig. S6). Therefore, some additional factor must contribute to the evolution of the decision variable encoded by the P3. We hypothesized that this additional factor could correspond to the concept of *urgency*.

Urgency is often conceived as an evidence-independent, time-growing signal that facilitates making a decision in various contexts by modulating the amount of evidence required for a decision as time elapses (Braunlich and Seger, 2016; Churchland et al., 2008; Ditterich, 2006; Thura, 2020), and it has been shown to depend on the basal ganglia (Bogacz et al., 2010; Thura and Cisek, 2017). Here, we followed studies showing a multiplicative interaction between evidence accumulation processes and urgency signals (Cisek et al., 2009; Ditterich, 2006; Thura and Cisek, 2014). In our task, conditions differed not only in the balance, but also in the amount of evidence available. Our best-fitting model suggested that participants' behaviour was best explained by a competing accumulation process with a condition-specific gain adjustment, which we postulated to be mediated by a growing and multiplicative urgency signal. When evidence was scarce (NL), the urgency signal was estimated to be highest (Fig. 5b). This effectively constitutes an adaptive strategy for reaching a decision when evidence is rare: by increasing the gain in such contexts, each new piece of evidence makes an enhanced contribution to the decision variable. Thus, fewer evidence items will therefore suffice to reach a decision threshold (see Fig. 2a). In turn, reducing gain in contexts where evidence is abundant allows the accumulation of more pieces of evidence, thus preventing potentially hasty and erroneous decisions. Thus, a multiplicative urgency parameter can naturally account for the behavioural adaptations observed in our task, and the neural dynamics tracked by the P3. If the P3 reflected urgency only, we would expect the P3 traces to be modulated by the time at which they were evoked, but not by which type of evidence evoked it. However, we show that the evolution of the P3 evoked by category-specific stimuli diverged according to participants' choices. Thus, the P3 seems to reflect not only the growing urgency signal, but *also* evidence accumulation. Our results are compatible with recent suggestions that decision-related activity may be modulated by urgency (Steinemann et al., 2018). It remains unclear whether such urgency effects occur directly at the evidence accumulation stage itself, or rather result indirectly from accumulation of early sensory representations that have been modulated by urgency at the encoding stage.

In sum, based on the neural data and our computational modelling, we suggest that the amplitude of the P3 evoked by any given stimulus depends on 1) the current state of an internal evidence accumulator 2) a modulating factor given by a context-dependant urgency signal growing over time. Fig. 7 illustrates a schematic of the cognitive mechanisms underlying the decision-making process that we believe the P3 tracks. We suggest that the P3 reflects the activity of category-specific neural populations processing accumulating evidence. This is consistent with the fact that distinct neural populations specifically represent different different decision alternatives in the monkey parietal cortex (Fitzgerald et al., 2011; Freedman and Assad, 2006; Swaminathan and Freedman, 2012). Finally, modelling suggests that the availability of evidence modulates the rate at which an internal urgency signal grows, which in turn modulates the gain ( $g$ ) of the evidence accumulators. As a decision-making process about the distribution underlying the stimuli evolves (i.e. are *bd* more frequent than *pq*?), so the accumulated evidence ( $\Sigma$ ) moves closer to threshold. Context-specific modulations of urgency signals have important implications: when evidence is abun-



**Fig. 7. Conceptual model of the cognitive mechanisms underlying the P3.** **a.** Suggested computational mechanisms driving evidence accumulation in our task. Evidence is processed sequentially, and each stimulus is processed with a certain gain ( $g$ ) in a category-specific accumulator. The accumulators for competing choices are mutually inhibitory. The gain in turn depends on a time-varying urgency signal. The rate at which urgency, and hence gain, increases over time is a function of the availability of relevant evidence. The more evidence is available, the smaller the steepness of the urgency signal, allowing the accumulation of more evidence items before reaching a threshold. Once one of the accumulators reaches a threshold a decision is made. **b.** Schematic representation of the evoked P3 responses in a condition where evidence is abundant (NH, *top*) or rare (NL, *bottom*). The evolution along the x axis illustrates the P3 amplitudes over time, which are a function of a competitive evidence external accumulation process modulated by a growing internal urgency signal. In turn, this enables crossing the threshold after fewer evidence items have been presented. The fact that the P3 amplitudes evoked by +Ev and -Ev are different despite the absolute numbers of evidence being the same shows the effect of lateral inhibition. **c.** Schematic of the multiplicative interaction between external evidence and context-dependant internal urgency signal to produce a composite decision variable. When evidence is abundant (NH), the urgency signal grows slowly and the evolution of the decision variable is mostly driven by external evidence. When evidence is scarce (NL), the urgency signal grows faster and its contribution to the evolution of the decision variable is greater.

dant, decision variables evolve up to threshold driven mostly by external information, even when this evidence is neutral. In turn, when evidence is scarce, decision variables would rise to threshold driven mostly by an internal urgency signal that boosts evidence processing. Can this model account for the observed behavioural patterns of symmetry-breaking?

