Implications of recursive Bayesian Sensory Inference

Erik Skjoldan Mortensen, Maud Eline Ottenheijm, and Mark Schram Christensen

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

Abstract

In this article, we examine how the balance between velocity-based and position-based proprioceptive feedback influences state estimation during motor control. We introduce a computational model of arm state inference grounded in Bayesian sensory integration and compare its behaviour to findings from three classical sensorimotor studies. The model allows us to contrast the predicted behaviour of two proprioceptive configurations: one relying primarily on position signals and another relying primarily on velocity signals, reflecting the distinct contributions of type II and type Ia muscle spindle afferents. Our simulations show that a system with strong reliance on velocity-based feedback tends to represent movement relative to previously estimated positions. In a simulated reaching task with briefly presented offset visual feedback, such an Agent produces systematic endpoint errors even after visual feedback is removed. In contrast, an Agent relying mainly on positional feedback is able to correctly update its inferred hand position after visual feedback is removed, thereby limiting this type of biased endpoint errors. When biased visual feedback remains continuously available, or when muscle vibration is simulated, the two configurations produce very similar behaviour. These results indicate that markedly different assumptions about the weighting of positional and velocity proprioceptive cues can yield similar observable behaviour in some of the often-used experimental setups designed to probe state inference processes. This underscores the importance of carefully considering the composition of proprioceptive signals when building computational models and interpreting human sensorimotor experiments. We highlight the task conditions under which our model predicts clear behavioural differences arising from the relative contribution of velocity versus positional feedback.

1 Author summary

During motor control, the central nervous system tracks both the position and movement of our limbs. Even with our eyes closed, we can bring the tips of our index fingers together, a feat that depends on sensory signals from specialised receptors in our muscles. One set of these receptors is most sensitive to the rate of muscle lengthening, providing information about movement, while another set signals the muscle's current length, giving us a sense of posture. There has so far been limited discussion of the extent to which these two feedback channels are used to augment one another. For example, signals about the rate of muscle lengthening reflect how our pose is changing from moment to moment, but it remains unknown to what extent the brain uses this information to update its estimate of limb position. In this study, we simulate three classical sensorimotor experiments while varying the precision of these two types of sensory feedback. Two simulations show that different assumptions can yield similar movement patterns, highlighting the need for caution when interpreting such experiments. The third suggests that movement-related feedback may contribute more to our sense of limb position than previously recognised.

2 Introduction

Moving our hand to a target requires feedback on the arm's initial state and how it evolves over the course of the reaching movement. When moving, we have multiple senses to guide us; vision may inform us of both

the target location and the location of our hand in a visual reference frame, while proprioception allows us to feel our hand and arm through feedback from muscle, skin, and joint sensors. In addition to these feedback signals, our intention to move forms a feed-forward signal that may likewise inform us of how these different sources of sensory feedback are expected to change over time. Yet it is a basic principle of imperfect and noisy biological sensors and neural communication that these different sources of information will at any given time never perfectly match each other, nor the ground truth, which they are attempting to represent (Faisal et al., 2008). Instead, such signals are at best simplified and noisy abstractions, which the central nervous system must apply a certain set of rules to in order to estimate the pose and movement of the body; as has been noted in various forms over the years, a map is not the territory it represents (Korzybski, 1933), but it may nevertheless be very useful for navigation.

2.1 Bayesian inference

This large array of available, though to some extent mismatching, information highlights the need for a method to integrate different sources of information into a coherent state estimate, which can then be used for motor control. Bayesian sensory integration has been proposed as a way to understand this phenomenon (Knill & Pouget, 2004; Franklin & Wolpert, 2011). A central state estimator may, over time, learn the expected level of noise present in a given input sensory channel and output motor channel and use this noise estimate to weigh both incoming and outgoing signals by their expected precision. Several studies have demonstrated behaviour indicative of such precision-weighted sensory integration, when experimenters have manipulated the availability, reliability, and bias of different types of sensory feedback (Van Beers et al., 1996, 1999; Ernst & Banks, 2002; Körding & Wolpert, 2004; Chancel, Blanchard, et al., 2016).

By linking such Bayesian sensory integration with a dynamics model, the theory may further describe how the perceived state of the arm is influenced not only by the currently incoming sensory feedback but also by the previously perceived state of the arm. The currently perceived state of the arm represents prior knowledge that new sensory data should be used to update, rather than attempt to construct the whole state of the arm "from scratch" using only the latest sensory feedback. For example, if the arm is believed to be stationary, then the currently perceived position represents a prior for the current state, which, along with expected noise in the various sensor channels, may then also become a prior for incoming sensory feedback to be weighted against. If a control signal to contract the biceps was recently sent, then the dynamic model can be used to predict both how the elbow should begin flexing in response and how downstream incoming sensory feedback is expected to change based on this action. Several studies have shown how efferent drive affects the perception of limb position, both when afferent feedback is blocked (Gandevia et al., 2006) and available (Laufer et al., 2001; Smith et al., 2009; Erickson & Karduna, 2012), demonstrating the role of intended action for state perception.

As such, recursive Bayesian sensory integration provides a framework for updating the internally represented state of the arm (or body, in a broader context) through precision-weighted incoming sensory feedback and outgoing motor control signals. The basic functioning of this proposed system relies on utilising predictions of incoming sensory feedback, which has further been proposed to serve the additional function of compensating for the delays associated with neural conduction and processing, which causes lag both when sending efferent commands and receiving afferent feedback. Outgoing signals affect the state of the arm sometime in the future, while incoming signals represent the past. Predictions made through a dynamics model are hypothesised to bridge this temporal gap (Miall et al., 1993; Miall & Wolpert, 1996; Franklin & Wolpert, 2011).

This model of Bayesian sensory integration forms the basis for state estimation in both Optimal Feedback Control (Wolpert et al., 1995; Todorov & Jordan, 2002) and Active Inference (Friston, 2010), two theories that have been central to the field of sensorimotor control over the last few decades. While the above description follows the formulation as used in Optimal Feedback Control more closely in terms of the use of efference copy signals, a similar principle applies in the context of Active Inference. In Active Inference, forward models provide sensory predictions of desired states, with actions serving to minimise prediction errors relative to these priors. While the computational implementations differ, both approaches share the principle of precision-weighted Bayesian

integration of state predictions and sensory feedback.

2.2 The constraints of proprioception

Proprioception is a broad term that encompasses the perception of the body's state in terms of posture, position, and movement. In particular, proprioception typically refers to the part of perception arising from a collection of diverse peripheral mechanoreceptors situated in the skin, muscles, and joints (Tuthill & Azim, 2018). To limit the scope of our introduction to the subject, we will focus here on two of the more well-studied proprioceptors: the type Ia and type II muscle afferents. These are stretch-sensitive muscle spindles found distributed in skeletal muscles, which provide state-dependent feedback to the central nervous system. They integrate directly with the motor control system at the spinal level through various local reflex circuits, providing low-latency feedback, as well as supplying ascending input to the brain to allow for multisensory central state estimation and movement planning (Proske & Gandevia, 2012). The type Ia afferent primarily signals the rate of change of muscle length, providing a velocity-based signal (Prochazka & Gorassini, 1998; Roll et al., 2004; Albert et al., 2006), while the type II afferents signal muscle length directly (Botterman & Eldred, 1982; Roll & Vedel, 1982), although the activity is dependent upon the movement task being performed. For example, movements made against gravity may cause a phase shift towards the receptors signalling acceleration (type Ia) and velocity (type II) respectively (Dimitriou & Edin, 2008). In practice, there is likely substantial overlap in the responses of each receptor type, such that they may both, to some extent, encode a position- and velocity-dependent signal, and in some instances also acceleration and further time-derivatives in the case of the type Ia afferents (Botterman & Eldred, 1982; Macefield & Knellwolf, 2018; Blum et al., 2020). They both encode their respective values in their firing frequency, such that type Ia afferents increase their firing frequency proportionally to the stretch-velocity of the parent muscle, while the type II afferents encode the length of the muscle through proportional firing frequency (Botterman & Eldred, 1982; Roll & Vedel, 1982; Grill & Hallett, 1995).

While absolute position sensors, such as the type II afferents, are of clear value for any attempt to estimate the current pose of the body, they also face principled challenges. An absolute state sensor must be able to encode an extensive range of possible states, inherently limiting the resolution of the signal. They are further susceptible to any potential drift of the signal, which must be accounted for to avoid induction of biased estimates; microneurographic studies of type II afferents reveal substantial decay of the firing frequency of type II afferents already within seconds of a ramp-and-hold stretch of a muscle (Roll & Vedel, 1982; Edin & Vallbo, 1990). Additionally, many factors in the internal environment affect the firing frequency of muscle afferents, including temperature (Mense, 1978), pH (Fischer & Schäfer, 2005), tissue pressure (Torell & Dimitriou, 2024), and arterial pulsation (Birznieks et al., 2012), all of which change to some degree in response to muscular work. These may all pose a challenge for extracting the position- and movement-related signals of interest; it may become challenging to disambiguate any change in posture from background noise or sensor drift due to the aforementioned factors. While velocity sensors, such as the type Ia afferent, are not immune to these effects, their phasic response to any change in joint configuration continues to provide a high signal-to-noise ratio signal of when and how the current state is changing. Nevertheless, it is seen that, if an imposed movement is kept slow enough to fall below the detection threshold of velocity sensors, then much larger changes in position are required for reliable detection, compared to faster movements (Hall & McCloskey, 1983; Taylor & McCloskey, 1992). This has led to the general interpretation that type II afferents are likely more involved with the estimation of slow and maintained changes in limb position and posture, while type Ia afferents are more central to the control of faster movements (Banks et al., 2021).

In a Bayesian filter, velocity and positional sensors may be used to augment each other in a reciprocal manner. Velocity sensors, beyond allowing for the direct measurement of movement speed, also provide a source of information regarding how position changes over time. By combining this with even a noisy position sensor, the absolute position error is prevented from drifting too far, while short-term relative position changes are represented well due to the velocity signal. Several findings support the idea that the CNS may be utilising these signals in this manner. For example, viewing the hand prior to reaching towards a target improves endpoint

precision (Elliott & Calvert, 1990; Rossetti et al., 1994; Vindras et al., 1998), while movements performed from an offset hand position estimate bias the endpoint position estimate by a similar amount (Brown et al., 2003b; Vindras et al., 2005; Patterson et al., 2017). This indicates that the perceived arm configuration and hand position are continuously updated in a manner that depends on a previously perceived position; the perception and execution of relative movements continues to function well, even when absolute position errors have accumulated due to drift or have been induced with offset visual feedback. Studies of proprioceptive drift demonstrate how substantial perceptual errors emerge over time as movements are performed without visual feedback (Brown et al., 2003b; Smeets et al., 2006; Patterson et al., 2017), indicating that absolute positional state estimation is relatively poor compared to the accuracy with which individual movements continue to be performed.

2.3 Study goals

Our aim with this study was to investigate the impact of the assumed relative contributions of position- and velocity-based proprioceptive feedback for state estimation. To this end, we simulate the behaviour of a Bayesian inference-based Agent controlling a 2-jointed planar arm. We compare the realised set of simulated data with three classic studies of sensorimotor control, and discuss how our simulation approach can provide further insight into the mechanisms behind the empirical results reported in each study, considering both the simulated motion and the Agent-perceived arm state. These three studies were chosen as classic examples of widely used techniques for investigating state inference.

We contrast two configurations of proprioceptive feedback: one that relies primarily on position-based feedback and the other mainly on velocity-based feedback. The simulated results demonstrate that they can produce very similar patterns of movement which match human behaviour, but that they require different interpretations of the results. Additionally, we also highlight an example experimental setup where the two configurations produce quite different results, which tends to favour a mainly velocity-dependent proprioceptive system. We further discuss how future experiments may be designed to better identify the types of feedback utilised during state estimation.

3 Results

In the following, we will present the results of a simulated Agent, which relies on recursive Bayesian Inference to estimate the state of a two-joint planar arm, as it performs various reaching tasks, set up to match three classical studies of sensorimotor control. In the following, we will first briefly describe the model, and then present the results of each of the three simulated tasks.

We simulate an Agent using recursive Bayesian Inference to infer the state of a planar 2-joint arm (see Fig. 1) in a continuous loop, as it moves towards a specified target. For our implementation of recursive Bayesian Inference, we use an Unscented Kalman Filter (Julier et al., 2000). The Unscented Kalman Filter (UKF) is a model-based state estimator that continually predicts and updates its point state estimate, with uncertainty in state estimation represented as Gaussians, while also tracking any correlations in uncertainty. To control the arm, we utilise a Linear Quadratic Regulator (LQR), a model-based control algorithm that balances the deviation from the desired state with a control cost, with motor control signals corrupted by signal-dependent noise (Harris & Wolpert, 1998). This combination of a Kalman filter-based estimator with LQR control and signal-dependent motor noise is closely analogous to the approach used by Todorov and Jordan (2002).

With this model, our focus is on exploring the behavioural effects of various induced sensory biases in the context of Bayesian Inference, while the formation of motor commands is beyond the scope of the model.

In pseudo-code, the simulation loop is defined in Fig. 2. On each simulated trial, this loop is run for N iterations such that $\Delta t = \frac{t_k}{N}$, with $\Delta t = 10$ ms and N set to match the time requirements of the simulated task. k thus represents each discretised time step iteration in the simulation. In the following, we will briefly outline this perception-action loop, to provide a high-level intuition of the functioning of the Agent. Further details of each step and function is provided in the Methods section.

