Inferring control objectives in a virtual balancing task in humans and monkeys

Mohsen Sadeghi[1,2*], Reza Sharif Razavian[3*], Salah Bazzi[1,2,4], Raeed Chowdhury[5], Aaron Batista[5], Patrick Loughlin[5], Dagmar Sternad[1,2,4,6]

[1] Departments of Biology, [2] Electrical and Computer Engineering, [6] Physics, and the [4] Institute for Experiential Robotics, Northeastern University, Boston, MA, USA.
[5] Department of Bioengineering, and Center for the Neural Basis of Cognition, University of Pittsburgh, PA, USA.
** Equal contribution*

Impact statement

A computational approach identifies control strategies in humans and monkeys to serve as basis for analysis of neural correlates of skillful manipulation.

Abstract

Natural behaviors have redundancy, which implies that humans and animals can achieve their goals with different control strategies. Given only observations of behavior, is it possible to infer the control strategy that the subject is employing? This challenge is particularly acute in animal behavior because we cannot ask or instruct the subject to use a particular control strategy. This study presents a three-pronged approach to infer an animal's control strategy from behavior. First, both humans and monkeys performed a virtual balancing task for which different control strategies could be utilized. Under matched experimental conditions, corresponding behaviors were observed in humans and monkeys. Second, a generative model was developed that identified two main control strategies to achieve the task goal. Model simulations were used to identify aspects of behavior that could distinguish which control strategy was being used. Third, these behavioral signatures allowed us to infer the control strategy used by human subjects who had been instructed to use one control strategy or the other. Based on this validation, we could then infer strategies from animal subjects. Being able to positively identify a subject's control strategy from behavior can provide a powerful tool to neurophysiologists as they seek the neural mechanisms of sensorimotor coordination.

Introduction

Almost all actions in daily life can be achieved in multiple ways that all can lead to the desired task goals. As an example, consider a driver steering a car on a curvy road. She may choose different paths depending on whether she wants to maintain a consistent distance from the median strip or whether she aims to minimize changes in velocity. Both strategies can take the driver to her destination, maybe even arriving at the same time, although the precise path taken by the car in both situations will differ. How could one identify the underlying control objectives from differences in observed behavior? A considerable number of studies in human movement neuroscience have aimed to identify the control strategies in a given task based on their kinematic manifestations (Braun et al., 2009; Izawa et al., 2008; Nagengast et al., 2009; Razavian et al., 2023; Uno et al., 1989; Wong et al., 2021). However, experimental tasks are often chosen to elicit consistent behavioral features across repetitions and individuals, not only to facilitate analysis, but also to constrain control to a single objective. Behavior in natural settings, however, is often complex and highly variable across repetitions, and individuals can employ a multitude of strategies to accomplish a task. To date, understanding of such variable behavior - let alone its neural bases - has posed formidable challenges (Croxson et al., 2009; Diedrichsen et al., 2010; Kawato, 1999; Scott, 2004).

Attempts to understand the neural underpinnings of control objectives have been pursued in research on both humans and non-human primates (Benyamini & Zacksenhouse, 2015; Cross et al., 2023; Croxson et al., 2009; Desrochers et al., 2016; Kao et al., 2021; Miall et al., 2007; Nashed et al., 2014; Omrani et al., 2016). Yet, these two lines of inquiry have remained largely parallel with few direct bridges: human behavioral and computational research has mainly focused on the analysis of behavior, while animal research has used invasive methods such as intracortical recordings to gain direct insights into the neural mechanisms of movement control. Experiments with humans tend to use detailed experimental manipulations to elicit features of motor behavior that afford insights into its governing principles. Using a wide range of tasks, from simple reaching to interacting with complex objects, mathematical models with specific control algorithms have been used to reproduce the salient features of behavior (Crevecoeur et al., 2019; Diedrichsen, 2007; Nagengast et al., 2009; Nayeem et al., 2021; Razavian et al., 2023; Yeo et al., 2016). However, understanding the neural underpinnings of movement control at the intracortical level in healthy humans has remained a challenge. On the other hand, animal research, in particular with non-human primates, allows sophisticated methods to directly record neural activity to afford insights into neural correlates of motor behavior. Ultimately, this knowledge should transfer to how the human brain functions (Badre et al., 2015), but those links must be built.

To achieve this objective, cooperative study designs between human and animal motor research are needed to understand the neural basis of human motor skill (Badre et al., 2015; Rajalingham et al., 2022). However, there are difficult challenges to overcome: First, cooperative design requires matching behavioral tasks that can be performed similarly and with the same conditions by both humans and animals. The most appropriate animal model for many human behaviors are monkeys. Second, the goals and constraints of behavioral studies with monkeys and humans are somewhat different, which can preclude a direct comparison. Behavioral tasks used with monkeys are typically simpler than those used with humans, due to the animals' more limited cognitive capacities. Also, studies with monkeys aim for highly repeatable behaviors to facilitate the examination of neural activity by aggregating it across trials or days. In contrast, studies of human behavior can push toward tasks that are more cognitively sophisticated and that capture the complexity that abounds in natural activities. This study bridges the

gap between human and monkey behavioral studies to build toward an understanding of the neural principles of human motor control.

