bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Stimulus reliability automatically biases temporal integration of discrete perceptual targets yielding suboptimal decisions

Short title: Reliability-weighted integrative decision-making

Authors: **Dragan Rangelov[1], Rebecca West[2], and Jason B. Mattingley[1,2,3]**

[1]Queensland Brain Institute, The University of Queensland, St Lucia, QLD Australia

[2]School of Psychology, The University of Queensland, St Lucia, QLD Australia

[3]Canadian Institute for Advanced Research (CIFAR), Toronto, Canada

Corresponding author: Dragan Rangelov

d.rangelov@uq.edu.au

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Abstract

Decision making is a ubiquitous cognitive process that determines choice behaviour. In recent years there has been increased interest in how information about multiple discrete sensory events are combined in support of single, integrated decisions. Previous studies have shown that integrative decision-making is biased in favour of more reliable stimuli. As reliability-weighted integration typically mimics statistically optimal integration, it remains unclear whether reliability biases are automatic or strategic. To dissociate reliability-weighting and optimal decisions, we developed a task that required participants to monitor two successive epochs containing brief, suprathreshold coherent motion signals which varied in their reliability. Rather than judging the individual target motion directions, however, participants had to reproduce the *average* motion direction of the two targets. Using mixture distribution modelling and linear regression to model behavioural data, we found robust biases in favour of the more reliable stimulus, despite the fact that unbiased responses were optimal in our paradigm. Using population-tuning modelling to characterise feature specific brain activity recorded using electroencephalography, we observed robust and sustained feature-specific responses to target signals in both epochs. Using the same method, we were able to capture the temporal dynamics of integrated decision-making by characterising tuning to the average motion direction. Critically, the tuning profiles to the average motion direction exhibited biases in favour of the more reliable signal, in keeping with the modelled behavioural responses. Taken together, our findings reveal that temporal integration of discrete sensory events is automatically and suboptimally weighted according to stimulus reliability.

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Introduction

Decision-making is a ubiquitous cognitive process involved in any form of choice behaviour, from choosing to cross a busy street to choosing a life partner. Sequential-sampling evidence-accumulation models have had great success in modelling both choice behaviour and the neural correlates of decision-making, from single cell firing to the activity of large neuronal populations (1–6). According to these models, evidence in support of one or more choices accumulates in time toward a decision criterion, and the decision is made once the accumulated value reaches a criterion threshold. This architecture can account for response-time distributions of correct and error responses in a variety of behavioural tasks, from simple detection of sensory events to memory retrieval (7–9). Importantly, the neural activity recorded in both humans (10–12) and in animal models (13–15) closely mimics the time-course of the hypothetical decision variable, thereby lending neurobiological support to the idea that the brain accumulates evidence toward a decision threshold.

In recent years there has been increased interest in the cognitive and neural mechanisms underpinning more complex decision making (15–19). One pertinent issue is how multiple sources of evidence might be combined in support of a single decision. For example, to safely cross a busy street, one should consider the traffic coming from both sides of the road. This decision presumably engages at least two evidence-accumulation processes, each of which should converge on the same decision, namely, whether to cross or to wait. While much has been learned about decision making in relation to single stimuli – in our example, monitoring cars on just one side of the street – little is known about how the brain integrates two (or more) distinct sources of evidence into one decision. Here we characterised the cognitive and neural mechanisms underpinning such 'integrative' decision-making.

To investigate integrative decision-making, most studies to date have used a variant of the 'redundant signals' paradigm (17,18,20–22) in which a stream of multisensory stimuli – typically auditory and visual pulses – are concurrently presented and observers have to discriminate whether the number of pulses in either stream is lower or higher than an arbitrary criterion. The pulse counts in the two streams vary independently, so that, at the end of trial, the streams can either support the same decision (i.e., congruent trials), or different decisions. Importantly, the streams are temporally jittered so that they cannot be integrated into a single multisensory event at the level of initial sensory encoding. Instead, each stream should engage a separate evidence accumulator and yield two decisions that are integrated at later stages of the processing hierarchy. Following this rationale, a comparison between single-signal trials (unimodal) and redundant-signals trials (multimodal) can reveal mechanisms of integrative decision-making. Typically, participants are more

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

accurate in *congruent* multimodal than in unimodal trials suggesting that the two decisions are indeed integrated at some stage. Critically, the higher accuracy in multimodal trials scales with stimulus reliability of the individual streams, suggesting that the integration process is biased in favour of the more reliable sensory stream. In fact, biases in integrative decision-making closely resemble statistically optimal signal integration (23,24) in which the contributions of individual signals are weighted by their reliability so as to yield statistically optimal decisions.

While previous research has demonstrated that integrative decision-making is subject to biases based on factors such as stimulus reliability, it remains unclear whether these biases are automatic or strategic. One ubiquitous finding in the literature on integrative decision-making is that some observers do not integrate decisions at all, but rather rely exclusively on signals of higher reliability (17,18,20,22). This finding suggests that integrating decisions may be subject to higher order influences such as a trade-off analysis of increased accuracy at the expense of increased effort. It is important to note that the redundant signals paradigm, in which deciding on the basis of a single stream already affords accurate decisions, is especially vulnerable to such higher order effects. While characterising higher order biases is an important open issue (25), it remains unclear whether integrating several discrete decisions automatically favours sources of higher reliability in tasks where the integration is essential, rather than opportunistic.

To address this important issue, we developed a task which required explicit integration of two simple visual decisions on brief periods of coherent motion in successive stimulus displays (Fig. 1a). On every trial, we presented two epochs of coherently moving dots separated by 1 s of randomly moving dots. At the end of the trial the task was to reproduce the direction of the ***average*** target motion. To illustrate, if a trial contained successive motion directions toward 10 o'clock and 2 o'clock, participants should indicate an *average* motion direction of 12 o'clock. Participants made their decisions without time constraints by adjusting the orientation of a response dial. To manipulate stimulus reliability, motion coherence in the first and second epoch could either be low (40% of coherently moving dots) or high (80%). Critically, as we explain below, these motion coherence values were deliberately chosen to be well above normal motion coherence thresholds. Target reliability across the two epochs of each trial was factorially combined so that different combinations (low/low, low/high, high/low, and high/high) were presented equally often and in a random order. The main goal was to characterise the effects of stimulus reliability on cognitive and neural mechanisms of integrative decision-making.

To adjudicate between automatic versus strategic reliability-weighted biases in integrative decision-making, it is important to demonstrate that unbiased integration is, in principle, the optimal strategy. If a motion direction is particularly difficult to discern at low coherence and very easy at

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

high coherence, for example, then unbiased averaging of a very noisy and a very precise representation would yield relatively poor performance. In other words, the specific combination of low and high coherences will determine the optimality of unbiased integration. For this reason, we used motion coherence levels that were well above threshold (40% and 80%) as opposed to threshold-level stimuli, which are typically around 6–7% for most observers (26). The supra-threshold coherence levels we used should afford the best possible motion discrimination for both low and high coherence signals, thus permitting unbiased source integration. Any bias in favour of signals of higher reliability would strongly suggest that reliability-weighted integration of discrete decisions is automatic. To ensure that different coherence levels afforded similar response precision, in Experiment 1 we compared error magnitudes for reproducing *single* motion directions of low and high coherence.

In contrast to commonly used forced-choice paradigms (1,2,4), which are optimised for characterising the temporal dynamics of decision-making, the reproduction task provided us with a continuous, feature-specific read-out of integrated decisions. In that sense, our paradigm is complementary to forced-choice paradigms; by measuring the difference between the expected and the reproduced average-motion direction on each trial, we could use mixture distribution modelling (27,28) and linear regression (16,29,30) to characterise behavioural biases in integrative decision-making. To characterise the neural correlates of decision-making, we recorded brain activity using electroencephalography (EEG). We were primarily interested in measuring feature-specific brain responses to presented motion signals using population-tuning modelling of brain activity (31–36). Critically, our experimental design also permitted us to use the same method to decode feature-specific brain responses to the *average* motion direction. This quantity must be internally computed by integrating the representations of the two target signals and, as such, must be closely related to the mechanisms of integrative decision-making. If stimulus reliability affects integration processes, then population-tuning responses to the average motion direction should depend on the combination of motion coherence levels across target epochs within a trial.

