When Velocity Becomes Position: Proprioceptive Contributions to State Estimation

Erik Skjoldan Mortensen and Mark Schram Christensen

Department of Psychology, University of Copenhagen
Øster Farimagsgade 2A, Copenhagen, Denmark

Abstract

In this study, we investigate the relative contributions of positional- and velocity-based proprioceptive feedback to state estimation when visual feedback of the hand is unavailable. We performed a virtual reality experiment in which healthy human participants (N = 22) completed two tasks: a continuous target-tracking task and a sequential reaching task. In both cases, visual feedback of the hand position was withheld, except for a brief moment, during which an offset to the left or right was sometimes applied. By analysing the movement path following the offset visual feedback and comparing it to the non-offset condition, we can infer the extent to which the brain relies on rate-of-change feedback compared to absolute positional feedback for estimating the hand position. We compare and contrast our empirical data with simulated movements from a computational model based on Bayesian sensory integration under varied assumptions of the reliance on positional and velocity-based cues. Our human and simulated data together indicate a heavy reliance on velocity-based proprioceptive feedback, as movement offsets induced by biased visual feedback decay slowly during both target tracking and sequential reaching tasks. The eventual convergence between offset and non-offset trials demonstrates that absolute positional information is incorporated, albeit at a slow rate.

Author summary

The brain continually estimates the current pose and movement of the body, supporting motor control and giving rise to the perception of the body's state. When visual feedback is lacking, we rely heavily on the sense of proprioception to infer the state of the body. Key among the proprioceptive receptors are the muscle spindle afferents, which are divided into two specialised subgroups: one that signals absolute muscle length, and the other that signals the rate-of-change of muscle length. In this study, we investigate the extent to which the brain utilises rate-of-change information to supplement its inference of the current body position. While such a function would fit well within several of the most popular theories of sensorimotor control, the investigation of this process has so far received limited attention. By having a group of research participants complete two virtual reality-based experiments, we demonstrate that the performed movements remain relative to a visually induced offset. By further simulating the same tasks in a computational model based on Bayesian sensory integration, we can show that this pattern of behaviour is well explained by assuming a relatively high reliance on rate-of-change versus absolute proprioceptive feedback.

Introduction

To effectively move our hand from one position to another requires some estimate of the initial state of the arm. Several sources of information are available to the central nervous system (CNS) to perform this estimation; some are based on afferent signals, key among these being visual and proprioceptive feedback, as well as internally generated signals such as an efferent motor copy or a prior of intended movement that is still being relayed to the muscles. In this study, we focus on the proprioceptive system and examine the relative reliance on position-

and velocity-based feedback for state estimation and how this may lead to drift between the actual state of the hand and its perceived state when visual feedback is lacking.

Proprioceptive sensory organs

Proprioception refers to the sense of body position and movement, which arises from a collection of peripheral mechanoreceptors (Tuthill & Azim, 2018). Central among these are the type Ia and type II muscle spindle afferents, which are sensitive to stretch of the parent muscle (Proske & Gandevia, 2012). Muscle/tendon vibration is one of the classical experimental techniques used to investigate how these sensors contribute to the perceived pose and movement of the body. Vibration of a muscle/tendon increases the firing rate of both of these types of receptors, leading to biasing effects on the perceived position and movement of the vibrated limb, corresponding to stretching of the vibrated muscle (Eklund, 1972; Goodwin et al., 1972; Roll & Vedel, 1982). Additionally, such vibration techniques have been used to replay microneurographically recorded signals from muscle spindle afferents, allowing for the re-induction of afferent firing patterns, which allowed the research participants to recognise complex movement patterns (Albert et al., 2005, 2006). Such studies clearly demonstrate that these sensors are at least sufficient for the perception of complex and specific changes in limb posture and velocity, and do not merely elicit vague sensations. The type Ia afferents are generally described as signalling muscle stretch velocity, while the type II afferents signal absolute muscle length (Proske & Gandevia, 2012), although they likely overlap substantially in their function. The activity of both sensor types depends on both muscle length and velocity to some extent, and may further shift towards signalling acceleration in some instances (Macefield & Knellwolf, 2018).

The fact that their functions overlap and the tendency of both types of afferents to be biased to some extent by muscle vibration has made it difficult to disentangle the relative contributions and reliance of absolute position and rate-of-change sensors to the final integrated perception of limb state. It is, however, generally assumed that type Ia sensors are particularly important for faster movements, while type II afferents are central for estimation of maintained posture and more gradual changes (Banks et al., 2021).

It should further be noted that our brief introduction to proprioceptive receptors here is not exhaustive, and that Golgi tendon organ afferents, skin stretch receptors, as well as joint receptors all contribute meaningfully to the collective proprioceptive sense. We have here focused on muscle spindle afferents, in order to limit the scope of our introduction, while exemplifying how position- and velocity-based sensory feedback enters the sensorimotor system.

Sensory integration and state estimation

Proponents of modern theories of sensorimotor control, such as Optimal Control Theory (and derived frameworks) and Active Inference, generally argue that sensory feedback is assumed by the CNS to arise from the same underlying hidden cause, which the CNS is attempting to infer the state of (Wolpert et al., 1995; Todorov & Jordan, 2002; K. J. Friston et al., 2010). I.e., when position- and velocity-based proprioceptive feedback as well as visual feedback is received by the CNS, it is integrated in a self-consistent manner, such that these signals ideally augment and filter each other, rather than give rise to multiple distinct perceptions of the same arm, whenever mismatching signals are received.

The presence of sensory and motor noise as well as transmission delays is fundamental to the nervous system (Van Beers et al., 2002; Faisal et al., 2008), underscoring the need for effective mechanisms to integrate signals across sensor types, sensory modalities, and time. Bayesian sensory integration has been proposed to meet this requirement (Knill & Pouget, 2004), providing a principled way to combine multiple sources of uncertain and noisy information, weighting each by its assumed state-sensor precision. Several studies in which research subjects have been required to combine different perceptual cues, and in which the reliability of some cues has been manipulated, have provided strong support for such Bayes-optimal sensory integration (Körding & Wolpert, 2004; Chancel et al., 2016; Limanowski & Friston, 2020; Chancel & Ehrsson, 2023). The perceived state may then further be integrated over time through learned laws of motion or generative models; e.g., the arm may be

expected to move smoothly through space, with congruent changes in position, velocity, and acceleration, rather than disappear in one spot and reappear at another (K. Friston, 2008; K. J. Friston et al., 2010).

With such a formulation, it becomes clear how velocity-based proprioceptive feedback can contribute to the perceived position of a limb, even if it does not communicate it directly; if the start position of the hand is well represented, e.g. due to visual feedback, then access to high-precision velocity information would allow for high-precision inference of any change in position, relative to the start position, over the course of a movement.

Reaching tasks are a central tool for investigating this issue. When indicating the position of a visual target without visual feedback of the hand, a research subject implicitly indicates that the hand is perceived to match the target's position. Any endpoint error, therefore, indicates an offset between the perceived and actual hand position, given that the movement is executed with a focus on precision rather than speed. Several studies have indicated that a currently perceived position may well be inferred relative to a previously perceived position; for example, it is seen that endpoint errors are reduced if vision of the hand is provided prior to movement onset, compared to not seeing the hand at the start position (Elliott & Calvert, 1990; Rossetti et al., 1994; Vindras et al., 1998). Additionally, repetitive movements remain of consistent length and direction across many repeated sequential reaches, even when absolute position errors have accumulated substantially (Brown et al., 2003a, 2003b; Patterson et al., 2017). Similarly, induced position offsets from briefly presented biased visual feedback lead to biased endpoint errors in the opposite direction, indicating that the perceived hand position is inferred as a change from the presented visual input, even when such feedback becomes out of date (Körding & Wolpert, 2004; Hewitson et al., 2018).

Bayesian integration

The above-mentioned findings indicate that the perceived state is history-dependent; motor behaviour appears to be based on an internal state that is updated over time, taking previous state estimates into account, rather than being constructed solely from the most recent sensory feedback. In the context of Bayesian sensory integration, this can be expressed as a recursive loop where the current state estimate is used to construct a prior for the next step through forward-prediction, which is then updated by incoming sensory feedback. Such models of sensory integration have a long history of use in computational modelling of state estimation (Wolpert & Ghahramani, 2000). We will here give a brief introduction to the basic steps used during such state estimation. In this study, we employ a computational model based on this type of recursive Bayesian filtering to compare the model's simulated behaviour with the observed behaviour of human research participants. The model is set up such that we can vary the relative reliance on position- and velocity-based proprioceptive feedback.

The full implementation of the simulation model is described in our previous paper (Mortensen et al., 2025). In brief, the simulated Agent attempts to infer the state of a two-jointed planar arm and move it towards a specified target, set up to match the tasks completed by our research participants. For the Agent to infer the current state of the arm, we implemented an Unscented Kalman Filter (Julier & Uhlmann, 2004), an approximation of Bayes-optimal recursive filtering in the context of the non-linear transition and observation models necessary for a multi-jointed arm model. At the end of each simulated time step, a Linear Quadratic Regulator was used to generate efferent motor commands, which are applied to the arm on the next step; the motor commands are subjected to signal-dependent motor noise (Harris & Wolpert, 1998). This simulation setup is conceptually similar to the model presented by Todorov and Jordan (Todorov & Jordan, 2002).

In the following, we will briefly outline the functioning of a Kalman Filter (Kalman, 1960), which describes Bayesian filtering under simple linear dynamics with Gaussian noise, in order to provide an intuition for the implied effects of varying the reliance on position- or velocity-based proprioceptive feedback.

Consider a simplified example of an Agent moving along a one-dimensional line, receiving noisy feedback about both its position along the line and the velocity at which it is moving:

$$\mathbf{x}_k = \begin{bmatrix} x_k \\ \dot{x}_k \end{bmatrix}, \quad \mathbf{z}_k = \mathbf{x}_k^{true} + \mathcal{N}(0, \mathbf{R}),$$

$\mathbf{x}_k$ here represents the internal state estimate of the Agent at time step k, consisting of its position (x_k) and velocity ($\dot{x}_k$), as it is attempting to infer its true state $\mathbf{x}_k^{true}$, through noisy sensory observations, $\mathbf{z}_k$. Bayes' rule describes the general framework for how the Agent may attempt to infer and update its position and velocity state over time as more sensory information is received, when expressed as a recursive filter:

$$\underbrace{p(\mathbf{x}_k \mid \mathbf{z}_{1:k})}_{posterior} \propto \underbrace{p(\mathbf{z}_k \mid \mathbf{x}_k)}_{likelihood} \cdot \underbrace{p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1})}_{prior}$$

It tells us that the current posterior probability of a hypothesised state ($\mathbf{x}_k$) given all the historical noisy sensory observations ($\mathbf{z}_{1:k}$), including the current set obtained at step k, is proportional to the probability of the current observations ($\mathbf{z}_k$) given the hypothesised state times the prior probability of the hypothesised state. The prior represents a forward prediction of the current state, given all the historical sensory data, up to, but not including, the current observations; it is the previous state estimate projected forward in time, predicting the current state. As the prior depends on the previous state, we can see how this process iteratively updates and refines its state estimate, while still considering previously received sensory observations. The Kalman Filter thus works in a continuous two-step process of prediction followed by correction when new sensory data is received.

