Temporal Scaling of Human Scalp-Recorded Potentials During Interval Estimation

Cameron D. Hassall, Jack Harley, Nils Kolling, and Laurence T. Hunt*

Oxford Centre for Human Brain Activity, Wellcome Centre for Integrative Neuroimaging,

Department of Psychiatry, University of Oxford, Oxford OX3 7JX, United Kingdom

*laurence.hunt@psych.ox.ac.uk

Abstract

Standard event-related potential analysis assumes fixed-latency responses relative to experimental events – yet recent single unit recordings have revealed neural activity *scales* to span different durations during behaviours demanding flexible timing. We use a novel approach to unmix fixed-time and scaled-time components in human electroencephalography, recorded across three tasks. A consistent and distinct scaled-time component is revealed, demonstrating temporal scaling can reliably be measured at the scalp.

Introduction

Action and perception in the real world require flexible timing. We can walk quickly or slowly, recognize the same piece of music played at different tempos, and form temporal expectations over long and short intervals. Flexible timing is critical in our lives, yet its neural correlates have proven difficult to study. One source of difficulty is disagreement over how the brain represents time. For example, the classic pacemaker-accumulator model (Treisman, 1963) relies on a dedicated timing mechanism. Other models represent time intrinsically through oscillatory alignment (Matell & Meck, 2004) or network population dynamics (Buonomano & Maass, 2009). However, it has recently been shown that brain activity at the level of individual neurons can be best explained by a *temporal scaling* mode (Wang et al., 2018).

When monkeys are cued to produce either a short or a long interval, medial frontal cortex (MFC) unit activity can be explained by a single response that is stretched or compressed according to the length of the produced interval – a temporally scaled response. This suggests that flexible motor timing is achieved by adjusting the speed of a common neural process. Temporal scaling of neural responses is also implicit in other settings, such as the process of evidence integration during decision making (O'Connell et al., 2018). Indeed, recent approaches to studying time-warped responses in neural populations have revealed time-warping as a common property across many different population recordings (Williams et al., 2020).

It is currently unclear how temporal scaling of neural responses may manifest at the scalp (if at all) using non-invasive recording in humans. Although electroencephalography (EEG) has played a prominent role in understanding the neural basis of timing (Macar & Vidal, 2004), the method commonly used to analyse such data has been the event-related potential (ERP), which averages event-locked EEG across multiple repetitions. This implicitly assumes

that neural activity occurs at *fixed-time* latencies with respect to experimental events. Sometimes these event-related potentials have been found to ramp at different speeds for different temporal intervals (Macar & Vidal, 2004), perhaps suggestive of temporal scaling - but crucially, they appear mixed at the scalp with fixed-time components, due to the *superposition problem* (Chapter 2 in Luck, 2014).

We therefore developed an approach to *unmix* scaled-time and fixed-time components in the EEG, which we first tested on simulated data (Fig. 1a). Our proposed method builds on existing regression-based approaches (Ehinger & Dimigen, 2019; Smith & Kutas, 2015a) that have proven useful in unmixing fixed-time components that overlap, e.g. stimulus-related activity and response-related activity. These approaches estimate the ERP using a general linear model (GLM) in which the design matrix is filled with time-lagged dummy variables (1s around the events of interest, 0s otherwise). Importantly, these 'stick functions' can overlap in time to capture overlap in the underlying neural responses (Fig. 1b); in situations without any overlap, the GLM would exactly return the conventional ERP. To reveal scaled-time responses, we allowed the duration of the stick function to vary depending upon the interval between stimulus and response, meaning that the same neural response could span different durations on different trials. As such, the returned scaled-time potential is no longer a function of real-world ('wall clock') time, but instead a function of the *percentage of time elapsed* between stimulus and response.

Results

As a proof of concept, we simulated data for at a single EEG sensor for an interval timing task, consisting of two *fixed-time components* (locked to cues and responses), and one *scaled-*

