



Original Articles

On-line confidence monitoring during decision making

Dror Dotan^{a,b,*}, Florent Meyniel^a, Stanislas Dehaene^{a,c}^a Cognitive Neuroimaging Unit, CEA DRF/I2BM, INSERM, Université Paris-Sud, Université Paris-Saclay, F-91191 Gif/Yvette, France^b Language and Brain Lab, School of Education and the Sagol School of Neuroscience, Tel Aviv University, Tel Aviv 69978, Israel^c Collège de France, 11 Place Marcelin Berthelot, 75005 Paris, France

ARTICLE INFO

Keywords:

Decision making
Confidence
Accumulation of evidence
Trajectory tracking
Speed-accuracy trade-off

ABSTRACT

Humans can readily assess their degree of confidence in their decisions. Two models of confidence computation have been proposed: post hoc computation using post-decision variables and heuristics, versus online computation using continuous assessment of evidence throughout the decision-making process. Here, we arbitrate between these theories by continuously monitoring finger movements during a manual sequential decision-making task. Analysis of finger kinematics indicated that subjects kept separate online records of evidence and confidence: finger deviation continuously reflected the ongoing accumulation of evidence, whereas finger speed continuously reflected the momentary degree of confidence. Furthermore, end-of-trial finger speed predicted the post-decisional subjective confidence rating. These data indicate that confidence is computed on-line, throughout the decision process. Speed-confidence correlations were previously interpreted as a post-decision heuristic, whereby slow decisions decrease subjective confidence, but our results suggest an adaptive mechanism that involves the opposite causality: by slowing down when unconfident, participants gain time to improve their decisions.

1. Introduction

Confidence is defined as our degree of belief that a certain thought or action is correct (Grimaldi, Lau, & Basso, 2015; Meyniel, Sigman, & Mainen, 2015). There is growing evidence that humans and other animals possess a sense of confidence in their decisions (Baranski & Petrusic, 1994; Grimaldi et al., 2015; Kepecs & Mainen, 2012; Kiani & Shadlen, 2009; Meyniel, Schlunegger, & Dehaene, 2015). Although confidence can be subject to various biases, the very fact that animals and humans are able to approximate the likelihood of a decision being correct is an impressive feat that fits with the increasingly influential view that the brain is able to compute with probabilities and their distributions (Beck et al., 2008; Kording & Wolpert, 2004; Pouget, Drugowitsch, & Kepecs, 2016). However, precise knowledge of how confidence is computed is still lacking. Two classes of models of confidence computation can be contrasted. One class emphasizes that confidence is computed in a post hoc manner, in order to retrospectively evaluate a recent decision (Balakrishnan & Ratcliff, 1996; Ferrell, 1995), using heuristics and post-decision variables (Kahneman & Tversky, 1982; Pleskac & Busemeyer, 2010; Resulaj, Kiani, Wolpert, & Shadlen, 2009). For instance, one model proposes that subjects use a summary of the decision process, namely, reaction time, as an index to confidence: trials that are responded fast are judged as more likely to be

correct, which is indeed a valid heuristic in many situations (Kiani, Corthell, & Shadlen, 2014). Another computational model proposes that confidence is based not only on the evidence accumulated to make the decision, but also on additional evidence accumulated after the decision (Pleskac & Busemeyer, 2010). In general, this approach tends to view confidence judgment as a slow and imperfect mechanism, that follows decision making and uses memory and heuristics to re-evaluate our decisions (Dunlosky & Metcalfe, 2008).

Another class of models, however, emphasizes that a sense of confidence can emerge from the decision-making process itself. According to this account, confidence is computed online throughout the decision-making process, in parallel to or even as part of the accumulation of evidence that supports the decision (Fetsch, Kiani, & Shadlen, 2014; Kiani & Shadlen, 2009; Meyniel et al., 2015). For example, one computational model proposes that the brain can process probability distributions and therefore, throughout the decision-making process, carries a full representation of the probability that a given inference is correct (Pouget et al., 2016). Such online monitoring of confidence could be helpful in regulating our decisions while they are being made, for instance in order to withhold decision and look for more information (Meyniel et al., 2015).

The online and post-decisional accounts of confidence are not mutually exclusive but complementary: even if confidence is computed

* Corresponding author at: School of Education, Tel Aviv University, Tel Aviv 69978, Israel.
E-mail address: dror.dotan@gmail.com (D. Dotan).

online, it can still be submitted to various post-decisional transformations and biases before one reaches a conscious, reportable level of subjective confidence in a decision. However, while evidence for post-decisional confidence processing is well-established (Pleskac & Busemeyer, 2010; Resulaj et al., 2009), the existence of online, pre-decisional confidence monitoring processes is still debated (Pouget et al., 2016).

Measuring pre-decision confidence poses methodological challenges. Most metacognitive paradigms are retrospective, asking participants to rate their subjective confidence in a past decision (Dunlosky & Metcalfe, 2008). Other paradigms, allowing the participant to opt out of the decision (Fetsch, Kiani, Newsome, & Shadlen, 2014; Kiani & Shadlen, 2009), provide behavioral information about the decision (when the participants do not opt out) or about confidence (by comparing opt-out and no-opt-out trials), but not about both on a given trial. Implicit measures of confidence derived from neural recordings (Charles, King, & Dehaene, 2014; Fetsch et al., 2014; Kepecs, Uchida, Zariwala, & Mainen, 2008; Kiani & Shadlen, 2009; Kiani et al., 2014) avoid these problems, but they rely on invasive electrophysiological or costly brain-imaging measures from which it remains difficult to disentangle decision and confidence signals. Here, we show how an elementary behavioral measurement – tracking the participants' finger movement during decision making – can be used to analyze the decision-making process and obtain separate implicit measures of a prospective decision and the associated confidence.

30 human adults performed a simple two-alternative forced-choice task on a touchscreen. On each trial, 1, 3, or 5 arrows, each pointing leftward or rightward, were presented sequentially, and participants were asked to decide whether most arrows pointed to the left or to the right. This paradigm is inspired by the classical Shadlen-Newsome motion direction detection task in which sensory evidence must be accumulated across time (Shadlen & Newsome, 2001). However, our stimuli were not continuous but employed few discrete bouts of evidence, thereby allowing for a precise analysis of changes in decision making (de Lange, Jensen, & Dehaene, 2010; de Lange, van Gaal, Lamme, & Dehaene, 2011; Yang & Shadlen, 2007). In the Discussion, we elaborate further on the similarities and differences between our paradigm and classical paradigms of perceptual decision making. Crucially, our participants responded by continuously moving their finger on the touchscreen from a fixed starting point to one of two response buttons, without ever stopping (Fig. 1). Previous studies showed that changes in finger direction reflect intermediate stages of decision making (Berthier, 1996; Erb, Moher, Sobel, & Song, 2016; Friedman, Brown, & Finkbeiner, 2013; Pinheiro-Chagas, Dotan, Piazza, & Dehaene, 2017; Resulaj et al., 2009). Here, given previous results on confidence and decision times (Baranski & Petrusic, 1994; Kiani et al., 2014), we propose that, additionally, the instantaneous finger speed reflects online fluctuations in the participant's prospective confidence that the final decision will be correct.