Using a similar analytical approach, previous studies (35,37) have successfully characterised feature-specific brain responses to both task-relevant and task-irrelevant stimuli. Although weaker than neural responses to relevant stimuli, robust responses to irrelevant stimuli suggest that the population-tuning analyses have a strong sensory component. To gauge the degree to which decoded motion-specific responses reflect sensory processing, on every trial of the motion-averaging task we presented two overlaid patches of dots in different, distinctive colours. At the beginning of every trial, a colour cue indicated the dot-patch that would contain the target-motion direction; the other dot-patch served as a concurrent distractor which could be ignored. Different,

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

uncorrelated motion signals were briefly (500 ms) and concurrently presented in both target and distractor patches. Introducing distractor signals enabled us to dissociate sensory responses – which should be comparable for target and distractor signals – and decision-related brain responses – which should be more prominent for targets than distractors.

Unlike commonly used decision-making paradigms in which all stimuli are task-relevant, contrasting target and distractor processing should elucidate the role of selective attention in integrative decision-making. On the basis of the vast literature on selective attention (38), we expected target motion signals to influence behavioural measures of integrative decision-making more strongly than distractor signals. What is less clear, however, is whether and in what way selective attention might affect feature-specific brain responses to distractor signals. While unlikely, it is possible that the brain represents both relevant and irrelevant sensory input with equal fidelity, predicting a comparable degree of population tuning to target and distractor signals. A more likely outcome (39,40) is that selective attention modulates the temporal dynamics of tuning to target and distractor motion, predicting suppression of distractor-related responses following an initial, sensory response to both target and distractor signals. By presenting both target and distractor signals and using population-tuning modelling, we could adjudicate between these two alternatives.

Results

In Experiment 1 (Fig. 1a, upper panel), we presented only one epoch containing a motion stimulus, and participants had to ascertain the target motion direction while ignoring a concurrently presented distractor motion event. By presenting a single epoch of coherent motion, we could test whether participants' ability to discern motion direction was comparable across low (40%) and high (80%) coherence levels. By presenting both target and distractor motion signals within each single-epoch trial, we could also ask whether selective attention operates differently at different coherence levels. Importantly, on every trial the strength of motion coherence of the concurrently presented target and distractor signals was the same, either both high coherence or both low coherence. To quantify overall task performance, we used mixture distribution modelling of error magnitudes (see Methods) to separately analyse the response precision (K) of noisy target responses and the proportion of random guesses (P_g). To quantify the degree to which individual target and distractor motion signals influenced the responses, we used linear regression (ordinary least squares; OLS) with complex-valued data. The absolute value of the regression coefficients associated with individual signals will reflect the degree to which target and distractor signals influenced the decision, or decision weights.

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

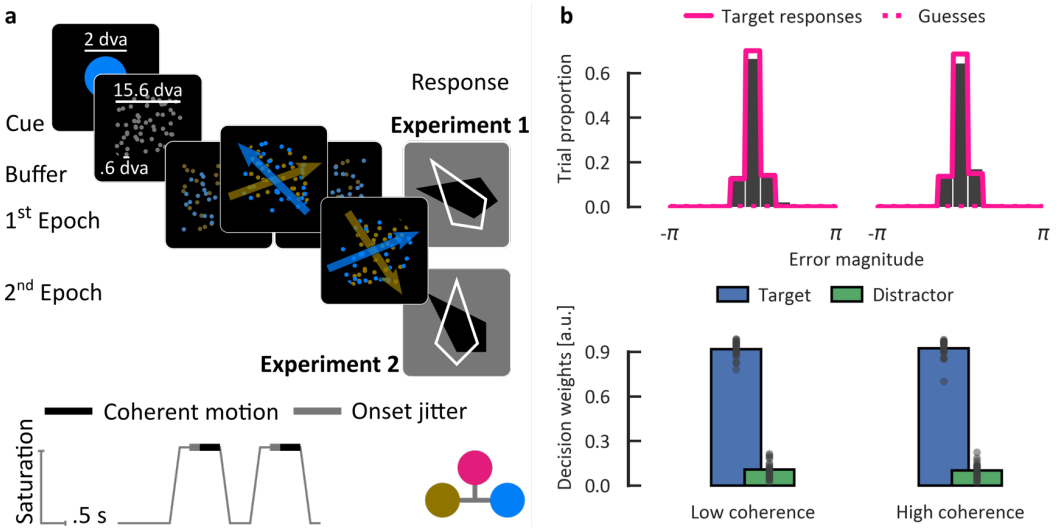

Figure 1. Overview of the experimental paradigm and the results of Experiment 1. (a) Upper panel shows a typical display sequence per trial. Lower panel shows the time-course of different events and possible colours. A coloured cue indicating the task-relevant colour (fixed per participant) was followed by a patch of grey dots moving randomly. After 1 s buffer periods, the colour saturation increased gradually to reveal two intermingled fields of distinctly coloured target and distractor dots. Coherent motion signals were presented for .5 s in both fields, jittered relative to the maximum saturation onset (.25–.5 s). In Experiment 1, only one epoch was presented and participants reproduced the target motion direction. In Experiment 2, two epochs were presented and participants reproduced the average motion direction of the two target motion stimuli. (b) Upper panel shows the histograms (bars) of the observed error magnitudes (expected – actual response) together with the predictions of the mixture-distribution modelling (pink lines) per motion coherence (low, high) in Experiment 1. Lower panel shows contributions of target- and distractor-motion signals to the response (i.e., the decision weights) per motion coherence (low, high) in Experiment 1. Dots represent individual participants.

Consistent with the supra-threshold levels of coherence we used, distributions of observed error magnitudes were unimodal and narrow (overall FWHM = 50°, Fig. 1b, upper panel) suggesting that participants performed the task with ease. The mixture distribution modelling yielded very close fits to the observed error magnitudes. The estimated response precision was high and comparable between low and high coherence motion targets ($K_{M/SEM}$ = 14.86/.71 and 13.73/.71, respectively, one-sample t_{22} = 1.12, $p_{one\text{-}tailed}$ = .137). Similarly, the estimated guessing rates were very low and comparable between coherences (both P_g = .03/.01, t_{22} < 1). Analyses of decision weights revealed that participants successfully focused on target signals and ignored concurrently presented distractors, as indicated by target weights which were around 9-times higher than concurrently presented distractors (.92 vs. .10, $F_{1,22}$ = 1,693, p < .001, η_p^2 = .99, Fig. 1b, lower panel). Critically,

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

both target and distractor weights were statistically indistinguishable across low and high coherence conditions (main effect of coherence $F_{1,22} < 1$, coherence x stimulus type $F_{1,22} = 1.34$, $p = .259$).

To summarise, Experiment 1 showed that participants were able to discern target motion direction with comparable precision at both low and high coherence values. Similarly, the efficiency of selective attention, as indexed by the difference between target and distractor weights, was comparable across low and high coherence trials. These results suggest that the low and high coherence levels we chose were both close to the asymptote of the psychometric curve. Consequently, unbiased integration of signals with low and high reliability should permit accurate integrative decisions.

In Experiment 2, we examined how temporally discrete perceptual decisions are integrated into a single decision. To this end, two epochs of coherent motion, rather than a single epoch as in Experiment 1, were presented in each trial. Participants had to reproduce the average motion direction of the two target signals while ignoring the concurrently presented distractors. We were primarily interested in characterising biases in behavioural decision weights as a function stimulus reliability. In addition, we recorded brain activity using EEG. To characterise the temporal dynamics of individual perceptual decisions, we quantified a well-documented neural correlate of decision-making (11,12,41), the centro-parietal positivity (CPP), time-locked to the onsets of both colour-saturation modulation and coherent motion signals (see Methods and Fig. 1a). Further, we characterised feature-specific neural responses to the presented motion using population-tuning modelling. Finally, using the same method, we quantified feature-specific brain activity related to the average motion direction. The average-motion responses should be closely related to the integration of the two target signals and, as such, are a good candidate for the neural correlate of integrative decision-making.

Similar to Experiment 1, the mixture distribution modelling of error magnitudes in Experiment 2 yielded close fits to the observed errors. Relative to Experiment 1, having to reproduce the average yielded lower response precision and higher guessing rates ($K = 8.16/.34$, $Pg = .09/.01$, Fig. 2a). Importantly, the variance of error magnitudes in the averaging task (Experiment 2) was equivalent to the simple linear sum of the variances estimated from the single target reproduction (Experiment 1) for all combinations of dot coherence across the two epochs (all $p_{FDR\text{-}corrected} > .05$). The fact that the variance of single signals was additive when averaging them indicates that participants performed the two tasks in a qualitatively similar way.