With the Agent possessing an internal model of its expected dynamics (here semi-explicit Euler integration) and a learned state-sensor precision mapping, it now has everything required to optimally update its internal state estimate with new sensory observations, while keeping track of how certain it is in its state estimate. Assuming Gaussian noise, to construct the prior for the next hypothesised state, given the current posterior state:

$$\dot{x}_k^{prior} = \dot{x}_{k-1}^{post} + u_{k-1}\Delta t \tag{1}$$

$$\sigma_{\dot{x},k}^{2\,prior} = \sigma_{\dot{x},k-1}^{2\,post} + Q_{\dot{x}} \tag{2}$$

$$x_k^{prior} = x_{k-1}^{post} + \dot{x}_k^{prior}\Delta t \tag{3}$$

$$\sigma_{x,k}^{2\,prior} = \sigma_{x,k-1}^{2\,post} + \sigma_{\dot{x},k}^{2\,prior}\,(\Delta t)^2 \tag{4}$$

Here u_{k-1} represents a control signal expected to affect the velocity since the last update; this could be thought of as an efferent motor copy generated by the Agent, which will accelerate it along the line, and change its velocity with some uncertainty, denoted $Q_{\dot{x}}$. In the context of signal-dependent motor noise, this quantity would scale with u.

The prior of velocity is now the previous posterior of velocity plus the expected change due to this control signal, with the prior of position being the sum of the previous posterior position and the expected change in position due to its estimated velocity.

σ^2 describes the variance of the Gaussian uncertainty of the various values. For example, the prior uncertainty of the Agent's position along the line at time step k, before seeing new data, corresponds to the sum of the previous posterior position uncertainty ($\sigma_{x,k-1}^{2\,post}$) and the uncertainty in velocity scaled by time squared ($\sigma_{\dot{x},k}^{2\,prior}$).

While the above describes how the prior of the state is estimated, it does not yet consider any expected noise associated with the observed sensory feedback.

$$\mathbf{z}_k = \mathbf{x}_k^{true} + \epsilon_k \tag{5}$$

$$\epsilon_k \sim \mathcal{N}(0, \mathbf{R}) \tag{6}$$

$$p(\mathbf{z}_k \mid \mathbf{x}_k) = \mathcal{N}(\mathbf{z}_k;\ \mathbf{x}_k,\ \mathbf{R}) \tag{7}$$

Here, $\mathbf{z}_k$ represents noisy sensory observations of $\mathbf{x}_k^{true}$ with Gaussian noise added (with $\mathbf{R}$ variance), which leads to the definition of the likelihood function. Of note, there is an inverse relationship between the two steps of the filter regarding uncertainty. During the forward-prediction (prior construction) step, variances sum, meaning uncertainty increases, compared the previous posterior estimate. However, when new evidence is integrated, the precisions of the likelihood and prior sum. Since precision is the reciprocal of variance ($J = \frac{1}{\sigma^2}$), this means the

uncertainty of the posterior will always be lower than either prior or the measurement alone.

Given this description of Bayesian filtering, we can consider the implied effects of varying the relative precision of positional and velocity feedback:

High-precision velocity, low-precision position (HVLP):
The Agent here estimates its velocity with high precision, which allows for a highly precise prior of position when predicted forward to the next time step, which the low-precision positional feedback will not drastically update. Moment-to-moment changes in position is smooth and will be well-estimated, but any small errors in velocity will accumulate, and any error in the inferred position will tend to be propagated forward in time.

Low-precision velocity, high-precision position (LVHP):
The Agent will struggle to estimate velocity, and therefore have wider (uncertain) priors for position; the comparatively precise position sensor will have a much greater effect on the posterior positional estimate. Moment-to-moment changes in position will be less well-represented, but substantial drift from the signalled position will be avoided.

With this study, we aimed to investigate the relative reliance on proprioceptive absolute positional sensors compared to velocity sensors through the lens of proprioceptive drift and Bayesian sensory integration. We designed a virtual reality (VR) based target tracking task, where the target follows a predictable and continuous circular path; this setup allows us to obtain a continuous readout of the perceived position of the hand, which we can then compare to the actual tracked position.

By providing brief, offset visual feedback of the hand position during the beginning of a trial, we can induce a positional offset in the perceived hand position, similar to the approach used by Kording and Wolpert (Körding & Wolpert, 2004). This allows us to compare how the continued movement is made following the removal of visual feedback between the offset and non-offset conditions, which enables us to infer to what extent the perceived hand position is updated based on relative versus absolute positional information. We can consider two extremes as examples of the type of results we are looking for: If the position of the hand is inferred mainly from velocity-based signals, with minimal absolute positional proprioceptive signals, then we would expect any visually induced offset to remain constant when averaged across trials. On the other hand, if the hand position is inferred mainly from absolute positional signals, we should expect fairly rapid convergence of the paths followed during offset and non-offset trials once visual feedback is removed.

We follow this with a conceptually similar sequential reaching task, where visual feedback is provided only briefly during the first outbound reach, as this type of task is more frequently employed for such long-duration movement tasks (Brown et al., 2003a, 2003b; Smeets et al., 2006; Patterson et al., 2017); we here aimed to confirm that any findings from the target-tracking task were similarly applicable.

We simulated the same task in a computational model based on Bayesian sensory integration. Specifically, we compare two different proprioceptive configurations, one that mainly relies on velocity-based feedback, while the other relies primarily on position-based feedback.

Overall, our results indicate that the perceived position of the hand is highly dependent on rate-of-change signals, as visual offsets continue to affect it long after they are removed.

Results

In order to investigate the relative reliance of position- and velocity-based proprioceptive feedback for inference of arm state and hand position, we designed a virtual reality (VR) based experimental setup, as seen in Fig 1. The research participants (N = 22) were seated at a table and outfitted with a VR headset, with the tabletop tracked into the VR environment to provide a stable visual reference of the working environment. They would grasp the right-hand VR controller, which was placed in a 3D-printed mount with a flat base. A felt slider attached to the base allowed the participants to slide the controller along the tabletop with minimal friction, effectively allowing for movement of the controller along the 2D plane of the tabletop. The height of the table

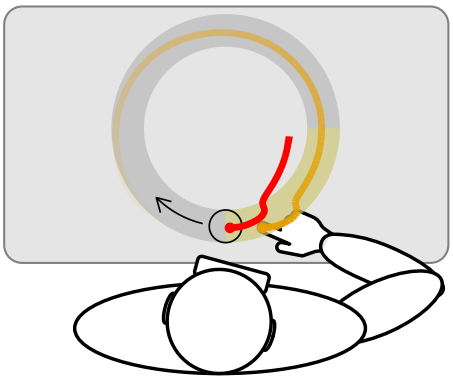

Figure 1: **Visualisation of the target tracking task.** The research participant is seated at a table and is outfitted with a VR headset. Their task is to track a moving target (small black circle) that follows a clockwise circular track (large grey circle). The target is followed for 4.5 revolutions, with each revolution lasting 4 seconds (total of 18 seconds). During the last quarter of the first revolution (shown as a yellow highlight), 1 second of visual feedback is provided, which might be offset 5 cm to either side (red trace, shown here is offset to the left, prompting a corrective movement to the right). The orange path shows the real movement of the tracked VR controller, which is not visible in the VR environment.

was adjusted so that when the grasped controller was resting on the table with the elbow bent 90°, the elbow would be approximately 5 cm above the table, preventing the arm from sliding across the table and providing additional sensory feedback. Within the virtual environment, a semi-transparent target sphere (radius = 2.5 cm) was displayed above the table at a height matching that of the top of the held controller. During each trial, the target would follow a clockwise circular path (15 cm radius) over the tabletop, with each revolution lasting 4 seconds. The hand starting point, and the near edge of the target path, was approximately 10 cm in front of the centre of the chest of the seated participant, with the full path of the target within easy reach of the right arm holding the controller. Each trial lasted 20 seconds from when the participant indicated readiness; visual feedback of the controller position (a smaller sphere of 1.25 cm radius, centred on the thumb button) would be removed, and the target would start moving from the far edge of the circle, intersecting the hand start position after 2 seconds, at which point the participant was instructed the follow the moving target as precisely as possible for the next 4.5 revolutions. A total of 40 trials per participant were completed. In three-quarters of the trials, visual feedback would be reinstated for 1 second in the 3-4 second interval after target tracking onset, corresponding to the last quarter of the first revolution (see yellow highlight in Fig 1). In trials with visual feedback, it was either correctly mapped or offset 5 cm to either side; that is, 10 trials each of: no visual feedback, correctly mapped visual feedback or offset 5 cm to either side. After each trial ended, visual feedback of the controller position would only be reinstated once the controller was returned to within 5 cm of the hand start position, to minimise any endpoint feedback provided to the participant.

Target tracking task

Fig 2 plots the averaged movement path across all 22 participants, for all three conditions with visual feedback. Fig 2A, subpanel '3-4 s' shows the initial response when (potentially offset) visual feedback is displayed, demonstrating the expected corrective movement of close to 5 cm in the opposite direction, such that the visual feedback is aligned to the target position. The following 4 subpanels show that the performed tracking movement tends to remain offset by the now-removed visual feedback for the remainder of the trial, although the offset and non-offset conditions do approach convergence of their respective paths towards the end of the tracked movement. This is similarly highlighted in Fig 2B, which plots the slow decay of the induced positional offset along the x-axis.

The long duration of the induced offset in the perceived position of the hand, following visual feedback, indicates that the perceived hand position is highly dependent on a signal regarding how the arm position is changing, rather than signals which directly relate to its absolute position; if the latter was the case, we should expect a fast convergence of the offset and non-offset paths following the removal of visual feedback. The fact that

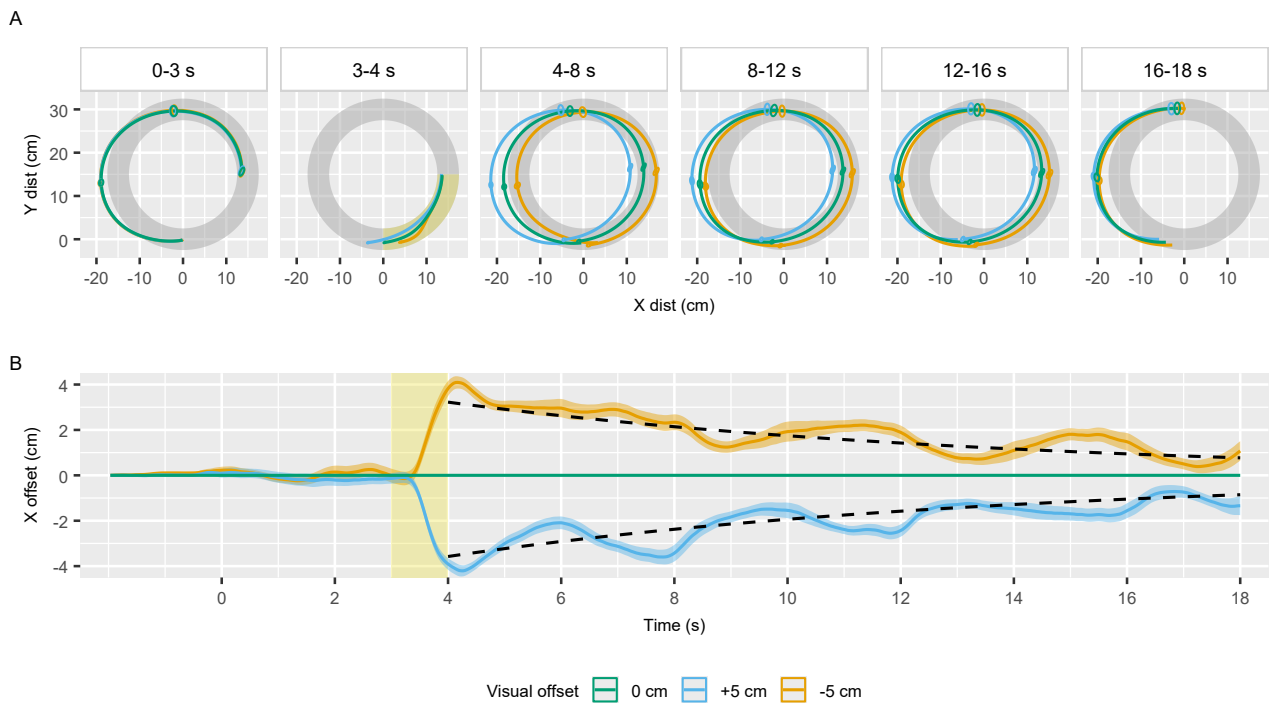

Figure 2: **Circular tracking task: Mean movement paths across participants.** Panel A (top row of subpanels) plots the averaged movement paths for trials with non-offset visual feedback (green), left offset (orange), and right offset (blue), across all participants. Movement direction was always clockwise. Ellipses showing the 1 SEM contour of subject means are plotted once per second (4 per revolution). Visual feedback was only provided in the 3-4 second interval (yellow highlight). Panel B plots the averaged induced offset along the X-axis, calculated by subtracting the mean X position of the control 0 cm offset condition from the other two conditions, done within-subject to account for any consistent accuracy errors. The ribbon shows the +/- 1 SEM of the induced offset. Dashed black lines indicate the fit of the group-mean exponential decay.

the three tested conditions eventually do converge back together, however, highlights that absolute positional information is available for sensory integration. The slowness of this convergence process can be interpreted as the available positional feedback being assigned a relatively low weight for sensory integration.