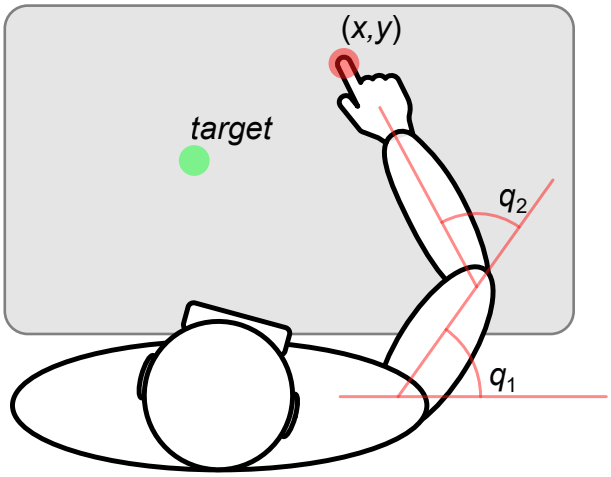

Figure 1: **Visualisation of the simulated setup.** The Agent is sitting by a table with the weight of its arm supported by air sledges, providing minimal friction to the table surface and allowing only planar movement. It is attempting to infer the state of its arm through noisy proprioceptive and visual sensory feedback. Proprioceptive information is received in the form of the angular positions of the shoulder and elbow ($\mathbf{q} = [q1, q2]$) and their corresponding angular velocities ($\dot{\mathbf{q}} = [\dot{q}1, \dot{q}2]$). The Agent is outfitted with Virtual Reality glasses, providing visual feedback of only the hand location, which is implemented as a set of 2D Cartesian coordinates, (x, y). The Agent is able to produce noisy torques at each joint ($\boldsymbol{\tau} = [\tau 1, \tau 2]$), in an effort to drive the hand towards a target location.

To summarise, the Agent continually attempts to estimate the state of its arm and produce motor commands to move its hand towards an indicated target location within a given maximum number of steps. The Agent lacks direct access to the true state of its arm ($\mathbf{x}_k^{\text{true}}$), and must instead rely on noisy sensory observations ($\mathbf{z}_k^{\text{obs}}$). The state of the arm is here fully defined in terms of the shoulder and elbow angular positions and velocities, with the position of the hand further defined by the forward kinematic mapping, given the angular positions and the lengths of the arm segments (see Eq. P.1-2 in Fig. 6). The sensory feedback that the Agent receives is of two types: partly proprioceptive feedback, consisting of noisy versions of the angular positions and velocities of each joint, meant to represent the information available from type II and type Ia muscle spindle afferents, respectively, and partly visual feedback, here implemented as a set of (x, y) Cartesian coordinates of the hand position, likewise with added noise.

Together with an internal model of the coupled dynamics of its arm (matching the actual dynamics function f_{disc}) and its previous efferent motor command, the Agent is able to form predictions about the evolution of the arm configuration ($\mathbf{x}_{k-1}^{\text{post}}$ to $\mathbf{x}_k^{\text{prior}}$), and how certain it is in this prediction ($\mathbf{P}_k^{\text{prior}}$). From this predicted state configuration, predictions of expected sensory feedback may be formed, based on both the certainty of the predicted state estimate and the expected level of noise present in each sensor channel; e.g., the Agent may believe visual feedback to be of high precision, and would, in that case, expect the incoming visual feedback to be correspondingly close to the predicted state.

By comparing the expected sensory feedback with the actual sensory feedback ($\mathbf{z}_k^{\text{obs}}$), the Agent updates its current state, the posterior $\mathbf{x}_k^{\text{post}}$ and the associated (un)certainty $\mathbf{P}_k^{\text{post}}$.

With this posterior estimate, the Agent forms an efferent motor command, here symbolised by $\boldsymbol{\tau}_k^{\text{efferent}}$, which is corrupted by signal-dependent noise before being applied to the physics simulation of the true arm on the next step.

3.1 Simulation 1: Single-joint reaching task with antagonist vibration

In the first simulation, the Agent is tasked to reach a target position 30° away, using only an elbow extension movement, with the shoulder joint locked in place, and the position of the elbow therefore remaining stationary.

Simulation loop

$$
\left.
\begin{array}{ll}
\textbf{Physics:} & \mathbf{x}_k^{\text{true}} \leftarrow f_{\text{disc}}\big(\mathbf{x}_{k-1}^{\text{true}},\, \boldsymbol{\tau}_{k-1}\big)\,, \quad \mathbf{z}_k^{\text{obs}} \leftarrow h\big(\mathbf{x}_k^{\text{true}}\big) + \mathbf{b}_k + \mathbf{v}_k \\
\textbf{Predict:} & (\mathbf{x}_k^{\text{prior}}, \mathbf{P}_k^{\text{prior}}) \leftarrow \text{UKF_predict}\big(\mathbf{x}_{k-1}^{\text{post}},\, \mathbf{P}_{k-1}^{\text{post}},\, \boldsymbol{\tau}_{k-1}^{\text{efferent}}\big) \\
\textbf{Update:} & (\mathbf{x}_k^{\text{post}}, \mathbf{P}_k^{\text{post}}) \leftarrow \text{UKF_update}\Big(\mathbf{x}_k^{\text{prior}},\, \mathbf{P}_k^{\text{prior}},\, \mathbf{z}_k^{\text{obs}}\Big) \\
\textbf{Act:} & \boldsymbol{\tau}_k^{\text{efferent}} \leftarrow \text{LQR}\big(\mathbf{x}_k^{\text{post}}, \mathbf{x}_k^{\text{target}}\big)\,, \quad \boldsymbol{\tau}_k \leftarrow \text{MotorNoise}\big(\boldsymbol{\tau}_k^{\text{efferent}}\big)
\end{array}
\right\} k
$$

Figure 2: **Simplified pseudo-code of the simulation loop.** The simulation loop is run for N iterations with each step k lasting 10 ms. **Physics:** $\mathbf{x}_k^{\text{true}}$ represents the true state of the Agent's arm at time t_k, defined by the shoulder and elbow angular positions and velocities. This state is updated by the discretised physics simulation function f_{disc}, based on the previous state and muscle-generated torque $\boldsymbol{\tau}_{k-1}$. A set of sensory observations ($\mathbf{z}_k^{\text{obs}}$) is generated from the updated arm state, based on the state-sensor mapping ($h(\mathbf{x}_k^{\text{true}})$) with added Gaussian sensor noise ($\mathbf{v}_k$) and possible sensor bias ($\mathbf{b}_k$, e.g., a visual offset). **Predict:** The Agent predicts the current state of its arm ($\mathbf{x}_k^{\text{prior}}$) and its uncertainty in this prediction (state covariance, $\mathbf{P}_k^{\text{prior}}$) by passing its previous posterior state estimate and uncertainty ($\mathbf{x}_{k-1}^{\text{post}}$, $\mathbf{P}_{k-1}^{\text{post}}$), and its associated efferent motor commands $\boldsymbol{\tau}_{k-1}^{\text{efferent}}$ through a dynamic prediction model. **Update:** The Agent updates its posterior state estimate and uncertainty by integrating sensory observations with its predictions. **Act:** The Agent calculates a set of efferent motor commands that will guide its hand towards a specified target state ($\mathbf{x}_k^{\text{target}}$). These motor commands are corrupted by signal-dependent motor noise before being applied to the physics simulation on the next step. See main text for a detailed explanation.

Visual feedback is withheld, so the Agent must rely solely on efferent motor output and proprioceptive feedback to estimate its state as it guides its hand towards the target. Here, we aimed to conceptually recreate the setup presented by Cody et al. (1990), although we simulate elbow rather than wrist movements. Movements were simulated with a duration of 1 second, matching the original setup. Cody et al. demonstrated the classically observed pattern of biased endpoint errors and reduced movement speed in response to various durations and amplitudes of antagonist muscle vibration during active movement. Through varying vibration amplitude, they found that increasing intensity caused successively increased undershooting of the target and larger decreases in movement speed. Varying the point at which the vibration motor was turned on also affected the degree of undershooting; starting vibration earlier resulted in greater undershooting for a given intensity. We test and compare two proprioceptive configurations on both the intensity- and duration-dependent conditions: one that relies mainly on velocity-based feedback, with relatively high-precision velocity feedback and low-precision position feedback (HVLP), and one that relies mainly on position-based feedback, with relatively low-precision velocity feedback and high-precision position feedback (LVHP). Each of these two configurations is matched with a different implementation of simulated antagonist vibration, which provides either a velocity-sampling bias (for the HVLP) or a position-sampling bias (for the LVHP); it is generally recognised that muscle vibration stimulates both type II and type Ia muscle spindle afferents (Burke et al., 1976); both carry a position- and velocity-dependent signal, and so it seems likely that both sensed position and velocity may be affected to some degree.

Fig. 3 shows the averaged elbow angular position plotted across a 1 s movement, for each of the described Agent configurations. These are split across two possible implementations of simulated muscle vibration as either a bias of the sampled angular position (Fig. 3 AII & BII) or velocity (Fig. 3 AI & BI). Here, we aim to highlight that both tested implementations of simulated muscle vibration produce very similar results, depending on the assumed relative reliance on position- and velocity-based proprioceptive cues. While the results here are not fully identical, either would be a good match to the observations made by Cody et al. (Cody et al., 1990); whether vibration is implemented as a velocity-bias for the mainly velocity-dependent HVLP Agent configuration or as a position-bias for the mainly position-dependent LVHP configuration, both producing similar movement patterns, leading to increased undershooting of the target for higher vibration intensities and longer vibration durations.

The baseline performance of the two Agent configurations are closely matched; the HVLP configuration had a mean endpoint of 29.0° (SD = 0.76), while the LVHP configuration had a mean endpoint of 29.1° (SD =

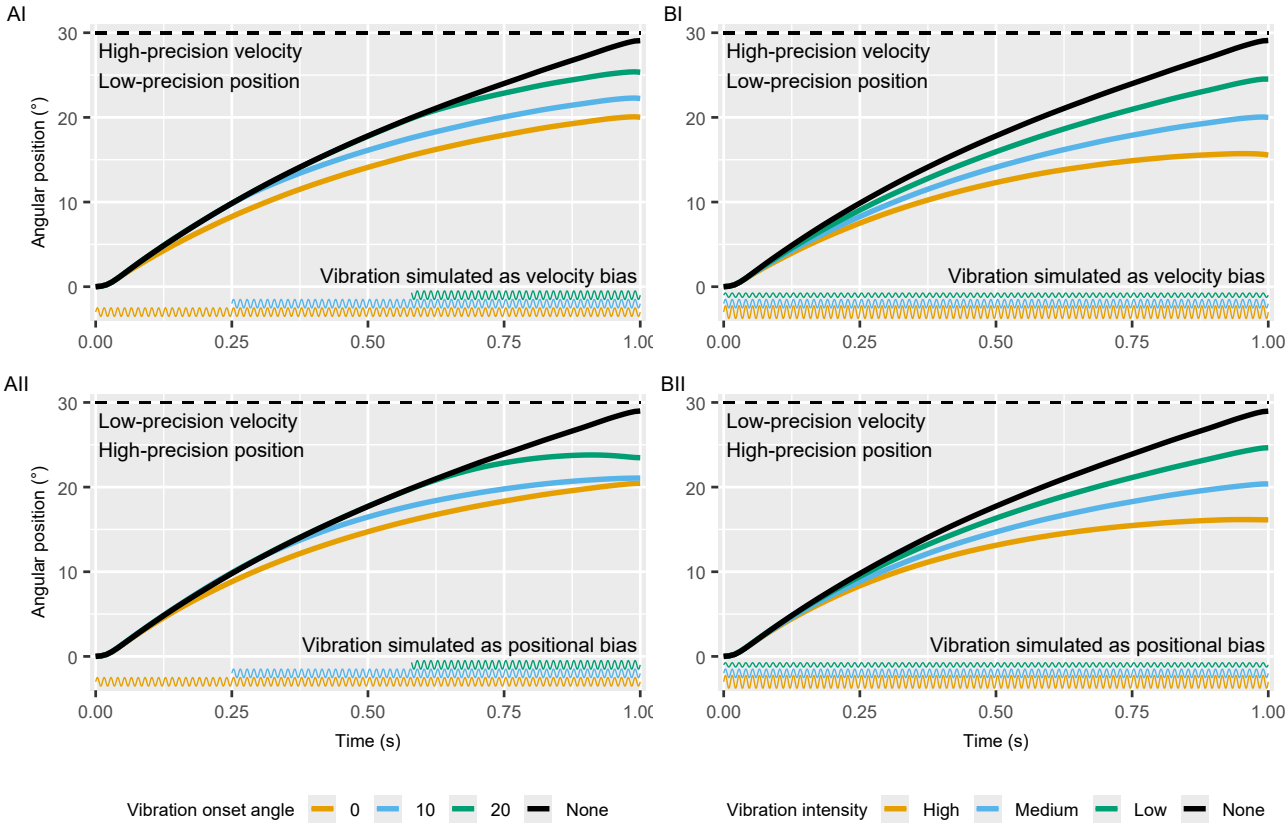

Figure 3: **Results of simulation 1.** Shows the angular position over a 1-second elbow extension movement towards a target 30° away, analogous to the setup presented by Cody et al. (1990). 0° refers to the start position, while the dashed line at 30° is the target position. AI and AII show the effect of simulated antagonist vibration onset partway through the movement, as illustrated by the corresponding coloured sinus traces along the bottom of each plot. BI and BII show the effect of varying the simulated vibration intensity across the whole movement, likewise illustrated by the height of the sinus traces. AI and BI are simulated with high-precision velocity and low-precision positional proprioceptive feedback (HVLP), and with simulated vibration implemented as a velocity bias. AII and BII are simulated with low-precision velocity and high-precision positional proprioceptive feedback (LVHP), and with simulated vibration implemented as a position bias. The vibration intensity used in AI and AII matches the Medium vibration intensity in BI and BII. Plotted paths represent a mean of simulated 500 trials each. See the Methods section for details of simulation.