We next analysed the effect of dot coherence in different epochs on response precision and guessing rates. Consistent with Experiment 1, response precision was comparable for low and high coherence in the first epoch ($K = 8.12$ and 8.19, respectively, $F < 1$). By contrast, response precision

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

was significantly lower for low than high coherence targets in the second epoch (7.61 and 8.70, respectively, $F_{1,21}$ = 14.27, p = .001, $\eta_p^2 = .40$). The interaction between coherence levels in the first and second epochs was not significant (F < 1). Independently of epoch, participants were more likely to guess for low coherence signals than high (.10 and .06, overall, main effect of coherence in the first epoch $F_{1,21}$ = 25.75, p < .001, $\eta_p^2 = .55$, in the second $F_{1,21}$ = 13.80, p = .001, $\eta_p^2 = .40$, interaction F < 1).

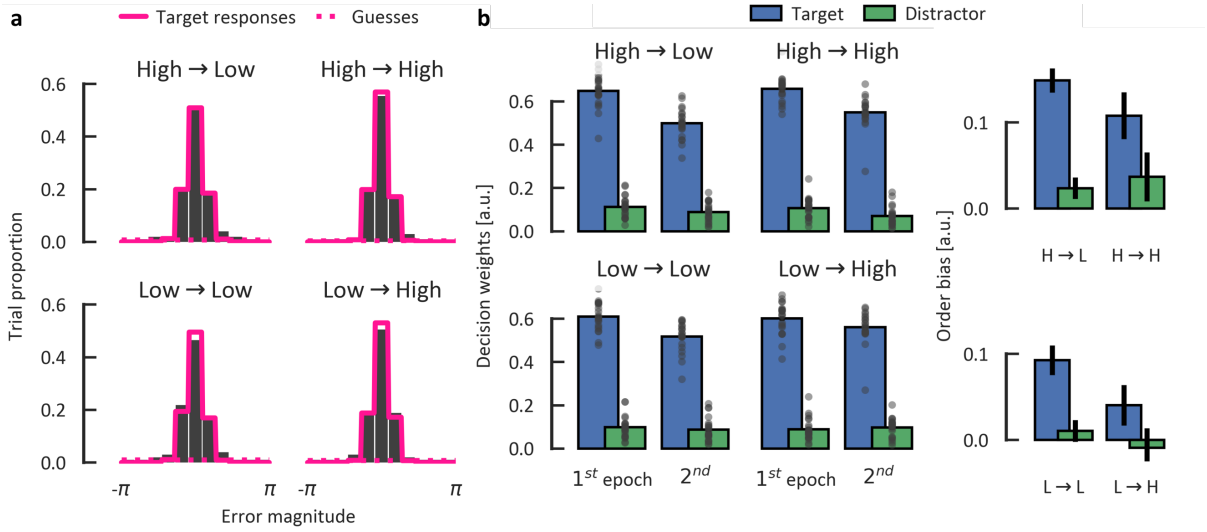

Figure 2. Behavioural results of Experiment 2. (a) Distributions of the observed error magnitudes (bars) in Experiment 2 together with the predicted distributions (pink lines) on the basis of mixture distribution modelling across different combinations of motion coherence (Low/High) in the first and second epoch (First epoch coherence → Second). (b) Left: Decision weights across different conditions of Experiment 2 separately per epoch and stimulus type. Right: Difference in respective decision weights for the first and second epoch separately per condition and stimulus type. Error bars denote within-participants standard errors of the mean (42). Conventions as in Figure 1.

Similar to Experiment 1, decision weights were significantly larger for target than distractor motion signals (.56 and .09, respectively, $F_{1,21}$ = 2,028, p < .001, $\eta_p^2 = .99$, Fig. 2b). Unexpectedly, however, both target and distractor weights were higher in the first epoch than in the second (.63 and .53 for targets, respectively, $F_{1,21}$ = 24.28, p < .001, $\eta_p^2 = .54$; .10 and .09 for distractors, $F_{1,21}$ = 4.60, p = .044, $\eta_p^2 = .18$). As the trials were relatively long (5 s), one might expect memory decay (29,43) to affect the first target more than the second, more recent target. The primacy bias we observed in Experiment 2 is opposite to what one might expect from simple memory decay, and might be indicative of how participants represented individual signals prior to integrating them. Most importantly, this order bias was further qualified by two-way interactions with motion coherence in both the first epoch and the second (Fig 2b, right panel). Specifically, whereas

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

increasing the coherence in the first epoch *increased* the order bias ($F_{1,21}$ = 16.30, p < .001, $\eta_p^2 = .44$), increasing the coherence in the second epoch *decreased* the bias ($F_{1,21}$ = 4.34, p = .049, $\eta_p^2 = .17$). The interaction of dot coherence with order bias indicates that averaging two signals was biased in favour of the signal of higher reliability. Importantly, this bias arose despite the fact that, consistent with the task instructions to ***average*** the two target signals, unbiased averaging should be the optimal response strategy. Moreover, the results of Experiment 1 confirmed that participants could perform equally accurate motion judgements on both high and low coherence signals when they were presented in single-epoch trials (see Fig. 1b).

To characterise the time-course of evidence accumulation in the motion-averaging task, we next analysed the CPP (Fig. 3), a positive deflection over centro-parietal electrodes that has been shown to closely mimic temporal dynamics of evidence accumulation (11,12,41). Visual inspection of the ERP topographies (Fig. 3a) revealed a pattern of central-medial positivity, consistent with the typical CPP topography. Visual inspection of the CPP time-course over the trial (Fig. 3b) revealed a phasic modulation that closely followed the trial sequence: during periods of grey, randomly moving dots, the CPP amplitude was at baseline. Two sharp deflections closely followed the onset of colour modulation in the first and the second epochs, and two broader deflections coincided with periods of coherent motion. This phasic modulation of the CPP suggests that several, temporally separable episodes of evidence accumulation took place within each trial. (By contrast, a single, continuous episode of evidence accumulation should have resulted in a sustained response across the trial, which was not observed.)

We next analysed shorter segments (500 ms) time-locked to the onsets of colour modulation and coherent motion (Fig. 3c and 3d, respectively). For both colour- and motion-locked epochs, the CPP deflection started at around 200 ms after the onset, consistent with the notion that the CPP is not merely a sensory-evoked response, but rather reflects higher level processes following sensory encoding (see the Intercept line in Fig. 3c and 3d, bottom panel). The colour-locked CPP was modulated only by the epoch, with a steeper rise and higher peak amplitude in the first epoch than the second (Fig. 3c). Analyses of the motion-locked CPP revealed a robust effect of motion coherence, with a steeper rise and higher peak amplitude for high coherence relative to low in both epochs (Fig. 3d). Importantly, only the coherence for the currently presented motion stimulus affected the CPP, as evidenced by a significant interaction between epoch (first/second) and the motion coherence per epoch (First coherence x Epoch and Second coherence x Epoch lines, bottom panel). This finding suggests that the motion-locked CPP reflects evidence accumulation in support of discriminating the currently presented motion target (i.e., a simple perceptual decision), rather than the averaging process (i.e., an integrated decision).