To quantify the time dynamics of this convergence process, as well as the between-subject variability, we fit a Bayesian hierarchical exponential decay model to the trial-level data shown in Fig 2B. The model had the general form (see Methods for further details of model fitting):

$$x_{offset}(t) \sim \mathcal{N}(exp(\alpha_k - p * t), \sigma)$$

α_k here represents the initial offset when visual feedback was removed, with a different parameter fitted depending on whether the visual offset was in the positive or negative x direction (k). The model revealed a group-mean decay exponent (p) of 0.10 (95% credible interval (CI) [0.08, 0.12]), corresponding to a half-life of the induced offset of 6.9 seconds (95% CI [5.8, 8.7]). The between-subject variability in the exponent was estimated as an SD of 0.05 (95% CI [0.04, 0.07]), ranging from individual p estimates of 0.02 and 0.21 (corresponding to estimated half lives between 3.3 and 40.6 seconds). The fit of the group-mean exponential decay function is plotted in 2.

It should be noted that while the group average movement paths give the impression that the induced offset is purely along the x-axis, this is to some extent a result of averaging across somewhat varied individual responses. The above-considered outcome of relative offset along the x-axis is therefore a simplification. We will expand on this point in the two following sections. Similarly, if it is assumed that the gradual reduction of the induced offset over time can be explained by Bayesian filtering processes, then the speed of the washout/decay would be expected to depend both on both the precision of the proprioceptive feedback as well as uncertainty in currently

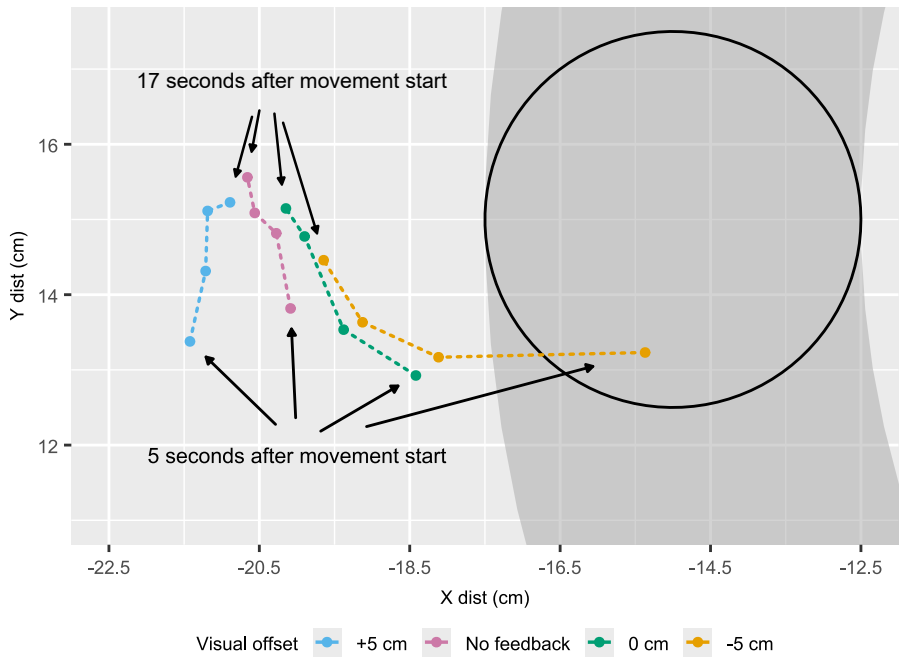

Figure 3: **Circular tracking task: Perceptual drift over time.** The figure shows the tendency toward positional drift during continuous target tracking, as the target repeatedly passes the same position over the course of each trial. Each trace represents how the group mean hand position changes over the course of the target tracking task for each of the four visual feedback conditions. All points represent the hand position when the target passes the nine o'clock position on each subsequent revolution of the target; the first point along each trace represents 5 seconds into the movement (1 second after visual feedback is removed), the next points 9, 13 and 17 seconds respectively. The outline of the target and its path is plotted for scale.

estimated state; it would be easy to imagine that the latter of these two changes, as time since last visual feedback increases. Immediately after visual feedback, the estimated state might still have high certainty, with uncertainty potentially increasing as movement continues over the next 14 seconds.

Perceptual drift

While we highlighted above how the conditions with visual offsets slowly decay towards the non-offset condition along the x-axis, this represents a simplified description of the overall pattern of drift that is apparent in our data. In Fig 3, we plot out how the position of the hand changes over the course of the target tracking task as the target passes the same position. To exemplify this, we have zoomed in on the left edge of the path followed by the target, and here plot the mean position of the hand when the target passes the nine o'clock position on each repeated revolution; the lower end of each trace represents five seconds after movement onset, the next one nine seconds after, the next 13 and the final point (arrow) 17 seconds. For the three conditions with visual feedback, the first point corresponds to 1 second after visual feedback was removed; additionally, we also plot the condition with no visual feedback.

This highlights how the proprioceptively perceived hand position appears to drift over the course of the tracking movement, with all four conditions appearing to converge towards a single point. Additionally, it may be highlighted how the second point along the green trace (correctly mapped feedback) is rather close to the first point along the purple trace (no feedback); these both represent one and a quarter revolutions of the target since visual feedback was last provided for the given condition.

Example individuals

To highlight the variability in individual responses, we have added Fig 4, which displays the movement paths of two example participants. With these two examples, we have chosen two subjects who display rather different patterns of movement and responses following visual offset. Example A in the top row of Fig 4 shows a very

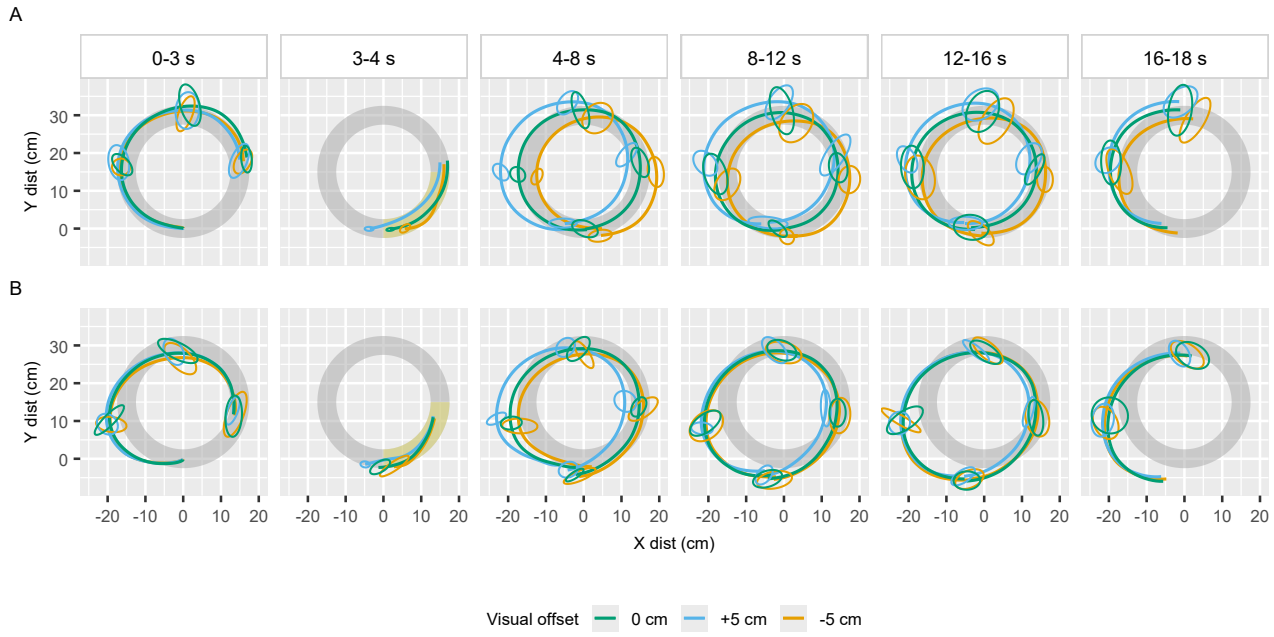

Figure 4: **Circular tracking task: Mean movement paths of two example participants.** Averaged movement paths of two example participants. Ellipses showing the 1-SD contour are plotted once per second (4 per revolution).

strong tendency to move relative to the visually induced offset, with clear differences between the three visual conditions remaining until the end of the trial. Additionally, the response of this participant highlights that induced offset is not necessarily limited to the Cartesian x-axis; in particular, it appears that the induced offset tends to gradually shift from along the x-axis to rather be along the y-axis by the end of the trial. One possible interpretation here is that the induced offset is in a proprioceptive/motor reference frame, rather than a visual one, which might result in movement paths such as these; we will return to this point in the next section regarding our simulation results.

Example B, on the other hand, appears to be minimally affected by the visual offset already around 4 seconds after visual feedback is removed. Additionally, this example also illustrates the tendency of some participants to make consistent errors in the shape of the circular movement path, potentially indicating distortions/offsets between their proprioceptive and visual reference frames.

Similarly, other participants tended to follow either a too-small or a too-large circle, or one that is distorted in some way (see Fig S1-3, where the remaining 20 subjects' data are plotted). These within-subject consistent errors tend to be averaged out in the group-average plotted in Fig 2.

Target tracking simulation

To further explore how the integration of rate-of-change and absolute positional signals may lead to the above-plotted movement paths, we simulated the same task in a computational model based on Bayesian sensory integration. The simulated Agent was set up to control a planar 2-link arm and to infer its state based on noisy angular position and velocity sensors for its shoulder and elbow joints, along with Cartesian 2D visual feedback. We tested two configurations of proprioceptive feedback for the current target-tracking task: one with high-precision velocity feedback and low-precision position feedback (HVLP) and one with low-precision velocity feedback and high-precision position feedback (LVHP). The results of the simulations are shown in 5.