2. Method

2.1. Participants and task

The participants were 30 university students (mean age = 25;11, SD = 4;0) and gave informed consent prior to participating. One participant rated almost all trials (97%) as “100% confident” and was excluded. On each trial, participants saw on a tablet computer a sequence of arrows that included one arrow (2 possible sequences, each presented 64 times), 3 arrows ($2^3 = 8$ sequences, 16 times each) or 5 arrows (32 sequences, 12 times each). The numbers of arrows were not disclosed to the participants. They were instructed to indicate where the majority of arrows pointed to by dragging their finger from a starting point at the bottom of the screen to a response button on the top-right or top-left corner of the screen (Fig. 1a). Touching the starting point triggered a central fixation dot on the top of the screen, where arrows appear, and finger movement (crossing $y = 50$ pixels from the bottom of screen) triggered the arrow sequence. We used an Apple iPad air with 1024×768 resolution (5.2 px/mm), black background, white arrows (150 × 50 px), and grey response buttons (200 × 100 px) and starting point (60 × 40 px). Lifting the finger in mid-trial, moving the finger backwards, or starting a trial with sideways (rather than upward) movement, aborted the trial. Trials were also aborted when the finger movement was too slow (excluding a grace period of the trial's first 300 ms): less than 3 s per trial or less than 1.5 s to reach the first third of the screen. Aborted trials were excluded from analysis and presented again later in the experiment. Immediately after each trial, participants rated retrospectively their subjective confidence about their decision (i.e. the probability of the decision being correct) on a continuous vertical scale (top = “I'm sure”; middle = “I have no idea”; bottom = “I'm sure I was wrong”). The scale was presented in the middle of the screen, i.e., to rate their confidence, the participants first had to move their finger from the top of the screen, where it was at the end of the trial, back to the middle of the screen. Statistical analyses were done with Matlab and R (R Core Team, 2015). In <http://trajtracker.com>, we provide our trajectory-tracking analysis tools as well as a Python-based experimentation software equivalent to the one that we used here. The trajectory raw data is enclosed as Supplementary online material.

2.2. Data processing and terminology

Evidence is the sum of all stimulus arrows (\rightarrow is +1, \leftarrow is -1). $|Evidence|$ is its absolute value. *Accuracy* is the fraction of correct responses. *Confidence rating* refers to the participant's post-decision subjective rating (0–100 scale). *Movement time* is the time from the first arrow onset (which is immediately after the finger started moving) until the finger reached a response button, and *average speed* is the inverse of movement time. *Time point* refers to a particular time within a trial,

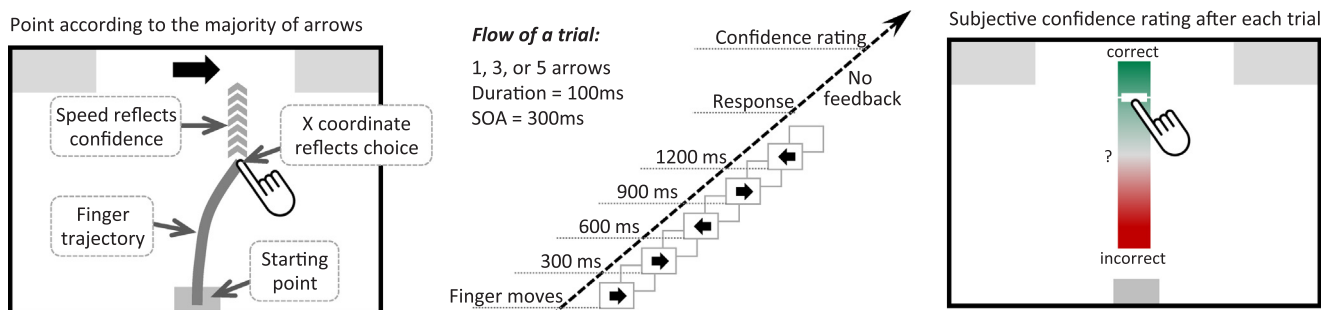


Fig. 1. Task and screen layout. On each trial, 1, 3, or 5 arrows, each pointing left or right, were presented sequentially. Participants dragged their finger on a touchscreen towards the response button corresponding to the majority of arrows. Their finger movement was continuously recorded. The onset of the first arrow was triggered by finger movement. After touching a response button, a slider appeared and participants rated their confidence about their decision, from “certainly correct” to “certainly incorrect”. A color version of this figure is available in the online version of the article.

specified relative to the onset of the first arrow.

Trajectory: The x , y coordinates were recorded at 16 Hz ($\sigma = 1$) using the screen resolution (1024×768 px), and transformed to 100 Hz using cubic spline interpolation. $x = 0$ is the middle of the screen. Momentary x , y speeds and accelerations per time point were obtained at the derivative of the x , y coordinates, after applying a Gaussian smoothing ($\sigma = 20$ ms) before derivation. *End speed* is the speed 100 ms before reaching the response button. We used a 100 ms gap to exclude the period in which the finger may sharply slow down as it approaches the response buttons; we note that the end-speed results reported below were replicated for a range of delays.

Speed, unless stated otherwise, refers to the finger's momentary speed along the y -axis, i.e. the speed with which it moves towards the top of the screen, where the response buttons are placed. Note that finger movement in the horizontal axis may be driven by the ongoing decision (aiming rightwards or leftwards), so that horizontal speed may reflect the buildup of the decision. We decided a priori to analyze speed only along the y -axis precisely to avoid this confound, and to show that speed specifically indexes confidence rather than the buildup of the decision. In the Results section, we bring several arguments to justify this decision from the data. Nevertheless, in the Supplementary online material (Fig. S5), we show that all critical results are replicated if one considers the xy -speed instead of speed along the y -axis only, and therefore that our conclusions are robust to this choice.

2.3. Regression analyses

Correlations among various parameters (confidence rating, accuracy, finger coordinates or speed, specific arrows, etc.) were examined using regression analyses. Specific regression models are detailed in the text below. In these regressions, predictors denoting the direction of a specific arrow were coded as $+1$ (right) and -1 (left). Accuracy was coded as 0, 1. In some regressions, the specific sequence of arrows was added as a covariate; in such cases, each sequence was coded as a binary predictor (i.e., for 5 arrows, which allow for 32 possible sequences, there were 32 such covariate predictors).

When the dependent variable and the predictors were trial-level measures, one regression was run per participant. When the dependent variable was a within-trajectory measures (e.g., x coordinate, y speed), one regression was run per participant and per time point, in 50 ms intervals. The significance of a specific predictor in a specific time point was assessed at the group level by comparing the participants' regression coefficients against 0 with one-tailed t -test.

2.4. The onset time of each arrow's effect

To identify when each arrow started affecting the finger movement, we examined how the finger x coordinates were affected by that arrow. On each time point, in 10 ms intervals, the finger x coordinates were submitted to a two-way repeated measures ANOVA with two within-participant factors: the arrow direction, and the specific sequence of previous arrow directions. The arrow's onset time was defined as the first time point in which the main effect of arrow direction was significant ($p < .05$) and remained so during at least 50 ms.

Using an identical method with y speed as the dependent variable, we identified when each arrow started affecting the finger speed.

2.5. Changes of mind

Another type of analysis aimed to examine whether certain situations promoted changes of mind (Resulaj et al., 2009). Such changes of mind may be revealed by situations where the finger changed its course, i.e., sudden increases/decreases in the horizontal speed. These "acceleration bursts" were defined as trajectory sections in which the unsigned acceleration exceeded 144 pixels/s^2 during at least 70 ms. Bursts that started before 100 ms were excluded, as they may result

from pre-stimulus default behavior (Dotan, 2017). Acceleration was the 2nd derivative of the x coordinates, after applying a Gaussian smoothing before each derivation ($\sigma = 100$ ms, but the effects described below were replicable with several different σ values). In the text below, we describe specific analyses that examined how the number of acceleration bursts per trial (denoted #Acc) was affected by various parameters. By applying the same analysis method to the y axis, we also examined the number of vertical speed fluctuations per trial.

A related measure is the number of times a trajectory bends to the right or to the left during a trial. This measure is useful to control for putative motor effects related with finger deviations. A *bend* in the trajectory was defined as a trajectory section of at least 100 ms during which the finger continuously changed its direction in either clockwise or counter-clockwise manner, excluding very small bends ($\Delta\theta < 5^\circ$ between start and end of a bend; but including these small bends yielded essentially the same results).

3. Results

3.1. End-of-trial measures

Analysis of end-of-trial measures – accuracy, subjective confidence rating, and average finger speed – indicated that they were all sensitive to trial difficulty. Trials with fewer arrows had overall higher accuracy, higher confidence ratings, and faster finger movements (Fig. 2a). This effect was confirmed by three separate repeated measures ANOVAs on the per-stimulus average accuracy, confidence ratings, or speed. In each of these three ANOVAs, the number of arrows (1, 3 or 5) was a within-subject factor and the participant was the random factor (all $F(2, 56) > 110$, $p < .001$).