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

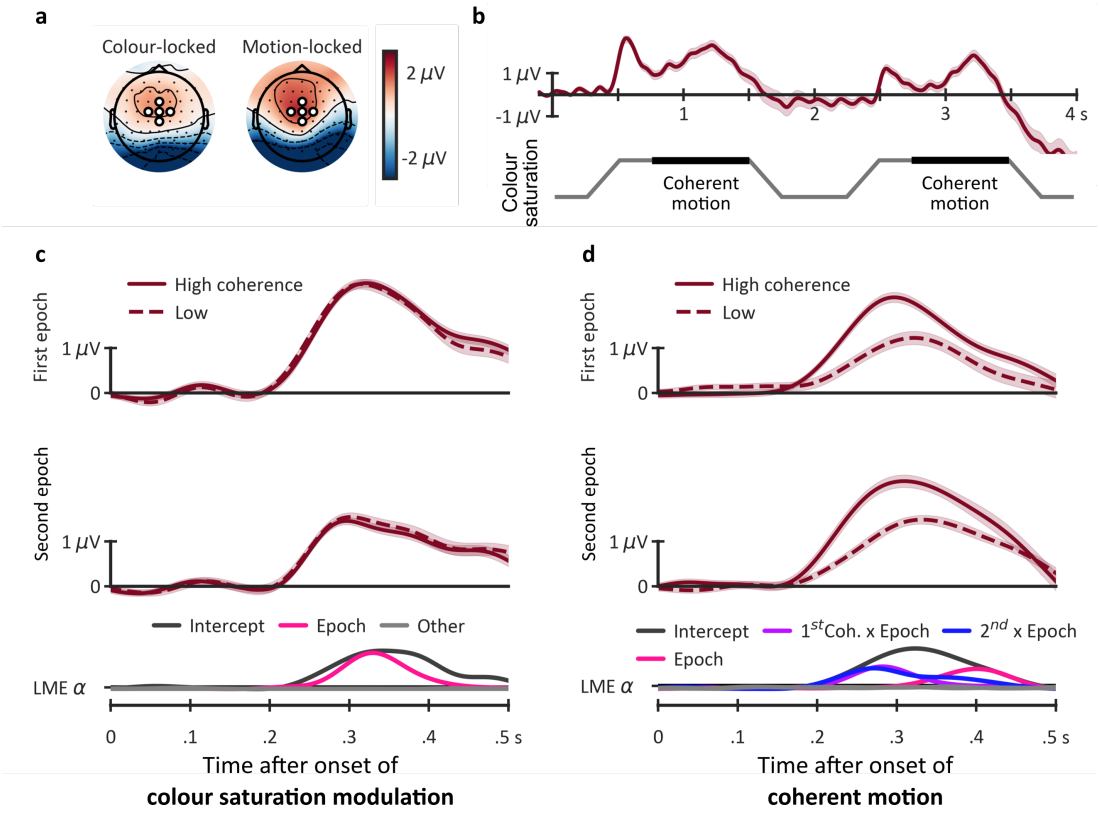

Figure 3. ERP results from Experiment 2. (a) ERP topography for colour-locked responses, showing the average of time-samples at which the colour saturation reached a maximum (.75 s and 2.75 s for the first and the second epoch, respectively); and for motion-locked responses, showing the average of time-samples at which the motion signal was most likely to be presented (at 1.25 and 3.25 s). Marked electrodes (white circles) were used for computing the CPP time-course. (b) The CPP time-course (upper panel) together with the time-course of different events within a trial (lower panel). (c) Colour-locked CPP during the first .5 s following the onset of colour modulation in the first epoch (upper panel) and the second epoch (central panel). Only random motion was presented during this period. The lower panel shows $-\log_{10}(p_{FDR\text{-}corrected})$ for different terms in the stepwise linear mixed effect model of colour-locked CPPs. (d) Motion-locked CPP. The colours were at maximum saturation during the analysed 500 ms. Conventions as in panel (b). For the purpose of presentation, all time-traces were low-pass filtered at 10 Hz using a Butterworth 4^{th} order infinite impulse-response filter. The analyses were performed on unfiltered data. Shaded areas denote ±1 within-participants SEM.

Whereas analyses of behavioural decision weights revealed robust interactions between the order bias and motion coherence across the two epochs, the ERP analyses suggest that the order effect and the coherence effects might have separable neural correlates. The order bias, evident in the colour-locked CPP, appears to reflect the process of selecting the target patch against the distractor patch – and, potentially, staying focused on the target patch – whereas the coherence effect, evident in the motion-locked CPP, appears to reflect the strength of the subsequently presented motion signals.

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

We next characterised time-resolved motion-specific responses to target and distractor signals (Fig. 4a) using population-tuning modelling (31–33) of the motion-locked EEG signals (see Methods for details). Inspection of tuning to target signals revealed a robust and sustained motion-specific response. The onset of significant motion-tuning coincided with the peak latency of the motion-locked CPP, suggesting that motion-tuning reflects the decision about the currently presented motion stimulus. There was no such tuning to distractor signals, which further supports the notion that motion tuning captures the dynamics of deciding about task-relevant signals, rather than stimulus-driven responses regardless of task relevance. Importantly, tuning to the target motion direction was sustained well after motion offset (indicated by dotted vertical lines), suggesting that representations of individual signals were maintained until both targets had been presented.

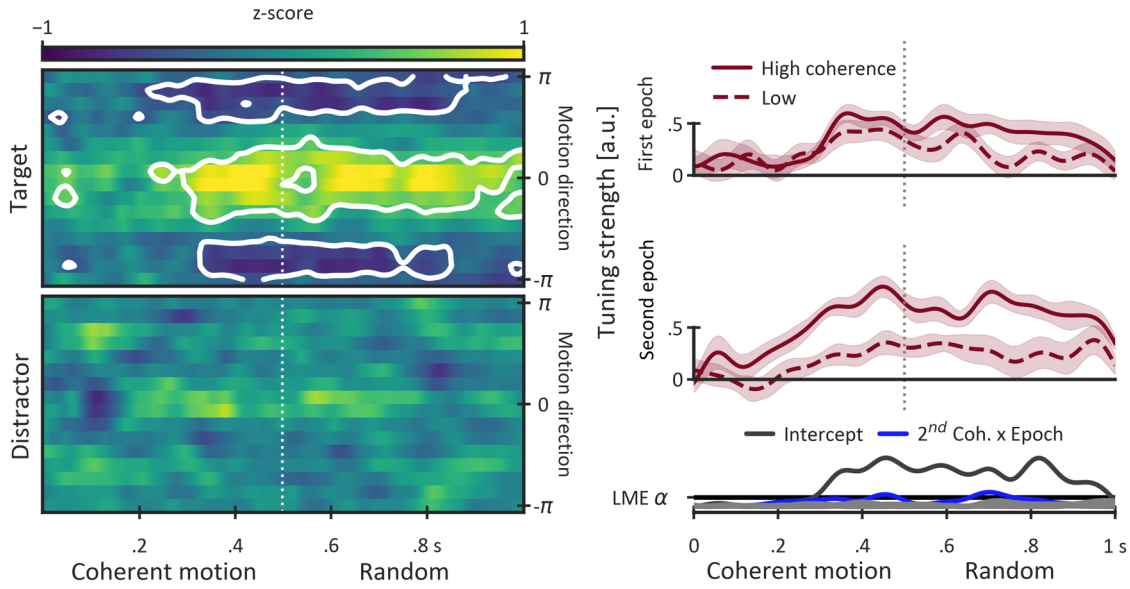

Figure 4. Population-tuning modelling of neural activity in response to targets and distractors in Experiment 2. (a) Time-resolved responses to target and distractor signals (upper and lower panels, respectively) of 16 hypothetical motion-specific channels spanning the full circle (from -π to +π) and centred on the actual presented motion signal. The trial-averaged profile of channel responses per participant and time-sample was z-scored across channels. White contours denote areas which were significantly different from 0 (FDR-corrected across all time-samples and channels). Significant modulation of channel-response profiles with yellow areas around 0 and blue areas around -π and +π indicate robust motion tuning to the presented signal. (b) Time-resolved tuning strength to *target signals* across first and second epochs (upper and lower panels, respectively). The tuning strength is an aggregate index of the channel-response profile (see Methods for details) with 0 indicating no tuning. Conventions as in Figure 3.

Inspection of the time-resolved tuning strength for *target signals* (Fig. 4b) across different epochs revealed comparable tuning to high and low coherence signals in the first epoch. This is

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

consistent with the observation of no effect of dot coherence in the first epoch on behavioural response precision, and it supports the notion that the two coherence levels afforded decisions of comparable precision. In keeping with the significant effect of dot coherence in the second epoch on behavioural response precision, the tuning strength in the second epoch was significantly higher for high coherence than for low coherence targets. These results, in conjunction with the lower decision weights and weaker colour-locked CPP responses in the second epoch relative to the first, suggest that the overall level of focused attention decreased from the first to the second epoch within a trial.