The HVLP configuration was set up with 16 times higher precision of sampled joint angular velocity feedback compared to angular position feedback, with the precision values swapped for the LVHP configuration, such that they were here 16 times for the position sensor. This leads to markedly different patterns of behaviour, with the HVLP configuration being comparatively more reliant on forward-prediction of position estimates, compared

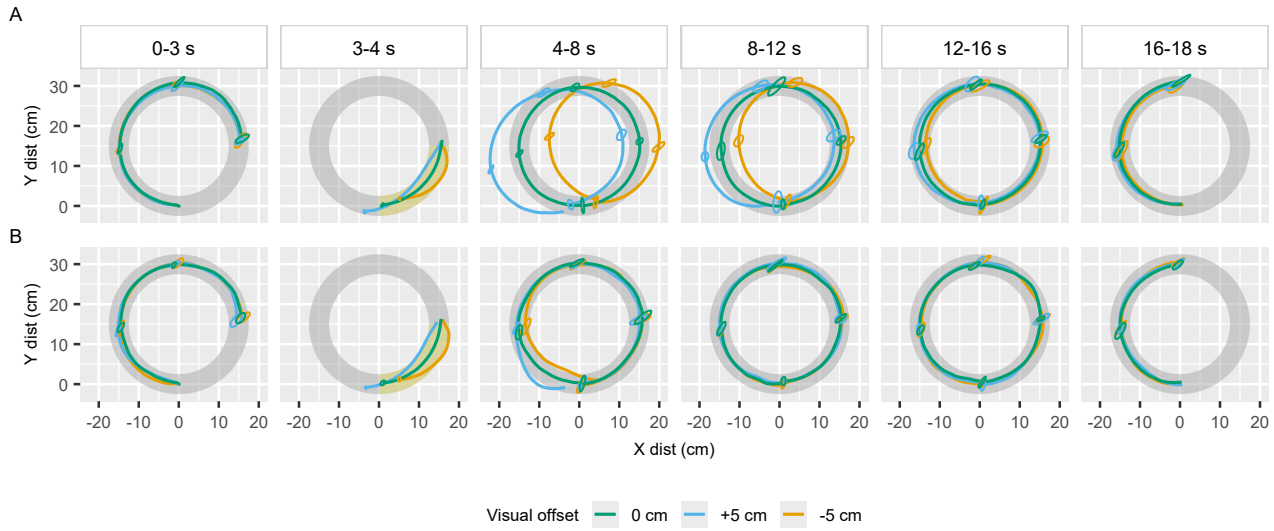

Figure 5: **Circular tracking task: Simulated movement paths.** A shows the averaged movement path of the high-precision velocity and low-precision position proprioceptive configuration (HVLP), and B shows the low-precision velocity and high-precision position configuration (LVHP). Ellipses plotted once per second (4 times per revolution) show the 1 standard deviation contour at that point.

to the LVHP configuration. Both configurations had identical and comparatively very high precision of visual feedback. This meant that, when visual feedback was available, it would quickly dominate the posterior of the perceived arm state. See the Methods section for further details.

Comparing the simulated data of the Bayesian Agent to the observed human data reveals a close similarity between the pattern of behaviour of HVLP configuration and our human participants, while the LVHP configuration shows markedly different movement following offset visual feedback. This modelling approach demonstrates how a high reliance on sampled angular velocity, with relatively low precision of angular position measurements, can lead to an offset in the perceived state which persists for a long duration after incorrect visual feedback. Additionally, the HVLP configuration reveals a somewhat similar tendency for the visual feedback to produce an offset that is not confined to the x-axis, as seen for example participant A shown in Fig 4. In the case of the simulated Agent, we can explicitly describe why this occurs. Because its internal state estimate is tracked in proprioceptive coordinates (angular positions and velocities of its joints), it also means that the visually induced state offset is represented in the proprioceptive reference frame. This leads to non-constant offsets along the x- and y-axis in Cartesian space for the different arm poses required to track the target around the full circle.

The LVHP configuration, in contrast, quickly corrects its trajectory following the removal of visual feedback, due to its reliance on high-precision positional feedback. From the point of view of both simulated Agents, they both perceive themselves to faithfully be following the target throughout each trial; corrections made in response to the offset visual feedback are mostly completed within approximately 500 ms after presentation.

Sequential reaching task

Investigations of proprioceptive drift are often conducted through testing of a sequential target-directed reaching, rather than continuous tracking tasks, such as the one presented above. In order to link our observations to this literature, we also completed a second task consisting of three sequential reaches (out, back and out again) with the same group of participants. This was done to validate that the observations made in the target tracking task were generalisable to sequential reaching tasks.

For this task, the participants were required to sequentially move between and indicate the position of two different targets, one placed directly in front of them (same spot as the start position for the circle tracking task), and the other positioned 38 cm in front of this target. They were instructed to indicate the position of each target as precisely as possible by aligning the tracked thumb button to the centre of each target, and then pressing the

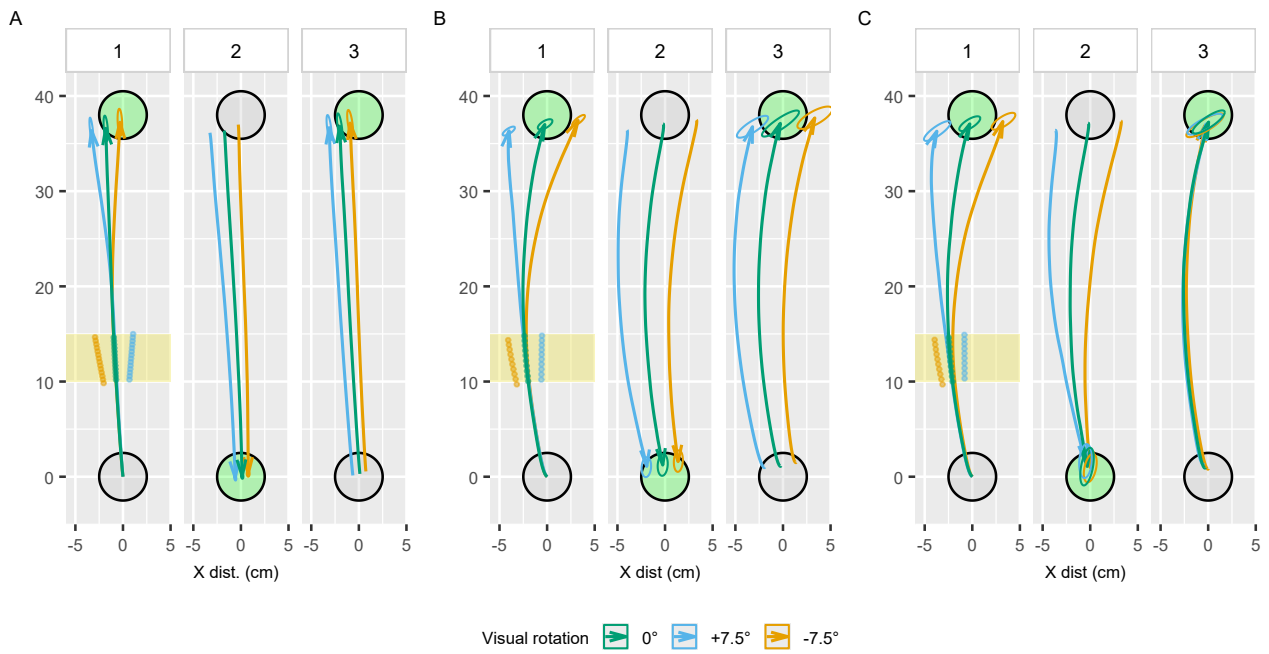

Figure 6: **Sequential reaching task: Human and simulated data.** The averaged movement path across three sequential reaching movements towards targets placed 38 cm apart. Visual feedback is withheld, except for 100 ms after passing the 10 cm mark during the first reach, illustrated with the yellow highlight. Visual feedback was rotated either 0 °, +7.5 °, or -7.5 ° around the starting position; the offset visual feedback is shown in the yellow highlight. A: Averaged movement path from the group of participants. B: Average of the simulated movement paths for the HVLP configuration. C: Average of the simulated movement paths for the LVHP configuration. Note that ellipses in A show the 1-SEM contour, while ellipses in B and C show the 1-SD contour instead; A is a mean of individual means across trials, while B and C are simply a mean of repeated trials (these may be considered more akin to plotting out the means of a single participant).

button. The current target was coloured green, while the previous target was grey; pressing the thumb button would cause them to switch colours. They were free to move at any speed they chose.

Visual feedback was withheld, except for a 100 ms interval following passing the 10 cm mark on the first outwards reach; it might either be correctly mapped, or offset to either side corresponding to a 7.5 ° rotation around the starting point (approx. offset of 1.3 cm at onset). A total of 60 trials per participant was completed, with a one-third split between correct visual feedback and offset to either side. The results of this task are plotted in Fig 6A, along the the simulated results of the same task, for the HVLP configuration (Fig 6B) and the LVHP configuration (Fig 6C). The task was simulated with one second per reach, corresponding approximately to the group mean time per reach, which was at 0.99 seconds (range of individual means was between 0.48 and 1.65).

The movement paths of the human participants plotted in A show much the same pattern as those seen in the circular tracking task in Fig 2, namely that the visually induced offset continues to affect the indicated target position for the subsequent reaches.

While the HVLP and the LVHP configuration produce almost identical movement paths during the first outbound reach, both showing a similar style of correction as the group of human participants, they are substantially different over the next two movements. Where the HVLP configuration keeps the induced offset across all three movements, similar to the human participants, the LVHP configuration is no longer affected by the visual feedback by the time the second target is reached.

Discussion

When moving the hand to a visual target, the executed motor commands represent the perceived difference between the hand's position and the target's position. By tasking research participants to move to/track the

perceived position of the target, we thus gain insight into how they are mapping between proprioceptive and visual reference frames, and how they are updating their perceived arm and hand position when lacking visual feedback. We approached this by analysing the hand's movement path following a visually induced offset in the perceived hand position, through the lens of Bayesian sensory integration. Access to high-precision proprioceptive positional feedback would, in this framework, predict a relatively fast decay of a visually induced state-estimate offset, after visual feedback is no longer available. If high-precision positional feedback is not available, then position may still be estimated by relying on rate-of-change-based signals, but such a method of updating the perceived position is necessarily more history-dependent, in that it will heavily depend on a previously perceived position.

Visually induced bias

The results of our target tracking task show that the perceived hand position remains influenced by visual feedback long after this source of information is no longer available. This indicates considerable reliance on signals encoding changes in arm posture, as opposed to feedback providing direct information about the current arm position. However, we do find that absolute positional feedback clearly does contribute to the perceived state, albeit at a fairly slow rate. This is seen as the offset of the perceived position does tend to decay away over the course of the continued movement; at the group level, the offset movement paths approach full convergence with the non-offset path at the end of the trial, 14 seconds after visual feedback is removed; specifically, we estimated the group mean half-life at 6.9 seconds, when we fit the induced x-offset to an exponential decay function. At the individual level, we found substantial variation in decay rates of the induced offsets, ranging from half-lives of 3.3 to 40.6 seconds; these may reflect individual variation in the relative reliance of position- and velocity-based feedback, as well as overall levels of sensory and motor noise.

It should be noted that the exponential decay function used here is a simplified approximation of the underlying assumed Bayesian filtering process. In Bayesian filtering, the rate at which the estimated positional state converges towards the mean of a noisy positional feedback channel depends on both the relative precision of the positional and velocity channels - here taken to reflect the proprioceptive feedback from type II and Ia afferents, respectively - and the precision of the current state estimate, which will evolve over the course of the movement. Immediately following the removal of visual feedback, the perceived hand position is likely to be represented more precisely compared to towards the end of the movement.