Within the 5-arrow trials (on which we focus from now on, but see Fig. S7 for replication with 3-arrow trials), accuracy, confidence rating, and speed were higher for trials with more evidence (repeated measures ANOVAs on the per-stimulus average accuracy, confidence ratings, or speed, all with a numeric factor of $|Evidence|$ = the absolute difference between the numbers of left and right arrows: all $F(2, 56) > 68$, $p < .001$). More specifically, both accuracy and confidence rating (averaged per stimulus) varied monotonically with the ratio between left and right evidence (Fig. S1).

To examine the contribution of each of the five consecutive arrows to the choice of a response button, we used logistic regression: for each participant, the response was regressed against the 5 arrow directions. The per-participant regression β values were then compared to 0 with t -test (excluding one participant whose regression did not converge). Each of the 5 arrows had a significant positive effect (all $t(27) > 5.8$, one-tailed $p < .001$, FDR corrected). As previously reported (de Lange et al., 2010), arrows were considered even when they were redundant: late arrows affected the decision even in trials where a decision could be reached based on the first 3 or 4 arrows. This was demonstrated with logistic regressions on the response against arrow direction, now limited to subsets of trials where an early decision was possible. One such regression included only trials where the sufficient information was provided at the 3rd arrow (e.g., $\rightarrow\rightarrow\rightarrow\leftarrow\leftarrow$). The response was regressed on the directions of the 4th and 5th arrows. A second regression included trials where sufficient information was provided at the 4th arrow (e.g., $\rightarrow\rightarrow\leftarrow\rightarrow\leftarrow$). The response was regressed on the 5th arrow direction. In both analyses, the per-participant regression β values of each arrow were significantly higher than zero (all $t(28) \geq 2.8$, one-tailed $p < .01$, FDR corrected). Still, late arrows were underweighted relatively to early arrows: in a repeated measures ANOVA with the all-trials logistic regression β values as dependent variable and the arrow position in the sequence as a within-subject numeric factor, the arrow position had a significant effect ($F(1, 27) = 46$, $p < .001$).

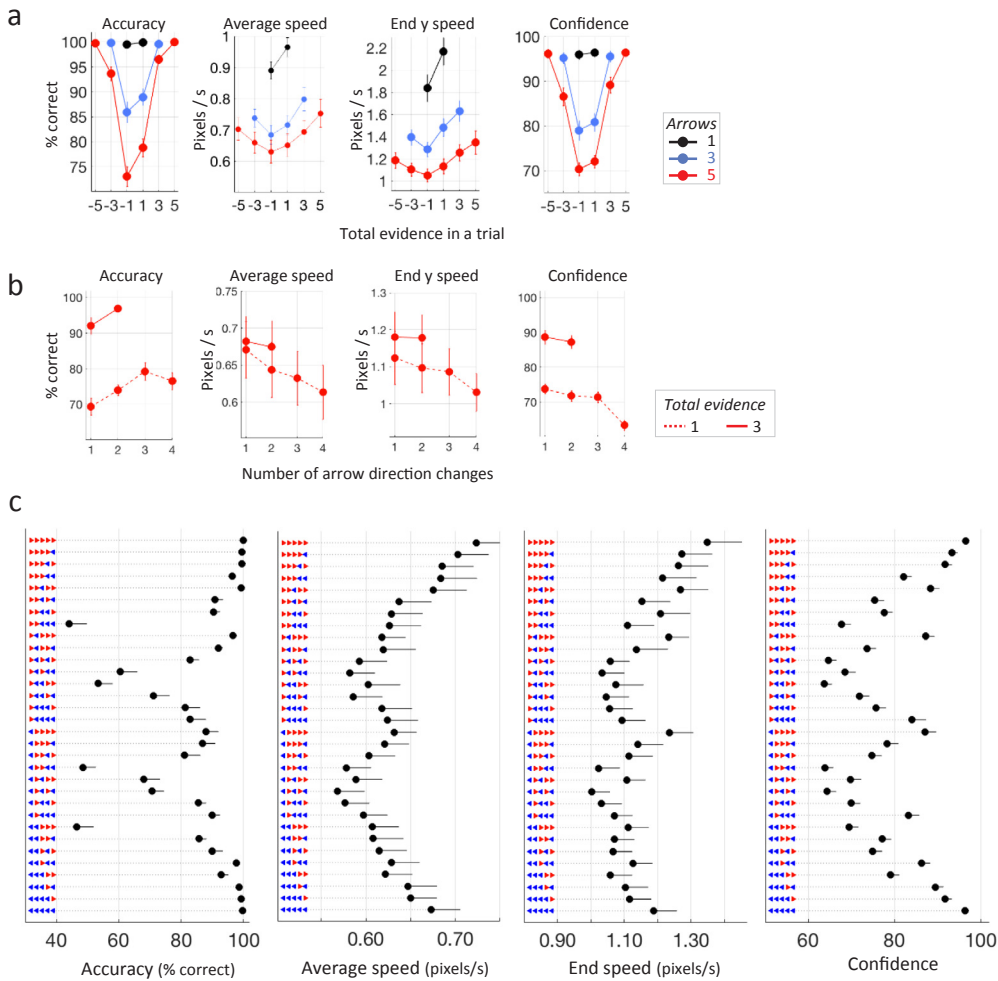


Fig. 2. Effect of the stimulus sequences on end-trial measures of accuracy, confidence, and speed. (a) All variables decreased for higher number of arrows and increased for more evidence. (b) A greater number of arrow-direction-changes reduced confidence and speed, but increased accuracy (only 5-arrow trials are included here). (c) The average accuracy, speed, and confidence for each stimulus. Error bars show one inter-subject standard error. A color version of this figure is available in the online version of the article.

3.2. Real-time accumulation of evidence

We next showed that finger movement tracked the pre-decision accumulation of evidence. Plotting the average finger position as a function of time yielded a clear tree structure that closely reflected the different sequences of arrows (Figs. 3a and S2a). This plot suggests that each arrow started affecting the finger x coordinate shortly after its presentation. The onset time of each arrow’s effect was calculated as described in Methods (Section 2.4; plotted in Fig. 3a as vertical lines). All arrows had a significant effect on the finger x coordinate, starting on average 406 ms after the arrow appeared on screen, and separated by 335 ms for consecutive arrows, in agreement with the actual stimulus-onset-asynchrony (SOA) of 300 ms.

We then examined the buildup of each arrow’s effect throughout the trial. At each time point, the finger horizontal position (x coordinate) was regressed per participant against the five arrows, coded as 5 predictors (Fig. 3b). All arrows had significant effects (per-participant β values significantly higher than 0 in a *t*-test), which started on average 417 ms after the arrow onset and remained significant in all time points until the end of the trial. These results were robust to the exclusion of error trials and the easiest trials (Fig. S3), as well as to the choice of different measures for finger movement (Fig. S4).

The separable effects of each arrow suggest that evidence is accumulated incrementally within a trial. Alternatively, it is possible that participants made only a single finger deviation during a given trial, and the effect of successive arrows would have resulted from averaging trials with different deviation latencies. However, our data ruled out this possibility, because several changes of mind were observed *within*

trials: we observed that trials with a higher number of arrow direction changes (#ADC) had a greater number of horizontal acceleration/deceleration bursts (#Acc) per trial (Figs. 3c and S6a, b; #Acc was calculated as defined in Section 2.5). This effect was statistically reliable: when #Acc was regressed for each participant against |Evidence| and #ADC, the per-participant β values of the #ADC predictor were significantly higher than zero ($t(28) = 4.49$, one-tailed $p < .001$). We note that the #Acc-by-#ADC slope was low (only ~ 0.15 horizontal acceleration/deceleration bursts per ADC). This could indicate either that participants did not change their mind on every single arrow direction change, or that our measure did not detect all changes of mind.