In neutral trials, participants were more likely to decide to act if the first target letter that they saw provided evidence for action, rather than inaction – irrespective of subsequent evidence (Fig. 2b). Competing accumulator models like the ones we explored predict this kind of primacy effect (Tsetso et al., 2012). Small differences at the beginning of an accumulation process are amplified over time thanks to the mutually inhibiting connections between competing accumulators. In the absence of strong countermanding evidence, decision outcomes tend to follow early hypotheses. In our study, these asymmetries depended on the time at which external stimuli were presented. However, early brain activity has been reported to contain information that enables decoding free decisions in humans (Soon et al., 2008), and it has been suggested that free decisions depend on accumulation processes similar to those operating in perceptual ones (Bode et al., 2013). Thus, early asymmetries of either endogenous or exogenous origins can effectively bias accumulation processes and thus act as symmetry-breakers.

In addition, we found that actions typically shortly followed the presentation of a +Ev item (Fig. S1). That is, the time of actions was dependant on immediate evidence to a certain extent in all conditions. This is to be expected, since participants' decisions were based on evidence. However, the influence of immediate evidence on the time of actions was strongest in the NL condition, where external evidence was particularly rare (Fig. 2d and Fig. S1). This finding can be explained by the condition-specific modulations of urgency that we modelled. The increased urgency estimated for the NL condition means that the impact of any individual piece of evidence on the decision variable is comparatively high. Therefore, the probability that any given evidence item will make the decision variable cross the threshold is highest in that condition, accounting for the recency effects.

We also investigated whether differences in the neural precursors of action could account for the observed behaviour. Classic interpretations of the RP suggest that it reflects endogenous neural activity leading to action. Research has typically measured the RP in experiments where the participant acts as and when they wish, without any imperative stimulus or explicit sensory evidence to trigger the decision to act (e.g. Libet et al., 1983). However, recent studies have suggested that its presence may depend on the arbitrariness of action rather than its intentionality (Maoz et al., 2019). Our task was a hybrid between classic

spontaneous action paradigms and cued-reaction tasks: participants had no strong time pressure to execute actions, and although their actions were informed by the evidence, they were not triggered by any specific individual stimulus. Readiness potentials were visible preceding actions in all conditions, and the amplitudes did not significantly differ across conditions (Fig. 6). The fact that an RP was also visible preceding actions in the pro-Action condition, where actions were deliberate and strongly determined by the environment is not compatible with the idea that the stochastic accumulation of internal noise that the RP reflects (Schurger et al., 2012) precedes actions *only* in contexts where actions are arbitrary (Maoz et al., 2019). Thus, while the RP may not act as a symmetry-breaking mechanism, the fact that it can be found preceding deliberate actions suggests it may be related to the precise timing of self-paced actions rather than their arbitrariness. Further, we found no differences in RP amplitude between our various evidence conditions. This suggests that the rate at which the RP builds up does not mirror the strength or amount of external evidence and may rather depend on purely internal signals. However, further research is required to clarify the potential relationship between external evidence strength and RP shape.

### Data and code availability

Code and data supporting the analyses reported in this study will be made available at [http://github.com/Elisabeth-Pares/NIMG\\_Pares-Pujolras\\_et\\_al\\_2021](http://github.com/Elisabeth-Pares/NIMG_Pares-Pujolras_et_al_2021)

### Acknowledgments

This work was partly supported by a joint grant from the John Templeton Foundation and the Fetzer Institute (to PH). The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation or the Fetzer Institute. Furthermore, this work was supported by European Research Council Advanced Grant, HUMVOL, awarded to PH.

### Authors' contributions

EPP and PH conceptualized and designed the task. EPP and YA conducted the experiment. EPP and ET analysed the data. EPP, ET, and PH wrote the manuscript.

### Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.neuroimage.2021.117863](https://doi.org/10.1016/j.neuroimage.2021.117863).