Additionally, it might be considered more likely that the induced offset is in a body-oriented/proprioceptive reference frame, rather than in a simple visual Euclidean one; a hint towards this in our data may both be seen in Fig 2B, where the induced offset along the x-axis appears to be somewhat cyclical, depending on the target (and hand) position. Similarly, in Fig 4A, we plot out an example participant, for whom the induced offset along the x-axis appears to gradually rotate in Cartesian space to eventually be more along the y-axis. Fig 5A displays a somewhat similar pattern in the simulated data in the first few revolutions after visual feedback; because the induced offset is in proprioceptive space, it does not result in a clean lateral translation of movement paths in Cartesian space. While the direction and skew of the Cartesian offsets in movement paths are different in the simulated example, such differences can likely, at least partially, be ascribed to differences in arm position (elbow-out versus elbow-down) as well as the differing number of degrees of freedom involved in the movement.

Furthermore, we see a similar pattern when the movement is broken into several distinct sequential reaches, as seen in the results of the sequential reaching task, in terms of both the persistence and slow decay of the induced offset.

These results are closely related to the classical study by Körding and Wolpert (Körding & Wolpert, 2004), who demonstrated that briefly presented offset visual feedback during a reaching movement is incorporated into the perceived hand position in a manner consistent with a Bayesian process. These effects were observed through the induction of biased endpoint errors, which indicated that the second part of the movement, following the removal of visual feedback, was executed relative to the provided visual feedback. Through the manipulation of the precision of visual feedback, they demonstrated that the movement was not merely completed relative to

the exact provided position of the visual feedback, but rather relative to the integrated perceived hand position; when the visual feedback was blurred, this would be some somewhere between the expected position of the hand prior to receiving visual feedback and the presented offset position, dependent on the level of blur. These results are similarly in agreement with a range of other studies, where it has been demonstrated how the availability of visual feedback of the hand position prior to movement onset increases endpoint precision (Elliott & Madalena, 1987; Rossetti et al., 1994; Vindras et al., 1998). Such findings indicate that the perceived position at the end of a movement depends on how well the position of hand is represented prior to onset of movement, which similarly points at a significant dependence on signals representing how position changed over the course of the movement.

Perceptual drift

While the findings of Körding and Wolpert and the related studies referenced above are specific to single and short-duration movements, our study builds on these by demonstrating effects during longer-duration continuous and sequential movements, providing a link to the existing literature on proprioceptive drift. This is a phenomenon that has been explored through tasks involving many sequential reaching movements between two or more targets, without visual feedback of the hand, as demonstrated by Brown et al. (Brown et al., 2003a, 2003b) and Patterson et al. (Patterson et al., 2017). In these studies, it is shown that the absolute perceived hand position may drift substantially across multiple sequential movements, while individual movements remain well-formed, with highly consistent direction and movement amplitude. In the current study we build upon on these findings, by showing a very similar tendency for a continuous tracking task; by looking at the averaged hand position when the target passes the same position over the course of the target tracking task, we see a tendency at the group level for a slight forwards and to the left drift over the course of the trial. Additionally, we here illustrate how the four different conditions, three with either correct or offset visual feedback, and one without, appear to converge towards a single point, as the influence of recent visual feedback (whether correct or not) gradually decays away.

Their results further show how this baseline offset is slowly approached over multiple reaches; the first few movements made after visual feedback is removed show only a fraction of the error, which is approached after 50 movements made without visual feedback. It seems likely that the apparent drift convergence point which our data points at in Fig 3 results from such a base offset.

Smeets et al. (Smeets et al., 2006) further demonstrated similar drifting behaviour over sequential reaches with randomised direction, while highlighting that individual drift direction and magnitude remained constant both within and across different test days. This finding strongly indicates that the observed drift is not merely due to the effect of random errors arising from general noise in the sensorimotor system, but may rather reflect a baseline and stable offset present in the mapping between visual and proprioceptive positional reference frames; as multiple movements are made without visual feedback, this offset is slowly incorporated into the perceived state of the arm in process well-approximated by a Bayes-optimal model, while highlighting that the observed drift in accuracy does not require drift of the state-sensor mapping (Smeets et al., 2006).

Together, these studies of long-duration movement without visual feedback point to the fundamental challenge posed by attempting to align the different reference frames provided by estimating position (of both the body and external targets) from the activation of photoreceptors in the retina to the proprioceptive signals received from mechanoreceptors situated in the skin, joints and muscles.

Inferring changes in position

The findings of the above-referenced studies, as well as the results we have here presented, may be explained by assuming that the CNS relies on rate-of-change signals to infer how the position of the hand is changing over the course of one or several movements, in addition to any proprioceptive signals regarding absolute position.

We investigated this as a possible explanation by simulating both the target tracking task and the sequential reaching task in a simplified model, based on Bayesian sensory integration. We here simulated each task with two different configurations of proprioceptive feedback, one which relies mainly on velocity-based feedback with only poor positional feedback, and one that, in contrast, relies mainly directly on positional feedback with poor

velocity-based feedback. We have previously demonstrated that the model can replicate human performance in response to both visual offsets and muscle/tendon vibration (Mortensen et al., 2025).

In both tasks, the simulated performance of the HVLP configuration showed a response to offset visual feedback that was markedly similar to that observed in our human participants. As this configuration relies heavily on velocity-based feedback with only low-precision positional feedback, it effectively integrates proprioceptively sensed velocity to estimate how its arm configuration is changing over time. The high precision of the (possibly offset) visual feedback means that the visual feedback becomes highly determining for its estimated state when it is available; any induced state offset is propagated forward due to the heavy velocity-based filtering.

The low-precision positional feedback prevents velocity-based errors from accumulating unchecked; despite its small Kalman gain, it gradually pulls the estimated state back toward the true state. This produces the slow realignment of movement paths in the offset versus non-offset visual conditions.

The simulation model here uses an ideal observation model to map between its perceived state and visual and proprioceptive observations. As such, the estimated state will tend to be pulled towards the true state by the available proprioceptive positional feedback, and it will not show the kind of spontaneous drifting behaviour towards a constant offset, as demonstrated by Smeets et al. (Smeets et al., 2006), as well as in the behaviour of our human participants here. One approach to expand on this work in the future could involve employing a non-perfect forward kinematic model to map between proprioceptive and visual reference frames.

Both position- and velocity-based proprioceptive feedback originates from the signalling of muscle spindle afferents, with velocity-signals mainly originating from type Ia afferents and positional signals from the II afferents, though there is some overlap in their function. Our findings here thus matches up well with previous research on the function of these two afferent types; a recent review by Banks et al. highlighted the conclusion that the type II afferents are of particular importance for signalling slow and maintaining changes in body position, while the type Ia afferents are a primary source of feedback regarding the sensation of movement (Banks et al., 2021).

Additionally, it should be noted that while we have here positioned our discussion of the results in terms of proprioceptive velocity-based feedback, efferent motor commands also present themselves as a possible source of information regarding how the state of the arm is changing over the course of a movement. This was clearly demonstrated by Gandevia et al. (Gandevia et al., 2006), where the ischaemia-paralysed arm was perceived to have changed position after an (ineffective) attempt was made to move it, showing that internally generated motor intentions can on their own contribute to the perception of arm state. This was further modelled by Smeets et al. (Smeets et al., 2006), who positioned their interpretation on the assumption that intended movements carry the totality of information regarding changes in hand position, while proprioceptive feedback is modelled as positional-only. In the case of the simulation model we used here, the efferent motor commands contributed equally to the estimate of arm movement velocity and position changes in the two compared configurations. While it is not clear how the relative contribution to perceived movement velocity can be cleanly attributed to efferent motor signals compared to afferent rate-of-change proprioceptive feedback during active movement, both notions support the interpretation that these signals can be translated into expected changes in position, which are then propagated forward in time. Findings based on applying antagonist muscle vibration during active movement have, however, shown that undershooting of targets during reaching tasks occurs together with movement slowdowns (Cody et al., 1990; Eschelmuller et al., 2023; Mortensen & Christensen, 2025); such muscle vibration is known to primarily (though not exclusively) bias the firing rate of type Ia afferents and affect the sense of limb motion (Proske & Gandevia, 2012). These observations would tend to support the notion that velocity feedback from type Ia afferents is integrated to augment the inference of limb position.

Conclusion

To summarise, we have demonstrated how offset visual feedback continues to affect the perceived position of the hand, long after it is no longer available, both within a single continuous movement as well as across several sequential movements. This clear history-dependence in how the currently perceived hand position is estimated is well in line with the predictions made by predictive frameworks of sensory integration. This appears to be the

case specifically if we assume relatively high reliance on rate-of-change signals, such as those originating from type Ia muscle spindle afferents, and relatively low reliance on absolute positional signals.

Methods

The study was conducted at the Department of Psychology, University of Copenhagen, at the Section for Cognition and Neuropsychology, and was approved by the departmental ethics review board (no. IP-EC-030723-01), in accordance with the Declaration of Helsinki (excepting prior registration in a public database). Twenty-two healthy participants (16 women, 6 men) were enrolled in the study, all of whom were right-handed, with a mean age of 23 (range 20-29). Informed and written consent was obtained from all participants before participation in the experiment. All data were collected during a single 3-hour session, and participants were compensated at a rate of 160 DKK/hour. The data for the two tasks described here were collected as part of a larger study investigating various aspects of sensorimotor control.

Experimental setup

The participant was seated at a table, with the table height adjusted so that when the right hand rested on the table with the upper arm in a vertical position, the elbow was slightly extended past 90°. This ensured that the lower arm would not slide across the tabletop during the two tasks, preventing any additional tactile feedback regarding the executed movement.

The participant was outfitted with a Meta Quest 3 virtual reality (VR) headset (Meta Platforms Technologies Ireland Limited, Dublin, Ireland), which was connected to the experiment PC via Quest Link. The experiment software was programmed in the Unity game engine (Unity Technologies, San Francisco, USA), using the Unity Experiment Framework (Brookes et al., 2020) for the trial and block structure. The VR headset was set up with 90 Hz refresh rate, which controlled both the refresh rate of the displays as well as the rate at which the system reported the tracked position of the controller. The right-hand VR controller was seated in a 3D-printed mount with a flat base with a felt slider attached underneath. This allowed the VR controller to slide across the tabletop with minimal friction, and participants were instructed to always let the controller rest on the table and not lift it. This was monitored for in the tracking data; no trials had to be rejected due to a lifted controller. The left-hand VR controller was attached to the tabletop, which allowed a virtual copy of the table to be tracked in into the virtual environment. This virtual table was always visible, ensuring stable visual feedback regarding the working environment.

When visual feedback of the controller location was provided in the virtual environment, it consisted of a simple sphere (1.25 cm radius), providing position but not orientation feedback. The height of this sphere was centred on the upper part of the controller, specifically on the button on which the thumb was resting; this button was also used as input during the two completed tasks. The target in both tasks was a sphere (2.5 cm radius), positioned at the same height; the centre of the controller sphere could always be matched to the centre of the target sphere with the controller resting on the table.

Two different tasks were completed, a target tracking task and a sequential reaching task. During the target tracking task, the target would move at a constant speed, following the perimeter of circle with a radius of 15 cm, with each full revolution lasting 4 seconds. The participant would indicate their readiness by placing the right VR controller at an indicated start location approximately 10 cm in front of their chest, and pressing the thumb button on the controller. At this point, the visual feedback of the controller would be removed, leaving only the table and the target in view. The target, positioned 30 cm in front of the start location would begin moving along the perimeter of the circle in a clockwise direction (when see from above), intersecting the start location after two seconds, at which point the participant was instructed to follow and track the target, specifically by keeping the thumb button as close to the centre of the target as possible. The target would continue for 18 seconds (4.5 revolutions after tracking onset). A total of 40 trials were completed per participant. In three-quarters of these,

visual feedback was reinstated for one second, in the last quarter of the first revolution (see 1). Visual feedback was assigned pseudorandomly to be correctly mapped or offset by 5 cm to either side.