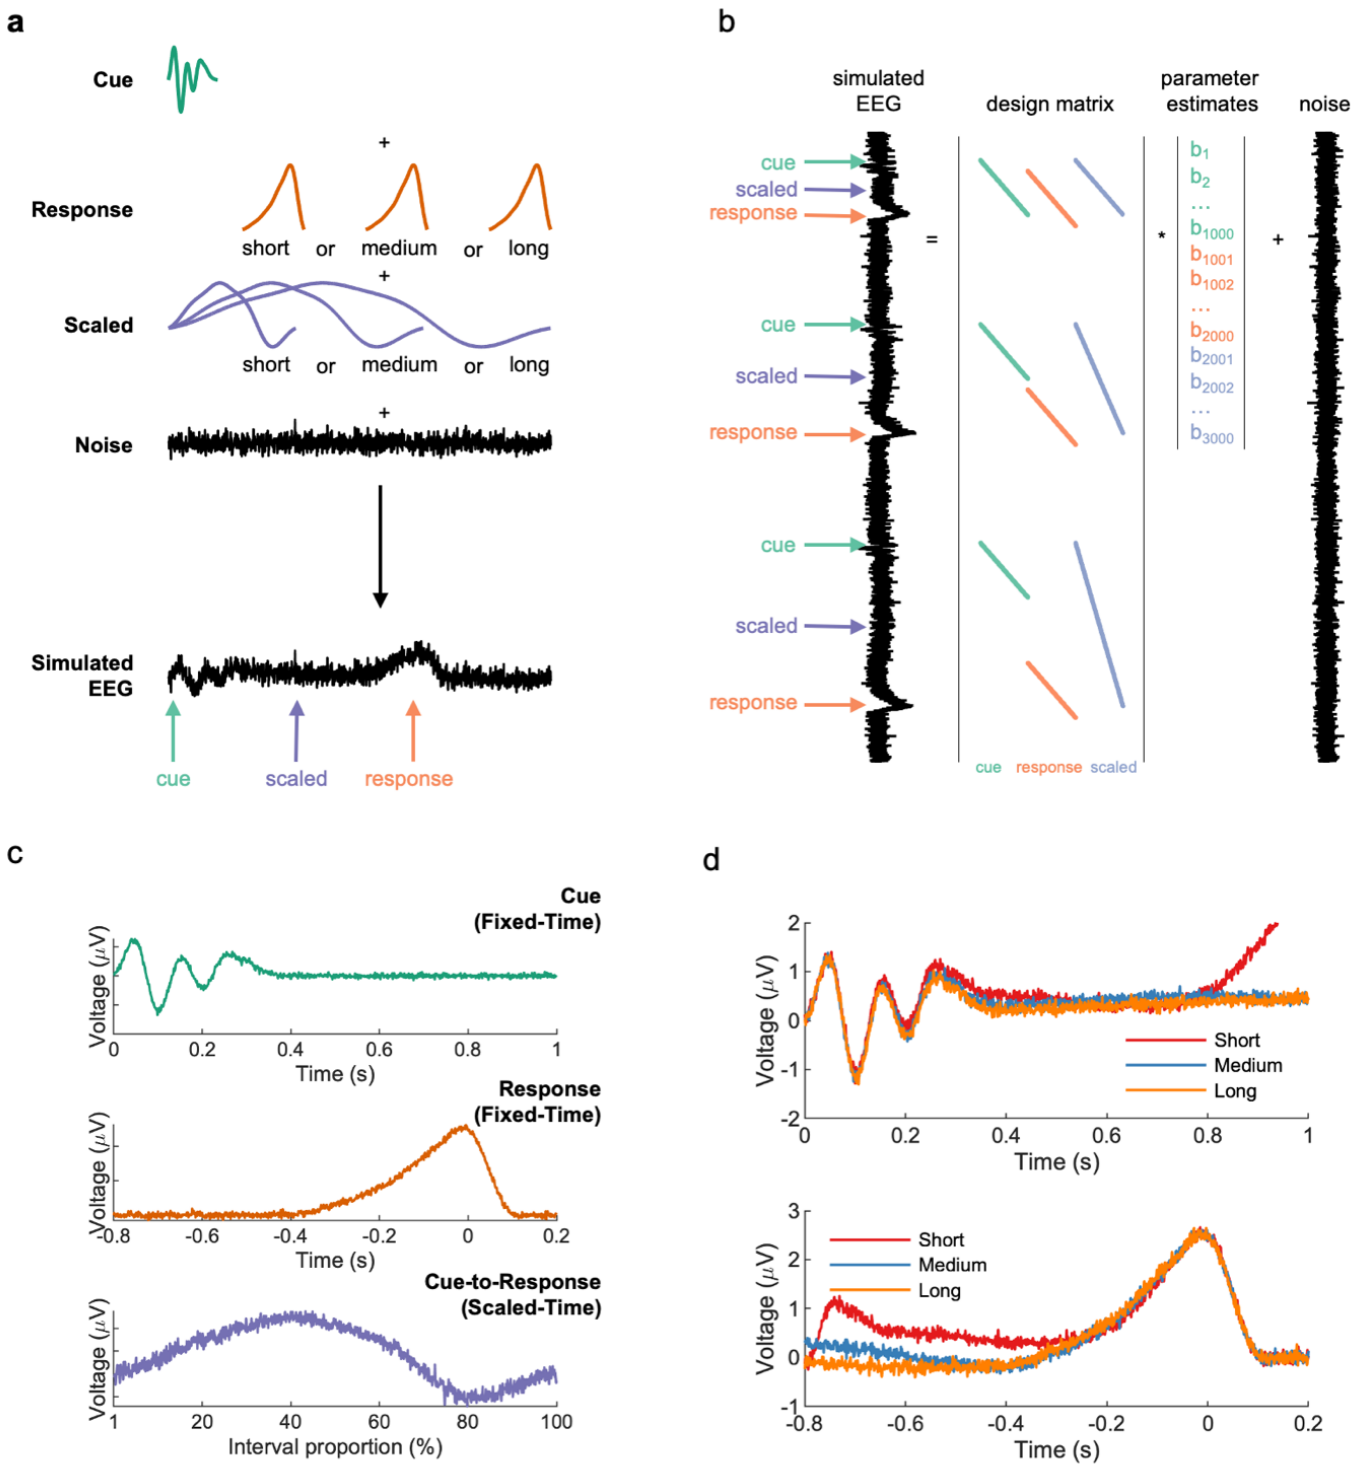

Figure 1. **Regression based unmixing of simulated data successfully recovers scaled-time and fixed-time components.** (a) EEG data were simulated by summing fixed-time components (cue and response), a scaled-time component with differing durations for different trials (short, medium, or long), and noise. (b) The simulated responses were unmixed via a GLM with stick basis functions: cue-locked, response-locked, and a single scaled-time basis spanning from cue to response. (c) The GLM successfully recovered all three components, including the scaled-time component. (d) A conventional ERP analysis (cue-locked and response-locked averages) of the same data obscured the scaled-time component.

time component spanning between cues and responses (Fig. 1a). Our proposed method was successful in recovering all three components (Fig. 1c), whereas a conventional ERP

approach obscured the scaled-time component (Fig. 1d). Crucially, in real EEG data we also repeat this approach across all sensors, potentially revealing different scalp topographies (and hence different neural sources) for fixed-time versus scaled-time components.

We used this approach to analyze EEG recorded during three interval timing tasks (Supplementary Fig. 1). In one task, participants produced a target interval (short, medium, or long) following a cue. Feedback was provided, and participants were able to closely match the target intervals. In a second, participants evaluated a computer-produced interval. The closer the produced interval was to the target interval, the more likely participants were to judge the response as 'on time'. In a third (previously analyzed (Breska & Deouell, 2017a, 2017b)) task, participants made temporal predictions about upcoming events based on rhythmic predictions.