We used an experimental paradigm, the Critical Stability Task (CST), that can be performed by both humans and monkeys (Quick et al., 2018). The CST requires the subject to balance an unstable virtual system governed by a very simple dynamical equation (see Met hods). Performing the task is akin to balancing a virtual pole. The CST has features that make it suitable for the study of more complex motor behaviors. First, while the goal remains the same, the difficulty of the task can be titrated. Second, it involves interactions with an object (albeit virtual in our case) so that continuous adjustments are required to succeed. Each trial evokes unique behavior that may reflect different control strategies to accomplish the task. In addition, even if the same control strategy is employed, each trial generates different behavior due to sensorimotor noise and the task's instability. As in the car driving analogy, the subjects might seek to optimize position, or they might seek to optimize velocity, and different behavioral strategies may lead to equal success.

Because of its complexity and redundancy, each trial of the CST is unique. The goal of the study is to infer the subject's control policy (i.e., optimize position or optimize velocity) from observations of their behavior. When the subjects are humans, it is possible to instruct them to employ a particular strategy or to ask them posthoc what strategy they adopted to succeed at the task. This explicit route is definitely not available with monkeys. As we are still quite far from 'reading out' strategies from neural activity, we need to start with behavior to infer the control strategies. Hence, this study adopted a computational approach based on optimal control theory to simulate behavior during the CST in various conditions. This approach allowed us to make predictions about the behavioral signatures associated with different control policies, which we then used to analyze the experimental data from both humans and monkeys.

In overview, this study investigated, through experimental data and model-based simulations, the sensorimotor origins of behavioral strategies in humans and non-human primates performing the CST. We developed the experimental paradigm such that humans and monkeys executed the task under matching conditions while recording movement kinematics in exactly the same way. An optimal control model was used to simulate different control objectives, through which we identified two different control strategies in the experimental data of humans and monkeys. We discuss how in the future these results could guide the analysis of neural data collected from monkeys to understand the neural underpinnings of different control policies in an interactive feedback-driven task with redundancy.

Results

The Critical Stability Task (CST) involved balancing an unstable system using horizontal movements of the hand to keep a cursor from moving off the screen (**Figure 1**A, C). This study collected data from human subjects performing the CST and compared it to previously collected data from monkeys performing the same task. The hand's displacements were recorded by 3D motion capture (Qualisys, Gothenburg), with a reflective marker attached to the hand. The cursor dynamics were generated by a linear first-order dynamical system, relating hand and cursor kinematics as described in Quick et al., 2018:

$$\dot{x} = \lambda(x + p) \qquad (1)$$

where x and $\dot{x}$ are the horizontal cursor position and cursor velocity on the screen, p is the horizontal hand position, and λ is a positive constant fixed at the beginning of each trial. The parameter λ sets the gain of the system. When λ is larger, the cursor would tend to move faster, making the task more difficult as faster and more precise hand movements were required to maintain balance. Correspondingly, success rates at the task decreased with increasing λ. To summarize the skill of human and monkey participants, we identified the value at which subjects succeeded at only 50% of the trials and defined that value as the "critical" λ.

The task goal was to keep the cursor within a range of space shown on the screen, i.e., $-c \leq x(t) \leq c$, where c was a positive constant. This created a redundancy in achieving the task goal as there were infinitely many ways in which one could balance the cursor inside the specified region. We examined movement kinematics to identify control strategies employed by different subjects, or across different trials.

In a previous study, two Rhesus monkeys were trained to perform the CST under increasing difficulty levels (Quick et al., 2018). Similarly, here 18 human subjects were recruited to perform the same task under comparable experimental conditions as the monkeys (see Methods). **Figure 1** illustrates the experimental setup for both monkeys and humans (**Figure 1**A and 1C) and shows examples of their behavior (**Figure 1**B and 1D). Overall, there were similarities in performance between humans and monkeys. To further quantify and compare this performance across humans and monkeys, we defined a set of control metrics to assess different aspects of control as detailed in the following.