During the sequential reaching task, the same start location was used, which was marked with a grey sphere, with another grey target placed 38 cm in front of this (both 2.5 radius). The trial was similarly initiated by pressing the thumb button when at the start location, which would again remove visual feedback of the controller location. After a random delay of 500-1000 ms, the far target would turn green, which functioned as the 'go' signal. From here, the participant was tasked to indicate the far target as precisely as possible, by sliding the controller there and pressing the thumb button again; this would switch the colours of the two targets, such that the start location would now be green, prompting a move to this target. This would be repeated for a total of three movements (out, back and out again). Visual feedback would be reinstated for 100 ms after passing the 10 cm mark on the first outbound reach; it would again be either correctly mapped, or offset to either side corresponding to a 7.5 ° rotation around the start position (approx. 1.3 cm lateral offset at at the 10 cm onset mark). A total of 60 trials were completed for 20 repetitions of each condition, in pseudo-random order.

Five practice trials with full visual feedback of the controller position was completed prior to starting each task to familiarise the participant with the procedure and the movement required.

Simulation model

We describe the simulation model's implementation in detail in our previous paper (Mortensen et al., 2025); we will briefly summarise it below. The model was set up as a two-jointed planar arm, with a fixed and known shoulder position. The controlling Bayesian Agent would attempt to infer the state of its arm from noisy visual and proprioceptive sensory feedback sources with an Unscented Kalman Filter (Julier & Uhlmann, 2004), while producing efferent motor commands with a Linear Quadratic Regulator; these commands would be subject to signal-dependent noise before being applied to each joint of the arm. The simulation model is as such conceptually similar to the model presented by Todorov and Jordan (Todorov & Jordan, 2002) in their seminal paper on Optimal Feedback Control.

The simulated Agent would attempt to infer the state of its arm, defined by the angular positions and velocities of its shoulder and elbow joints. Visual feedback was provided as a set of sampled 2D Cartesian position coordinates. Proprioceptive feedback was provided through sampled identity values of the angular positions and velocities of each joint. Both visual and proprioceptive feedback were corrupted by white noise. We here use the same two configurations of proprioceptive noise as we previously presented with the model implementation (Mortensen et al., 2025). These are each configured with high-precision velocity feedback and low-precision position feedback (HVLP) or, alternatively, low-precision velocity feedback and high-precision position feedback (LVHP); these allow us to compare a primarily velocity- or position-reliant proprioceptive system.

Sensory noise was implemented as discretised white noise, allowing us to specify the noise amplitude independently of the chosen Δt. The simulation was run with a time step size of 10 ms. Both configurations had identical visual noise of $1\,\mathrm{mm\,s^{-1/2}}$, equivalent to per-step sampling SD of 10 mm. Signal-dependent motor noise was set at $0.05\,\mathrm{s^{-1/2}}$, corresponding to a per-step sampling SD of 50 % of the efferent motor command.

The two proprioceptive configurations differed only in the relative amplitude of proprioceptive sensory noise: the HVLP configuration had a 1-to-4 relative sampling SD of velocity versus positional sensors (leading to 16-to-1 relative precision), with these values flipped for the LVHP configuration. Specifically, the HVLP was configured with velocity noise of $0.015\,\mathrm{rad\,s^{-3/2}}$ and positional noise of $0.06\,\mathrm{rad\,s^{-1/2}}$ (leading to per-step sampling SD of $0.15\,\mathrm{rad\,s^{-1}}$ and 0.6 rad respectively), while the LVHP configuration had $0.06\,\mathrm{rad\,s^{-3/2}}$ velocity noise and $0.015\,\mathrm{rad\,s^{-1/2}}$ positional noise (per-step sampling SD of $0.6\,\mathrm{rad\,s^{-1}}$ and 0.15 rad).

Statistics

In order to quantify the time scale of how the visually induced offset decreased over time in the target tracking task, we fit a Bayesian exponential decay model. This model type was chosen to obtain an easily interpretable parameter describing the time course of the decay. We used a hierarchical structure with partial pooling to

account for the nested nature of the data (trials repeated within participants), and to model the inter-individual variation in initial offset (expected to be around 5 cm) as well as the rate of decay. The model was fit with the brms package (Bürkner, 2017) in R (R Core Team, 2024). Prior to model fitting, we computed the per-trial offset along the x-axis by subtracting the mean of the x-position at each time point from the corresponding x-position during the trials with a visual offset. This was to account for any static offset along the x-axis relative to the target, and was done within-individual. Furthermore, the resulting x-offset values from trials with a positive visual x-offset (prompting a negative-direction correction) were multiplied by -1. This resulted in a positive-towards-zero direction of decay for both offset conditions, meaning that the same decay parameter could be fit for both conditions; this matches the assumption that the decay rate for the offset is assumed to be independent of the direction of the induced offset, and simplifies the model by only estimating a single decay rate parameter. We include data from immediately after removal of visual feedback, which we here set to $t = 0$, until the end of each trial.

For each observation, we assumed:

$$x_offset_{i,j,k} \sim \mathcal{N}\big(\mu_{i,j,k}, \sigma_j^2\big)$$

$$\mu_{i,j,k} = exp(\alpha_{j,k} - p_j * t_i)$$

Where $x_offset_{i,j,k}$ denotes the calculated x-offset for participant j, at time t_i, and visual condition k. The parameter $\alpha_{j,k}$ is the condition- and participant-specific initial log-transformed offset at $t_i = 0$, while p_j controls the temporal decay and σ_j is the residual standard deviation for participant j. We modelled the hierarchical structure as:

$$\alpha_{jk} = \alpha_k + b_{jk}, \qquad b_{jk} \sim \mathcal{N}(0, \tau_\alpha^2),$$

$$p_j = \mu_p + u_{p,j}, \quad u_{p,j} \sim \mathcal{N}(0, \tau_p^2),$$

$$\sigma_j = \mu_\sigma + u_{\sigma,j}, \quad u_{\sigma,j} \sim \mathcal{N}(0, \tau_\sigma^2),$$

With $k \in \{+5\,\text{cm}, \text{-}5\,\text{cm}\}$ and $j = 1, \ldots, 22$. Here α_k are condition-level means of the (log-transformed) amplitude parameter α, and μ_p, μ_σ are population-level means of the decay and noise parameters, respectively. The hyperparameters $\tau_\alpha, \tau_p, \tau_\sigma$ are group-level standard deviations that capture between-participant variability in α_{jk}, p_j, and σ_j. Priors on all population-level parameters and standard deviations were weakly informative regularising priors (exact values can be found in the code repository).

Author contributions

ESM and MSC both conceptualised the work. ESM collected and analysed the data and drafted the manuscript. ESM and MSC both revised and approved the final manuscript.

Data availability statement

The following links will be made available at the time of peer-reviewed publication. The model code is available at https://github.com/CoInAct-group/BayesianInferenceArm and archived at https://osf.io/kvtcu. Experimental and simulated data sets are made available at https://osf.io/xz5jt.

Disclosures

The authors have no conflicts of interest to disclose.

References

Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, *82*(1), 35–45. https://doi.org/10.1115/1.3662552

Eklund, G. (1972). Position sense and state of contraction; the effects of vibration. *Journal of Neurology, Neurosurgery & Psychiatry*, *35*(5), 606–611. https://doi.org/10.1136/jnnp.35.5.606

Goodwin, G. M., McCloskey, D. I., & Matthews, P. B. C. (1972). Proprioceptive Illusions Induced by Muscle Vibration: Contribution by Muscle Spindles to Perception? *Science*, *175*(4028), 1382–1384. https://doi.org/10.1126/science.175.4028.1382

Roll, J. P., & Vedel, J. P. (1982). Kinaesthetic Role of Muscle Afferents in Man, Studied by Tendon Vibration and Microneurography. *Experimental Brain Research*, *47*(2), 177–190. https://doi.org/10.1007/BF00239377

Elliott, D., & Madalena, J. (1987). The Influence of Premovement Visual Information on Manual Aiming. *The Quarterly Journal of Experimental Psychology Section A*, *39*(3), 541–559. https://doi.org/10.1080/14640748708401802

Cody, F. W., Schwartz, M. P., & Smit, G. P. (1990). Proprioceptive guidance of human voluntary wrist movements studied using muscle vibration. *The Journal of Physiology*, *427*(1), 455–470. https://doi.org/10.1113/jphysiol.1990.sp018181

Elliott, D., & Calvert, R. (1990). The influence of uncertainty and premovement visual information on manual aiming. *Canadian Journal of Psychology / Revue canadienne de psychologie*, *44*(4), 501–511. https://doi.org/10.1037/h0084263

Rossetti, Y., Stelmach, G., Desmurget, M., Prablanc, C., & Jeannerod, M. (1994). The effect of viewing the static hand prior to movement onset on pointing kinematics and variability. *Experimental Brain Research*, *101*(2), 323–330. https://doi.org/10.1007/BF00228753

Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, *269*(5232), 1880–1882. https://doi.org/10.1126/SCIENCE.7569931

Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, *394*(6695), 780–784. https://doi.org/10.1038/29528

Vindras, P., Desmurget, M., Prablanc, C., & Viviani, P. (1998). Pointing errors reflect biases in the perception of the initial hand position. *Journal of Neurophysiology*, *79*(6), 3290–3294. https://doi.org/10.1152/JN.1998.79.6.3290

Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, *3*(11), 1212–1217. https://doi.org/10.1038/81497

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*(11), 1226–1235. https://doi.org/10.1038/NN963

Van Beers, R. J., Baraduc, P., & Wolpert, D. M. (2002). Role of uncertainty in sensorimotor control. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *357*(1424), 1137–1145. https://doi.org/10.1098/RSTB.2002.1101

Brown, L. E., Rosenbaum, D. A., & Sainburg, R. L. (2003a). Limb Position Drift: Implications for Control of Posture and Movement. *Journal of Neurophysiology*, *90*(5), 3105–3118. https://doi.org/10.1152/JN.00013.2003

Brown, L. E., Rosenbaum, D. A., & Sainburg, R. L. (2003b). Movement speed effects on limb position drift. *Experimental Brain Research*, *153*(2), 266–274. https://doi.org/10.1007/s00221-003-1601-7

Julier, S. J., & Uhlmann, J. K. (2004). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, *92*(3), 401–422. https://doi.org/10.1109/JPROC.2003.823141

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719. https://doi.org/10.1016/j.tins.2004.10.007

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247. https://doi.org/10.1038/nature02169

Albert, F., Ribot-Ciscar, E., Fiocchi, M., Bergenheim, M., & Roll, J. P. (2005). Proprioceptive feedback in humans expresses motor invariants during writing. *Experimental Brain Research*, *164*(2), 242–249. https://doi.org/10.1007/S00221-005-2246-5/TABLES/1

Albert, F., Bergenheim, M., Ribot-Ciscar, E., & Roll, J. P. (2006). The Ia afferent feedback of a given movement evokes the illusion of the same movement when returned to the subject via muscle tendon vibration. *Experimental Brain Research*, *172*(2), 163–174. https://doi.org/10.1007/s00221-005-0325-2

Gandevia, S. C., Smith, J. L., Crawford, M., Proske, U., & Taylor, J. L. (2006). Motor commands contribute to human position sense. *Journal of Physiology*, *571*(3), 703–710. https://doi.org/10.1113/JPHYSIOL.2005.103093

Smeets, J. B., Van Den Dobbelsteen, J. J., De Grave, D. D., Van Beers, R. J., & Brenner, E. (2006). Sensory integration does not lead to sensory calibration. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(49), 18781–18786. https://doi.org/10.1073/pnas.0607687103

Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, *9*(4), 292–303. https://doi.org/10.1038/nrn2258