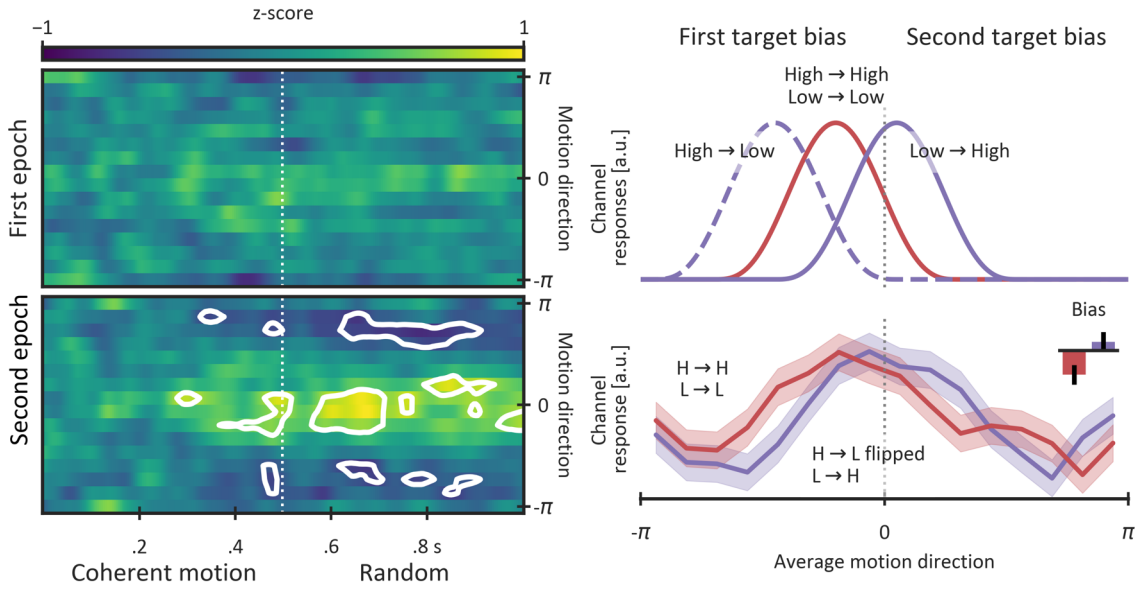

Figure 5. Population-tuning modelling of neural activity associated with the average motion direction in Experiment 2. (a) Time-resolved average-motion tuning in the first and the second epoch (upper and lower panels, respectively). Conventions as in Figure 4. (b) *Upper panel*: Expected shifts in the tuning profile reflecting the observed behavioural order bias. The motion channels were sorted relative to the first- and the second-presented motion direction so that shifts to the left from 0 reflect a first target bias. *Lower panel*: Observed shifts in the tuning profiles for the same-coherence trials (High→High and Low→Low, red line) and different-coherence trials (High→Low and Low→High, purple line). The profiles were computed by averaging channel responses in the .5–1 s interval during which there was robust tuning to the average motion direction. To match the number of trials in the same-coherence and different-coherence trials and to maximise the signal-to-noise ratio in the different-coherence conditions, the channels were flipped for the High→Low trials so that the direction of a potential order bias would be the same as for the Low→High trials. *Small inset panel*: the average shift of the tuning profiles computed as the difference between the mean response for channels tuned to -π –0 and 0 –π intervals. Whiskers denote ±1 within-participants SEM.

In a final analysis, we characterised the temporal dynamics of integrated decision-making by estimating neural tuning to the *average* motion direction (Fig. 5a). For the first epoch, there was no

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

significant motion tuning. This finding was expected, as neural representations of average motion direction can only be determined *after* presentation of the second motion target within the trial. By contrast, there was robust and sustained tuning to the average motion direction in the second epoch, starting from the offset of coherent motion. Note that during this period only random motion was actually presented on the screen, so tuning to the average motion direction could not have been stimulus-driven.

To investigate how the brain integrates two discrete decisions, we next quantified the tuning profiles for the average motion direction (Fig. 5b). We focused on quantifying potential shifts in the profile that might reflect the first and the second target bias. To do so, the motion channels on each trial were re-coded so that negative channels were closer to the first presented target, and positive channels were closer to the second. Thus, a leftward shift would indicate a first-target bias, and a rightward shift would indicate a second-target bias. The channels' responses were then averaged across trials separately for different combinations of stimulus reliability across epochs. Motivated by the first-target bias we observed in behavioural decision weights (Fig. 2b), we expected to see a leftward shift for trials with the same coherence across epochs (Fig. 5b, upper panel). To increase signal-to-noise ratio, we averaged the tuning profiles for High→High and Low→Low coherence trials. The observed tuning profiles (Fig. 5b, lower panel) confirmed our expectation, as we observed a statistically significant leftward shift consistent with a first-target bias (M/SEM = -.16 /.06, t_{21} = 2.32, $p_{FDR\text{-}corrected}$ = .046, Fig. 5b, inset panel). This result mimics the first-target bias in behavioural decision weights, and it suggests that the brain integrates two target signals in a biased way, with the first signal contributing more strongly than the second.

In contrast to the same-coherence trials (High→High, Low→Low), for the different-coherence trials we expected to observe effects of signal reliability on the tuning profile shifts. For the High→Low sequence, in which the order bias and the reliability bias both favoured the first target, we expected to see an even stronger leftward shift relative to the same-coherence trials (Fig. 5b, upper panel). For the Low→High sequence, on the other hand, in which the order bias and the reliability bias favoured different targets, we expected to see weaker shifts in the tuning profile relative to the same-coherence trials (Fig. 5b, upper panel). To match the numbers of trials in the same- and different-coherence conditions and to increase signal-to-noise ratio in the different-coherence condition, we flipped the tunning profiles for High→Low trials so that the expected shift direction was the same for High→Low and Low→High trials. Confirming our predictions, the first-target bias for different-coherence trials was not significantly different from zero (.05/.06, t_{21} = .66, $p_{FDR\text{-}corrected}$ = .519, Fig. 5b, inset panel), and it was significantly smaller than the bias for same-

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

coherence trials (t_{21} = 2.35, $p_{FDR\text{-}corrected}$ = .046). Taken together, analyses of shifts in tuning profiles revealed that the brain relies more on signals with high reliability than low when integrating the two.

Across two experiments in which we combined behavioural testing and whole-brain recording, we have shown that behavioural decision weights and population-tuning profiles to the average motion direction exhibit qualitatively similar biases in favour of stimuli of higher reliability. These findings demonstrate that temporal integration of discrete perceptual decisions is biased even when the bias is suboptimal and unbiased, optimal integration is possible. Our findings suggest that the brain encodes the reliability of sensory inputs and that the encoded reliability is used automatically to weight respective inputs during integrative decision-making.

The relatively narrow distributions of error magnitudes and high response precision suggest that participants were able to successfully select target signals and ignore concurrently presented distractors. This finding was further corroborated by very low decision weights for distractors, which were five to nine times lower than the respective target weights. Perhaps most interestingly, and contrary to expectations, the population-tuning modelling of distractor motion signals revealed no distractor-specific neural activity. As participants were given some preparation time (500–750 ms) before the motion onset, it is likely that this time was sufficient for attentional resources to be engaged exclusively on the target-motion stimulus. By contrast, the motion decoding for target signals was robust and sustained well after signal offset, suggesting that the population-tuning modelling primarily captured decision-making processes as opposed to purely sensory-evoked activity patterns. With this in mind, it is likely that decoding of the average motion direction also reflected the dynamics of integrated decision making. One might ask whether it is possible that tuning to the average motion direction simply reflects tuning to the second target, given that the two directions were not entirely uncorrelated. This seems unlikely, however, as comparable tuning to the average motion direction should also have been observed in the first epoch (for which target signals were likewise not uncorrelated with the average), but this was clearly not the case. Moreover, if tuning to the average motion direction was driven by the second target then the time course of average motion tuning should have been similar to that of the second target, which, again, clearly was not the case. We therefore conclude that the robust tuning to the average target motion observed in Experiment 2 reflects the temporal dynamics of integrative decision making.

An unexpected finding in Experiment 2 was a reliable order bias, with higher decision weights for the first epoch than the second. Additionally, the colour-locked CPP had a steeper slope and a higher peak in the first epoch than the second. By contrast, the motion-locked CPPs did not differ much between epochs, at least prior to the peak of the CPP response. These findings suggest that the order bias originates from processes related to selecting the task-relevant dot patch, rather

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

than from processing of the motion signals themselves. Perhaps most compellingly, the order bias was also evident in tuning to the average motion direction, with robust shifts in favour of the first target-motion direction. This order effect was independent of reliability-weighted source integration, as we observed the shift in population tuning for epochs that were matched in coherence. Additionally, as the stimuli across the two epochs were matched for low-level properties, the order bias cannot be stimulus-specific. Finally, the first-target bias speaks against a simple memory decay explanation, which instead would predict a recency effect (29,43). Taken together, the most parsimonious explanation for the primacy bias is that the effectiveness of attentional selection decreases from the first epoch to the second. Characterising attentional selection dynamics in relation to integrated decision-making was not in the focus of the present study, and follow-up studies will be needed to address this issue in more detail. For the interested reader, we have recently conducted a study (44) that focused on the relationship between selective attention and decision-making using a similar experimental paradigm but a different analytical approach.