In all three tasks, we observed a scaled-time component that was distinct from the preceding and following fixed-time components (Fig. 2), which resembled conventional ERPs (Supplementary Fig. 3). Typically, ERP components are defined by their polarity, scalp distribution, and latency (Luck, 2014). The observed scaled-time components shared a common polarity (negative) and scalp distribution (central), which is notably consistent with the medial frontal recording site where temporally scaled single-unit responses were previously identified (Wang et al., 2018). Although our scaled-time components were estimated by time-warping a common signal so that they could span a variable delay period, their 'latency' was nevertheless consistent, in that the scaled-time signal grew and appeared to peak later in the timed interval. This is again reminiscent of the time course of scaled time

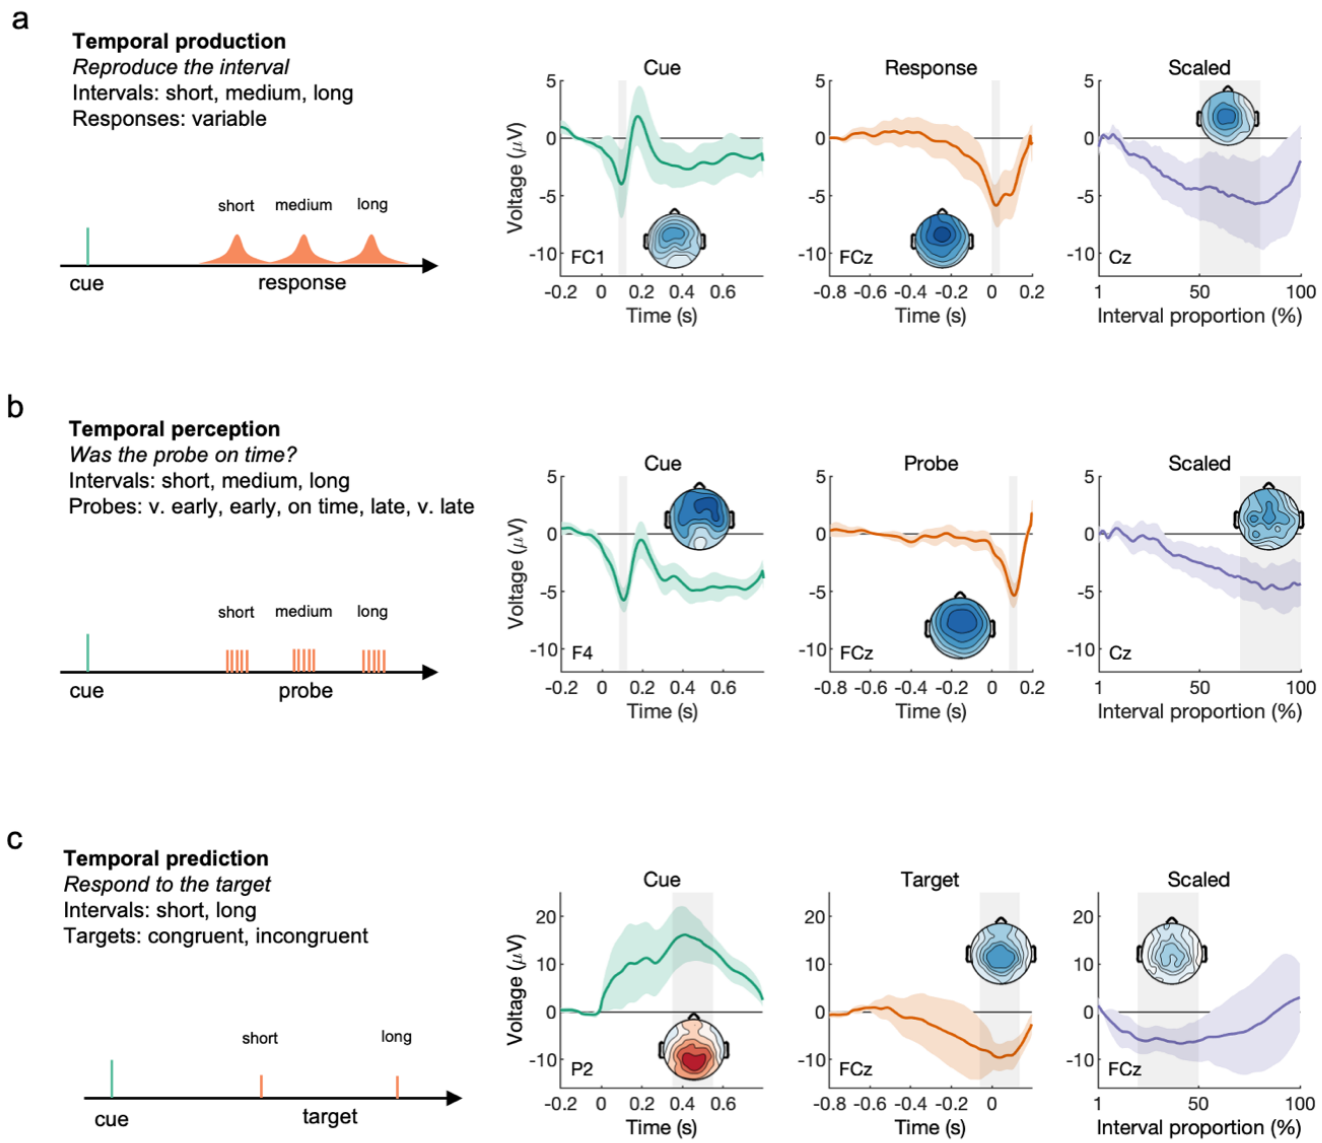

Figure 2. **Scaled-time components were consistently observed across three time-estimation paradigms, with distinct scalp topographies from fixed-time components.** Data were analysed from: (a) a temporal production task; (b) a temporal judgement task; (c) a temporal prediction task[11,12]. All had distinct fixed-time components relative to task-relevant events (left/middle columns), and a common negative scaled-time component over central electrodes, reflecting interval time (right column). Error bars represent 95% confidence intervals.