0.86), in the non-vibrated condition. In both configurations and all vibration conditions, the Agent correctly guides its posterior perceived hand to the target position, i.e., from the point of view of the Agent, it reaches the target (see Fig. S1). It may also be noted that if the simulated vibration bias were applied to the non-dominant source of proprioceptive information, e.g., a position bias to the HVLP configuration, little to no effect would be observed, as one might expect. We chose to report the here-presented combinations to highlight how quite different configurations of proprioceptive precision and simulated vibration can produce very similar motor behaviour.

3.2 Simulation 2: Reaching task with rotated visual feedback

In Simulation 2, the Agent was required to make a reaching movement to a target 22.2 cm in front of its start hand position, requiring movement of both the shoulder and elbow joints. Visual feedback of the hand location was supplied throughout the whole movement, with a visual rotation of up to 10° in either direction applied after passing the 5.7 cm mark, matching the experimental setup by Fourneret and Jeannerod (1998). The induced visual offset is gradually applied, with a final offset of approximately 2.9 cm for a 10° rotation once the target is reached. Trials were simulated with a duration of 2.5 s, approximating the reported duration by Fourneret and Jeannerod. They demonstrated how their participants corrected their movements in response to the visual

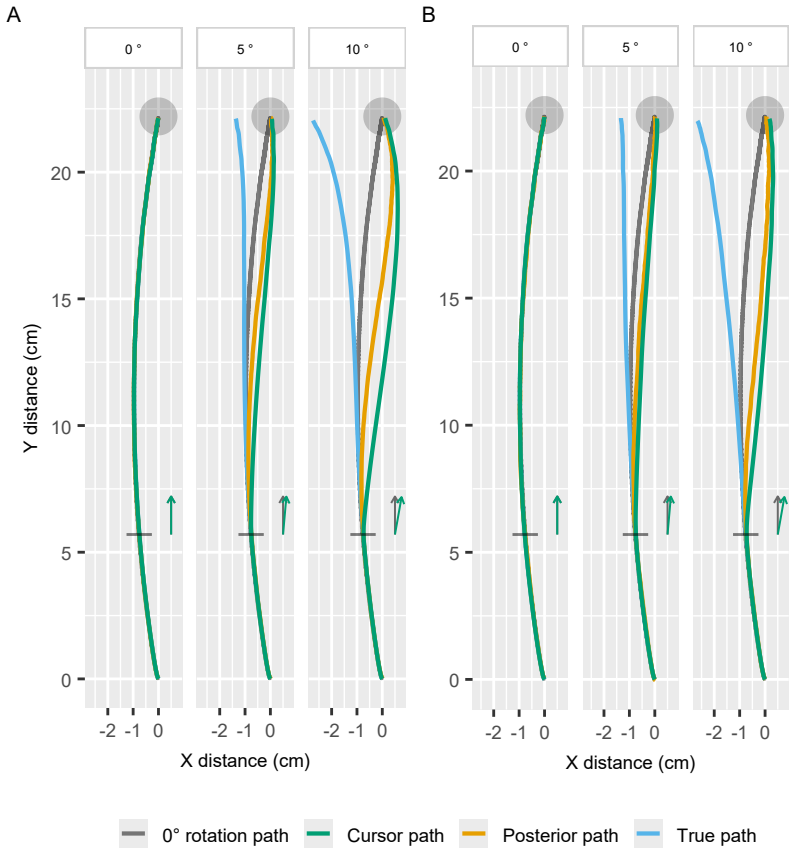

Figure 4: **Results of simulation 2.** Shows the simulated effect of applying a clockwise visual rotation of 0, 5, or 10° starting 5.7 cm into a reaching movement towards a target at a distance of 22.2 cm, mirroring the setup of Fourneret and Jeannerod (1998). Panel A shows an Agent with proprioceptive high-precision elocity feedback and low-precision position feedback, while panel B shows one with proprioceptive low-precision velocity feedback and high-precision position feedback. The green arrows indicate the rotation applied in each plot. In each condition, the Agent guides the posterior perceived position of its hand to the target location. The true path of the unperturbed (0° rotation) condition is added in grey to facilitate comparison of how the rotated visual feedback affects the reaching movement. The paths shown here represent the averages of 500 simulated trials per condition. The cursor path is shown without added visual noise.

feedback rotation, such that their movement paths stayed approximately linear towards the target location, but appeared to be unaware of their own motor performance. They tended to report that their perception of their executed movement deviated in the same direction as the visual rotation, rather than the opposite direction, which would correspond to their actual executed movement.

In Fig. 4, we plot the averaged path of the visual cursor (green), the Agent-perceived posterior path (orange), and the true path of the hand (blue) as the Agent reaches towards the target, for three different visual clockwise rotations of 0°, 5° and 10°. We also plot the averaged path during the 0° condition, as a point of comparison for the remaining conditions. While Fourneret and Jeannerod instructed their participants to follow a straight-line path, we did not implement an explicit straight-line-preference function in our simulation; the slightly curved path seen in Fig. 4 emerges from the implemented LQR motor control function.

The figure shows almost identical movement paths and sensory integration for both the HVLP configuration (Fig. 4A) and the LVHP configuration (Fig. 4B). Both configurations appear to match the perceptual experience reported by the research participants in the study by Fourneret and Jeannerod (1998); the Agent-perceived posterior state deviates in the same direction as the visual feedback, despite the fact that both efferent motor commands and proprioceptive feedback would drive the posterior towards the true hand location. From the Agent's point of view, it observes that it is moving unintendedly too far in the direction indicated by the visual feedback, and it corrects for this mistake by moving slightly more in the opposite direction.

Three different sources of information influence the posterior arm state: partly forward prediction based on efferent motor commands and the previous state estimate, partly visual feedback of hand location, and finally, proprioceptive feedback, which can be further broken down into angular positions and velocities. As the Agent assumes the same causal origin of all three sources of information, the posterior perceived path will tend towards a weighted average of each source of information; when sensory feedback is not biased, this allows for filtering of sensory feedback and efferent output, such that the perceived arm state tracks the true arm state as well as possible, given the present sensory and motor noise (the two left-most sub-panels of A and B, titled "0°"). However, when incongruency is introduced, the posterior will then tend to be dominated by the most trusted source of information, which in this case is visual feedback. As rotated visual feedback is received after the 5.7 cm mark, the posterior state estimate of the Agent begins to show a deviation in the same direction as the offset visual feedback. Reacting in response to this perceived deviation, the Agent begins to modify its efferent motor commands, so as to reach the target from its current posterior state estimate, here moving slightly more to the left than for the unperturbed trials.

Our model thus predicts that in cases such as this, where visual feedback is continuously available, it makes little difference whether proprioceptive feedback is mainly in the form of velocity-based (Fig. 4A) or position-based (Fig. 4B) information, as the visual feedback ends up dominating the posterior state estimation in either case, producing very similar outcomes.

3.3 Simulation 3: Reacting to brief offset visual feedback during reaching

In Simulation 3, we (partly) replicated the classic experiment performed by Körding and Wolpert (Körding & Wolpert, 2004). Here, the Agent was required to reach towards a target positioned 20 cm in front of its start hand position. Visual feedback of the hand position was withheld during the movement, except for 100 ms following the hand passing the halfway (10 cm) mark, forcing the Agent to rely on proprioceptive feedback and efferent output for state estimation for most of the movement. When visual feedback was presented, it could be offset either 0, 0.5, or 1 cm laterally to the left or right, and might either have no, medium, or high blur added to it, decreasing the Agent-assumed precision of the visual feedback (see Methods for details of simulation approach). While Körding and Wolpert employed both a training period to adapt their research participants to a baseline 1 cm offset, followed by a testing period, we forego the baseline offset and adaptation period and simulate the testing phase directly. They demonstrated that subjects in this setup would correct their movement so that the implied cursor, rather than their true hand position, was guided to the target, showing a movement correction similar to that shown by Fourneret and Jeannerod in Simulation 2, despite the brief period of visual feedback. When the precision of the displayed visual feedback was decreased, the correction would be less pronounced, indicating a still-present, but less influential, integration of blurred visual feedback. Our goal with this simulation is, first, to show how the precision of brief visual feedback affects the posterior state estimate, and second, to show the differences between the HVLP and LVHP configurations in how the posterior evolves after visual feedback is again removed.

In Fig. 5 AI and AII, we plot out the averaged true hand path, the posterior Agent-perceived path, and the cursor path, here for an example condition where the cursor is offset 1 cm to the right. Three different visual conditions are plotted, each showing no blur, medium blur, and high blur of the visual feedback that the Agent receives. Here we compare the same two configurations of the Agent as in Simulation 1, the HVLP (Fig. 5 AI and BI) and the LVHP configuration (Fig. 5 AII and BII).

For both Agent types, the initial 10 cm of the movement appears identical. When an offset to the position of visual feedback is added, the forward kinematic mapping is no longer valid, resulting in a discrepancy between the predicted visual feedback and the observed values. A consistent offset in the provided visual feedback will gradually cause the Agent to update its state estimate to one that would be more likely to produce the visual observations; it may be seen how the posterior path is "pulled" towards the cursor during the brief presentation window. The rate at which the posterior is pulled in this direction depends upon the relative precision of the different sources of information; this is seen clearly when Medium and High visual blur are compared to the no

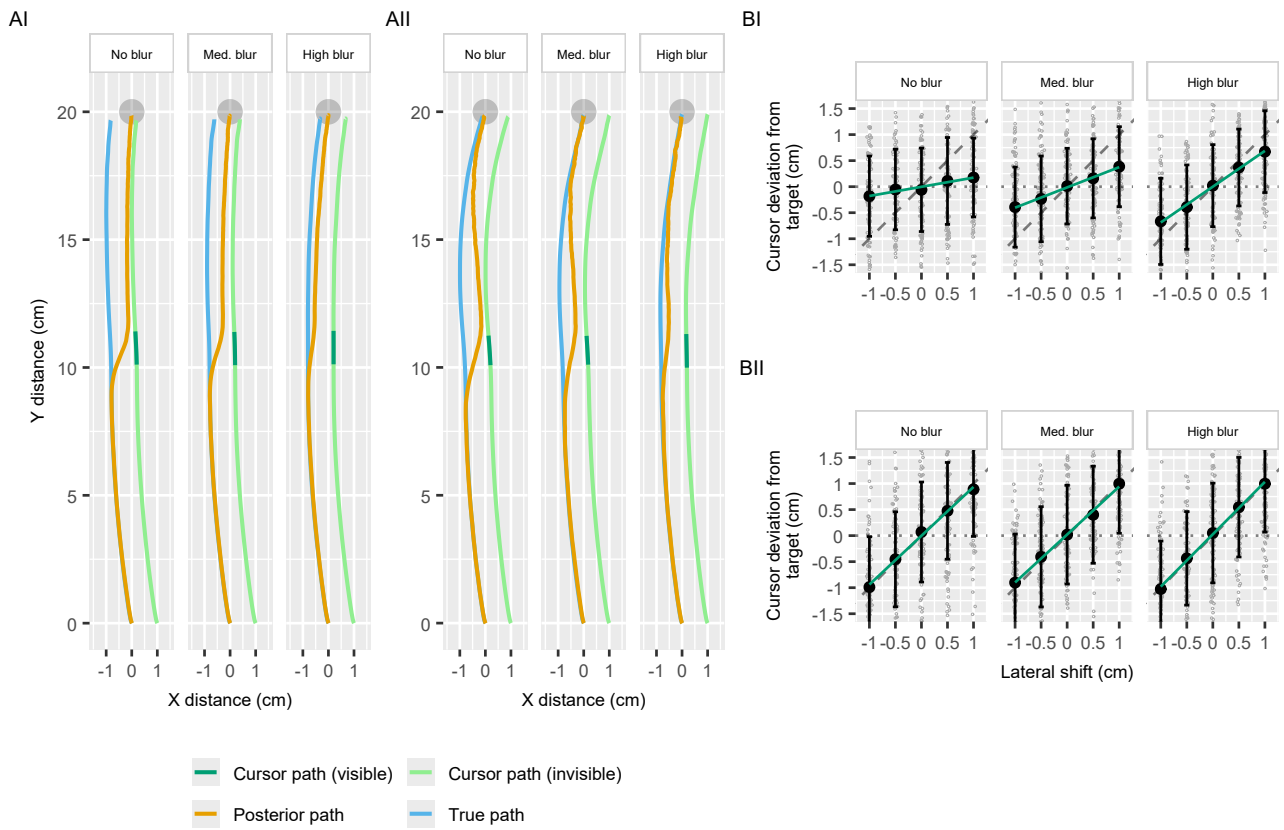

Figure 5: **Results of simulation 3.** Simulation results of a setup inspired by the experiment performed by Körding and Wolpert (2004). Panel AI shows an Agent with proprioceptive high-precision velocity feedback and low-precision position feedback, while panel AII shows one with proprioceptive low-precision velocity feedback and high-precision position feedback. Both AI and AII display a condition where a 1 cm visual offset to the right is applied to the cursor, which remains visible for only 100 ms after passing the 10 cm mark. Three sub-panels are shown, one each for no visual blur, medium visual blur, and high visual blur. Plotted movement paths represent the means across 500 simulated trials per condition. BI and BII each correspond to AI and AII conditions, respectively, and show the endpoint cursor deviation (X distance) from the target plotted for five different lateral shifts of the cursor (negative offsets are to the left). Small grey dots are individual trials (subset of 100 per condition is shown), large black dots and error bars are mean +/- SD, and green lines are linear ordinary least squares regression lines. The diagonal dashed line corresponds to no endpoint correction by visual feedback, while the horizontal dotted line corresponds to full correction (no cursor deviation from the target). See the Methods section for details of the simulation approach.

blur condition.