The reproduction task we employed enabled us to probe the nature of the representations underlying integrated decision-making. In typical decision-making paradigms (4), the response is a categorical decision, for example, whether motion direction is to the left or to the right. While forced-choice paradigms lend themselves to speeded responding, and permit use of computational modelling to characterise different aspects of decision making, they do not capture the precision of the sensory and memory representations that underlie evidence accumulation processes. A well-known property of the brain's responses to sensory input (18,45–47) is that they are graded, forming a probabilistic stimulus representation in feature space. In the case of motion signals, the large-scale neural representation of a given motion direction should resemble a bell-shaped curve with a peak over the actual direction which gradually decreases for motion directions further away from the peak. Classical decision-making paradigms would be sensitive to the location of the peak, but would have difficulty characterising the variability of the probabilistic representation – and that variability seems to play a critical role in integrated decision-making. Using mixture distribution modelling for behavioural measures, and population tuning modelling for neural measures, we were able to characterise both the peak and the variance of the underlying probabilistic representations.

A key question regarding reliability-weighted integrated decision-making concerns how two discrete representations get combined in support of an integrated decision. Using the redundant signals paradigm, previous research on signal integration at both sensory (23,24) and decision-making (17,22,48) stages has suggested that a simple multiplication of two probabilistic representations could drive reliability-weighted integration. The multiplication, however, predicts *lower* variability (18,23) of the integrated representation relative to the variability of individual

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

sources. By contrast, the variability of integrated decisions in our study was *higher* than that of single decisions, following the summation rule for two random variables. Therefore, it appears that for tasks such as the one employed here the two probabilistic representations are summed rather than multiplied.

As both the behavioural and neural results of Experiment 2 depended more strongly on high-reliability events than on low, the summation process appears to be biased, with higher weights for high-reliability representations. It is unlikely that these weights reflect learning processes, as in our paradigm the coherence of individual signals was unpredictable. Rather, the summation weights were likely encoded in parallel with encoding of the motion direction signals. One possibility is that the weights directly reflect the *variability* of the probabilistic signal representations: a simple combination of two probabilistic representations of differing variances would be shifted in favour of the representation of higher reliability. Another possibility is that the weights reflect a *belief* about the accuracy of the respective representations. In this scenario, even though the two reliability levels afforded comparable accuracy (as confirmed in Experiment 1), the strong perceptual differences between low- and high-reliability signals would have resulted in different beliefs about signals of different coherence. While at present we cannot adjudicate between the two potential correlates of the summation weights, the absence of strong coherence effects in our study suggests that the representations of the high and low reliability signals were comparably accurate, speaking in favour of the latter, *beliefs-as-weights* alternative. Further studies, most likely in combination with hierarchical computational modelling (49), would be necessary to address this issue conclusively. At present, computational models of decision-making in reproduction tasks are just beginning to appear (50,51), and more research will be needed before applying these models to integrative decision-making tasks.

In summary, here we have shown that combining two discrete, temporally separated signals in support of a single, integrated decision is biased in favour of higher reliability signals. Unlike previous studies in which reliability-weighted integration was statistically optimal, in the present study biased integration was suboptimal. These findings suggest that reliability-weighted integrated decision-making is automatic, taking place even when it is detrimental for performance.

Methods

Participants. 25 neurotypical adults (mean age 22 years, 14 females) participated in Experiment 1. All had normal or corrected-to-normal visual acuity and normal colour vision confirmed by Ishihara colour plates. The sample size was selected with the aim to achieve high power (β = .9 at α = .05) to detect a medium to large effect size (Cohen's d_z = .65) for a one-tailed, one-sample t-test between response error magnitude for low- and high-motion coherence. Due to

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

corrupted data acquisition, one participant was immediately removed from further analyses. Another group of 25 neurotypical adult humans (mean age 21, 15 females) took part in Experiment 2. Based on behavioural performance (see below for details), one and two participants were identified as outliers in Experiments 1 and 2, respectively, and removed from further analyses. Based on the EEG signals, two participants were identified as outliers in Experiment 2, and removed from further analyses. The final sample comprised 23 and 22 participants in Experiments 1 and 2, respectively. The study was approved by the Human Research Ethics Committee of The University of Queensland (approval nr 2016001247), and was conducted in accordance with the Human Subjects Guidelines of the Declaration of Helsinki. All participants provided written informed consent prior to experimental testing.

Stimuli, task, and procedure. In both experiments, every trial started with a coloured cue, indicating the target colour (Fig. 1a). Two out of three easily discernible colours (pink, HSL values of 0, 75, 50; yellow, 90, 75, 50; and cyan, 270, 75, 50) served as target and distractor colours. The target-distractor colour pairs (e.g., pink target and blue distractor) were fixed per participant and counterbalanced between participants. After the cue, a circular patch (15.6 dva diameter) of 160 grey randomly-moving dots (.6 dva diameter, speed 2.5 dva/s, infinite dot life) appeared. To prevent a stimulus onset-evoked response from influencing electrophysiological measures of decision-making, the grey patch remained on screen for 1 s. Thereafter, the dot saturation increased gradually (over .25 s), revealing two overlapping patches of coloured dots (80 dots per patch), in the target and distractor colours. The colour saturation remained at maximum for 1 s and then gradually returned to grey. During maximum saturation, coherent motion signals were presented briefly (.5 s) in both patches. The onset of coherent motion was jittered (.25–.5 s) relative to the maximum colour saturation. The motion coherence was pseudo-randomly selected for every epoch of coloured dots, with low (40%) and high (80%) coherences presented equally often. Participants had to monitor for target motion signals and ignore distractors. A feedback stimulus was presented after every trial indicating response accuracy in that trial. Response accuracy rather than speed was emphasised, and participants were given ample time to respond (max. 6 s).

The aim of Experiment 1 was to investigate the ability to discern target motion signals at different (40% and 80%) coherence levels. Only one epoch of coloured dots was presented per trial and participants had to reproduce the target motion direction by adjusting the orientation of a response dial. Per participant, two pairs of motion directions (i.e., four in total) were selected as target motions from a range of directions (0–360 degrees in 15-degree steps). Within a pair, the directions were 30 degrees apart from each other and the two pairs were 180 degrees apart (e.g., 15 and 45, 195 and 225 degrees). Different combinations of directions were counterbalanced across

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

participants. The distractor motion was within a ±30–150 degree range relative to the presented target motion. Only behavioural data were recorded.

The aim of Experiment 2 was to investigate integration of target signals across two epochs. Experiment 2 was identical to Experiment 1, with two exceptions. First, two epochs of coloured dots, rather than one, were presented in every trial and participants had to reproduce the average motion direction of the two target signals while ignoring distractor motion events. Second, target motion in the first epoch was selected randomly from 0–360 degrees range in 1 degree steps. The target motion in the second epoch was selected from a ±30–150 degree range relative to the first target. The dot coherence across the two epochs was selected pseudo-randomly so that all four combinations (low/high in the first epoch × low/high in the second) were presented equally often. Both behavioural and electroencephalography data were recorded.

Apparatus. The experiments were conducted in a dark, acoustically and electromagnetically shielded room. The stimuli were presented on a 24" monitor with 1920×1080 resolution and a refresh rate of 144 Hz. The experimental software was custom-coded in Python using the PsychoPy toolbox (52,53). EEG signals were recorded using 64 Ag-AgCl electrodes (BioSemi ActiveTwo) arranged in the 10-20 layout, and sampled at 1,024 Hz.

Behavioural analyses. To identify outlier participants, the distributions of error magnitudes (i.e., the angular difference between the response and the correct answer) were compared to a uniform distribution (i.e., pure guessing) using the Kolmogorov-Smirnov test. Participants for whom the probability of the null hypothesis (i.e., a uniform distribution of error magnitudes) exceeded .001 were removed from further analyses. The remaining distributions per experimental condition and per participant were fitted to a theoretical model (54), and responses were separated into noisy target responses and random guesses. To quantify decision weights, a multiple-regression (OLS) model with a term for each of the presented motion directions, expressed as complex numbers, was fitted to the responses, separately per participant and experimental condition. The absolute value of the resulting regression coefficients reflects the influence of each of the presented coherent motion signals on the response, i.e., its decision weight.