components across the neural population in medial frontal recordings (Wang et al., 2018). Single-sample *t*-tests of the mean voltage in the shaded regions in Fig. 2 revealed a significant scaled-time signal in the production task ($t(9) = -3.19$, $p = .01$, Cohen's $d = -1.01$), the perception task ($t(9) = -4.79$, $p = .001$, Cohen's $d = -1.52$), and the prediction task ($t(18) = -4.03$, $p < .001$, Cohen's $d = -0.92$). In many cases, scaled-time components were reliably observed at the single-subject level (Supplementary Figs. 4-6).

We then examined how the scaled-time component related to behavioural variability: does the latency of the scaled-time component predict participants' response time? To measure component latency, we applied an approach developed in Hunt et al. (2015) and Möcks (1986), using principal component analysis (PCA) to model delay activity over central electrodes in the temporal production task, after first regressing out fixed-time components from the data. PCA was applied separately to each of three produced intervals. This consistently revealed a first principal component that matched the shape of the scaled-time component identified in Fig. 2a, and a second principal component that matched its *temporal derivative*. This analysis not only provides an additional way of confirming the presence of the scaled-time component in our data (as it is the first principal component of the residuals after removing fixed-time components), but crucially adding or subtracting PC2 captures variation in the *latency* of this scaled-time component (Fig. 2b). Across response time quantiles, we found that PC2 scores were significantly related to response times (Fig. 3c; $F(2,18) = 9.05$, $p = .002$). This implies that the earlier in time that the scaled-time component peaked, the faster the subject would respond on that trial. Note, however, that within-condition behaviour was highly consistent in Task 1 (Supplementary Fig. 2); we would therefore expect the relationship between scaled-time component latency and behaviour to be even stronger in tasks with greater response variability between trials.