After visual feedback is again removed, the simulation results show substantial differences in how the posterior state, and consequently the performed movement, is affected, depending on the configuration of the proprioceptive feedback. The HVLP Agent completes the movement with essentially the same constant offset (Fig. 5AI), due to its reliance on high-precision velocity-based proprioceptive feedback. This configuration of proprioceptive feedback allows for high-precision estimation of changes in joint angles, while only a weak signal is available for sensing absolute joint angles. The visual-induced offset in perceived joint angles is therefore mostly kept constant over the second half of the movement. This produces a pattern of endpoint errors as a function of lateral offset and visual precision (Fig. 5 BI) that closely matches those observed by Körding and Wolpert (2004). The LVHP Agent, on the other hand, has relatively precise position sensors, while velocity feedback and, therefore, changes in joint angles, are only weakly contributing to the perceived angular positions. This has the effect of "pulling" the posterior arm state back towards the true arm state during the second half of the movement. When the target is reached, the perceptual offset induced by the visual feedback is almost washed out (Fig. 5 BII), a pattern of results that is substantially different from those reported by Körding and Wolpert (2004).

The endpoint precision of the two Agent configurations is approximately matched, with a lateral error distri-

bution of 0.71 cm SD for the HVLP Agent and 0.77 cm SD for the LVHP Agent (for the baseline condition with no blur and no lateral offset).

4 Discussion

Theories based on the concepts of Bayesian sensory integration and state estimation have become central to the field of sensorimotor control over the last few decades (Todorov & Jordan, 2002; Knill & Pouget, 2004; Friston, 2010; Franklin & Wolpert, 2011). Here, we set out to investigate the impact of two contrasting configurations of proprioceptive feedback in the context of Bayesian state estimation: one that relies primarily on velocity-based feedback (high-precision velocity, low-precision position: HVLP) and one that primarily relies on position-based feedback (low-precision velocity, high-precision position: LVHP). Including proprioceptive sensors beyond purely position-based inputs, and specifying their precision, is an important modelling choice that is often omitted for the sake of simplicity and readability. For example, the computational models of Active Inference presented by Friston et al. (2010) is one of the most cited such articles in the field and utilises a position-only sense of proprioception for their oculomotor and arm-reaching examples. While Todorov and Jordan (2002) do include velocity sensors in their simulation model, which do augment the positional estimation of their simulated Agent to some extent, this function is not discussed.

In Simulation 1, we demonstrate that both the HVLP and LVHP proprioceptive feedback configurations can result in human-like adjustments in movement speed and endpoint errors when the effects of antagonist muscle vibration are simulated. Both of these configurations predict the observations made by Cody et al. (1990), showing progressively greater undershooting of a target location with increasing intensity and duration of the simulated vibration, with corresponding slowing of the movement.

These responses depended on how muscle vibration was implemented; for the HVLP configuration, vibration was implemented as a joint velocity sampling bias, whereas for the LVHP configuration, it was implemented as a joint position sampling bias. Muscle/tendon vibration acts to some degree on both type II and type Ia muscle spindle afferents, each primarily signalling position and velocity values, though with some overlap, with the effect on each driven by both vibration frequency and amplitude (Burke et al., 1976; Roll & Vedel, 1982; Roll et al., 1989). Any effects observed in a given muscle vibration study must at least be considered to be partly caused by biasing of a positional and/or velocity feedback signal. While we simulate and show the individual effects of either a position or velocity offset, it seems likely that actual muscle vibration would cause both at the same time; implementing such dual-action in the two presented proprioceptive configurations would amplify the effects slightly, with the majority of the effect still mediated by the most precise of the two modes of proprioceptive feedback. Our results highlight the need for careful design and interpretation when utilising muscle vibration to investigate the sensori-motor system, as fundamentally different effect pathways are here predicted to cause very similar observed patterns of behaviour in the context of Bayesian sensory integration. Additionally, researchers using similar simulation models should carefully consider their implementation of proprioceptive feedback, as both velocity-based and position-based implementations of proprioception can produce human-like behaviour under simulated interventions such as these. While the addition of either feedback stream can augment the inference of the other, there are fundamental differences between whether information enters the inference process as rate-of-change or absolute values.

Similarly, Simulation 2 shows negligible differences in the simulated results between the HVLP and LVHP proprioceptive configurations, as visual feedback of the hand location tends to dominate the posterior state estimate in both cases. Fourneret and Jeannerod found that most participants in their study reported experiencing movement in the same direction as the offset visual feedback, rather than in the direction corresponding to the actual movement of their hand as they corrected for the offset visual feedback. For both configurations in our simulation, the posterior state estimate is updated in much the same way when visual feedback is rotated, leading to similar motor corrections. If we consider the posterior state estimate from the here-presented simulations as what might be available for conscious perception and report, then our simulations would predict the same pattern of reports. While both proprioceptive feedback and efferent output would tend to pull the posterior state

estimate in the same direction as the actual movement of the hand, the high precision of visual feedback used in our simulation results in the posterior path of the hand being very close to the path of the rotated visual cursor. It has generally been found that in instances with conflicting proprioceptive and visual feedback, only a single, combined state estimate is experienced, rather than one based on visual feedback and one based on proprioceptive feedback; when available, visual feedback tends to dominate the posterior state estimation (Lackner & Taublieb, 1984; Körding & Wolpert, 2004; Chancel, Blanchard, et al., 2016; Chancel, Brun, et al., 2016; Chancel & Ehrsson, 2023). In the case of our simulations, the initial deviation caused by biased visual feedback drives the correction to ensure the target is reached. Both visual and proprioceptive feedback and efferent motor output contribute to the final posterior state estimate, while the relatively high precision of visual feedback ensures that the posterior perceived arm state stays close to that indicated by the visual feedback. The motor corrections ensure that the perceived path deviates only slightly from the straight-line path.

In contrast, Simulation 3 reveals fundamental differences in the behaviour predicted by velocity-dominated proprioceptive feedback compared to position-dominated proprioceptive feedback. Here we attempted to replicate the setup presented by Körding and Wolpert (2004). They elegantly demonstrated how such briefly presented visual feedback could induce an offset in the perceived position of the hand, as indicated by the effects on the endpoint lateral errors. By further manipulating the precision of the visual feedback, they illustrated how such visual feedback appears to be integrated in a Bayes-optimal fashion, with increasingly blurred visual feedback inducing smaller offsets to the perceived position of the hand. We can additionally conclude that the induced offset in the inferred position of the hand following the integration of offset visual feedback is at least somewhat temporally stable; after visual feedback is removed again, the executed movement appears to be made relative to what the inferred position was when visual feedback was removed. This would tend to indicate that whatever the precision of proprioceptive positional feedback is, it is in this case insufficient to effectively correct for the induced offset over the second half of the reaching movement; likewise, we can also infer that the positional precision of unblurred visual feedback is high enough that even 100 ms is sufficient to allow for close to full correction of the movement trajectory.

Our simulations of the HVLP proprioceptive configuration explain this pattern of relative movement through its heavy reliance on velocity-based proprioceptive feedback. This setup leads to movements that are effectively performed relative to the state offset induced by the offset visual feedback, as the high-precision velocity feedback allows high-precision estimation of changes in joint angles. Once the offset is induced through the biased visual feedback, it is effectively maintained; the relatively low-precision positional feedback is insufficient to correct for the induced offset following removal of visual feedback. In contrast, the high-precision positional feedback of the LVHP configuration is sufficient for the induced state offset to quickly drift back towards the true state of the arm after the biasing visual feedback is removed; by the time the target is reached, the state estimate is no longer affected by the out-of-date offset visual feedback.

We further manipulated the Agent-perceived precision of the visual feedback in order to approximate the usage of blurred visual feedback employed by Körding and Wolpert. While this approach of manipulating the Agent-perceived precision of visual feedback is necessarily a simplification of the method used by Körding and Wolpert to reduce reliance on visual feedback, Limanovski and Friston (2020) have previously employed a similar simulation approach to affect the relative attendance to proprioceptive and visual cues, which resulted in a good match to their empirical observations; we use it here to decrease the attendance to visual feedback, indicating the belief that blurred visual input is trusted less. This simulation approach produces results that closely resemble the observations by Körding and Wolpert (2004), as well as a more recent replication (Hewitson et al., 2018). Here, our simulation further reveals that decreasing the perceived precision of the visual feedback reduces the rate of information integration from the visual source. This duration dependence of information is central to the formulation of our simulation model; being exposed to a given sensory channel for longer is here inherently more informative. Such duration dependence of visual precision has previously been indicated in terms of the localisation accuracy of visual stimuli at various durations (Adam et al., 1993, 2007; Zimmermann et al., 2013); the precision of visual stimulus localisation steadily increases over time, plateauing at around 300-500 ms. A future study investigating the interaction effects of visual blur and exposure duration during a motor task could

further explore this point, as our model predicts that a longer exposure time with a given visual blur can produce the same endpoint error deviation as a lesser blur, shown for a shorter duration.

Here, it is worth noting that the endpoint deviation in this simulated task also depends on the Agent's movement speed. The results presented here are based on simulated reaches with a duration of 2 s. While Körding and Wolpert do not report the movement duration characteristics in their study, 2 s may well be longer than the average duration in their study. Simulating the same task with a higher movement speed would tend to decrease the differences between the HVLP and LVHP configuration as they are here defined: partly, visual feedback would remain available until the hand is closer to the target, and partly, there would be less time to accumulate proprioceptive sensory evidence after visual feedback is removed and before the target is reached. As such, the LVHP configuration could also match the results reported by Körding and Wolpert quite well, given different assumptions about the utilised movement speed.

We employed what might turn out to be fairly extreme differences in the relative reliance on position- and velocity-based proprioceptive feedback, in order to highlight the different movement characteristics that would be expected of each. It seems likely that reality reflects a more even mixture, which may further depend upon the task demands. It has, for example, been indicated that it is possible to attend differentially to position- and velocity-based proprioceptive cues (Sittig et al., 1985), and similarly to focus either on visual or proprioceptive feedback (Limanowski & Friston, 2020), both affecting the perceived state of the arm.

These simulation results should therefore not be taken as strong evidence in favour of the human proprioceptive system more closely matching the HVLP configuration. However, studies using a longer series of subsequent movements without visual feedback have demonstrated that while absolute proprioceptive position estimation is prone to substantial drift over time, the execution of individual relative movements continues to be well-formed, with minimal variation in movement length and direction (Brown et al., 2003a, 2003b; Smeets et al., 2006; Patterson et al., 2017). Similarly, viewing the hand start location before a reaching movement increases the endpoint precision (Elliott & Calvert, 1990; Rossetti et al., 1994; Vindras et al., 1998). Such results would tend to support the notion that relative movements are well represented by the nervous system, whereas precisely estimating absolute end-effector position appears to be a substantial challenge as the time since the last visual feedback increases.

Future experiments that combine the approach of Körding and Wolpert (2004) with long-duration movement tasks (e.g. Brown et al., 2003b; Smeets et al., 2006; Patterson et al., 2017) would seem well positioned to further probe the relative reliance on rate-of-change versus absolute position signals. Such experiments could examine the time course of endpoint errors following brief initial offset visual feedback; quickly discarding visual offsets would tend to indicate minimal reliance on velocity-based feedback, while retaining such an induced offset for longer durations would support the notion of substantial reliance on velocity-based feedback for state inference.

4.1 Limitations

A key limitation of this simulation setup is that we do not include any delay of afferent feedback and efferent output; all sensory samples are generated in real time during each time step, and are immediately available for inference, and the computed torque is likewise immediately applied to the arm. This approach was chosen to limit the complexity of the model, but it has some implications which are worth considering. Much like a velocity estimate may be used to predict changes in position, an estimate of acceleration due to intended motor torque and expectation of friction is likewise useful for predicting changes to both velocity and position. This is relevant, as an extension to the model to include afferent and efferent delays would make the model considerably more reliant on its forward predictions. For example, if a 50 ms delay was added to both proprioceptive feedback and efferent motor commands, then the filter would be required to perform a 100 ms forward prediction, rather than the current single step of 10 ms, in order to bridge the temporal gap between the out-of-date sensory feedback and the point in the future when the next efferent command can affect the arm. This would affect the relative reliance on sensory feedback and efferent output for state estimation, and would necessitate different tuning of the sensory and/or motor noise parameters. Similarly, it would, for example, be very reasonable to include force-

based measurements, analogous to what might be signalled by Golgi tendon organ afferents, which would add additional state-dependent information, again requiring different tuning of sensory/motor precision to achieve similar behaviour as in the current implementation. This again highlights that the specific values chosen for the simulations here are necessarily very specific to the model implementation presented here; there are many valid modelling choices that would require different parameter tuning. This means that while the parameters are to some extent biologically interpretable, their interpretation should be done exclusively in context.