EEG analyses. EEG signals were analysed using the MNE-Python toolbox (55). The data were offline re-referenced to the average electrode, low-pass filtered at 99 Hz and notch-filtered at 50 Hz to eliminate line noise. The recorded signal was pre-processed using the FASTER algorithm for automated artefact rejection (56). The pre-processed signal was down-sampled to 256 Hz, segmented into 4 s periods between the onset of the first epoch and the response-display onset, baseline-corrected relative to -.1–0 s pre-trial and linearly de-trended. Outlier trials and participants were identified using the FASTER algorithm and removed from further analyses.

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Next, whole-trial time-traces were further segmented into shorter (.5 s) periods time-locked to the onset of colour-saturation increase and the onset of coherent motion in the first and second epoch, and baseline-corrected relative to .1–0 s pre-onset interval. To characterise the temporal dynamics of evidence accumulation, we quantified an ERP known as the central-parietal positivity (CPP). Previous research has shown that the time-course of the CPP closely resembles the time-course of evidence accumulation: specifically, its amplitude builds up gradually, the build-up slope is proportional to stimulus quality, and it is observed even in the absence of overt responses (11,12,41). Visual inspection of the ERP topographies revealed a positive deflection in a cluster of central-medial electrodes (FCz, C1, Cz, C2, and CPz) consistent with the CPP ERP – to improve signal-to-noise-ratio, the average of these electrodes was used in further analyses. Next, the CPP voltage per time-sample, trial and participant was submitted to a stepwise linear mixed-effects model with epoch (first/second) and motion coherence (low/high) per epoch as fixed effects, and participant as a random effect. A likelihood ratio test between models of higher and lower complexity was used to assess the significance of each main effect and interaction terms. To control for multiple comparisons, the p-values for all time-samples and all model terms were jointly corrected using the false discovery rate algorithm (57).

To recover feature-specific information about motion signals from the EEG signals (Fig. 6), we used a population tuning curve model (31–33). To that end, the first and the second epochs (1 s segments time-locked to the onset of coherent motion) from all trials were concatenated (ca 700 segments per participant), temporally smoothed by convolving the time series with a gaussian window (SD = 16 ms), shuffled and split into 10 testing sets ($B'_{64\ electrodes \times n\ epochs}$). For each testing set, the remaining epochs were used as the training set (B). We modelled motion tuning as a set of 16 half-wave-rectified sinusoid channels raised to the 15th power and centred on equidistant motion directions spanning 0–360°. The motion directions presented in each epoch were then convolved with the motion channels to yield a matrix of channel responses to motion signals ($C_{16\ channels \times n\ epochs}$) in the training set. To model the mean EEG amplitude across training epochs, a fixed intercept term was added to the matrix of channel responses. The training EEG data were then modelled using the following linear model: $B = WC$, where W represents the weight matrix ($W_{64\ electrodes \times 16\ channels}$) relating the EEG data and the population tuning model. The W matrix was estimated using ordinary least squares: $W = BC^T(CC^T)^{-1}$ using Moore-Penrose inverse. Finally, the responses of the population tuning model in the testing set were estimated using the following equation: $C'_{16\ channels \times n\ epochs} = (W^TW)^{-1}W^TB'$. The profile of channel responses would reflect motion tuning: a uniform (i.e., flat) response profile would correspond to no tuning, whereas a prominent peak at channels close to the presented motion would reflect strong tuning.

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

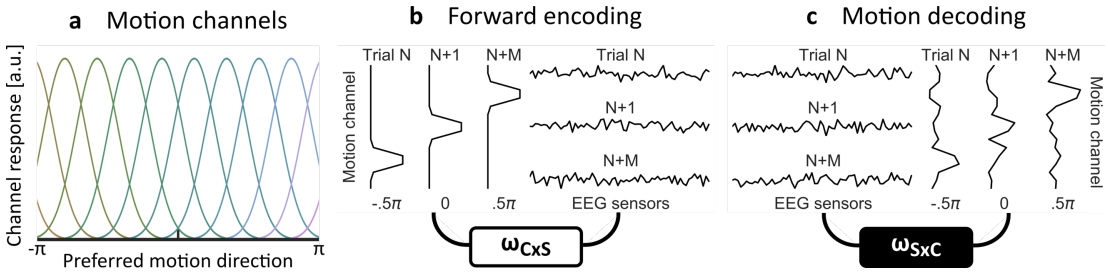

Figure 6. Illustration of the population-tuning analysis pipeline. (a) A set of hypothetical motion-selective channels is created, each of which is preferentially tuned to one of 16 equidistantly spaced motion directions spanning the full circle (-π to π). (b) The array of presented motion directions per trial in the training set of trials is convolved with the response profiles of the motion channels to yield trial-specific responses (left panel). The trial-specific responses are used as predictors to model the EEG time-traces from the respective trials (right). The forward encoding yields a weight matrix (ω_{CxS}) relating channel responses to the EEG. (c) The inverted weight matrix (ω_{SxC}) is applied to the EEG time-traces from the testing set of trials to retrieve the profiles of the motion channel responses.

The motion tuning analyses were conducted per time sample per participant. The tuning to target and distractor motion signals was analysed separately. We also estimated tuning to the average motion direction (i.e., the expected, average response) using the first and the second epoch data in two separate analyses. To characterise the overall tuning strength across different conditions, the vectors of channel responses were centred on the actual presented motion direction, temporally smoothed using a gaussian window (SD = 16 ms) and averaged across trials. As an aggregate index of the response profile, we computed the dot product (θ) between the channel responses and the complex-valued preferred motion directions for respective channels. To quantify the tuning strength, we used the following equation: $f_\theta = abs(\theta)cos(\theta)$. This equation yields a good, non-parametric descriptor of the overall tuning shape: (i) the $abs(\theta)$ yields 0 when there is no tuning (i.e., the distribution of channel responses is flat) and (ii) the $cos(\theta)$ reflects a mismatch between the expected preferred motion direction (which is normalized to 0) and the empirical peaks in the tuning profile. Similar to the ERP analyses, the tuning strength time series were analysed using stepwise mixed-effects general linear models, separately per time sample.