### References

Balsdon, T., Wyart, V., Mamassian, P., 2020. Confidence controls perceptual evidence accumulation. *Nat. Commun.* 11. doi:10.1038/s41467-020-15561-w.

Bates, D., Mächler, M., Bolker, B.M., Walker, S.C., 2015. Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* 67 (1), 251–264. doi:10.18637/jss.v067.i01.

Benjamini, Y., Heller, R., Yekutieli, D., 2009. Selective inference in complex research. *Philosoph. Trans. R. Soc. A Math. Phys. Eng. Sci.* 367 (1906), 4255–4271. doi:10.1098/rsta.2009.0127.

Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc.* 57 (1), 289–300.

Blankertz, B., Acqualagna, L., Dähne, S., Haufe, S., Schultze-Kraft, M., Sturm, I., Ušćumlić, M., Wenzel, M.A., Curio, G., Müller, K.R., 2016. The Berlin brain-computer interface: progress beyond communication and control. *Front Neurosci.* (NOV) 10. doi:10.3389/fnins.2016.00530.

Bode, S., Bogler, C., Haynes, J.-D., 2013. Similar neural mechanisms for perceptual guesses and free decisions. *Neuroimage* 65, 456–465. doi:10.1016/j.neuroimage.2012.09.064.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., Cohen, J.D., 2006. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113 (4), 700–765. doi:10.1037/0033-295X.113.4.700.

Bogacz, R., Wagenmakers, E.J., Forstmann, B.U., Nieuwenhuis, S., 2010. The neural basis of the speed-accuracy tradeoff. *Trends Neurosci.* 33 (1), 10–16. doi:10.1016/j.tins.2009.09.002.

Brainard, D.H., 1997. The psychophysics toolbox. *Spat Vis* 10 (4), 433–436. doi:10.1163/156856897X00357.

Braunlich, K., Seger, C.A., 2016. Categorical evidence, confidence, and urgency during probabilistic categorization. *Neuroimage* 125, 941–952. doi:10.1016/j.neuroimage.2015.11.011.

Brown, S.D., Heathcote, A., 2008. The simplest complete model of choice response time: linear ballistic accumulation. *Cogn. Psychol.* 57, 153–178. doi:10.1016/j.cogpsych.2007.12.002.

Burnham, K.P., Anderson, D.R., 2001. Kullback-Leibler information as a basis for strong inference in ecological studies. *Wildlife Res.* 111–119 [print] 28(2).

Carland, M.A., Thura, D., Cisek, P., 2019. The urge to decide and act: implications for brain function and dysfunction. *Neuroscientist* 25 (5), 491–511. doi:10.1177/1073858419841553.

Churchland, A.K., Kiani, R., Shadlen, M.N., 2008. Decision-making with multiple alternatives. *Nat. Neurosci.* 11 (6), 693–702. doi:10.1038/nn.2123.

Cisek, P., Puskas, G.A., El-Murr, S., 2009. Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* 29 (37), 11560–11571. doi:10.1523/JNEUROSCI.1844-09.2009.

de Lafuente, V., Jazayeri, M., Shadlen, M.N., 2015. Representation of accumulating evidence for a decision in two parietal areas. *J. Neurosci.* 35 (10), 4306–4318. doi:10.1523/JNEUROSCI.2451-14.2015.

Deecke, L., Kornhuber, H.H., 1978. An electrical sign of participation of the mesial “supplementary” motor cortex in human voluntary finger movement. *Brain Res.* doi:10.1016/0006-8993(78)90561-9.

Delorme, A., Makeig, S., 2004. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J. Neurosci. Methods* 134 (1), 9–21. doi:10.1016/j.jneumeth.2003.10.009.

Ditterich, J., 2006. Evidence for time-variant decision making. *Eur. J. Neurosci.* 24 (12), 3628–3641. doi:10.1111/j.1460-9568.2006.05221.x.

Duncan-Johnson, C.C., Donchin, E., 1977. On quantifying surprise: the variation of event-related potentials with subjective probability. *Psychophysiology* 14 (5), 456–467.

Fitzgerald, J.K., Freedman, D.J., Assad, J.A., 2011. Generalized associative representations in parietal cortex. *Nat. Neurosci.* 14 (8), 1075–1079. doi:10.1038/nn.2878.

Forstmann, B.U., Ratcliff, R., Wagenmakers, E., 2016. Sequential sampling models in cognitive neuroscience: advantages, applications, and extensions. *Annu. Rev. Psychol.* 67, 641–666. doi:10.1146/annurev-psych-122414-033645.