3.3. Confidence

3.3.1. End-of-trial measures

We now turn to the analysis of confidence, starting with end-of-trial measures. Post-decision confidence ratings correlated with the participants’ objective accuracy: regressing confidence against accuracy yielded a mean coefficient $\beta = 0.48$ (S.E.M. = 0.03), significantly higher than zero ($t(28) = 18.9$, $p < .001$). This accuracy-confidence correlation was partly, but not solely, driven by stimulus properties (see Fig. 2c): it survived the inclusion of sequence type as covariate ($\beta = 0.38$, $t(28) = 12.2$, $p < .001$), indicating that, even for a fixed stimulus, participants tracked their trial-by-trial fluctuations in the decision accuracy.

However, we also found specific cases in which confidence rating and accuracy dissociated: these variables were differentially affected by the number of arrow direction changes (#ADC). For a fixed amount of

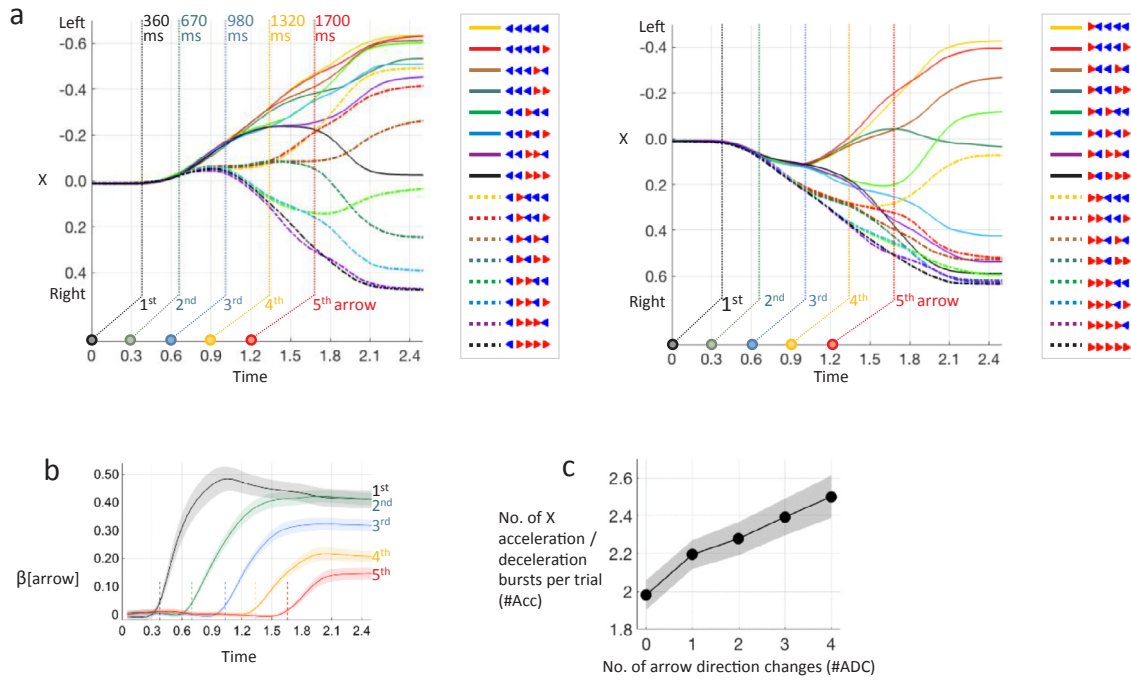


Fig. 3. The finger horizontal position indexes the online accumulation of evidence. (a) Average x coordinate as a function of time for each of the 32 arrow sequences. The vertical lines show the time when the effect of each arrow became significant. (b) Time course of the arrow effects. For each subject and each time point, across trials, the x coordinates were regressed against five predictors coding the direction of each arrow. The plot shows the regression weights (β) averaged across subjects, with their standard error. The dashed vertical lines indicate the onset of each arrow effect (when average β reached 10% of its asymptotic value). The sequential effect of arrows suggests a gradual accumulation of evidence. (c) The number of left-right deviations within a trial increased with the number of arrow direction changes, indicating within-trial changes of mind. A color version of this figure is available in the online version of the article.

evidence ($|Evidence| = 1$), a greater #ADC reduced confidence ratings, but increased accuracy (Fig. 2b). To confirm this accuracy-confidence dissociation, accuracy and confidence were averaged per sequence type and participant, Z-scored separately, and submitted together to repeated-measures ANOVA whose factors were the measure type (accuracy, confidence rating) and #ADC. A significant interaction confirmed the dissociation ($F(1, 28) = 72.9, p < .001$).

Larger #ADC also resulted in lower finger speed (Fig. 2b), suggesting that finger speed tracked subjective confidence even when confidence deviated from the objective accuracy. The decreasing effect of #ADC on confidence ratings and finger speed was observed even within objectively correct trials (comparing the per-participant confidence-by-#ADC or speed-by-#ADC regression slopes against 0 with t -test, $t(28) > 7.6, p < .001$), i.e., this effect was not caused by slower movement and lower confidence on error trials. In summary, end-of-trial measures indicated that average speed varied more closely with subjective confidence than with objective accuracy.

3.3.2. Within-trial confidence

We next examined whether the *instantaneous* finger speed tracked in real time the participant’s prospective confidence that the final decision will be correct. Plotting the momentary y-speed as function of time again yielded a clear tree structure (Fig. 4a), now reflecting fluctuations in confidence: the finger accelerated after arrows that increased the amount of evidence ($\Delta|Evidence| > 0$) and decelerated after arrows that decreased it (see green circles in Fig. 4a). These speed changes are only partially visible in this figure, because the averaging over trials smoothed the speed. For a more detailed statistics of these effects, see Fig. S2b. The effect of $\Delta|Evidence|$ on speed was confirmed by regressing speed on $\Delta|Evidence|$ (coded as 1 or -1) of the 2nd to the 5th arrows: the per-participant coefficients were significantly higher than 0 for each arrow (Fig. 4b).

If finger speed reflects momentary confidence, then arrow direction changes (ADC), which decrease the post-decision confidence (see

above), should also decrease the momentary speed. This was indeed the case: when speed was regressed against ADC and $\Delta|Evidence|$ in a multiple regression, both effects were significant (Fig. 4c; the predictors were $\Delta|Evidence|$ of the 2nd–5th arrows as before, and direction changes in the 3rd to 5th arrows, each coded as 0, 1). The ADC effect on speed cannot be reduced to a pure motor effect (e.g., slowing down whenever the finger changes its direction following an ADC): the regression results survived the inclusion of the trajectory curvature as an additional covariate (measured as $|\frac{d\theta}{dt}|$ after smoothing θ with a Gaussian, $\sigma = 20$ ms).

Importantly, neither the $\Delta|Evidence|$ effect nor the ADC effect can be explained as an artifact of the other, because each is observed independently of the other when inspecting specific subsets of the data. First, the ADC effect was observed even when $\Delta|Evidence|$ was fixed (this happens whenever previous arrows sum to $|Evidence| = 0$, so that a new arrow necessarily increases $|Evidence|$ regardless of its direction; see grey circles in Fig. 4a). Second, conversely, an effect of $\Delta|Evidence|$ was observed even when it went against an arrow direction change, such as when a repetition of the same arrow direction decreased $|Evidence|$ and a direction change increased $|Evidence|$ (e.g., the 4th arrow in $\rightarrow\rightarrow\leftarrow\leftarrow$ vs. $\rightarrow\rightarrow\leftarrow\leftarrow$; see red circles in Fig. 4a). The independent effect of $\Delta|Evidence|$ shows that finger speed depends on the evidence accumulated throughout the trial and not just on local factors such as arrow direction changes.

Speed fluctuations did not result merely from averaging across trials, because the number of within-trial speed fluctuations correlated with the number arrow direction changes (Figs. 4d and S6c, d): the number of within-trial speed fluctuations (#Acc, see Section 2.5) was regressed for each participant against $|Evidence|$ and the number of arrow direction changes (#ADC), and the per-participant β values of the #ADC predictor were significantly higher than zero ($t(28) = 4.86$, one-tailed $p < .001$).