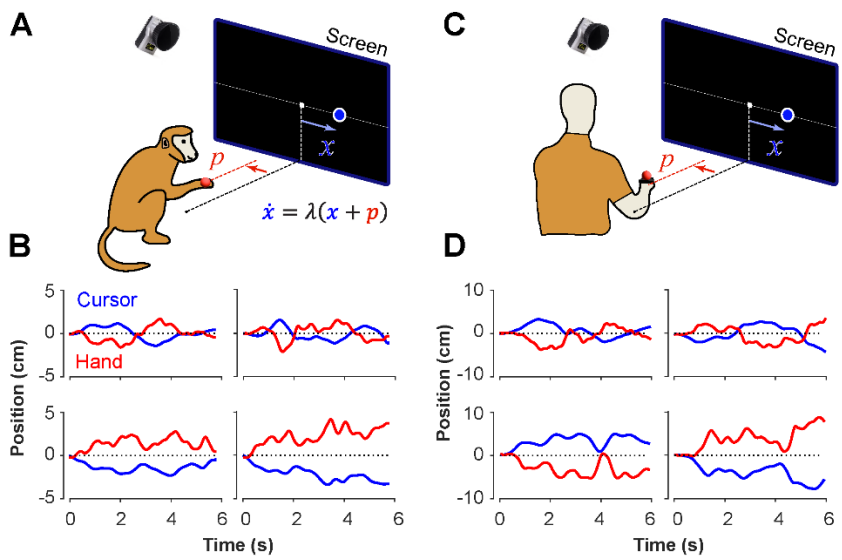

Figure 1: Experimental setup for monkeys and humans performing the CST*. Monkeys (**A**) and humans (**C**) controlled an unstable cursor displayed on a screen using lateral movements of their right hand. The hand movements were recorded using motion capture; the data were used in real-time to solve for the cursor position and velocity through the CST dynamics equation. Timeseries of the hand (red) and cursor (blue) movements shown for four example trials from monkeys (**B**) and humans (**D**).*

Experiment 1: CST performance without instructed strategy

In the first experiment, six human subjects performed the CST with the only instruction to "perform the task without failing to the best of your ability". Failure occurred if the cursor escaped the boundaries of the screen (±10cm from the center) within the trial duration of 6s. Subjects received categorical feedback about the outcome at the end of each trial in a text appearing on the screen reading "Well

done!" for success, and "Failed!" for failure. The degree of difficulty, set by λ, was increased stepwise across trials until the subject could no longer perform the task (see Methods for the specifics about the setting of λ values).

We first sought to examine the main characteristics of behavior in CST performance and how it compared between humans and monkeys. To quantify the overall behavior, four main metrics were employed as described and motivated below. To begin, we considered the overall success rate in the task among different individuals, before focusing on the kinematics of task performance. **Figure 2**A illustrates the success rates and how they dropped as the task difficulty increased. Both humans and monkeys showed a similar pattern of decrease in success rate which was well-captured with a sigmoidal function. Expectedly, individuals varied in their ability to achieve high difficulty levels as a measure of skillful performance, indicated by their "critical λ value", that is, the value of λ when the success rate drops below 50%. To investigate the performance in more detail, the kinematics of movement were examined, specifically the hand and cursor position during each trial. As indicated in equation (1), the hand position p was the control input to the system which aimed to control the cursor position x as the variable of interest. Due to the unstable nature of the task, drifting of the cursor towards the edge of the screen demanded a response by a hand movement to avoid failure. As such, two simple metrics characterized control, one quantifying how the movement of hand and cursor correlated, and a second one to what degree the hand response lagged cursor displacements. **Figure 2**B shows the correlation between the cursor movement and the hand movement as a function of task difficulty. The strength of the correlation increased as trials became more challenging in both monkeys and humans, asymptoting towards –1. According to equation (1), this behavior was equivalent to reducing the sum $(p + x)$ when λ increased, so as to prevent rapid changes in cursor velocity $\dot{x}$, and, hence, reduce the chance of failure.

The response lag from the cursor movement (observed feedback) to the hand movement (control response) is an important characteristic of a control system. As shown in **Figure 2**C, by increasing the task difficulty λ, the lag decreased for all subjects, meaning subjects generated faster corrective responses to cursor displacements in more difficult trials. A possible reason for such behavior is that higher λ values meant increased instability of the system, which required faster responses to avoid failure. Whereas in easy trials, due to slower dynamics of the system, subjects could afford delayed responses to cursor displacements (and hence, larger lags) and still manage to succeed.

As the fourth metric, we also calculated the control gain by measuring the ratio of root mean squared (RMS) of hand position to the RMS of cursor position for each trial. This measure determined to what extent the control signal (hand movement) compared in magnitude to the cursor movement. A large gain meant that on average across a trial, the hand exhibited larger movements than necessary to correct for cursor deviations. **Figure 2**D illustrates the calculated gain as a function of task difficulty for humans and monkeys. As shown, except for *Monkey J*, the gain showed a gradual decrease as the task difficulty increased for most individuals. Such decrease could be due to larger cursor movements at higher difficulty levels, and perhaps more efficient corrective hand responses to cursor displacements. To the latter, it is worth noting that for high λ values, small hand movements could cause large cursor displacements, which was detrimental to the task success. Therefore, pruning any task-irrelevant hand movements, consistent with promoting efficiency, seemed essential to succeed in more difficult trials. Overall, the control metrics presented in **Figure 2** give insight into how the CST was performed: as the task difficulty increased, subjects tended to respond to cursor displacements faster (that is, with lower lag), more precisely (seen in the stronger hand-cursor correlation), and more efficiently (with lower

gain). Behavior was comparable between humans and monkeys, which suggests that there were underlying control strategies used in common by both species. Next, we sought to detect those control strategies.