Friston, K. (2008). Hierarchical Models in the Brain. *PLOS Computational Biology*, *4*(11), e1000211. https://doi.org/10.1371/JOURNAL.PCBI.1000211

Friston, K. J., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: A free-energy formulation. *Biological Cybernetics*, *102*(3), 227–260. https://doi.org/10.1007/s00422-010-0364-z

Proske, U., & Gandevia, S. C. (2012). The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force. *Physiol Rev*, *92*, 1651–1697. https://doi.org/10.1152/physrev.00048.2011

Chancel, M., Blanchard, C., Guerraz, M., Montagnini, A., & Kavounoudias, A. (2016). Optimal visuotactile integration for velocity discrimination of self-hand movements. *Journal of Neurophysiology*, *116*(3), 1522–1535. https://doi.org/10.1152/jn.00883.2015

Bürkner, P. C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, *80*(1), 1–28. https://doi.org/10.18637/jss.v080.i01

Patterson, J. R., Brown, L. E., Wagstaff, D. A., & Sainburg, R. L. (2017). Limb position drift results from misalignment of proprioceptive and visual maps. *Neuroscience*, *346*, 382–394. https://doi.org/10.1016/J.NEUROSCIENCE.2017.01.040

Hewitson, C. L., Sowman, P. F., & Kaplan, D. M. (2018). Interlimb Generalization of Learned Bayesian Visuomotor Prior Occurs in Extrinsic Coordinates. *eNeuro*, *5*(4), 183–201. https://doi.org/10.1523/ENEURO.0183-18.2018

Macefield, V. G., & Knellwolf, T. P. (2018). Functional properties of human muscle spindles. *Journal of Neurophysiology*, *120*(2), 452–467. https://doi.org/10.1152/jn.00071.2018

Tuthill, J. C., & Azim, E. (2018). Proprioception. *Current Biology*, *28*(5), R194–R203. https://doi.org/10.1016/j.cub.2018.01.064

Brookes, J., Warburton, M., Alghadier, M., Mon-Williams, M., & Mushtaq, F. (2020). Studying human behavior with virtual reality: The Unity Experiment Framework. *Behavior Research Methods*, *52*(2), 455–463. https://doi.org/10.3758/s13428-019-01242-0

Limanowski, J., & Friston, K. (2020). Active inference under visuo-proprioceptive conflict: Simulation and empirical results. *Scientific Reports*, *10*(1), 1–14. https://doi.org/10.1038/s41598-020-61097-w

Banks, R. W., Ellaway, P. H., Prochazka, A., & Proske, U. (2021). Secondary endings of muscle spindles: Structure, reflex action, role in motor control and proprioception. *Experimental Physiology*, *106*(12), 2339–2366. https://doi.org/10.1113/EP089826

Chancel, M., & Ehrsson, H. H. (2023). Proprioceptive uncertainty promotes the rubber hand illusion. *Cortex*, *165*, 70–85. https://doi.org/10.1016/j.cortex.2023.04.005

Eschelmuller, G., Szarka, A., Gandossi, B., Inglis, J. T., & Chua, R. (2023). The effects of periodic and noisy tendon vibration on a kinesthetic targeting task. *Experimental Brain Research*, *242*(1), 59–66. https://doi.org/10.1007/S00221-023-06727-1

R Core Team. (2024). R: A Language and Environment for Statistical Computing. https://www.R-project.org/

Mortensen, E. S., & Christensen, M. S. (2025). Proprioceptive integration in motor control. *bioRxiv*. https://doi.org/10.1101/2025.07.18.665517

Mortensen, E. S., Ottenheijm, M. E., & Christensen, M. S. (2025). Implications of recursive Bayesian Sensory Inference. *bioRxiv*. https://doi.org/10.64898/2025.12.30.696730