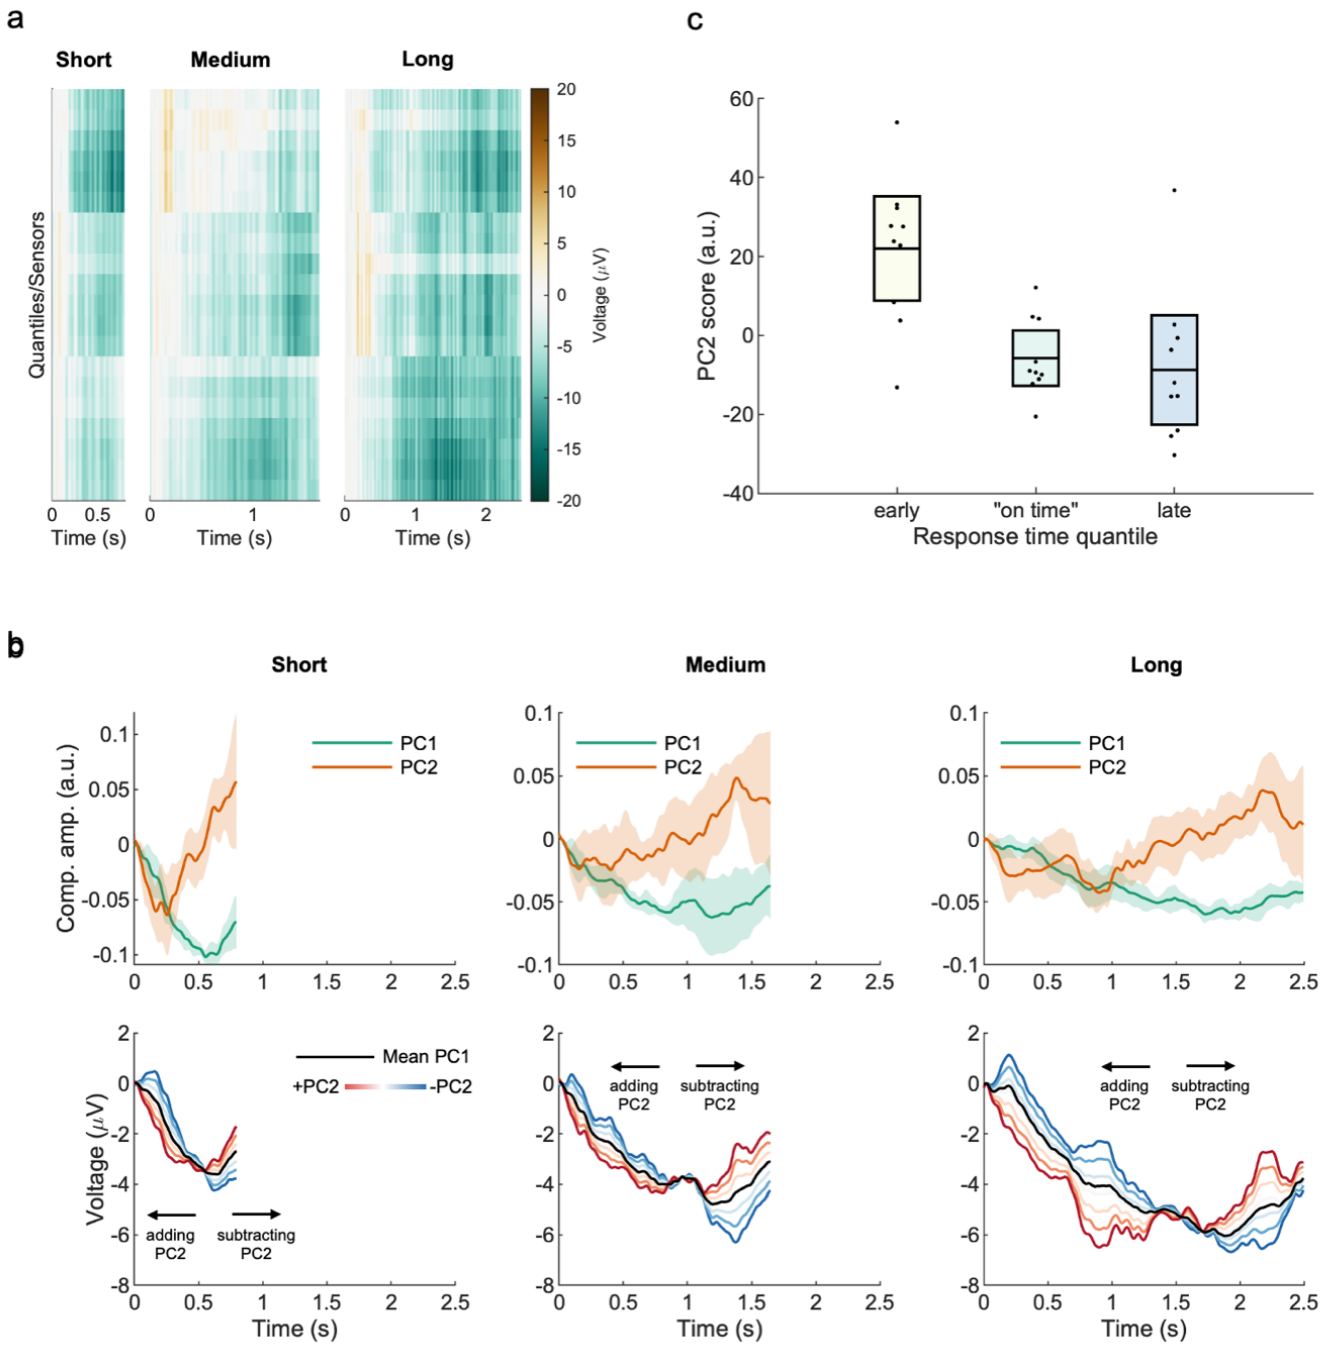

Figure 3. **Variation in scaled-time components predict behavioural variation in time estimation.** (a) Cue-locked EEG over central electrodes (FC1, FCz, FC2, Cz, CP1, CPz, CP2) was grouped by response time (early, on time, or late), averaged, and stacked for each target interval (short, medium, or long). Data for one participant is shown. A separate PCA was run for each target interval and participant. (b) The first two principal components for each target interval represent the amplitude (PC1) and first derivative (PC2) of the time-scaled component (top panel). Adding or subtracting different amounts of PC2 to PC1 shifted the peak earlier or later in time (bottom panel). (c) PC2 scores depended on response time, implying the scaled-time component peaked earlier for fast responses and later for slow responses. Error bars represent 95% confidence intervals.

Discussion

Our results provide a general method for recovering temporally scaled signals in human EEG, where scaled-time components are mixed at the scalp with conventional fixed-

time ERPs. We focused here on interval production and perception, but we anticipate other temporally scaled EEG and MEG signals will be discovered for cognitive processes known to unfold over varying timescales. For example, the neural basis of flexible sequential behaviours (such as speech) is still unknown, but may involve a form of temporal scaling (Remington et al., 2018). Flexible timing is also important in decision-making tasks, where evidence accumulation can proceed quickly or slowly depending on the strength of the evidence (O'Connell et al., 2018). Flexible timing can help facilitate a range of adaptive behaviours via *temporal attention* (Nobre & van Ede, 2018), while disordered timing characterizes several clinical disorders (Allman & Meck, 2012), underscoring the importance of characterising temporal scaling of neural responses in human participants.

Acknowledgements

This research was funded by a Natural Sciences and Engineering Research Council of Canada (NSERC) Postdoctoral Fellowship to C.D.H., a BBSRC AFL Fellowship (BB/R010803/1) to N.K., a Sir Henry Dale Fellowship from the Royal Society and Wellcome (208789/Z/17/Z) to L.T.H., a NARSAD Young Investigator Award from the Brain and Behavior Research Foundation to L.T.H., and supported by the NIHR Oxford Health Biomedical Research Centre. The Wellcome Centre for Integrative Neuroimaging was supported by core funding from Wellcome Trust (203139/Z/16/Z). For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

Author Contributions

N.K. and L.T.H. conceived the experiments and methodology. C.D.H. and L.T.H. designed the experiments and developed the methodology. C.D.H. and J.H. performed the experiments (except for the prediction task, where data was downloaded from a previous publication (Breska & Deouell, 2017a, 2017b)). C.D.H. and L.T.H. analyzed the data. C.D.H. and L.T.H. wrote the manuscript with input from the other authors.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper.

Correspondence and requests for materials should be addressed to L.T.H.

Methods

Simulations

We simulated cue-related and response-related EEG in a temporal production task using MATLAB 2020a (Mathworks, Natick, USA). Cue and response were separated by either a short, medium, or long interval. During the delay period, we simulated a scaled response that stretched or compressed to fill the interval. All three responses (cue, response, scaled) were summed together at appropriate lags (short, medium, or long), with noise – see Fig. 1a. In total, we simulated 50 trials of each condition (short, medium, long).