Freedman, D.J., Assad, J.A., 2006. Experience-dependent representation of visual categories in parietal cortex. *Nature* 443 (7107), 85–88. doi:10.1038/nature05078.

Hubbard, J., Kikumoto, A., Mayr, U., 2019. EEG decoding reveals the strength and temporal dynamics of goal-relevant representations. *Sci. Rep.* 9 (1), 1–11. doi:10.1038/s41598-019-45333-6.

Huk, A.C., Shadlen, M.N., 2005. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* 25 (45), 10420–10436. doi:10.1523/JNEUROSCI.4684-04.2005.

Jahanshahi, M., Jenkins, H., Brown, R.G., Marsden, C.D., Passingham, R.E., Brooks, D.J., 1995. Self-initiated versus externally triggered movements I. An investigation using measurement of regional cerebral blood flow with PET and movement-related potentials in normal and Parkinson's disease subjects. *Brain* 118, 913–933.

Jasper, H.H., 1958. The Ten-Twenty Electrode System of the International Federation. *951 Electroencephal. Clin. Neurophys.* 10, 371–375.

Keller, I., Heckhausen, H., 1990. Readiness potentials preceding spontaneous motor acts: voluntary vs. involuntary control. *Electroencephalogr. Clin. Neurophysiol.* 76 (4), 351–361. doi:10.1016/0013-4694(90)90036-J.

Kelly, S.P., O'Connell, R.G., 2013. Internal and external influences on the rate of sensory evidence accumulation in the human Brain. *J. Neurosci.* 33 (50), 19434–19441. doi:10.1523/jneurosci.3355-13.2013.

Khalighinejad, N., Schurger, A., Desantis, A., Zmigrod, L., Haggard, P., 2018. Precursor processes of human self-initiated action. *Neuroimage* 165, 35–47. doi:10.1016/j.neuroimage.2017.09.057, (July 2017)https://doi.org/.

Kiani, R., Hanks, T.D., Shadlen, M.N., 2008. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci. Offic. J. Soc. Neurosci.* 28 (12), 3017–3029. doi:10.1523/JNEUROSCI.4761-07.2008.

Kutas, M., McCarthy, G., Donchin, E., 1977. Augmenting mental chronometry: the p300 as a measure of stimulus evaluation time. *Science* 197 (4305), 792–795. doi:10.1126/science.887923.

Libet, B., Gleason, C.A., Wright, E.W., Pearl, D.K., 1983. Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). *Brain* 106 (3), 623–642. doi:10.1093/brain/106.3.623.

Maoz, U., Yaffe, G., Koch, C., Mudrik, L., 2019. Neural precursors of decisions that matter – an ERP study of deliberate and arbitrary choice. *Elife* 1–32. doi:10.1101/097626, https://doi.org/https://doi.org/.

Maris, E., Oostenveld, R., 2007. Nonparametric statistical testing of EEG- and MEG-data. *J. Neurosci. Methods* 164 (1), 177–190. doi:10.1016/j.jneumeth.2007.03.024.

Murakami, M., Shteingart, H., Loewenstein, Y., Mainen, Z.F., 2017. Distinct sources of deterministic and stochastic components of action timing decisions in rodent frontal cortex. *Neuron* 94, 908–919. doi:10.1016/j.neuron.2017.04.040.

Murakami, M., Vicente, M.L., Costa, G.M., Mainen, Z.F., 2014. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat. Neurosci.* 17 (11), 1574–1582. doi:10.1038/nn.3826.

Murphy, P.R., Boonstra, E., Nieuwenhuis, S., 2016. Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nat. Commun.* 7 (May), 1–14. doi:10.1038/ncomms13526.

Neville, H., Snyder, E., Woods, D., Galambos, R., 1982. Recognition and surprise alter the human visual evoked response. *Pnas* 79 (March), 2121–2123.