Instead, we aim to highlight how the relative reliance on position- and velocity-based proprioceptive feedback alters the pattern of inference and the resulting motor behaviour.

5 Conclusion

In summary, we presented an implementation of a computational model of sensory integration during active movements of a 2D planar arm, building on the principles of recursive Bayesian sensory integration. We exemplify the behaviour of the simulated Agent by comparing it with the reported results of three classical studies of sensorimotor control, demonstrating that the model can produce human-like motor behaviour under various visual and proprioceptive interventions.

With this model, we particularly highlight the importance of carefully considering the type and relative precision of the various channels of sensory feedback that a computational model is given access to, as different sets of assumptions can produce similar expected behaviour, matching human data, but requiring different interpretations. In Simulation 1, we demonstrate how different assumed effects of muscle/tendon vibration can produce very similar observed movement behaviour, matching empirical data, depending on the assumptions of the relative precision of position- and velocity-based proprioceptive feedback. This emphasises the care required when interpreting human data from such experiments. In Simulation 2, we show how the posterior arm state evolves in response to continuously displayed biased visual feedback, and discuss how the posterior of the Bayesian sensory integration process fits with the reported observations by Fourneret and Jeannerod (1998). In Simulation 3, we present how a setup using briefly presented offset visual feedback shows significant differences in the expected distribution of endpoint errors, depending on the relative weighting of position- and velocity-based proprioceptive cues. In particular, we highlight how high-precision velocity feedback allows for perception and execution of relative movements, as is seen in the study by Körding and Wolpert (2004), whose experimental setup was the basis for this simulation.

In conclusion, our modelling highlights that quite different assumptions of the type of feedback received by the proprioceptive system can, in some instances, produce very similar predicted behaviour, while in other situations, quite distinct patterns of results are produced, depending on the relative reliance on position- and velocity-based feedback. While our simulated results tend to support the contribution of velocity-based proprioceptive feedback for positional state inference, further research is needed to corroborate this claim.

6 Methods

The simulation was implemented in `Python` v3.11.5 (Python Software Foundation, 2023), using the `Numpy` v2.3.1 and `SciPy` v1.13.1 packages. All plotting of the results was done in R v4.5.1 (R Core Team, 2024), making use of the `tidyverse` v2.0.0 (Wickham et al., 2019).

6.0.1 Physics simulation

The arm is set up to move in the horizontal plane and should be considered as supported by a frictionless tabletop, such that gravity is not acting on any joints. The state of the arm is defined by $\mathbf{x}^{\text{true}}$ (Eq. P.1), consisting of the angular positions and velocities of the shoulder $(q_1, \dot{q}_1)$ and elbow $(q_2, \dot{q}_2)$ joints. These are related to the position of the hand as described by the forward kinematic mapping $\mathbf{P}_h(\mathbf{q})$ (Eq. P.2). The movement of the arm is simulated as dynamically coupled through the mass term $\mathbf{M}(q)$ and the Coriolis/centrifugal term $\mathbf{C}(q, \dot{q})$, with

Physics simulation

State and torque vectors

$$\mathbf{x}_k^{\text{true}} = [q_1, q_2, \dot{q}_1, \dot{q}_2]^\top, \quad \boldsymbol{\tau}_k = [\tau_1, \tau_2]^\top \tag{P.1}$$

Forward kinematic mapping

$$\mathbf{p}_h(\mathbf{q}) = \mathbf{p}_{\text{shoulder}} + \begin{bmatrix} L_1 \cos(q_1) + L_2 \cos(q_1 + q_2) \\ L_1 \sin(q_1) + L_2 \sin(q_1 + q_2) \end{bmatrix} \tag{P.2}$$

Dynamic joint coupling

$$\mathbf{M}(\mathbf{q})\,\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{D}\,\dot{\mathbf{q}} = \boldsymbol{\tau}, \quad \mathbf{D} = \text{diag}(d_1, d_2) \tag{P.3}$$

$$\mathbf{M}(\mathbf{q}) = \begin{bmatrix} m_1 l_{c1}^2 + m_2\big(L_1^2 + l_{c2}^2 + 2L_1 l_{c2} \cos q_2\big) + I_1 + I_2 & m_2\big(l_{c2}^2 + L_1 l_{c2} \cos q_2\big) + I_2 \\ m_2\big(l_{c2}^2 + L_1 l_{c2} \cos q_2\big) + I_2 & m_2 l_{c2}^2 + I_2 \end{bmatrix} \tag{P.4}$$

$$h(q_2) = -\,m_2 L_1 l_{c2}\, \sin q_2 \tag{P.5}$$

$$\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) = \begin{bmatrix} h(q_2)\,\big(2\dot{q}_1 \dot{q}_2 + \dot{q}_2^2\big) \\ h(q_2)\,\big(-\,\dot{q}_1^2\big) \end{bmatrix} \tag{P.6}$$

Discretisation

$$\dot{\mathbf{q}}_k = \dot{\mathbf{q}}_{k-1} + \ddot{\mathbf{q}}_{k-1}\,\Delta t \tag{P.7}$$

$$\mathbf{q}_k = \mathbf{q}_{k-1} + \dot{\mathbf{q}}_k\,\Delta t \tag{P.8}$$

Figure 6: **Physics simulation.** The true state vector, $\mathbf{x}_k^{\text{true}}$, represents the angular position and velocity of the elbow and shoulder joint, while $\boldsymbol{\tau}_k$ is the (noisy) muscle torque acting on the joints (Eq. P.1). The forward kinematic model $\mathbf{p}_h(\mathbf{q})$ represents the mapping from joint space to planar Cartesian space where visual feedback is received (Eq. P.2). The dynamic coupling between the joints is governed by the Mass ($\mathbf{M}$) and Coriolis/Centrifugal ($\mathbf{C}$) terms, with angular acceleration further affected by velocity-dependent viscous damping, $\mathbf{D}$ (P.3-6). These dynamics are discretised as semi-implicit Euler integration (Eq. P.7-8).

the generated muscle torques $\boldsymbol{\tau}$ acting on the joints (Eq. P.3-6). Additionally, friction is simulated as viscous damping in the joints with the variable $\mathbf{D}$, leading to exponential decay of angular velocity in the absence of applied muscle torques. The dynamics are discretised for simulation using semi-implicit Euler integration (Eq. P.7-8).

6.0.2 Measurement model and motor noise

At each time step k, the Agent receives a set of sensory observations $\mathbf{z}_k^{\text{obs}}$. These are based on the direct state-sensory mapping, $h(\mathbf{x}_k^{\text{true}})$, with two added terms (Eq. N.1). Partly, a time-varying bias may be added to simulate the various experimental interventions through $\mathbf{b}_k$, e.g., to add a visual offset to the returned visual feedback. Additionally, Gaussian noise is added, given by $\mathbf{v}_k$, with the noise of each sensor implemented as time-invariant and independent, given by $\mathbf{R}^{\text{true}}$ (Eq. N.1-2).

Motor noise is added to the efferent motor command ($\tau_{k-1}^{\text{efferent}}$) in a signal-dependent manner (Harris & Wolpert, 1998), scaling proportionally with the efferent torque (Eq. N.4-6).

Both sensory and motor noise are implemented as continuous-time white noise with a specified power spectral density, which is then discretised for sampling during simulation. This implementation means that noise injection per unit time is independent of the choice of Δt.

During forward prediction to form the prior for the current step, the Agent relies upon the efferent motor command and the previous state posterior estimate. Any uncertainty in the applied torque (Eq. N.7 leads

Measurement and motor noise

Sensory model and noise

$$\mathbf{z}_k^{\text{obs}} = \underbrace{\begin{bmatrix} \mathbf{p}_h(\mathbf{q}_k) \\ q_{1,k} \\ q_{2,k} \\ \dot{q}_{1,k} \\ \dot{q}_{2,k} \end{bmatrix}}_{h(\mathbf{x}_k^{\text{true}})} + \mathbf{b}_k + \mathbf{v}_k, \qquad \mathbf{b}_k \in \mathbb{R}^6, \qquad \mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}^{\text{true}}) \tag{N.1}$$

$$\mathbf{R}^{\text{agent}} = \text{diag}\Big((\tfrac{\sigma_{\text{vis}}^{\text{agent}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{\text{vis}}^{\text{agent}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{q}^{\text{agent}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{q}^{\text{agent}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{\dot{q}}^{\text{agent}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{\dot{q}}^{\text{agent}}}{\sqrt{\Delta t}})^2 \Big) \tag{N.2}$$

$$\mathbf{R}^{\text{true}} = \text{diag}\Big((\tfrac{\sigma_{\text{vis}}^{\text{true}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{\text{vis}}^{\text{true}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{q}^{\text{true}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{q}^{\text{true}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{\dot{q}}^{\text{true}}}{\sqrt{\Delta t}})^2, (\tfrac{\sigma_{\dot{q}}^{\text{true}}}{\sqrt{\Delta t}})^2 \Big) \tag{N.3}$$

Motor noise

$$\sigma_{\tau,k-1}^{\text{motor}} = \sigma_{\tau,\text{prop}} \, |\tau_{k-1}^{\text{efferent}}| \tag{N.4}$$

$$\epsilon_{k-1} \sim \mathcal{N}\Big(0, \, (\frac{\sigma_{\tau,k-1}^{\text{motor}}}{\sqrt{\Delta t}})^2\Big) \tag{N.5}$$

$$\tau_k = \tau_{k-1}^{\text{efferent}} + \epsilon_{k-1} \tag{N.6}$$

Agent-assumed process noise

$$Q_{c,\tau,k-1} = \big(\sigma_{\tau,\text{const}} + \sigma_{\tau,k-1}^{\text{motor}}\big)^2 \tag{N.7}$$

$$Q_{c,\alpha,k-1}\big(\mathbf{q}_{k-1}^{\text{post}}\big) = \frac{Q_{c,\tau,k-1}}{I\big(\mathbf{q}_{k-1}^{\text{post}}\big)^2}, \qquad I(\mathbf{q}_{k-i}^{\text{post}}) = [\mathbf{M}(\mathbf{q}_{k-i}^{\text{post}})]_{ii} \tag{N.8}$$

$$\text{Var}(\delta q)_k = \tfrac{1}{3} \, Q_{c,\alpha,k-1} \, \Delta t^3, \tag{N.9}$$

$$\text{Var}(\delta \dot{q})_k = Q_{c,\alpha,k-1} \, \Delta t, \tag{N.10}$$

$$\text{Cov}(\delta q, \, \delta \dot{q})_k = \tfrac{1}{2} \, Q_{c,\alpha,k-1} \, \Delta t^2 \tag{N.11}$$

$$\mathbf{Q}_k = \begin{bmatrix} \frac{1}{3} \, Q_{c,\alpha,k-1} \, \Delta t^3 & \frac{1}{2} \, Q_{c,\alpha,k-1} \, \Delta t^2 \\ \frac{1}{2} \, Q_{c,\alpha,k-1} \, \Delta t^2 & Q_{c,\alpha,k-1} \, \Delta t \end{bmatrix} \tag{N.12}$$

Figure 7: **Measurement and motor noise.** $h(\mathbf{x}_k)$ defines the state-sensor mapping, which together with added Gaussian noise $\mathbf{V}_k$ produces the observed sensory samples $\mathbf{z}_k^{\text{obs}}$ (Eq. N.1). $\mathbf{b}_k \in \mathbb{R}^6$ is an additive measurement bias (e.g., a visual Cartesian offset); the default is $\mathbf{b}_k = 0$. Sensory noise amplitude $\mathbf{R}^{\text{true}}$ is implemented as continuous time white noise with fixed densities; this formulation allows for the specification of noise-injection that is independent of the choice of Δt. $\mathbf{R}^{\text{agent}}$ is the assumed sensory noise used by the Agent for inference; by default, $\mathbf{R}^{\text{agent}} = \mathbf{R}^{\text{true}}$. Signal-dependent motor noise is added to the efferent command, and is likewise implemented as white noise (Eq. N.4-6). Due to how the loop is ordered (see Fig. 2), the torque to be applied at t_k is determined at t_{k-1}. The Agent-assumed process noise accounts for the expected motor noise and further adds a constant (Eq. N.7), which represents the extent to which the Agent believes that unanticipated forces may act on the arm, so as not to be overconfident in its predictions. Eq. N.8 shows how continuous time torque noise is expected to be transformed to acceleration noise; note that $\mathbf{q}_{k-1}^{\text{post}}$ is extracted from $\mathbf{x}_{k-1}^{\text{post}}$. Eq. N.9-12 show how the process noise is discretized. Process noise is treated identically for both joints, with no specified cross-joint correlation.

to uncertainty in the angular accelerations (Eq. N.8) applied to each joint, ultimately resulting in increased uncertainty in the angular positions and velocities (Eq. N.12). The Agent-assumed process noise matches the signal-dependent noise, reflecting the true motor noise, with an additional constant term ($\sigma_{\tau,const}$) meant to represent any unexpected external disturbances. Setting $\sigma_{\tau,const}$ to zero would result in the Agent perfectly representing the true noise magnitude affecting torques at each joint, and therefore making it highly confident in its predicted priors, while increasing it above zero increases the uncertainty of its predictions. This effectively means it will increase the relative weight of sensory feedback. In a simulation such as this, the actual injected torque noise can be controlled and matched perfectly by the Agent if programmed as such. However, it seems likely that a human performing a motor task should remain somewhat sceptical of their forward predictions from

efferent commands or predicted sensory effects of actions to allow for unexpected external perturbations. This is a simple way to implement this effect in a tunable manner. However, it should be noted that setting the value of this parameter (and other true and agent-assumed noise parameters) is inherently somewhat arbitrary and can significantly affect the Agent's behaviour and performance.