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

References

1. Ratcliff R, Smith PL, Brown SD, McKoon G. Diffusion Decision Model: Current Issues and History. Trends Cogn Sci. 2016 Apr 1;20(4):260–81.
2. Forstmann BU, Ratcliff R, Wagenmakers E-J. Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions. Annu Rev Psychol. 2016;67(1):641–66.
3. Summerfield C, Blangero A. Chapter 12 - Perceptual Decision-Making: What Do We Know, and What Do We Not Know? In: Dreher J-C, Tremblay L, editors. Decision Neuroscience [Internet]. San Diego: Academic Press; 2017 [cited 2018 Aug 30]. p. 149–62. Available from: http://www.sciencedirect.com/science/article/pii/B9780128053089000129
4. O'Connell RG, Shadlen MN, Wong-Lin K, Kelly SP. Bridging Neural and Computational Viewpoints on Perceptual Decision-Making. Trends Neurosci. 2018 Nov 1;41(11):838–52.
5. Philiastides MG, Auksztulewicz R, Heekeren HR, Blankenburg F. Causal Role of Dorsolateral Prefrontal Cortex in Human Perceptual Decision Making. Curr Biol. 2011 Jun 7;21(11):980–3.
6. Liu T, Pleskac TJ. Neural correlates of evidence accumulation in a perceptual decision task. J Neurophysiol. 2011 Aug 17;106(5):2383–98.
7. Ratcliff R. A theory of memory retrieval. Psychol Rev. 1978;85(2):59–108.
8. Smith PL, Ratcliff R. An integrated theory of attention and decision making in visual signal detection. Psychol Rev. 2009 Apr;116(2):283–317.
9. White CN, Ratcliff R, Starns JJ. Diffusion models of the flanker task: Discrete versus gradual attentional selection. Cognit Psychol. 2011 Dec 1;63(4):210–38.
10. Philiastides MG, Sajda P. Temporal Characterization of the Neural Correlates of Perceptual Decision Making in the Human Brain. Cereb Cortex. 2006 Apr 1;16(4):509–18.
11. Kelly SP, O'Connell RG. Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain. J Neurosci. 2013 Dec 11;33(50):19434–41.
12. O'Connell RG, Dockree PM, Kelly SP. A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. Nat Neurosci. 2012 Dec;15(12):1729–35.
13. Shadlen MN, Newsome WT. Motion perception: seeing and deciding. Proc Natl Acad Sci. 1996 Jan 23;93(2):628–33.
14. Shadlen MN, Newsome WT. Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. J Neurophysiol. 2001 Oct 1;86(4):1916–36.
15. Churchland AK, Kiani R, Shadlen MN. Decision-making with multiple alternatives. Nat Neurosci. 2008 Jun;11(6):693–702.
16. Spitzer B, Waschke L, Summerfield C. Selective overweighting of larger magnitudes during noisy numerical comparison. Nat Hum Behav. 2017 Aug;1(8):0145.
17. Raposo D, Sheppard JP, Schrater PR, Churchland AK. Multisensory Decision-Making in Rats and Humans. J Neurosci. 2012 Mar 14;32(11):3726–35.
18. Kanitscheider I, Brown A, Pouget A, Churchland AK. Multisensory decisions provide support for probabilistic number representations. J Neurophysiol. 2015 Mar 5;113(10):3490–8.
19. Leite FP, Ratcliff R. Modeling reaction time and accuracy of multiple-alternative decisions. Atten Percept Psychophys. 2010 Jan 1;72(1):246–73.
20. Raposo D, Kaufman MT, Churchland AK. A category-free neural population supports evolving demands during decision-making. Nat Neurosci. 2014 Dec;17(12):1784–92.
21. Brunton BW, Botvinick MM, Brody CD. Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. Science. 2013 Apr 5;340(6128):95–8.
22. Boyle SC, Kayser SJ, Kayser C. Neural correlates of multisensory reliability and perceptual weights emerge at early latencies during audio-visual integration. Eur J Neurosci. 2017 Nov 1;46(10):2565–77.

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

23. Ernst MO, Banks MS. Humans integrate visual and haptic information in a statistically optimal fashion. Nature. 2002 Jan;415(6870):429–33.

24. Seilheimer RL, Rosenberg A, Angelaki DE. Models and processes of multisensory cue combination. Curr Opin Neurobiol. 2014 Apr 1;25:38–46.

25. Churchland AK, Kiani R. Three challenges for connecting model to mechanism in decision-making. Curr Opin Behav Sci. 2016 Oct 1;11:74–80.

26. Scase MO, Braddick OJ, Raymond JE. What is Noise for the Motion System? Vision Res. 1996 Aug;36(16):2579–86.

27. Zhang W, Luck SJ. Discrete fixed-resolution representations in visual working memory. Nature. 2008 May 8;453(7192):233–5.

28. Bays PM, Husain M. Dynamic Shifts of Limited Working Memory Resources in Human Vision. Science. 2008 Aug 8;321(5890):851–4.

29. Wyart V, de Gardelle V, Scholl J, Summerfield C. Rhythmic Fluctuations in Evidence Accumulation during Decision Making in the Human Brain. Neuron. 2012 Nov 21;76(4):847–58.

30. Cheadle S, Wyart V, Tsetsos K, Myers N, de Gardelle V, Herce Castañón S, et al. Adaptive Gain Control during Human Perceptual Choice. Neuron. 2014 Mar 19;81(6):1429–41.

31. Brouwer GJ, Heeger DJ. Decoding and Reconstructing Color from Responses in Human Visual Cortex. J Neurosci. 2009 Nov 4;29(44):13992–4003.

32. Myers NE, Rohenkohl G, Wyart V, Woolrich MW, Nobre AC, Stokes MG. Testing sensory evidence against mnemonic templates. eLife. 2015;4:e09000–e09000.

33. Kok P, Mostert P, Lange FP de. Prior expectations induce prestimulus sensory templates. Proc Natl Acad Sci. 2017 Sep 12;201705652.

34. Mostert P, Kok P, de Lange FP. Dissociating sensory from decision processes in human perceptual decision making. Sci Rep [Internet]. 2015 Dec 15 [cited 2017 Oct 27];5. Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4678878/

35. Smout CA, Tang MF, Garrido MI, Mattingley JB. Attention promotes the neural encoding of prediction errors. PLOS Biol. 2019 Feb 27;17(2):e2006812.

36. Tang MF, Smout CA, Arabzadeh E, Mattingley JB. Prediction error and repetition suppression have distinct effects on neural representations of visual information. eLife [Internet]. 2018 [cited 2019 May 6];7. Available from: https://europepmc.org/articles/PMC6312401/

37. Garcia JO, Srinivasan R, Serences JT. Near-Real-Time Feature-Selective Modulations in Human Cortex. Curr Biol. 2013 Mar 18;23(6):515–22.

38. Wolfe JM, Horowitz TS. Five factors that guide attention in visual search. Nat Hum Behav. 2017 Mar 8;1(3):1–8.

39. Reynolds JH, Desimone R. Interacting Roles of Attention and Visual Salience in V4. Neuron. 2003 Mar 6;37(5):853–63.

40. Desimone R, Duncan J. Neural mechanisms of selective visual attention. Annu Rev Neurosci. 1995;18:193–222.

41. Twomey DM, Kelly SP, O'Connell RG. Abstract and Effector-Selective Decision Signals Exhibit Qualitatively Distinct Dynamics before Delayed Perceptual Reports. J Neurosci. 2016 Jul 13;36(28):7346–52.

42. Morey RD. Confidence intervals from normalized data: A correction to Cousineau. In: Tutorials in Quantitative Methods for Psychology , 4. 2008. p. 61–64.

43. Wyart V, Myers NE, Summerfield C. Neural Mechanisms of Human Perceptual Choice Under Focused and Divided Attention. J Neurosci. 2015 Feb 25;35(8):3485–98.

44. Rangelov D, Rangelov D, Mattingley JB. Selective attention modifies weights of distinct visual signals in human decision making. 2019; Available from: http://doi.org/10.1101/537910

bioRxiv preprint doi: https://doi.org/10.1101/841353; this version posted November 13, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

45. Hubel DH, Livingstone MS. Complex-unoriented cells in a subregion of primate area 18. Nature. 1985 May 23;315(6017):325–7.

46. Livingstone MS, Hubel DH. Segregation of form, color, movement, and depth: anatomy, physiology, and perception. Science. 1988 May 6;240(4853):740–9.

47. Livingstone MS, Hubel DH. Anatomy and physiology of a color system in the primate visual cortex. J Neurosci. 1984 Jan 1;4(1):309–56.

48. Sheppard JP, Raposo D, Churchland AK. Dynamic weighting of multisensory stimuli shapes decision-making in rats and humans. J Vis. 2013 May 1;13(6):4–4.

49. Mathys CD, Lomakina EI, Daunizeau J, Iglesias S, Brodersen KH, Friston KJ, et al. Uncertainty in perception and the Hierarchical Gaussian Filter. Front Hum Neurosci [Internet]. 2014 [cited 2019 Oct 25];8. Available from: https://www.frontiersin.org/articles/10.3389/fnhum.2014.00825/full

50. Smith PL. Diffusion theory of decision making in continuous report. Psychol Rev. 2016;123(4):425–51.

51. Ratcliff R. Decision making on spatially continuous scales. Psychol Rev. 2018;125(6):888–935.

52. Peirce JW. PsychoPy—Psychophysics software in Python. J Neurosci Methods. 2007 May 15;162(1–2):8–13.

53. Peirce JW. Generating Stimuli for Neuroscience Using PsychoPy. Front Neuroinformatics [Internet]. 2009 Jan 15 [cited 2012 Dec 5];2. Available from: http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2636899/

54. Schneegans S, Bays PM. No fixed item limit in visuospatial working memory. Cortex. 2016 Oct 1;83(Supplement C):181–93.

55. Gramfort A, Luessi M, Larson E, Engemann DA, Strohmeier D, Brodbeck C, et al. MEG and EEG data analysis with MNE-Python. Front Neurosci [Internet]. 2013 Dec 26 [cited 2017 Dec 20];7. Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3872725/

56. Nolan H, Whelan R, Reilly RB. FASTER: Fully Automated Statistical Thresholding for EEG artifact Rejection. J Neurosci Methods. 2010 Sep 30;192(1):152–62.

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