- O'Connell, R.G., Dockree, P.M., Kelly, S.P., 2012. A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat. Neurosci.* doi:10.1038/nn.3248, October.
- Oostenvelde, R., Fries, P., Maris, E., Schoffelen, J.M., 2011. FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput. Intell. Neurosci.* doi:10.1155/2011/156869.
- Papa, S.M., Artieda, J., Obeso, J.A., 1991. Cortical activity preceding self-initiated and externally triggered voluntary movement. *Movement Disord.* 6 (3), 217–224. doi:10.1002/mds.870060305.
- Ratcliff, R., Smith, P.L., Brown, S.D., McKoon, G., 2016. Diffusion decision model: current issues and history. *Trends Cogn. Sci. (Regul. Ed.)* 20 (4), 260–281. doi:10.1016/j.tics.2016.01.007.
- Roitman, J.D., Shadlen, M.N., 2002. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22 (21), 9475–9489. doi:10.1523/jneurosci.22-21-09475.2002.
- Schulze-Kraft, M., Birman, D., Rusconi, M., Allefeld, C., Görden, K., Dähne, S., Blankertz, B., Haynes, J.-D., 2016. The point of no return in vetoing self-initiated movements. *Proc. Natl. Acad. Sci. U.S.A.* 113 (4), 1080–1085. doi:10.1073/pnas.1513569112.
- Schurger, A., Sitt, J.D., Dehaene, S., 2012. An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proc. Natl. Acad. Sci. U.S.A.* 109 (42), E2904–E2913. doi:10.1073/pnas.1210467109.
- Shadlen, M., Newsome, W., 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* 86 (4), 1916–1936.
- Shadlen, N., Britten, H., Newsome, T., Movshon, Anthony, J., 1996. A computational analysis of the relationship between neuronal and behavioural responses to visual motion. *J. Neurosci.* 16 (4), 1486–1510.
- Soon, C.S., Brass, M., Heinze, H.-J., Haynes, J.-D., 2008. Unconscious determinants of free decisions in the human brain. *Nat. Neurosci.* 11 (5), 543–545. doi:10.1038/nn.2112.
- Steinemann, N.A., O'Connell, R.G., Kelly, S.P., 2018. Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. *Nat. Commun.* 9 (1), 3627. doi:10.1038/s41467-018-06117-0.
- Swaminathan, S.K., Freedman, D.J., 2012. Preferential encoding of visual categories in parietal cortex compared with prefrontal cortex. *Nat. Neurosci.* 15 (2), 315–320. doi:10.1038/nn.3016.
- Swaminathan, S.K., Masse, N.Y., Freedman, D.J., 2013. A comparison of lateral and medial intraparietal areas during a visual categorization task. *J. Neurosci.* 33 (32), 13157–13170. doi:10.1523/JNEUROSCI.5723-12.2013.
- Thura, D., 2020. Decision urgency invigorates movement in humans. *Behav. Brain Res.* 382, 112477. doi:10.1016/j.bbr.2020.112477, September 2019.
- Thura, D., Beauregard-Racine, J., Fradet, C.W., Cisek, P., 2012. Decision making by urgency gating: theory and experimental support. *J. Neurophysiol.* 108 (11), 2912–2930. doi:10.1152/jn.01071.2011.
- Thura, D., Cisek, P., 2014. Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* 81 (6), 1401–1416. doi:10.1016/j.neuron.2014.01.031.
- Thura, D., Cisek, P., 2017. The Basal Ganglia do not select reach targets but control the urgency of commitment. *Neuron* 95 (5), 1160–1170.e5. doi:10.1016/j.neuron.2017.07.039.
- Thura, D., Cos, I., Trung, J., Cisek, P., 2014. Context-dependent urgency influences speed-accuracy trade-offs in decision-making and movement execution. *J. Neurosci.* 34 (49), 16442–16454. doi:10.1523/JNEUROSCI.0162-14.2014.
- Tsetsos, K., Gao, J., McClelland, J.L., Usher, M., 2012. Using time-varying evidence to test models of decision dynamics: bounded diffusion vs. The leaky competing accumulator model. *Front. Neurosci.* 6 (JUN), 1–17. doi:10.3389/fnins.2012.00079.
- Turner, B.M., Sederberg, P.B., 2014. A generalized, likelihood-free method for posterior estimation. *Psychonom. Bull. Rev.* 21 (2), 227–250. doi:10.3758/s13423-013-0530-0.
- Twomey, D.M., Murphy, P.R., Kelly, S.P., O'Connell, R.G., 2015. The classic P300 encodes a build-to-threshold decision variable. *Eur. J. Neurosci.* 42 (1), 1636–1643. doi:10.1111/ejn.12936.
- Usher, M., McClelland, J.L., 2001. The Leaky, Competing Accumulator Model.pdf. *Psychol. Rev.* 108 (3), 550–592. <https://pdfs.semanticscholar.org/1c53/f6fdcef2afa88afe77332a9aa0ba2eb347a8.pdf>.
- Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Vázquez-Baeza, Y., 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* 17 (3), 261–272. doi:10.1038/s41592-019-0686-2.