6.0.3 State estimation

The agent estimates its arm state in two steps each loop. First, it forms a prior by propagating the previous posterior through its internal dynamics model together with the previous efferent motor command (Eq. K.7). The covariance $\mathbf{P}_k^{\text{prior}}$ includes the agent-assumed process noise $\mathbf{Q}_k$ (Eq. N.12). $\mathbf{P}_k^{\text{prior}}$ quantifies both marginal uncertainty of each element of $\mathbf{x}_k^{\text{prior}}$ and their expected correlations (e.g., elbow angle and angular velocity may be correlated, so underestimating angle tends to co-occur with underestimating velocity during a movement that increases the angle, and so on).

From this prior state estimate, predictions of the expected sensory input are formed (Eq. K.10-13). These are based on the state-sensory mapping ($h(\mathbf{x})$), producing both visual (via the forward kinematic mapping, Eq. P.2) and proprioceptive expected observations, given the expected noise of each sensor, given by $\mathbf{R}^{\text{agent}}$. Note that the Agent-assumed sensory noise ($\mathbf{R}^{\text{agent}}$) is equal to the true sensory noise $\mathbf{R}^{\text{true}}$, except for Simulation 3 for the application of visual blur.

The second main step consists of integrating new sensory observations, $\mathbf{z}_k^{\text{obs}}$. How much new measurements are trusted compared to the predicted state is governed by the Kalman gain $\mathbf{K}_k$; when sensory observations are expected to have low noise and correlate well with the true state, then observations are weighted heavily during the update, while high noise sensory feedback will have a smaller impact compared to the prior (Eq. K.14-16).

For a simple 1D example, this corresponds to precision-weighted averaging of the prior and the likelihood (a sensory observation) such that posterior μ and J (precision, the inverse of variance) would be given by:

$$\mu^{post} = \frac{J^{prior}\mu^{prior} + J^{likelihood}\mu^{likelihood}}{J^{prior} + J^{likelihood}}, \quad J^{post} = J^{prior} + J^{likelihood}, \quad J = \frac{1}{\sigma^2}$$

Here $\mu^{likelihood}$ is a sensory identity-measurement, while $J^{likelihood}$ is the expected precision of this measurement; these would each correspond to $\mathbf{z}_k^{obs}$ and $\mathbf{R}^{agent}$.

A key emergent feature of forward prediction is that as long as there is any uncertainty regarding velocity, and how it is changing, then uncertainty about both position and velocity estimates will increase; the Agent will be more unsure about the prior for the next time step than it was about its posterior on the previous step. Conversely, when new information is added in the form of noisy sensory feedback, certainty again increases from the prior to the posterior. Across time, the posterior state certainty will approach some equilibrium where the uncertainty gained from the forward prediction is balanced by the information contained in the incoming sensory feedback. This equilibrium will depend on the precision of the available sensory information and the uncertainty associated with the state evolution. For example, in the case of the model here, the precision of proprioceptive positional feedback puts an effective upper limit on how uncertain the Agent can ever be about its position; as seen in the 1D example above, precisions sum, such that the posterior will always have a greater precision than either the likelihood or the prior.

6.0.4 Motor control

Once the agent has estimated its current state, it computes joint torques to drive the hand toward a target using a Linear Quadratic Regulator (LQR) (Athans, 1971). At each step, we locally linearise the arm about the current estimate ($\mathbf{x}_k^{\text{post}}$) to obtain a discrete model (Eq. L.1-4) and minimise a quadratic cost (Eq. L.5). The finite-horizon problem is solved via the discrete-time backward Riccati recursion to produce feedback gains $\mathbf{K}_t^{\text{LQR}}$ (Eq. L.6-7). We then return the first torque $\boldsymbol{\tau}_0^{\textbf{efferent}}$ (Eq. L.9), enforcing per-joint rate of force development (Eq. L.10) and torque limits (Eq. L.11). This is repeated at each step in a receding-horizon fashion. To guide the intended movement time, we specify a time-to-target T_{plan} such that the planned horizon is $N = \frac{T_{plan}}{\Delta t}$, with

Unscented Kalman Filter

Sigma point distribution and weighting:

$$\lambda = \alpha^2(L+\kappa) - L, \qquad \gamma = \sqrt{L+\lambda} \tag{K.1}$$

$$\mathbf{W}_0^{(m)} = \frac{\lambda}{L+\lambda}, \qquad \mathbf{W}_0^{(c)} = \frac{\lambda}{L+\lambda} + (1-\alpha^2+\beta) \tag{K.2}$$

$$\mathbf{W}_i^{(m)} = \frac{1}{2(L+\lambda)}, \qquad \mathbf{W}_i^{(c)} = \frac{1}{2(L+\lambda)}, \quad i = 1,\dots,2L \tag{K.3}$$

Sigma points from $\mathbf{x}_{k-1}^{\text{post}}, \mathbf{P}_{k-1}^{\text{post}}$:

$$\mathbf{S} = \text{chol}\big(\mathbf{P}_{k-1}^{\text{post}}\big) \tag{K.4}$$

$$\mathbf{X}_{k-1}^{(0)} = \mathbf{x}_{k-1}^{\text{post}} \tag{K.5}$$

$$\mathbf{X}_{k-1}^{(i)} = \mathbf{x}_{k-1}^{\text{post}} + \gamma\,\mathbf{S}_{:i}, \quad \mathbf{X}_{k-1}^{(i+L)} = \mathbf{x}_{k-1}^{\text{post}} - \gamma\,\mathbf{S}_{:i} \quad (i = 1,\dots,L) \tag{K.6}$$

State prediction:

$$\mathbf{X}_k^{f,(i)} = f_{\text{disc}}\Big(\mathbf{X}_{k-1}^{(i)},\, \boldsymbol{\tau}_{k-1}^{\text{efferent}}\Big) \tag{K.7}$$

$$\mathbf{x}_k^{\text{prior}} = \sum_{i=0}^{2L} \mathbf{W}_i^{(m)}\,\mathbf{X}_k^{f,(i)} \tag{K.8}$$

$$\mathbf{P}_k^{\text{prior}} = \sum_{i=0}^{2L} \mathbf{W}_i^{(c)}\Big(\mathbf{X}_k^{f,(i)} - \mathbf{x}_k^{\text{prior}}\Big)(\cdot)^\top + \mathbf{Q}_k \tag{K.9}$$

Sensory prediction:

$$\mathbf{Z}_k^{(i)} = h\Big(\mathbf{X}_k^{f,(i)}\Big) \tag{K.10}$$

$$\bar{\mathbf{z}}_k = \sum_{i=0}^{2L} \mathbf{W}_i^{(m)}\,\mathbf{Z}_k^{(i)} \tag{K.11}$$

$$\mathbf{P}_{zz} = \sum_{i=0}^{2L} \mathbf{W}_i^{(c)}\Big(\mathbf{Z}_k^{(i)} - \bar{\mathbf{z}}_k\Big)(\cdot)^\top + \mathbf{R}^{\text{agent}} \tag{K.12}$$

$$\mathbf{P}_{xz} = \sum_{i=0}^{2L} \mathbf{W}_i^{(c)}\Big(\mathbf{X}_k^{f,(i)} - \mathbf{x}_k^{\text{prior}}\Big)\Big(\mathbf{Z}_k^{(i)} - \bar{\mathbf{z}}_k\Big)^\top \tag{K.13}$$

Update:

$$\mathbf{K}_k = \mathbf{P}_{xz}\,\mathbf{P}_{zz}^{-1} \tag{K.14}$$

$$\mathbf{x}_k^{\text{post}} = \mathbf{x}_k^{\text{prior}} + \mathbf{K}_k\big(\mathbf{z}_k^{\text{obs}} - \bar{\mathbf{z}}_k\big) \tag{K.15}$$

$$\mathbf{P}_k^{\text{post}} = \mathbf{P}_k^{\text{prior}} - \mathbf{K}_k\mathbf{P}_{zz}\mathbf{K}_k^\top \tag{K.16}$$

Figure 8: **Bayesian inference.** α, β and κ are UKF tuning parameters (Eq. K.1-3). γ is a factor used to control the spread of sigma points representing uncertainty $\mathbf{X}_{k-1}^{(i=1:9)}$ around the mean estimate sigma point $\mathbf{X}_{k-1}^{(i=0)}$ (Eq. K.4-6). L is the length of the state vector to be estimated, 4 in this case ($[q_1, q_2, \dot{q}_1, \dot{q}_2]$). A total of 9 ($2L+1$) sigma points are used here, one for the mean estimate and the remainder to represent uncertainty. Sigma points and efferent torque are passed through the discretized dynamics function (Eq. P.3-8) to produce the prior state estimate $\mathbf{x}_k^{\text{prior}}$ (Eq. K.8) and the associated uncertainty $\mathbf{P}_k^{\text{prior}}$, which is further affected by the Agent-assumed process noise (Eq. K.9). These together form the state prediction. From these sigma points, sensory predictions are formed through the state-sensor mapping (Eq. K.10-11). These, along with the measurement noise $\mathbf{R}^{agent}$, are used to predict the measurement covariance (Eq. K.12) and the state-measurement cross covariance (Eq. K.13). The ratio of these two matrices produces the Kalman gain $\mathbf{K}_k$, which scales the measurement innovation of the state (Eq. K.14-16).

T_{plan} reduced by Δt by each step, giving the remaining steps N_k at step k. The Agent then attempts to reach the target state ($\mathbf{x}^{target}$; target pose and zero velocity) through an inflated final state deviation weight $\mathbf{Q}_f^{\text{LQR}}$. $\mathbf{Q}^{\text{LQR}}$, $\mathbf{Q}_f^{\text{LQR}}$, and $\mathbf{R}^{\text{LQR}}$ are all tunable cost parameters (state and motor, respectively) that will shape the motor behaviour of Agent. Note that those three parameters, as well as $\mathbf{K}^{\text{LQR}}$, are not related to the previously mentioned process noise ($\mathbf{Q}_k$), sensory noise ($\mathbf{R}^{\text{true/agent}}$) or Kalman gain $\mathbf{K}$; they are reused here to match standard LQR notation.

Linear Quadratic Regulator

Locally linearised dynamics:

$$\dot{\mathbf{x}} = \mathbf{f}_c(\mathbf{x}, \boldsymbol{\tau}) = \begin{bmatrix} \dot{\mathbf{q}} \\ \mathbf{a}(\mathbf{q}, \dot{\mathbf{q}}, \boldsymbol{\tau}) \end{bmatrix} \tag{L.1}$$

$$\mathbf{a}(\mathbf{q}, \dot{\mathbf{q}}, \boldsymbol{\tau}) := \mathbf{M}(\mathbf{q})^{-1} \left(\boldsymbol{\tau} - \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) - \mathbf{D}\dot{\mathbf{q}}\right) \tag{L.2}$$

$$\mathbf{A}_d = \mathbf{I} + \Delta t \frac{\partial \mathbf{f}_c}{\partial \mathbf{x}}(\mathbf{x}_0, \mathbf{0}), \quad \mathbf{B}_d = \Delta t \frac{\partial \mathbf{f}_c}{\partial \boldsymbol{\tau}}(\mathbf{x}_0, \mathbf{0}) \tag{L.3}$$

$$\mathbf{x}_{t+1} = \mathbf{A}_d \mathbf{x}_t + \mathbf{B}_d \boldsymbol{\tau}_t, \quad \mathbf{x}_0 = \mathbf{x}_k^{\text{post}} \tag{L.4}$$

Quadratic cost and Riccati backward sweep:

$$J^{\text{LQR}} = \sum_{t=0}^{N_k - 1} \left(\tilde{\mathbf{x}}_t^\top \mathbf{Q}^{\text{LQR}} \tilde{\mathbf{x}}_t + \boldsymbol{\tau}_t^\top \mathbf{R}^{\text{LQR}} \boldsymbol{\tau}_t\right) + \tilde{\mathbf{x}}_{N_k}^\top \mathbf{Q}_f^{\text{LQR}} \tilde{\mathbf{x}}_{N_k} \tag{L.5}$$

$$\mathbf{K}_t^{\text{LQR}} = (\mathbf{R}^{\text{LQR}} + \mathbf{B}_d^\top \mathbf{P}_{t+1}^{\text{LQR}} \mathbf{B}_d)^{-1} \mathbf{B}_d^\top \mathbf{P}_{t+1}^{\text{LQR}} \mathbf{A}_d \tag{L.6}$$

$$\mathbf{P}_t^{\text{LQR}} = \mathbf{Q}^{\text{LQR}} + \mathbf{A}_d^\top \mathbf{P}_{t+1}^{\text{LQR}} (\mathbf{A}_d - \mathbf{B}_d \mathbf{K}_t), \quad \mathbf{P}_{N_k} = \mathbf{Q}_f^{\text{LQR}} \tag{L.7}$$

Receding horizon control at step k:

$$\tilde{\mathbf{x}}_0 = \mathbf{x}_k^{\text{post}} - \mathbf{x}_{\text{target}} \tag{L.8}$$

$$\boldsymbol{\tau}_0^{\text{ideal}} = -\mathbf{K}_0^{\text{LQR}} \tilde{\mathbf{x}}_0 \tag{L.9}$$

$$\Delta \boldsymbol{\tau}_k = \text{clip}\left(\boldsymbol{\tau}_0^{\text{ideal}} - \boldsymbol{\tau}_{k-1}^{\text{efferent}}, \pm \mathbf{rfd}_{\max}\right) \tag{L.10}$$

$$\boldsymbol{\tau}_k^{\text{efferent}} = \text{sat}\left(\boldsymbol{\tau}_{k-1} + \Delta \boldsymbol{\tau}_k, \pm \boldsymbol{\tau}_{\max}\right) \tag{L.11}$$

Figure 9: **Motor control.** The Agent controls its arm using a Linear Quadratic Regulator. At each time step k, we linearise the nonlinear dynamics around the current state estimate to obtain $\mathbf{A}_d$ and $\mathbf{B}_d$ (Eq. L.1-4), the discretised partial derivatives with respect to state and torque changes, respectively. Then the finite-horizon LQR is solved via the discrete-time backward Riccati recursion for N_k steps to obtain feedback gains $\mathbf{K}_t^{\text{LQR}}$ (Eq. L.5-7). N_k is the number of remaining steps at step k. We apply only the first computed torque in a receding-horizon manner (Eq. L.9). The weights $\mathbf{Q}^{\text{LQR}}$ and $\mathbf{R}^{\text{LQR}}$ are fixed, except for the final step, where the terminal weight $\mathbf{Q}_f^{\text{LQR}}$ increases the state penalty. Torques are constrained by per-step rate of force development limits (Eqs. L.10) and a hard saturation (Eq. L.11).