Does speed reflect online confidence, or could it solely reflect the momentary amount of evidence? The finding reported above, that

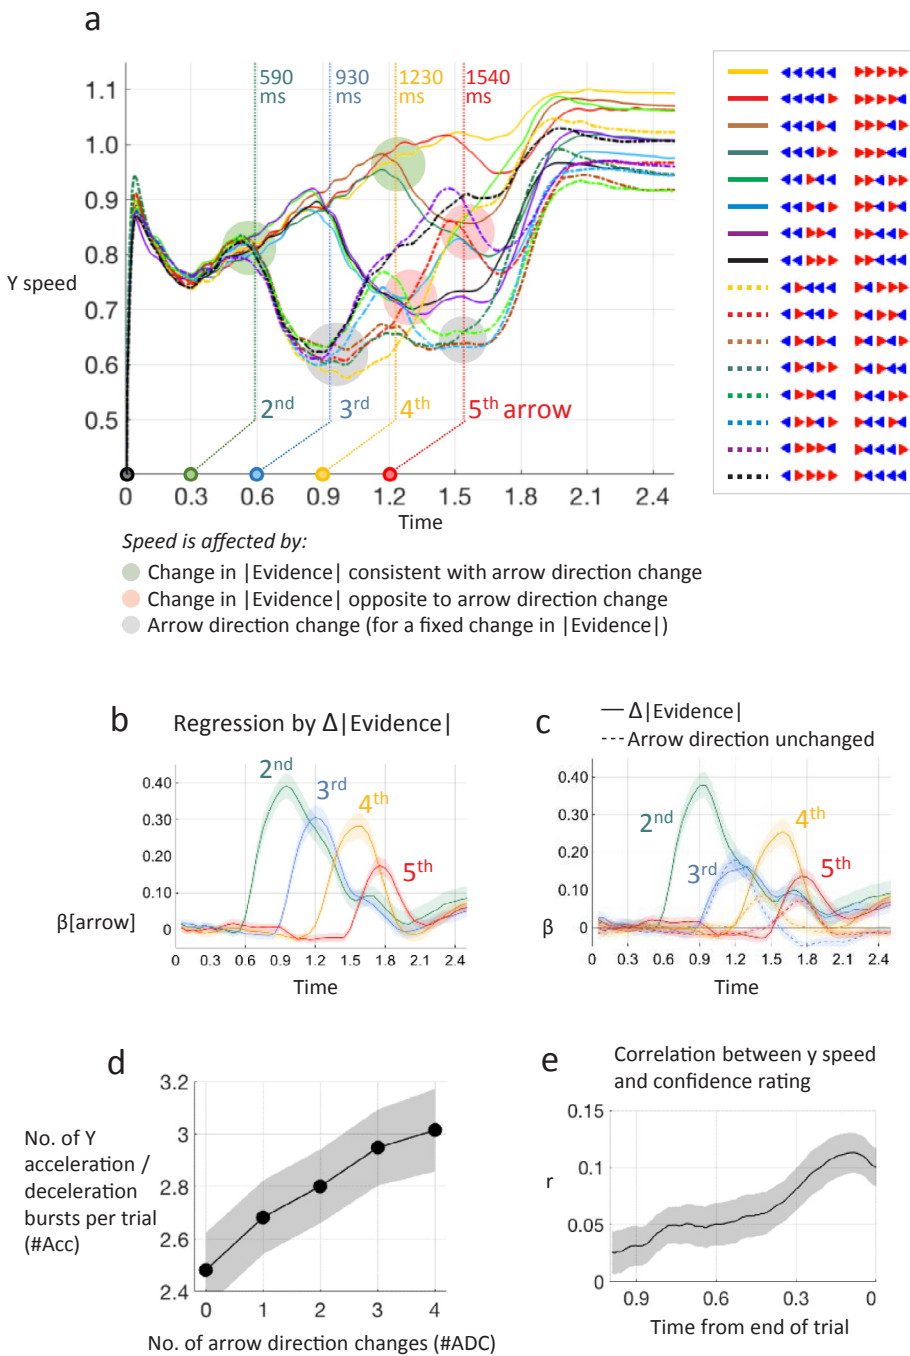


Fig. 4. The finger speed indexes the online computation of confidence. (a) Average y speed as a function of time for each sequence type (pooling together each pair of mirror sequences). The vertical lines show the time when an arrow started having a significant effect on the y speed (this time was calculated as described in Section 2.4). Two factors determined the effect of a given arrow: [1] The arrow's contribution to the absolute amount of evidence: increasing |Evidence| caused higher speed. [2] The arrow's agreement with the direction of the previous arrow: a change of direction decreased speed. The colored circles highlight examples of different combinations of these two factors. Either factor alone cannot explain the data: arrow direction changes affect speed even for fixed $\Delta|Evidence|$ (grey circles), and $\Delta|Evidence|$ affects speed even when its effect is opposite to that of the arrow direction change (red circles). (b) The time course of the effect of each arrow on finger speed. For each participant and each time point, the momentary y speed was regressed on the $\Delta|Evidence|$ provided by the 2nd to 5th arrows. (c) An extended regression model shows that, similarly to the post-decision confidence ratings, speed was not only affected by the accumulation of evidence, but also decreased locally following changes in the arrow direction. (d) The number of per-trial acceleration/deceleration bursts increased with the number of arrow direction changes, indicating within-trial fluctuations in speed. (e) The momentary y speed correlates with the post-decision confidence rating (r calculated per participant and then averaged), and the correlation increases towards the end of a trial. A color version of this figure is available in the online version of the article.

speed is affected not only by momentary evidence, but on top of it also by arrow direction changes, already suggests that speed is not merely about evidence. We now further tested whether trials with more speed fluctuations also had lower subjective confidence. We regressed, per participant, the confidence ratings against the number of y accelerations and decelerations per trial (#Acc, defined in Section 2.5). We also included as covariates the response accuracy (coded as 0 or 1), the specific sequence of arrows (coded as 32 binary predictors), and the trial serial number and its square (to control for drifts in confidence during the experiment session). The per-participant regression coefficients of #Acc were significantly negative ($t(28) = 3.13$, one-tailed $p = .001$). We verified that this result cannot be reduced to a pure motor effect: conceivably, the finger could slow down whenever it changes its direction, which may have been sufficient to explain the correlation between confidence and #Acc. However, the confidence-

#Acc correlation survived the inclusion of the number of finger direction changes (number of bends, described in Section 2.5) as an additional covariate in the regression ($t(28) = 2.64$, one-tailed $p = .003$).

A second finding that ties instantaneous speed with confidence is rooted in the hypothesis that a participant's post-decision confidence rating is derived from the online confidence shortly before the time of rating (near the end of the trial). This hypothesis predicts that the correlation between post-decisional confidence ratings and online prospective confidence (indexed as y-speed) would build up during the trial and culminate near the end of the trial. This was indeed the case (Fig. 4e). The correlation between end-speed and confidence rating was partly driven by the properties of the sequence type and the participants' accuracy, but did not reduce to it, because it survived the inclusion of sequence type and accuracy as covariates: for each participant, we regressed confidence ratings against the end-speed (y speed

100 ms before the end of the trial), the accuracy, the sequence type as 32 binary predictors, and the trial number and its square (to account for drift in confidence during the experiment). The speed regression coefficients were significantly higher than zero ($t(28) = 6.4$, one-tailed $p < .001$, mean $\beta = 0.08$). We acknowledge that although this correlation is significant across participants, it is low. This indicates other factors (e.g., accuracy, sequence type) also strongly account for the confidence ratings; it also suggests that our measure may suffer from noise and that additional factors may be at play. Nevertheless, end-speed predicted confidence rating better than several other speed measures, such as average speed; movement time (which is simply the inverse of average speed); peak y speed; peak y acceleration; Δ speed around the peak acceleration time (defined, for peak acceleration time T_{acc} , as $yspeed_{T_{acc} + 50ms} - yspeed_{T_{acc} - 50ms}$); the number of y acceleration/deceleration bursts per trial (see Section 2.5); and the absolute value of the end-of-trial x speed. None of these factors had a significant effect when added to the regression model (for all of them, $t(28) < 1.16$, one-tailed $p > .13$), and none decreased the weight of the end-speed factor. The finding that the speed effect cannot be reduced to an effect of average speed or movement time ($= \frac{1}{\text{average speed}}$) is especially important: it refutes the alternative explanation of the speed-confidence correlation as resulting from a post-decision process that heuristically uses movement time as an index of confidence (Kiani et al., 2014).