To unmix fixed-time and scaled-time components, we used a regression-based approach (Ehinger & Dimigen, 2019; Smith & Kutas, 2015a, 2015b) in which the continuous EEG at one sensor Y is modelled as a linear combination of the underlying event-related responses β, which are unknown initially. The model can be written in equation form as:

$$Y = X\beta + \varepsilon$$

where X is the design matrix and ε is the residual EEG not accounted for by the model. X contains as many rows as EEG data points, and as many columns as predictors (that is, the number of points in the estimated event-related responses). In our case, X was populated by 'stick functions' – non-zero values around the time of the modelled events, and zeros otherwise. We included in X two fixed-time components, the cue and the response, as stick functions of set EEG duration (with variables set to 1). In other words, the height of the fixed-time stick function was constant across events of the same type and equal to its width. To model a temporally-scaled response, we used the MATLAB *imresize* function (Image Processing Toolbox, R2020b) with 'box' interpolation to stretch/compress a stick function so that it spanned the duration between cue and response (other interpolation methods were tried – see Supplementary Fig. 7 – but this choice had little effect on the results). Thus, the duration of the scaled stick function varied from trial to trial (Fig. 1b). The goal here was to

estimate a single scaled-time response to account for EEG activity across multiple varying delay periods. For the fixed-time responses, each column of X represents a latency in ms before/after an experimental event; by contrast, for the scaled-time responses, each column of X represents the *percentage* of time that has elapsed between *two* events (stimulus and response). Simulation code is available at https://git.fmrib.ox.ac.uk/chassall/temporal-scaling.

Production and Perception Tasks

Participants

Participants completed both the production and perception tasks within the same recording session. We tested ten university-aged participants, 5 male, 2 left-handed, M_{age} = 23.40, 95% CI [21.29, 25.51]. Participants had normal or corrected-to-normal vision and no known neurological impairments. Participants provided informed consent approved by the Medical Sciences Interdivisional Research Ethics Committee at the University of Oxford. Following the experiment, participants were compensated £20 (£10 per hour of participation) plus a mean performance bonus of £3.23, 95% CI [2.92, 3.55].

Apparatus and Procedure

Participants were seated approximately 64 cm from a 27-inch LCD display (144 Hz, 1 ms response rate, 1920 by 1080 pixels, Acer XB270H, New Taipei City, Taiwan). Visual stimuli were presented using the Psychophysics Toolbox Extension (Brainard, 1997; Pelli, 1997) for MATLAB 2014b (Mathworks, Natick, USA). Participants were given written and verbal instructions to minimize head and eye movements. The goal of the production task was to produce a target interval and the goal of the perception task was to judge whether or not a computer-produced interval was correct.

The experiment was blocked with ten trials per block. There were 18 production blocks and 18 perception blocks, completed in random order. Prior to each block, participants listened to five isochronic tones indicating the target interval. Beeps were 400 Hz sine waves

of duration 50 ms and an onset/offset ramping to a point 1/8 of the length of the wave (to avoid abrupt transitions). The target interval was either short (0.8 s), medium (1.65 s), or long (2.5 s).

In production trials, participants listened to a beep then waited the target time before responding. Feedback appeared after a 400-600 ms delay (uniform distribution) and remained on the display for 1000 ms. Feedback was a 'quarter-to' clockface to indicate 'too early', a 'quarter-after' clockface to indicate 'too late', or a checkmark to indicate an on-time response. Feedback itself was determined by where the participant's response fell relative to a window around the target duration. The response window was initialized to +/- 100 ms around each target, then changed following each feedback via a staircase procedure: increased on each side by 10 ms following a correct response and decreased by 10 ms following an incorrect response (either too early or too late).

In perception trials, participants heard two beeps, then were asked to judge the correctness of the interval, that is, whether or not the test interval matched the target interval. Test intervals (very early, early, on time, late, very late) were set such that each subsequent interval was 25% longer than the previous (see Supplementary Table 1). Participants were then given feedback on their judgement – a checkmark for a correct judgement, or an 'x' for an incorrect judgement.

For each task, participants gained 2 points for each correct response and lost 1 point for each incorrect response. At the end of the experiment points were converted to a monetary bonus at a rate of £0.01 per point.

Data Collection

In the perception task we recorded participant response time from cue, trial outcome (early, late, on time), and staircase- response window. In the production task, we recorded trial 'on time' judgements (yes/no), and trial outcome (correct/incorrect).

We recorded 36 channels of EEG, referenced to AFz. Data were recorded at 1000 Hz using a Synamps amplifier and CURRY 8 software (Compumetrics Neuroscan, Charlotte, USA). The electrodes were sintered Ag/AgCl (EasyCap, Herrsching, Germany). 31 of the electrodes were laid out according to the 10-20 system. Additional electrodes were placed on the left and right mastoids, on the outer canthi of the left and right eyes, and below the right eye. The reference electrode was placed at location AFz, and the ground electrode at Fpz.

Prediction Task

In this previously published (Breska & Deouell, 2017a, 2017b) experiment, 19 participants responded to the onset of a visual target following a visual warning cue. The delay between cue and target was either short (700 ms) or long (1300 ms) and, in some conditions, congruent with a preceding stimulus stream. Only these predictable trials were included in the current analysis (i.e., the 'valid' trials in the 'Rhythmic' and 'Repeated-Interval' conditions – see Breska and Deouell, 2017a, 2017b for more detail).

Data Analysis

Behavioural data

For the perception task, we computed mean window size and mean produced interval for each participant. For the production task, we computed mean likelihood of responding yes to the 'on time' prompt, for each condition (short, medium, long) and interval (very early, early, on time, late, very late).

EEG Preprocessing

For all three tasks, EEG was preprocessed in MATLAB 2020b (Mathworks, Natick, USA) using EEGLAB (Delorme & Makeig, 2004). We first down-sampled the EEG to 200 Hz, then applied a 0.1-20 Hz bandpass filter and 50 Hz notch filter. The EEG was then re-referenced to the average of the left and right mastoids (and AFz recovered in the production/perception tasks). Ocular artifacts were removed using independent component

208 analysis (ICA). The ICA was trained on 3-second epochs of data following the appearance of
209 the fixation cross at the beginning of each trial. Ocular components were identified using the
210 *iclabel* function and then removed from the continuous data.