6.1 Model parameters and simulated tasks

As described in the preceding sections, the Agent relies on several parameters, the configuration of which affects various aspects of the physics simulation, state inference, and motor control. We will here provide an overview of those parameters which have some physical/biological interpretability, with all parameters defined in Appendix table S1. For the physical setup, the length of the upper arm was set to 0.33 m (L_1), and the length of the lower arm (and hand) to 0.45 m (L_2), with their respective masses set to 2 kg (m_1) and 1.7 kg (m_2). The shoulder joint is limited to movement between 30° horizontal extension and 135° flexion, while the elbow is limited to 0° extension and 170° flexion. For simplicity, both joints experience the same viscous damping of 1.5 N m s, which results in approximate angular velocity half-times of 0.25 s and 0.05 s respectively for the shoulder and elbow joints, meaning that the arm would very quickly slow down in the absence of applied muscle force (these are for elbow 45 ° flexed; shoulder half-time is posture-dependent). These length, mass, and friction values all represent

parameters that a human also would likely need to estimate and update as appropriate; for our simulations, we will assume that they are all constant and that the Agent knows and uses the correct values. The time step Δt is set to 10 ms for all simulations.

Visual noise was set to $1\,\mathrm{mm\,s^{-1/2}}$ per planar dimension, leading to per step sampling standard deviation of 10 mm. Proprioceptive sensory noise is implemented with identical noise densities for both joints. We test and compare two configurations of proprioceptive feedback with matched total Fisher information, but distributed differently across position and velocity sensors: one configuration relies mainly on velocity-based feedback, with relatively high-precision velocity feedback and low-precision position feedback (HVLP), and one that relies mainly on position-based feedback, with relatively low-precision velocity feedback and high-precision position feedback (LVHP). The HVLP configuration has a velocity noise of $0.015\,\mathrm{rad\,s^{-3/2}}$ (corresponding to $0.015\,\mathrm{rad/s/\sqrt{s}}$) and a position noise of $0.06\,\mathrm{rad\,s^{-1/2}}$. To transform these noise densities for per-step sampling, they are divided by $\sqrt{\Delta t}$; e.g. the above definition indicates that velocity noise here would have a standard deviation of $0.015\,\mathrm{rad\,s^{-1}}$ for intervals of 1 second, while per 10 ms step it is $0.15\,\mathrm{rad\,s^{-1}}$. The LVHP configuration has the two noise densities flipped, such that it has a velocity noise of $0.06\,\mathrm{rad\,s^{-3/2}}$ and a position noise of $0.015\,\mathrm{rad\,s^{-1/2}}$. These values were chosen to highlight both similarities and differences of velocity-dominated and position-dominated configurations of proprioceptive feedback in the chosen simulated tasks, while keeping the performance in terms of endpoint accuracy and precision similar in the non-perturbed control conditions. We address the sensibility of these choices in the discussion.

The signal-dependent motor noise scalar is set to $0.05\,\mathrm{s^{-1/2}}$, which leads to a standard deviation of the actual torque corresponding to 5 % of the efferent torque across a 1 second interval (per-step sampling standard deviation of 50 %, similar to the values used by Todorov and Jordan (2002)). Efferent torque values were limited by both max torque values (50 N m for the shoulder and 30 N m for the elbow), as well as rate of force development limits such that a minimum of 250 ms would be required to reach max torque (max Δ torque of 2 N m and 1.2 N m for the shoulder and elbow respectively per time step).

All simulations were initialized with the Agent perceived state matching the true state ($\mathbf{x}_{init}^{\mathrm{post}} = \mathbf{x}_{init}^{\mathrm{true}}$), and its state covariance matrix ($\mathbf{P}_{init}^{\mathrm{post}}$) set with sigmas of 1.0° and $1.0\,^{\circ}\,\mathrm{s^{-1}}$ respectively for position and velocity uncertainty for both joints (all non-diagonals initialized to 0).

6.1.1 Simulation 1

Here, we set up a simulation to match the experiment described by Cody et al. (1990). Their results, along with a range of related studies (e.g. Eklund, 1972; Goodwin et al., 1972; Roll and Vedel, 1982; Gilhodes et al., 1986; Inglis and Frank, 1990) demonstrate that the perceived limb pose and/or velocity may be biased through vibration of the muscles/tendons acting on the limb. The perceived limb is biased such that it corresponds to the vibrated muscle being stretched out further than it is, perceived as a positional and/or velocity offset. Both type II and Ia afferent firing rates in the passively lengthening muscle have been found to contain both position and velocity information, although the former emphasises position and the latter velocity. The Golgi tendon organ is a muscle tension sensor, and we have not included an equivalent (torque) sensor in our model.

In the simulated setup, the shoulder joint is locked in place, with a target requiring 30° of elbow extension to reach it. The allowed movement time is set to 1 s, matching the setup presented by Cody et al. (1990). We test and compare two proprioceptive configurations of the Agent as described in the previous section, one that relies primarily on velocity-based proprioceptive feedback (HVLP) and one that relies primarily on position-based proprioceptive feedback (LVHP). For each configuration, we test a different possible effect of muscle vibration, each targeting that configuration's dominant form of proprioceptive feedback (position/velocity). For the HVLP Agent, we test simulated muscle vibration implemented as a bias of sampled velocity, and for the LVHP Agent, we test it implemented as a position sampling bias. This is implemented through the relevant elbow position and velocity offset being applied to $\mathbf{b}_k$ (see Eq. N.1). We test low, medium, and high vibration intensities for each of the two implementations: a positional offset of -5 °, -10 °, or -15 °, and a velocity offset of -5 °/s, -10 °/s, or -15 °/s, respectively. Additionally, we test the medium vibration intensity with onset at 0 °, 10 °, and 20 ° into the movement. The direction of the applied bias is set to match the effect of antagonist vibration, such that it

produces a signal that implies greater-than-actual elbow extension/extension velocity.

6.1.2 Simulation 2

In Simulation 2, we base our setup on the study by Fourneret and Jeannerod (1998). Subjects were tasked with making forward reaches towards a target at a distance 22.2 cm with continuous cursor feedback; upon passing the 5.7 cm mark, a slight visual rotation of 0-10 ° to either side was applied to the cursor. This caused corrections to the reaching movements in the opposite direction, such that the cursor was correctly guided to the target location; after each trial, subjects were asked to report their actual movement direction by choosing from an array of arrows. Here, the majority of participants indicated the direction that corresponded to the observed rotated visual feedback, rather than the direction their hand travelled as they compensated for the incorrect visual feedback, and the reported movement direction was always less than the actual induced offset. Our goal with this simulation is to expand on their findings by considering the posterior of the proposed Bayesian process as what might be available for conscious report, rather than the raw data originating from either visual, proprioceptive, or efferent motor command sources.

We match the experimental setup by applying a visual rotation after the hand passes 5.7 cm, which is implemented through an offset to $\mathbf{b}_k$ dependent on the hand's position. Trials are simulated with a target duration of 2.5 s, which approximately matches the reach durations reported by Fourneret and Jeannerod (1998).

6.1.3 Simulation 3

In Simulation 3, we test and compare various levels of reduced visual precision for the HVLP and LVHP proprioceptive configurations, to show how each responds to offset visual feedback and how their perception continues to evolve after visual feedback is removed. Visual blur is implemented through increasing Agent-assumed visual noise in $\mathbf{R}^{\text{agent}}$; this effectively means that the Agent treats blurred visual observations as less informative, decreasing the visual Kalman gain and therefore relying relatively more on its state prior (and through it efferent motor commands) and proprioceptive feedback (see Eq. K.12-16). Specifically, we test for a "No blur" condition where $\sigma_{\text{vis}}^{\text{agent}}$ matches the true visual noise of $1.0\,\text{mm}\,\text{s}^{-1/2}$, a "Medium blur" condition where $\sigma_{\text{vis}}^{\text{agent}}$ is increased to $1.5\,\text{mm}\,\text{s}^{-1/2}$, and a "Large blur" condition where it is increased to $3.0\,\text{mm}\,\text{s}^{-1/2}$.

While the original experiment included both a training phase to adapt to a baseline 1 cm offset to the right, followed by a testing phase, we do not include a similar training phase in our simulation. Instead, we directly simulate the testing phase with five different lateral visual conditions of no offset, or 0.5 cm or 1 cm to either side; these are applied through $\mathbf{b}_k$, and our focus is on comparing the HVLP and LVHP configurations to the original pattern of results. Simulations are here made with a movement duration of 2 s; while Körding and Wolpert do not report the movement duration characteristics, a more recent (successful) replication used a maximum allowed time of 5.0 s (Hewitson et al., 2018).

7 Author contributions

ESM, MEO, and MSC all conceptualised the work. ESM and MEO both developed the modelling framework. ESM drafted the manuscript. ESM, MEO, and MSC all revised and approved the final manuscript.

8 Data availability statement

The following links will be made publicly available at the time of peer-reviewed publication. The simulation code is available at https://github.com/CoInAct-group/BayesianInferenceArm and archived at https://osf.io/kvtcu, where simulated datasets and figure-generation code are also available.

9 Funding

This project was supported by a DATA+ grant from UCPH. MSC was supported by DFF-FKK (0132-00141B) and the Carlsberg Foundation (CF22-0941).

10 Disclosures

The authors report no conflicts of interest.

References

Korzybski, A. (1933). *Science and sanity: An introduction to non-Aristotelian systems and general semantics.* International Non-Aristotelian Library Publishing Company.