A third finding that ties instantaneous speed with confidence is that end-speed correlated with subjective confidence more than with objective accuracy: for each participant, we regressed end-speed either against confidence or against accuracy, in both cases with the sequence type (32 binary predictors), the trial number, and its square as covariates. The confidence-based regression model yielded better end-speed estimates than the accuracy-based regression model (comparing the per-participant regression MSEs, paired $t(28) = 2.09$, one-tailed $p = .02$). With xy-speed as the speed measure, this last finding did not reach significance (the two models had similar MSEs, paired $t(28) = 0.9$, one-tailed $p = .18$). Note, however, that this is the only finding that was not replicated with xy-speed.

We conclude with a note on the use of y-speed rather than the tangential speed (“xy-speed”). We made this choice a priori, assuming that subject’s choice translates into movements along the x-axis whereas subject’s confidence translates into faster movements towards the top of the screen where the response buttons are located. Restricting the speed analyses to the y-axis was therefore intended to avoid confounding effects between choices and confidence. However, we acknowledge that y-speed may not be independent from movement along the x-axis: y-speed could be lower when the finger bends sideways, thereby creating an artificial negative correlation between y-speed and sideways movement. To address this concern, we included movements along the x-axis as covariates in our regression analyses, as described above. We also tested directly that x-speed did not negatively correlate with y-speed: we regressed y-speed against $|x\text{-speed}|$, with the specific arrow sequence as covariate (coded as 32 binary predictors) to allow for stimulus-specific effects. This regression was run per participant and per time point in 50 ms intervals. The per-participant regression coefficients of x speed were not significantly lower than zero in any time point. Adding the confidence rating as another covariate yielded essentially the same results. Last, we show in Fig. S5 that all major speed-related results were replicated when using xy-speed instead of y-speed.

4. Discussion

4.1. Pre-decision computation of confidence

Our decision-making paradigm presented participants with short sequences of discrete bouts of evidence (left and right arrows), and recorded their pointing trajectories and post-decision confidence ratings. The results indicate that finger direction and finger speed convey

distinct information, respectively about choice and confidence. The observation of systematic fluctuations in the momentary speed during the trial therefore implies the existence of a continuous, on-line, pre-decisional computation of confidence.

The conclusion that the finger momentary speed reflects confidence is supported by several findings. First, whereas the finger directions continuously reflected the accumulation of evidence presented, speed – similarly to subjective confidence ratings – was additionally modulated by reversals in the arrow directions within the sequence of stimuli: finger movement was faster when an arrow direction agreed with the previous arrow, and slower when the arrow direction changed; and a larger number of arrow reversals yielded a larger number of finger accelerations and decelerations per trial. These speed fluctuations clearly reflected confidence, because more within-trial acceleration/deceleration bursts correlated with lower post-decisional subjective confidence ratings. Importantly, these findings cannot be attributed to differences between the effects of arrows in different positions within the arrow sequence, because the findings are based solely on comparisons between arrows in a given position. Second, the end-of-trial speed reliably predicted the subjective confidence ratings on top of the stimulus, accuracy, and movement time, indicating that confidence is affected not only by evidence and decision accuracy but also by additional factors, indexed by the momentary speed. Last, the end-of-trial speed correlated with confidence rating more than with the decision accuracy, indicating that speed reflects confidence rather than accuracy. These findings support the dissociation between confidence and decision accuracy, and point to the finger speed as a unique index of confidence.

Importantly, our findings support the existence of pre-decision on-line prospective confidence, but they do not rule out, or even deny, the existence of processes that further shape the retrospective subjective confidence after a choice was made. The existence of such post-decisional processes was well-established by other studies (Pleskac & Busemeyer, 2010; Resulaj et al., 2009). The paradigm we proposed here may potentially serve as a method to assess the relationship between prospective and retrospective confidence.

4.2. Measuring decision and confidence

Several previous studies have shown that movement trajectories can reveal how decisions fluctuate in real time according to the evidence stream (Resulaj et al., 2009; van den Berg, Anandalingam, et al., 2016) and its changes (Kiani et al., 2014). Our paradigm, however, offers three major enhancements: first, the nearly-continuous measure of decision allowed us to track several changes of mind per trial (Resulaj et al., 2009) and relate them to the precise timing with which discrete bouts of evidence are accrued. Second, finger tracking can monitor not only the unfolding decision process (de Lange et al., 2010, 2011), but also the associated fluctuations in confidence. Third, finger tracking measures confidence implicitly, without requiring any additional instruction or response. This is a major advantage over explicit methods such as asking participants to report confidence and decision simultaneously (van den Berg, Anandalingam, et al., 2016) or allowing them to opt out of decision when unconfident (Fetsch et al., 2014; Kiani & Shadlen, 2009). The measure of finger speed can be applied to any decision task, without any substantial modification, instruction, or training, so it should find a broad range of applications (including in young children or animals). A caveat is that, by nature, the implicit measure proposed here is only indirectly related to subjective confidence, and the tightness of this relationship could be questioned. However, we reported several features of this measure that strongly argue for a tight relationship with subjective confidence.

Studies of decision making often used continuous stimuli – e.g., detect the predominant motion direction in a cloud of moving dots (Shadlen & Newsome, 2001). A major source of uncertainty in this paradigm is perceptual: the task is hard because most dots move

randomly, so subjects must accumulate evidence over time to average out this noise. In contrast, our paradigm (previously used by de Lange et al., 2010, 2011) presented discrete bouts of evidence as clearly visible arrows, so the perceptual uncertainty was presumably low. Indeed, the participants were nearly flawless in the single-arrow trials, indicating that they could easily discriminate the arrow directions. Nevertheless, in trials with multiple arrows they made many errors, indicating that our task involves uncertainty, yet from other, non-perceptual sources. Arguably, to tap the central (as opposed to sensory) decision mechanisms, a task with low perceptual uncertainty may even be an advantage, as it may help narrowing the source of uncertainty to the decision mechanisms (rather than to perceptual mechanisms). Conceivably, there are two different sources of non-perceptual uncertainty in our task. One source pertains to the process that accumulates evidence. Indeed, a recent study showed that the accumulation of perceptual evidence was the main source of noise in tasks that involve such accumulation (Drugowitsch, Wyart, Devauchelle, & Koehlin, 2016; but see Brunton, Botvinick, & Brody, 2013 for an opposite view). Another source of uncertainty pertains to the termination of the accumulation process: in any given trial, our participants did not know in advance the number of arrows, i.e., when the accumulation of evidence would end. Possibly, these two types of uncertainty may determine two different dimensions of confidence: retrospective (the evidence accumulated so far) and prospective (also taking into account uncertainty regarding the end of the accumulation process).

In spite of the methodological differences between our task and decision-making tasks with continuous stimuli, our data replicated several patterns previously observed in continuous-stimuli tasks. First, increasing the amount of evidence increased the accuracy, confidence, and speed (Fig. 2a; Kiani & Shadlen, 2009; Kiani et al., 2014). Second, for a constant sum of evidence ($E_{\text{evidence}_{\text{left}}} - E_{\text{evidence}_{\text{right}}}$), confidence levels were lower for larger total $|\text{left evidence}| + |\text{right evidence}|$ – more arrows in our case (Figs. 2a and S1), and longer stimuli in previous studies (Kiani et al., 2014, Fig. 6). Third, similarly to Kiani et al. (2014), we observed that higher confidence was associated with faster decision. These similarities suggest that our paradigm and the random-dot kinematogram task tap overlapping mechanisms.