211 ***ERPs***

212 To construct conventional event-related potentials (ERPs), we first created epochs of
213 EEG around cues (all tasks), responses (perception task), probes (production task), and
214 targets (prediction task). Cue-locked ERPs extended from 200 ms pre-cue to either 800, 1650,
215 or 2500 ms post-cue (the short, medium, and long targets) in the perception/production tasks
216 and 700 or 1300 ms in the prediction task (the short and long targets). Epochs were baseline-
217 corrected using a 200 ms pre-cue window. We also constructed epochs from 800, 1650, or
218 2500 ms prior to the response/probe in the production/perception tasks and 700 or 1300 ms
219 prior to the target in the prediction task to 200 ms after the response/probe/target. A baseline
220 was defined around the event of interest (mean EEG from -20 to 20 ms) and removed. We
221 then removed any trials in which the sample-to-sample voltage differed by more than 50 μV
222 or the voltage change across the entire epoch exceeded 150 μV. Cue and
223 response/probe/target epochs were then averaged for each participant, task (production,
224 perception, prediction), and condition (short, medium, long). Finally, participant averages
225 were combined into grand-average waveforms at electrode FCz, a location where timing-
226 related activity has been previously observed (Macar & Vidal, 2004). See Supplementary Fig.
227 3.

228 ***rERPs***

229 To unmix fixed-time and scaled-time components in our EEG data, we estimated
230 regression-ERPs (rERPs) following the same GLM procedure we used with our simulated
231 data, but now applied to each sensor. We used a design matrix consisting of a regular stick
232 functions for cue and response/probe/target and a stretched/compressed stick function

spanning the interval from cue to response/probe/target. In particular, we estimated cue-locked responses that spanned from 200 ms pre-cue to 800 ms post-cue. The response/probe/target response interval spanned from -800 to 200 ms. Each fixed-time response thus spanned 1000 ms, or 200 EEG sample points. The scaled-time component, as described earlier, was modelled as a single underlying component (set width in *X*) that spanned over multiple EEG durations (varying number of rows in *X*). Thus, our method required choosing how many scaled-time sample points to estimate (the width in *X*). For the production/perception tasks, we chose to estimate 330 scaled-time points, equivalent to the duration of the 'medium' interval. For the prediction task, we chose to estimate 200 scaled-time points, equivalent to the mean of the short and long conditions (700 ms, 1300 ms). Unlike the conventional ERP approach, this analysis was conducted on the continuous EEG. To identify artifacts in the continuous EEG, we used the *basicrap* function from the ERPLAB (Lopez-Calderon & Luck, 2014) toolbox with a 150 μV threshold (2000 ms window, 1000 ms step size). A sample was flagged if it was 'bad' for any channel. Flagged samples were excluded from the GLM (samples removed from the EEG and rows removed from the design matrix). Additionally, we removed samples/rows associated with unusually fast or slow responses in the production task (less than 0.2 s or more than 5 s). On average, we removed 2.17 % of samples in the production task (95% CI [1.49, 2.86]), 3.75 % of samples in the perception task (95% CI [2.39, 5.10]), and 1.03% of samples in the prediction task (95% CI [0.95, 1.10]).

To impose a smoothness constraint on our estimates, we used a first-derivative form of Tikhonov regularization. Tikhonov regularization reframes the GLM solution as the minimization of:

$$\|X\beta - \mathrm{Y}\|^2 + \lambda\|\mathrm{L}\beta\|^2$$

where L is the regularization operator and λ is the regularization parameter. In other words, we aimed to minimize a penalty term in addition to the usual residual. This has the solution

$$(X^T X + \lambda \mathrm{L})^{-1} X^T Y$$

In our case, L approximated the first derivative as a scaled finite difference (Reichel & Ye, 2008):

$$L = \frac{1}{2}\begin{bmatrix} 1 & -1 & 0 & \dots & 0 & 0 \\ 0 & 1 & -1 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 & -1 \end{bmatrix}$$

We then chose regularization parameters for each participant using 10-fold cross validation. Our goal here was to minimize the mean-squared error of the residual EEG at electrode FCz, our electrode of interest. The following λs were tested on each fold: 0.001, 0.01, 0, 1, 10, 100, 1000, 10000, 100000. An optimal λ was chosen for each participant corresponding to the parameter with the lowest mean mean-squared error across all folds. See Supplementary Table 3 for a summary.

Inferential Statistics

We quantified the amplitude of the scaled-time component in two ways. First, we computed a 95% confidence interval at each 'timepoint' in the scaled-time signal. Next, in line with conventional ERP analyses (Luck, 2014), we computed the mean signal around the apparent peak. We then confirmed the existence of the signal using a single-sample repeated-measures *t*-test. These analyses were done at the scalp location where the mean signal was greatest, i.e. electrode Cz in the production/perception tasks and FCz in the prediction task.

PCA

To explore the link between the scaled-time component and behaviour, we regressed out the fixed-time components from the EEG in the temporal production task – that is, we reconstructed the preprocessed data using only the scaled-time regressors plus residuals. Only mid-frontal electrodes were considered: FC1, FCz, FC2, Cz, CP1, CPz, and CP2. We then

constructed epochs starting at the cue and ending at the target interval (800 ms, 1650 ms, or 2500 ms). Epochs within each condition (short, medium, long) were further grouped into three equal-sized response-time bins (early, on time, late) and averaged for each electrode and participant (Fig. 3a). We then conducted a PCA for each condition (short, medium, long) and participant. See Supplementary Table 2 for amount of variance explained by PC1 and PC2. To visualize the effect of PC2, we computed the mean PC2 across all participants. We then added more or less of the mean PC2 to the mean PC1 projection and applied a 25-point moving-mean window for visualization purposes (Fig. 3b). In order to choose a reasonable range of PC2 scores, we examined the average minimum and maximum PC2 score for each participant and condition (short, medium, long). The PC2 score ranges were -21 to 15 (short), -41 to 38 (medium), and -40 to 55 (long). To assess the relationship between PC2 score and behaviour, we binned PC2 scores according to our response time bins (early, on time, late) and collapsed across conditions (short, medium, long). This gave us as single mean PC2 score for each participant and response time bin (early, on time, late), which we analyzed using a repeated-measures ANOVA (Fig. 3c).