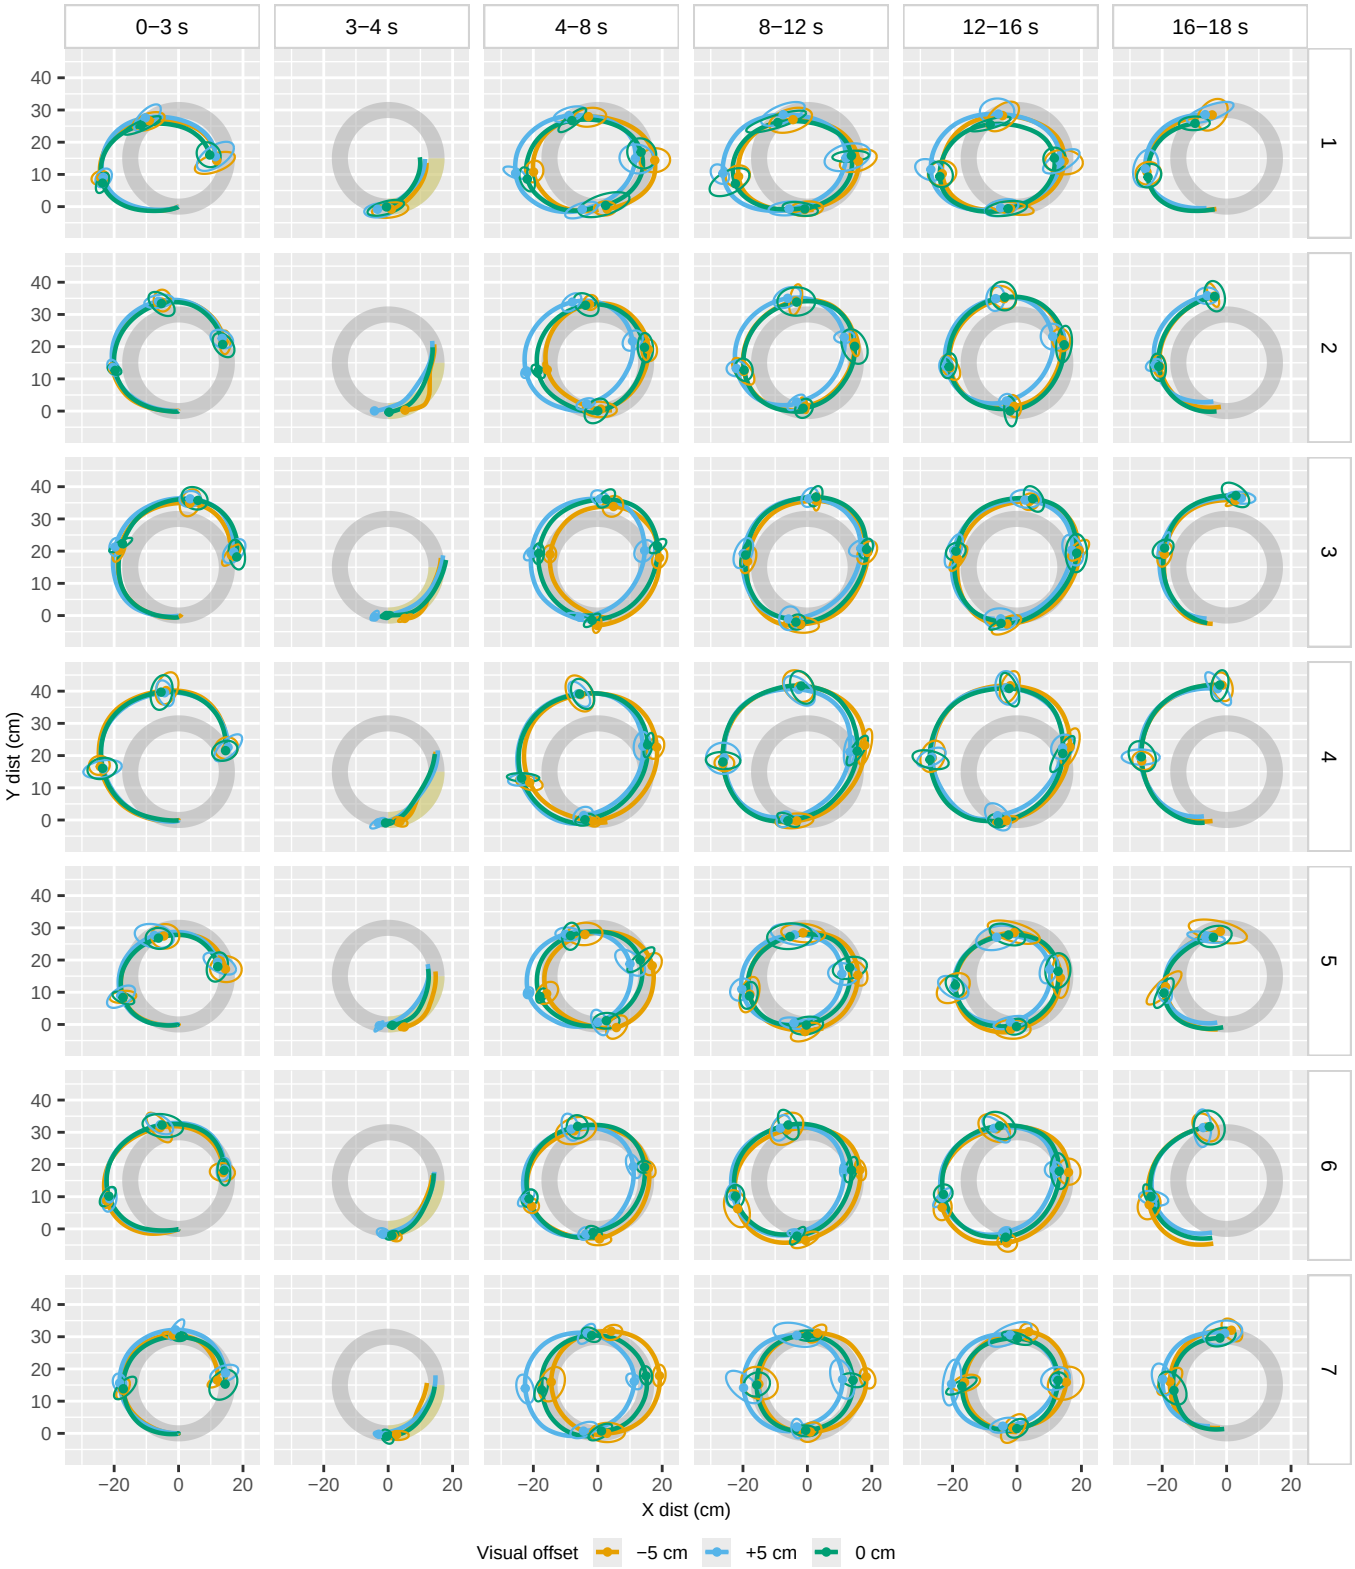

Figure S1: **Circular tracking task: Example participants 1:7**

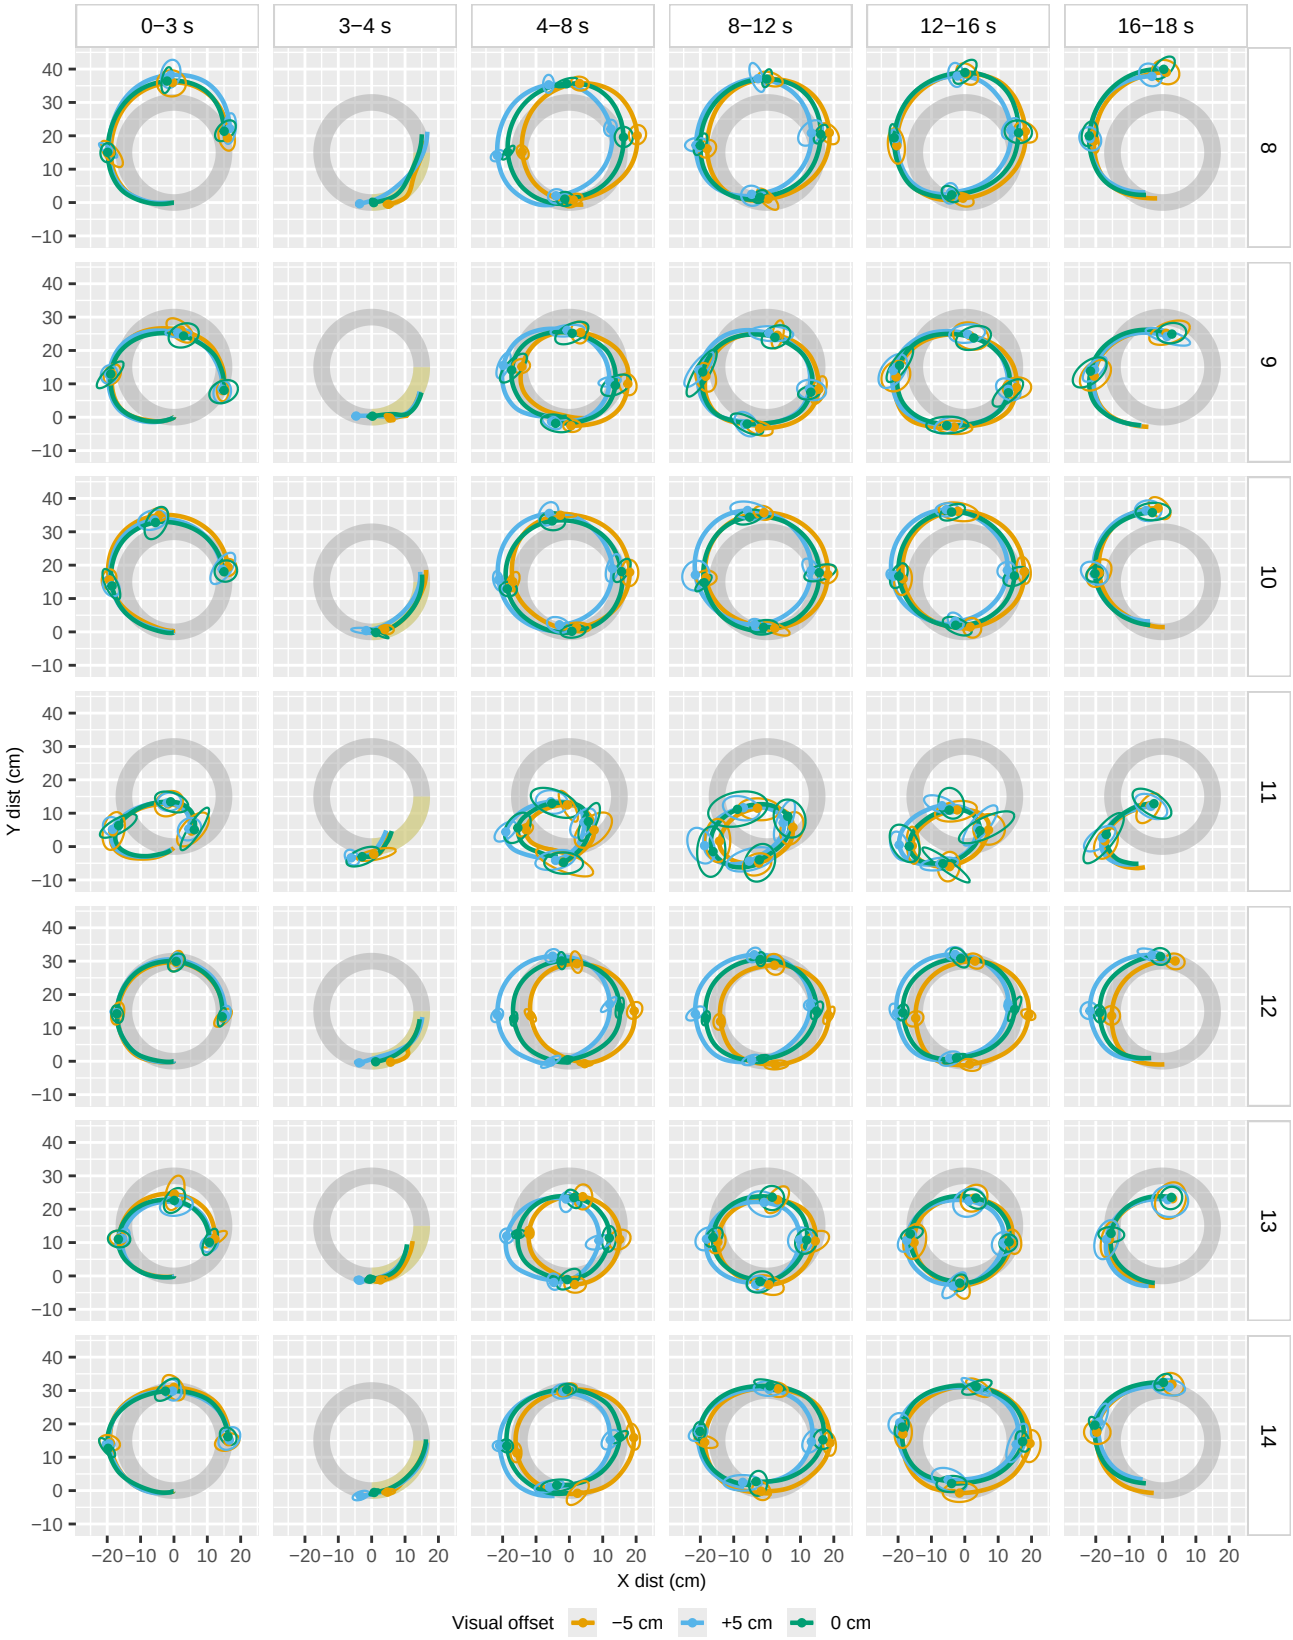

Figure S2: **Circular tracking task: Example participants 8:14**

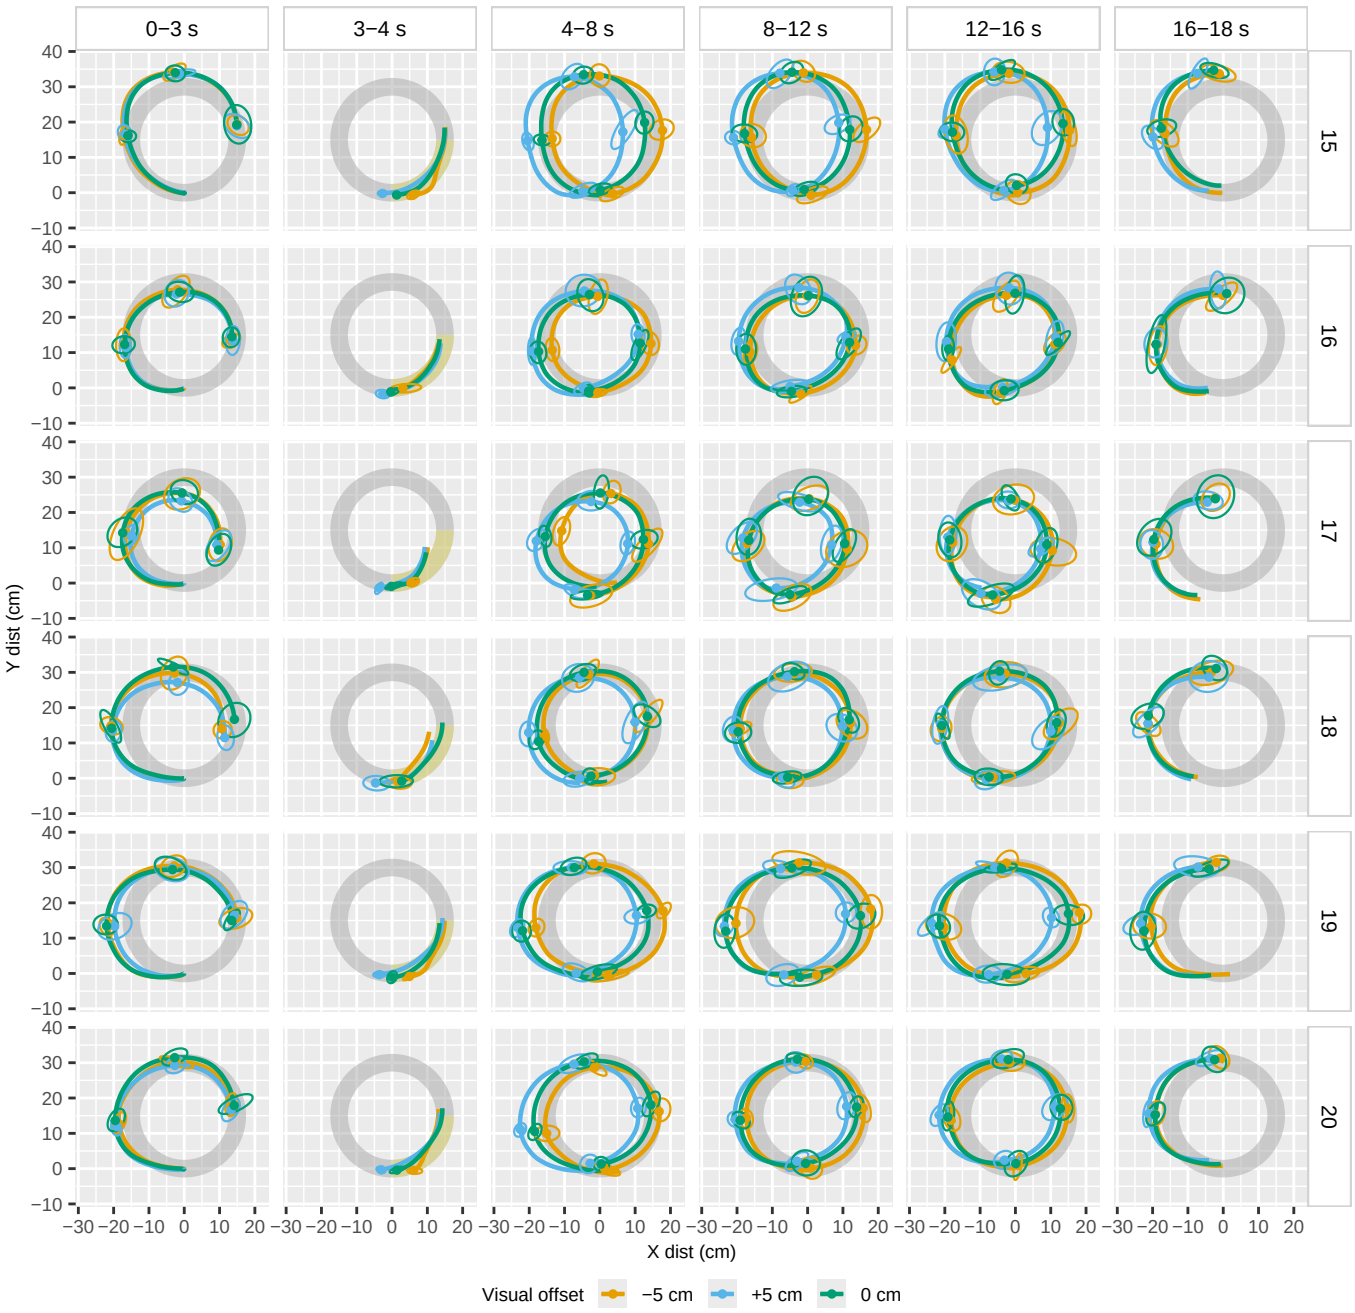

Figure S3: **Circular tracking task: Example participants 15:20**