4.3. The computation of online confidence

Our findings imply that confidence is computed on-line, but also that this computation is imperfect. In our task, it was affected by the accumulated evidence, but it was also biased, as changes in arrow direction reduced confidence. Analogous results were observed during MEG recordings in the same task, where arrow direction changes caused additional activation, irrespectively of the amount of evidence collected (de Lange et al., 2010). Importantly, this increased activation occurred not only in sensory areas, but also in fronto-central areas, in accord with our finding that direction changes affected decision. Although this bias remains to be fully understood, a possible explanation is that the changes in perceptual evidence strain the neural mechanisms of decision making due to the need to inhibit and/or reverse the current decision. Such mechanism would agree with studies showing that the processing of incoming evidence is distorted towards previously-accumulated evidence (Russo, Meloy, & Wilks, 2000), and could also explain why the effect of late arrows is smaller than that of earlier arrows (Fig. 3b).

An open question, not resolved by our data, is the exact mechanism that calculates online confidence. One possibility is that the momentary online confidence derives from current evidence but also from the participant's expectation about future observations that may reverse the current decision. For example, if an arrow sequence starts with $\rightarrow\leftarrow$, an ideal observer should expect the decision to occur after the next arrow in case this is a 3-arrow trial, and after up to 3 more arrows in a 5-arrow trial. But if the sequence starts with $\rightarrow\rightarrow$, the expected decision time is shorter. These expectations can be used to compute online confidence,

because they are tightly correlated with the probability that the already-accumulated evidence would agree with the final decision. Nevertheless, this mechanism, in which both evidence and expectations can be directly derived from the arrows presented, cannot fully account for confidence in our task: first, we observed the speed-confidence correlation even when controlling for the specific arrow sequence. Second, even when the arrow sequence allowed an early decision (e.g., sequences starting with $\rightarrow\rightarrow\rightarrow$), the late arrows were not ignored but continued affecting speed.

Another open question, also not resolved by our data, is whether the online confidence computation is continuous or proceeds in discrete steps (Latimer, Yates, Meister, Huk, & Pillow, 2015; Shadlen et al., 2016). An extreme possibility is that not only confidence, but even the decision is calculated in discrete steps, i.e., the participant's response in a given trial is preceded by several interim decisions to deviate left or right (Fishbach, Roy, Bastianen, Miller, & Houk, 2007; Friedman et al., 2013; van den Berg, Zylberberg, Kiani, Shadlen, & Wolpert, 2016). Under this view, what we measured here was not a pre-decision confidence, but the post-decision confidence of interim decisions. We note, however, that this view seems impossible to refute behaviorally, because its loose definition of "decision" implies by definition that no pre-decision behavior can ever be measured. Neurophysiological measures may be more appropriate to tackle the issue of continuous versus discrete computation (Latimer et al., 2015).

4.4. Does speed affect subjective confidence or vice versa?

While several aspects of our data indicate that online confidence correlates with the momentary finger speed, this correlation remains compatible with either direction of causality. One view, promulgated here, is that online confidence determines the momentary speed. Such causality is useful, and thus provides a principled reason for the speed-confidence correlation: slowing down when unconfident gives more time for decision making, and this strategy may therefore optimize the reward rate, much like the setting of optimal decision bounds does (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006). In this respect, our finding is reminiscent of speed optimization in motor control, where movement time in reaching tasks depends on the target's angular size (Fitts, 1954). Fitts' law indicates that movement speed can be finely tuned as a function of anticipated difficulty. Our data show that it is also modulated on-line by the difficulty of the ongoing decision. This view is compatible with studies showing that confidence is used as input to subsequent decisions (Dayan & Daw, 2008; Middlebrooks & Sommer, 2012; Purcell & Kiani, 2016; van den Berg, Zylberberg, et al., 2016; Vickers, 1979) and to the selection of high-level problem-solving strategies (Ball, Onarheim, & Christensen, 2010).

The opposite view is that finger speed affects the post-trial ratings of subjective confidence. This is essentially the proposition of Kiani et al. (2014): they proposed that subjective confidence varies with decision duration ($= \frac{1}{\text{average speed}}$) (see also Patel, Fleming, & Kilner, 2012). Our analyses refute the specific assertion that the relevant factor is the decision duration: the confidence ratings of our participants correlated with the end-of-trial speed more than with the decision duration. More importantly, even if Kiani et al.'s specific model was amended to consider the end-of-trial speed rather than the decision duration, their model would still remain with the assumption that confidence depends on speed rather than vice versa. Under this assumption, a mechanism would be needed to explain why the fluctuations in momentary speed closely track the absolute evidence in the stimulus stream. As we saw earlier (Fig. 4a, red circles), local factors such as changes in arrow direction indeed matter but they do not suffice to explain these fluctuations, as the tracked variable has all the characteristics of a confidence measure. Conversely, our hypothesis offers a potential alternative interpretation for Kiani et al.'s (2014) correlation between decision duration and confidence. Their critical finding was that experimentally

increasing the decision duration, without changing the amount of visual evidence, resulted in lower confidence reports. However, they lengthened the decision by adding reversals in stimulus movement direction. If these reversals were perceived similarly to our arrow direction changes, this may explain the lower confidence reports without calling for a speed → confidence causality.

4.5. Conclusion

We demonstrated that finger speed during decision making provides an implicit online measure of instantaneous confidence that is predictive of future explicit confidence ratings. Our findings indicate that confidence is computed continuously, online, throughout the decision-making process, thus lending support to models of the brain as a device that computes with probabilistic estimates and probability distributions (Meyniel et al., 2015, 2015; Pouget et al., 2016). In the future, the simple behavioral method presented here should provide a more direct access into this computation and its biases.

Acknowledgements

We thank Fosca Al-Roumi for her help in designing the experiment and Moria Lahis for her technical help in administering the experiments. This research was supported by INSERM, CEA, Collège de France, Université Paris-Sud, an ERC grant “NeuroSyntax” to S.D., and the Fondation Bettencourt Schueller. D.D. is grateful to the Azrieli Foundation for the award of an Azrieli Fellowship.

Conflict of interest

The authors declare that no competing interests exist.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.cognition.2017.11.001>.