Data Availability

Raw and preprocessed EEG for the production and perception tasks will be made publicly available at the time of publication. Raw data for the prediction task is available at https://doi.org/10.5061/dryad.5vb8h.

Code Availability

Simulation and analysis scripts are available at https://git.fmrib.ox.ac.uk/chassall/temporal-scaling.

References

Allman, M. J., & Meck, W. H. (2012). Pathophysiological distortions in time perception and timed performance. *Brain*, *135*(3), 656–677. https://doi.org/10.1093/brain/awr210

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

Breska, A., & Deouell, L. Y. (2017a). *Data from: Neural mechanisms of rhythm-based temporal prediction: delta phase-locking reflects temporal predictability but not rhythmic entrainment* (Version 1, p. 19333644881 bytes) [Data set]. Dryad. http://datadryad.org/stash/dataset/doi:10.5061/dryad.5vb8h

Breska, A., & Deouell, L. Y. (2017b). Neural mechanisms of rhythm-based temporal prediction: Delta phase-locking reflects temporal predictability but not rhythmic entrainment. *PLOS Biology*, *15*(2), e2001665. https://doi.org/10.1371/journal.pbio.2001665

Buonomano, D. V., & Maass, W. (2009). State-dependent computations: Spatiotemporal processing in cortical networks. *Nature Reviews. Neuroscience*, *10*(2), 113–125. https://doi.org/10.1038/nrn2558

Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21. https://doi.org/10/bqr2f2

Ehinger, B. V., & Dimigen, O. (2019). Unfold: An integrated toolbox for overlap correction, non-linear modeling, and regression-based EEG analysis. *PeerJ*, *7*, e7838. https://doi.org/10.7717/peerj.7838

Hunt, L. T., Behrens, T. E., Hosokawa, T., Wallis, J. D., & Kennerley, S. W. (2015). Capturing the temporal evolution of choice across prefrontal cortex. *ELife*, *4*, e11945. https://doi.org/10.7554/eLife.11945

Lopez-Calderon, J., & Luck, S. J. (2014). ERPLAB: An open-source toolbox for the analysis of event-related potentials. *Frontiers in Human Neuroscience*, *8*. https://doi.org/10.3389/fnhum.2014.00213

Luck, S. J. (2014). *An introduction to the event-related potential technique* (Second edition). The MIT Press.

Macar, F., & Vidal, F. (2004). Event-Related Potentials as Indices of Time Processing: A Review. *Journal of Psychophysiology*, *18*(2/3), 89–104. https://doi.org/10.1027/0269-8803.18.23.89

Matell, M. S., & Meck, W. H. (2004). Cortico-striatal circuits and interval timing: Coincidence detection of oscillatory processes. *Cognitive Brain Research*, *21*(2), 139–170. https://doi.org/10.1016/j.cogbrainres.2004.06.012

Möcks, J. (1986). The Influence of Latency Jitter in Principal Component Analysis of Event-Related Potentials. *Psychophysiology*, *23*(4), 480–484. https://doi.org/10.1111/j.1469-8986.1986.tb00659.x

Nobre, A. C., & van Ede, F. (2018). Anticipated moments: Temporal structure in attention. *Nature Reviews Neuroscience*, *19*(1), 34–48. https://doi.org/10.1038/nrn.2017.141

O'Connell, R. G., Shadlen, M. N., Wong-Lin, K., & Kelly, S. P. (2018). Bridging Neural and Computational Viewpoints on Perceptual Decision-Making. *Trends in Neurosciences*, *41*(11), 838–852. https://doi.org/10.1016/j.tins.2018.06.005

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442. https://doi.org/10.1163/156856897X00366

Reichel, L., & Ye, Q. (2008). Simple Square Smoothing Regularization Operators. *Electronic Transactions on Numerical Analysis*, *33*, 63–83.

Remington, E. D., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics. *Neuron*, *98*(5), 1005-1019.e5. https://doi.org/10.1016/j.neuron.2018.05.020

Smith, N. J., & Kutas, M. (2015a). Regression-based estimation of ERP waveforms: I. The rERP framework. *Psychophysiology*, *52*(2), 157–168. https://doi.org/10.1111/psyp.12317

Smith, N. J., & Kutas, M. (2015b). Regression-based estimation of ERP waveforms: II. Nonlinear effects, overlap correction, and practical considerations. *Psychophysiology*, *52*(2), 169–181. https://doi.org/10.1111/psyp.12320

Treisman, M. (1963). Temporal discrimination and the indifference interval: Implications for a model of the 'internal clock'. *Psychological Monographs: General and Applied*, *77*(13), 1–31. https://doi.org/10.1037/h0093864

Wang, J., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. *Nature Neuroscience*, *21*(1), 102–110. https://doi.org/10.1038/s41593-017-0028-6

Williams, A. H., Poole, B., Maheswaranathan, N., Dhawale, A. K., Fisher, T., Wilson, C. D., Brann, D. H., Trautmann, E. M., Ryu, S., Shusterman, R., Rinberg, D., Ölveczky, B. P., Shenoy, K. V., & Ganguli, S. (2020). Discovering Precise Temporal Patterns in Large-Scale Neural Recordings through Robust and Interpretable Time Warping. *Neuron*, *105*(2), 246-259.e8. https://doi.org/10.1016/j.neuron.2019.10.020