Athans, M. (1971). The Role and Use of the Stochastic Linear—Quadratic—Gaussian Problem in Control System Design. *IEEE Transactions on Automatic Control*, *16*(6), 529–552. https://doi.org/10.1109/TAC.1971.1099818

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

Burke, D., Hagbarth, K. E., Löfstedt, L., & Wallin, B. G. (1976). The responses of human muscle spindle endings to vibration of non-contracting muscles. *The Journal of Physiology*, *261*(3), 673–693. https://doi.org/10.1113/JPHYSIOL.1976.SP011580

Mense, S. (1978). Effects of temperature on the discharges of muscle spindles and tendon organs. *Pflügers Archiv European Journal of Physiology*, *374*(2), 159–166. https://doi.org/10.1007/BF00581297

Botterman, B. R., & Eldred, E. (1982). Static stretch sensitivity of Ia and II afferents in the cat's gastrocnemius. *Pflügers Archiv European Journal of Physiology*, *395*(3), 204–211. https://doi.org/10.1007/BF00584811

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

Hall, L. A., & McCloskey, D. I. (1983). Detections of movements imposed on finger, elbow and shoulder joints. *The Journal of Physiology*, *335*(1), 519–533. https://doi.org/10.1113/JPHYSIOL.1983.SP014548

Lackner, J. R., & Taublieb, A. B. (1984). Influence of vision on vibration-induced illusions of limb movement. *Experimental Neurology*, *85*(1), 97–106. https://doi.org/10.1016/0014-4886(84)90164-X

Sittig, A. C., van der Gon, J. J., & Gielen, C. C. (1985). Separate control of arm position and velocity demonstrated by vibration of muscle tendon in man. *Experimental Brain Research*, *60*(3), 445–453. https://doi.org/10.1007/BF00236930

Gilhodes, J., Roll, J., & Tardy-Gervet, M. (1986). Perceptual and motor effects of agonist-antagonist muscle vibration in man. *Experimental Brain Research*, *61*(2), 395–402. https://doi.org/10.1007/BF00239528

Roll, J. P., Vedel, J. P., & Ribot, E. (1989). Alteration of proprioceptive messages induced by tendon vibration in man: a microneurographic study. *Experimental Brain Research*, *76*(1), 213–222. https://doi.org/10.1007/BF00253639

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

Edin, B. B., & Vallbo, A. B. (1990). Muscle afferent responses to isometric contractions and relaxations in humans. *Journal of Neurophysiology*, *63*(6), 1307–1313. https://doi.org/10.1152/JN.1990.63.6.1307

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

Inglis, J. T., & Frank, J. S. (1990). The effect of agonist/antagonist muscle vibration on human position sense. *Experimental Brain Research*, *81*(3), 573–580. https://doi.org/10.1007/BF02423506

Taylor, J. L., & McCloskey, D. I. (1992). Detection of slow movements imposed at the elbow during active flexion in man. *The Journal of Physiology*, *457*(1), 503–513. https://doi.org/10.1113/JPHYSIOL.1992.SP019390

Adam, J. J., Ketelaars, M., Kingma, H., & Hoek, T. (1993). On the time course and accuracy of spatial localization: Basic data and a two-process model. *Acta Psychologica*, *84*(2), 135–159. https://doi.org/10.1016/0001-6918(93)90024-L

Miall, R. C., Weir, D. J., Wolpert, D. M., & Stein, J. F. (1993). Is the cerebellum a smith predictor? *Journal of Motor Behavior*, *25*(3), 203–216. https://doi.org/10.1080/00222895.1993.9942050

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

Grill, S. E., & Hallett, M. (1995). Velocity sensitivity of human muscle spindle afferents and slowly adapting type II cutaneous mechanoreceptors. *The Journal of Physiology*, *489*(2), 593–602. https://doi.org/10.1113/JPHYSIOL.1995.SP021075

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

Miall, R. C., & Wolpert, D. M. (1996). Forward Models for Physiological Motor Control. *Neural Networks*, *9*(8), 1265–1279. https://doi.org/10.1016/S0893-6080(96)00035-4

Van Beers, R. J., Sittig, A. C., & Denier Van Der Gon, J. J. (1996). How humans combine simultaneous proprioceptive and visual position information. *Exp Brain Res*, *111*(2), 253–261.

Fourneret, P., & Jeannerod, M. (1998). Limited conscious monitoring of motor performance in normal subjects. *Neuropsychologia*, *36*(11), 1133–1140. https://doi.org/10.1016/s0028-3932(98)00006-2

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

Prochazka, A., & Gorassini, M. (1998). Models of ensemble firing of muscle spindle afferents recorded during normal locomotion in cats. *The Journal of Physiology*, *507*(1), 277–291. https://doi.org/10.1111/J.1469-7793.1998.277BU.X

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

Van Beers, R. J., Sittig, A. C., & Denier Van Der Gon, J. J. (1999). Integration of Proprioceptive and Visual Position-Information: An Experimentally Supported Model. *Journal of Neurophysiology*, *81*(3). https://doi.org/10.1152/jn.1999.81.3.1355

Julier, S., Uhlmann, J., & Durrant-Whyte, H. F. (2000). A new method for the nonlinear transformation of means and covariances in filters and estimators. *IEEE Transactions on Automatic Control*, *45*(3), 477–482. https://doi.org/10.1109/9.847726

Laufer, Y., Hocherman, S., & Dickstein, R. (2001). Accuracy of reproducing hand position when using active compared with passive movement. *Physiotherapy research international : the journal for researchers and clinicians in physical therapy*, *6*(2), 65–75. https://doi.org/10.1002/PRI.215

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433. https://doi.org/10.1038/415429a

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

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

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

Roll, J. P., Albert, F., Ribot-Ciscar, E., & Bergenheim, M. (2004). "Proprioceptive signature" of cursive writing in humans: A multi-population coding. *Experimental Brain Research*, *157*(3), 359–368. https://doi.org/10.1007/s00221-004-1853-x

Fischer, M., & Schäfer, S. S. (2005). Effects of changes in pH on the afferent impulse activity of isolated cat muscle spindles. *Brain Research*, *1043*(1-2), 163–178. https://doi.org/10.1016/j.brainres.2005.02.059

Vindras, P., Desmurget, M., & Viviani, P. (2005). Error parsing in visuomotor pointing reveals independent processing of amplitude and direction. *Journal of Neurophysiology*, *94*(2), 1212–1224. https://doi.org/10.1152/JN.01295.2004

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

Adam, J. J., Davelaar, E. J., Van Der Gouw, A., & Willems, P. (2007). Evidence for attentional processing in spatial localization. *Psychological Research*, *72*(4), 433. https://doi.org/10.1007/S00426-007-0126-2

Dimitriou, M., & Edin, B. B. (2008). Discharges in human muscle spindle afferents during a key-pressing task. *The Journal of Physiology*, *586*(22), 5455–5470. https://doi.org/10.1113/JPHYSIOL.2008.160036

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

Smith, J. L., Crawford, M., Proske, U., Taylor, J. L., & Gandevia, S. C. (2009). Signals of motor command bias joint position sense in the presence of feedback from proprioceptors. *Journal of Applied Physiology*, *106*(3), 950–958. https://doi.org/10.1152/japplphysiol.91365.2008

Friston, K. J. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787

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

Franklin, D. W., & Wolpert, D. M. (2011). Computational Mechanisms of Sensorimotor Control. *Neuron*, *72*(3), 425–442. https://doi.org/10.1016/J.NEURON.2011.10.006

Birznieks, I., Boonstra, T. W., & Macefield, V. G. (2012). Modulation of Human Muscle Spindle Discharge by Arterial Pulsations - Functional Effects and Consequences. *PLOS ONE*, *7*(4), e35091. https://doi.org/10.1371/JOURNAL.PONE.0035091

Erickson, R. I., & Karduna, A. R. (2012). Three-dimensional repositioning tasks show differences in joint position sense between active and passive shoulder motion. *Journal of Orthopaedic Research*, *30*(5), 787–792. https://doi.org/10.1002/JOR.22007

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

Zimmermann, E., Morrone, M. C., & Burr, D. C. (2013). Spatial Position Information Accumulates Steadily over Time. *Journal of Neuroscience*, *33*(47), 18396–18401. https://doi.org/10.1523/JNEUROSCI.1864-13.2013

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

Chancel, M., Brun, C., Kavounoudias, A., & Guerraz, M. (2016). The kinaesthetic mirror illusion: How much does the mirror matter? *Experimental Brain Research*, *234*(6), 1459–1468. https://doi.org/10.1007/S00221-015-4549-5

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

Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T., Miller, E., Bache, S., Müller, K., Ooms, J., Robinson, D., Seidel, D., Spinu, V., … Yutani, H. (2019). Welcome to the Tidyverse. *Journal of Open Source Software*, *4*(43), 1686. https://doi.org/10.21105/joss.01686

Blum, K. P., Campbell, K. S., Horslen, B. C., Nardelli, P., Housley, S. N., Cope, T. C., & Ting, L. H. (2020). Diverse and complex muscle spindle afferent firing properties emerge from multiscale muscle mechanics. *eLife*, *9*. https://doi.org/10.7554/eLife.55177

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

Python Software Foundation. (2023). Python. https://www.python.org

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

Torell, F., & Dimitriou, M. (2024). Local muscle pressure stimulates the principal receptors for proprioception. *Cell Reports*, *43*(9). https://doi.org/10.1016/j.celrep.2024.114699

11 Supporting information

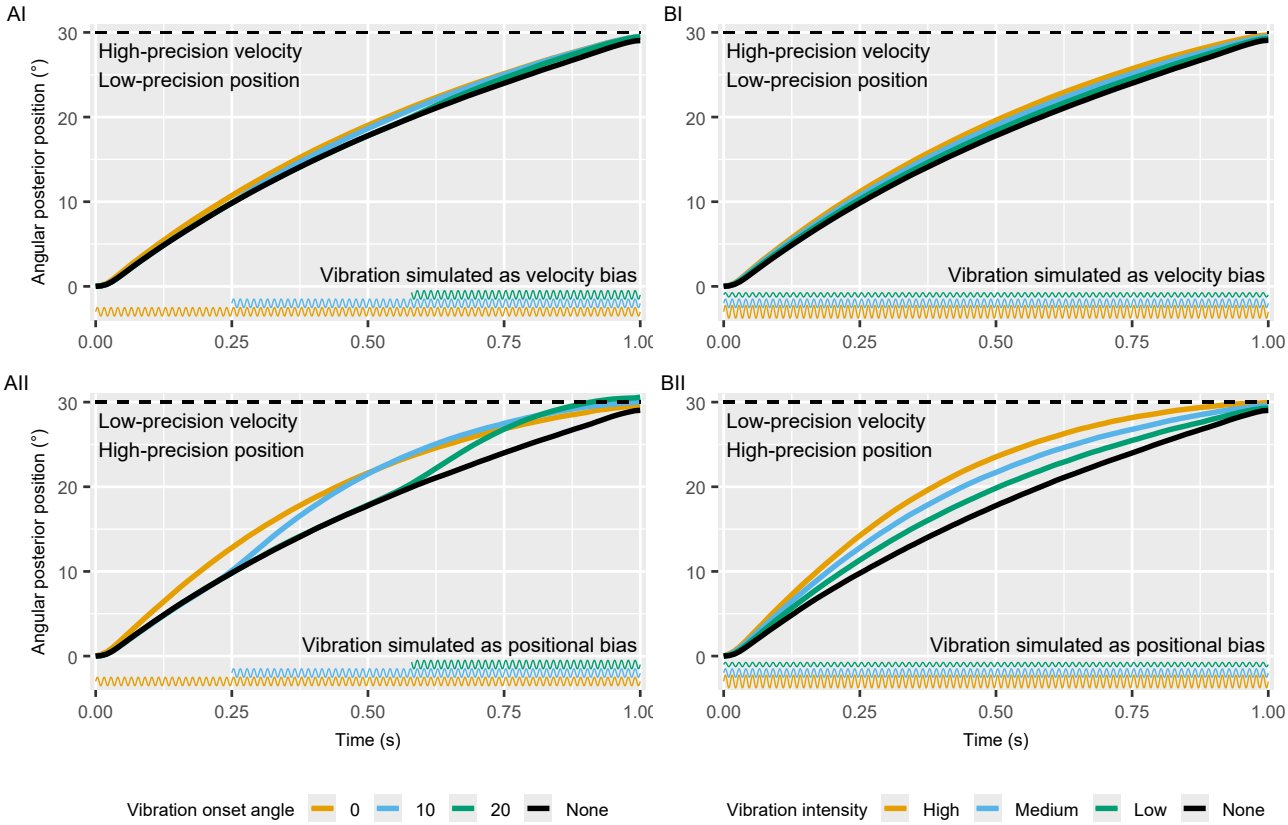

Figure S1: **Extended results of Simulation 1.** Shows the same plot as Fig. 3, but displays the Agent-perceived posterior values rather than true arm values; across all conditions, the Agent correctly guides its posterior hand to the target location, despite perturbations.

Parameter	**Symbol**	**Value**
Time step	Δt	$0.01\,\mathrm{s}$
Upper arm mass	m_1	$2.0\,\mathrm{kg}$
Lower arm mass	m_2	$1.7\,\mathrm{kg}$
Upper arm length	L_1	$0.33\,\mathrm{m}$
Lower arm length	L_2	$0.45\,\mathrm{m}$
Max torque shoulder	$\tau_{1,\max}$	$50\,\mathrm{N\,m}$
Max torque elbow	$\tau_{2,\max}$	$30\,\mathrm{N\,m}$
Max Δ torque shoulder	$\mathrm{rfd}_{1,\max}$	$200\,\mathrm{N\,m\,s^{-1}}$
Max Δ torque elbow	$\mathrm{rfd}_{2,\max}$	$120\,\mathrm{N\,m\,s^{-1}}$
Viscous damping (both joints)	D	$1.5\,\mathrm{N\,m\,s}$
Visual noise density (both dimensions)	σ_{vis}	$1\,\mathrm{mm\,s^{-1/2}}$
Prop. pos. noise density (both joints)	σ_q	$0.06\,\mathrm{rad\,s^{-1/2}}$
Prop. vel. noise density (both joints)	$\sigma_{\dot{q}}$	$0.015\,\mathrm{rad\,s^{-3/2}}$
Motor noise prop (multiplier)	$\sigma_{\tau,\mathrm{prop}}$	$0.05\,\mathrm{s^{-1/2}}$
UKF assumed additional process noise	$\sigma_{\tau,\mathrm{const}}$	$0.1\,\mathrm{N\,m\,s^{-1/2}}$
LQR state deviation cost	$\mathbf{Q}^{\mathrm{LQR}}$	$\mathrm{diag}(25, 25, 10, 10)$
LQR final state deviation cost	$\mathbf{Q}_f^{\mathrm{LQR}}$	$\mathrm{diag}(25000, 25000, 100, 100)$
LQR control cost	$\mathbf{R}^{\mathrm{LQR}}$	$\mathrm{diag}(0.1, 0.1)$
UKF alpha	α	0.001
UKF beta	β	2
UKF kappa	κ	0

Table S1: Table of simulation parameters. Noise densities are continuous-time; per-step SDs use division by $\sqrt{\Delta t}$. Propriopceitve pos. and vel. noise densities are here specified for the HVLP configuration; the amplitudes are flipped for the LVHP configuration.