References

- Balakrishnan, J. D., & Ratcliff, R. (1996). Testing models of decision making using confidence ratings in classification. *Journal of Experimental Psychology: Human Perception and Performance*, 22(3), 615–633. <http://dx.doi.org/10.1037/0096-1523.22.3.615>.
- Ball, L. J., Onarheim, B., & Christensen, B. T. (2010). Design requirements, epistemic uncertainty and solution development strategies in software design. *Design Studies*, 31(6), 567–589. <http://dx.doi.org/10.1016/j.destud.2010.09.003>.
- Baranski, J. V., & Petrusic, W. M. (1994). The calibration and resolution of confidence in perceptual judgments. *Perception & Psychophysics*, 55(4), 412–428. <http://dx.doi.org/10.3758/BF03205299>.
- Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., ... Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron*, 60(6), 1142–1152. <http://dx.doi.org/10.1016/j.neuron.2008.09.021>.
- Berthier, N. E. (1996). Learning to reach: A mathematical model. *Developmental Psychology*, 32(5), 811–823. <http://dx.doi.org/10.1037/0012-1649.32.5.811>.
- 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. *Psychological Review*, 113(4), 700–765. <http://dx.doi.org/10.1037/0033-295X.113.4.700>.
- Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, 340(6128) 95 LP-98.
- Charles, L., King, J. R., & Dehaene, S. (2014). Decoding the dynamics of action, intention, and error detection for conscious and subliminal stimuli. *The Journal of Neuroscience*, 34(4), 1158–1170. <http://dx.doi.org/10.1523/JNEUROSCI.2465-13.2014>.
- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience*, 8(4), 429–453. <http://dx.doi.org/10.3758/CABN.8.4.429>.
- de Lange, F. P., Jensen, O., & Dehaene, S. (2010). Accumulation of evidence during sequential decision making: The importance of top-down factors. *Journal of Neuroscience*, 30(2), 731–738. <http://dx.doi.org/10.1523/JNEUROSCI.4080-09.2010>.
- de Lange, F. P., van Gaal, S., Lamme, V. A. F., & Dehaene, S. (2011). How awareness changes the relative weights of evidence during human decision-making. *PLoS Biology*, 9(11), e1001203. <http://dx.doi.org/10.1371/journal.pbio.1001203>.
- Dotan, D. (2017). Tracking the dynamics of Bayesian priors. In D. Dotan (Ed.). *Multi-digit number processing: Cognitive mechanisms and their impairments (PhD dissertation)* (pp. 119–129). Tel Aviv University.
- Drugowitsch, J., Wyart, V., Devauchelle, A.-D., & Koehlin, E. (2016). Computational precision of mental inference as critical source of human choice suboptimality. *Neuron*, 92(6), 1398–1411. <http://dx.doi.org/10.1016/j.neuron.2016.11.005>.
- Dunlosky, J., & Metcalfe, J. (2008). *Metacognition*. Thousand Oaks, CA: Sage Publications.
- Erb, C. D., Moher, J., Sobel, D. M., & Song, J. H. (2016). Reach tracking reveals dissociable processes underlying cognitive control. *Cognition*, 152, 114–126. <http://dx.doi.org/10.1016/j.cognition.2016.03.015>.
- Ferrell, W. R. (1995). A model for realism of confidence judgments: Implications for underconfidence in sensory discrimination. *Perception & Psychophysics*, 57(2), 246–254. <http://dx.doi.org/10.3758/BF03206511>.
- Fetsch, C. R., Kiani, R., Newsome, W. T., & Shadlen, M. N. (2014). Effects of cortical microstimulation on confidence in a perceptual decision. *Neuron*, 83(4), 797–804. <http://dx.doi.org/10.1016/j.neuron.2014.07.011>.
- Fetsch, C. R., Kiani, R., & Shadlen, M. N. (2014). Predicting the accuracy of a decision: A neural mechanism of confidence. *Cold Spring Harbor Symposia on Quantitative Biology*, 79, 185–197. <http://dx.doi.org/10.1101/sqb.2014.79.024893>.
- Fishbach, A., Roy, S. A., Bastianen, C., Miller, L. E., & Houk, J. C. (2007). Deciding when and how to correct a movement: Discrete submovements as a decision making process. *Experimental Brain Research*, 177, 45–63. <http://dx.doi.org/10.1007/s00221-006-0652-y>.
- Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47(6), 381–391. <http://dx.doi.org/10.1037/h0055392>.
- Friedman, J., Brown, S., & Finkbeiner, M. (2013). Linking cognitive and reaching trajectories via intermittent movement control. *Journal of Mathematical Psychology*, 57, 140–151. <http://dx.doi.org/10.1016/j.jmp.2013.06.005>.
- Grimaldi, P., Lau, H., & Basso, M. A. (2015). There are things that we know that we know, and there are things that we do not know we do not know: Confidence in decision-making. *Neuroscience & Biobehavioral Reviews*, 55, 88–97. <http://dx.doi.org/10.1016/j.neubiorev.2015.04.006>.
- Kahneman, D., & Tversky, A. (1982). Variants of uncertainty. *Cognition*, 11(2), 143–157. [http://dx.doi.org/10.1016/0010-0277\(82\)90023-3](http://dx.doi.org/10.1016/0010-0277(82)90023-3).
- Kepecs, A., & Mainen, Z. F. (2012). A computational framework for the study of confidence in humans and animals. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1594), 1322–1337. <http://dx.doi.org/10.1098/rstb.2012.0037>.
- Kepecs, A., Uchida, N., Zariwala, H. A., & Mainen, Z. F. (2008). Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455(7210), 227–231. <http://dx.doi.org/10.1038/nature07200>.
- Kiani, R., Corthell, L., & Shadlen, M. N. (2014). Choice certainty is informed by both evidence and decision time. *Neuron*, 84(6), 1329–1342. <http://dx.doi.org/10.1016/j.neuron.2014.12.015>.
- Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, 324(5928), 759–764. <http://dx.doi.org/10.1126/science.1169405>.
- Kording, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247. <http://dx.doi.org/10.1038/nature02169>.
- Latimer, K. W., Yates, J. L., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244), 184–187. <http://dx.doi.org/10.1126/science.aaa4056>.
- Meyniel, F., Schlunegger, D., & Dehaene, S. (2015). The sense of confidence during probabilistic learning: A normative account. *PLoS Computational Biology*, 11(6), e1004305. <http://dx.doi.org/10.1371/journal.pcbi.1004305>.
- Meyniel, F., Sigman, M., & Mainen, Z. F. (2015). Confidence as Bayesian probability: From neural origins to behavior. *Neuron*, 88(1), 78–92. <http://dx.doi.org/10.1016/j.neuron.2015.09.039>.
- Middlebrooks, P. G., & Sommer, M. A. (2012). Neuronal correlates of metacognition in primate frontal cortex. *Neuron*, 75(3), 517–530. <http://dx.doi.org/10.1016/j.neuron.2012.05.028>.
- Patel, D., Fleming, S. M., & Kilner, J. M. (2012). Inferring subjective states through the observation of actions. *Proceedings of the Royal Society B: Biological Sciences*, 279(1748), 4853–4860. <http://dx.doi.org/10.1098/rspb.2012.1847>.
- Pinheiro-Chagas, P., Dotan, D., Piazza, M., & Dehaene, S. (2017). Finger tracking reveals the covert stages of mental arithmetic. *Open Mind*. http://dx.doi.org/10.1162/OPMI_a.00003.
- Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, 117(3), 864–901. <http://dx.doi.org/10.1037/a0019737>.
- Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: Distinct probabilistic quantities for different goals. *Nature Neuroscience*, 19(3), 366–374. <http://dx.doi.org/10.1038/nn.4240>.
- Purcell, B. A., & Kiani, R. (2016). Hierarchical decision processes that operate over distinct timescales underlie choice and changes in strategy. *Proceedings of the National Academy of Sciences*, 113(31), E4531–E4540. <http://dx.doi.org/10.1073/pnas.1524685113>.
- R Core Team (2015). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, 461(7261), 263–266. <http://dx.doi.org/10.1038/nature08275>.
- Russo, J. E., Meloy, M. G., & Wilks, T. J. (2000). Predecisional distortion of information by auditors and salespersons. *Management Science*, 46(1), 13–27. <http://dx.doi.org/10.1287/mnsc.46.1.13.15127>.
- Shadlen, M. N., Kiani, R., Newsome, W. T., Gold, J. I., Wolpert, D. M., Zylberberg, A., ... Roitman, J. (2016). Comment on “Single-trial spike trains in parietal cortex reveal

- discrete steps during decision-making” *Science*, 351(6280), 1406. <http://dx.doi.org/10.1126/science.aad3242>.
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86(4), 1916–1936.
- van den Berg, R., Anandalingam, K., Zylberberg, A., Kiani, R., Shadlen, M. N., & Wolpert, D. M. (2016). A common mechanism underlies changes of mind about decisions and confidence. *eLife*, 5, e12192. <http://dx.doi.org/10.7554/eLife.12192>.
- van den Berg, R., Zylberberg, A., Kiani, R., Shadlen, M. N., & Wolpert, D. M. (2016). Confidence is the bridge between multi-stage decisions. *Current Biology*, 26(23), 3157–3168. <http://dx.doi.org/10.1016/j.cub.2016.10.021>.
- Vickers, D. (1979). *Decision processes in visual perception*. London: Academic Press.
- Yang, T., & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature*, 447(7148), 1075–1080. <http://dx.doi.org/10.1038/nature05852>.