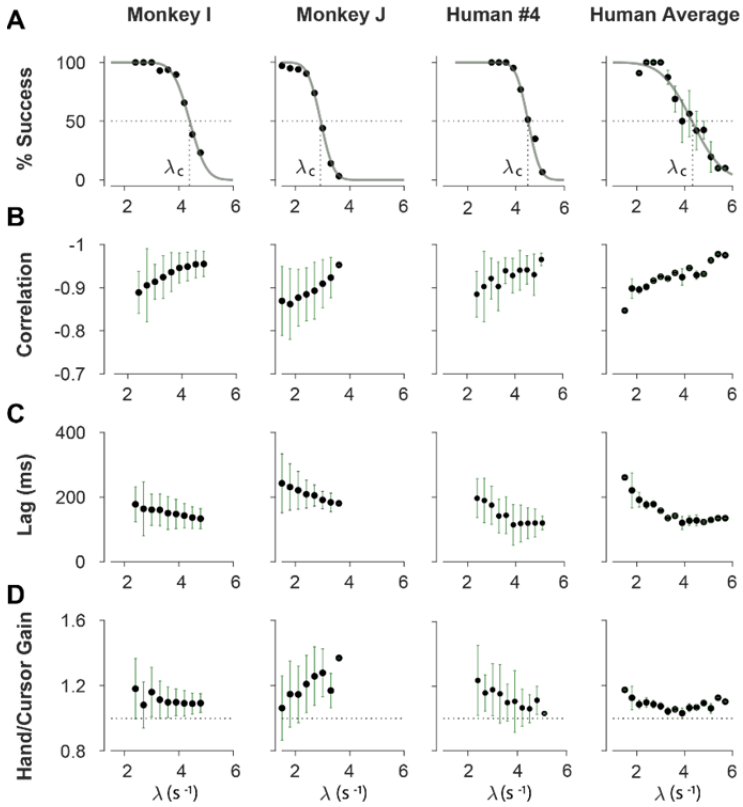

Figure 2: Overall behavioral characteristics of CST performance as a function of task difficulty (λ). *Data is shown for two individual monkeys (first two columns from left) from a previous study* (Quick et al., 2018), *as well as an example human individual (third column from left) and the average across human subjects (right-most column). For the individual subjects, each data point and its corresponding error bars represent the mean±SD across trials for any given difficulty level, respectively. For the human average plot, the data points and their corresponding error bars represent the mean±SE across individuals for each difficulty level.* ***A****. Psychometric curves for success rate (%) as a function of task difficulty (λ) the difficulty level at which the success rate crossed 50% was considered as the critical stability point (λ_c), indicating the individual's skill level in task.* ***B****. Correlation between the hand and cursor movement during CST.* ***C****. Sensorimotor lag between the cursor and the hand movements.* ***D****. Ratio of hand RMS over the cursor RMS calculated for each trial, representing the gain of the response.*

Redundancy of control strategies in CST performance

The CST, as described earlier, affords redundancy in the control strategies that could result in task success. Although covert in aggregate level of performance (i.e., **Figure 2**), single trial observations of hand and cursor movements suggested that different underlying control objectives might be at play. Two types of behavioral patterns appeared recognizable in the data. In one case, the cursor seemed to be always balanced around the center of the screen, and any deviations from the center induced a response to bring the cursor back to the center. This was reflected in the oscillatory movements of the cursor around the center, shown in example trials in Figure 1B and D (first row). In other trials, the cursor either exhibited a slow drift from the center or remained relatively still anywhere within the boundaries of the screen, with only limited attempts to bring the cursor back to the center (for example, **Figure 1**B and D, second row). We hypothesized that these patterns of behavior arise from different control objectives, each focused on a different state variable in the state-space of the cursor movement.

In the former case, the position of the cursor appeared to be the primary control variable. Under this strategy, subjects might pursue the objective of keeping the cursor near the center of the screen. We refer to this strategy as the Position Control strategy. In the latter case, the cursor velocity seemed to be of primary importance for control, with the objective to slow down cursor velocity regardless of its position in the workspace. We refer to this strategy as the Velocity Control strategy.

Can we distinguish between different control strategies by examining behavior? To test this idea, we took a computational approach by developing a generative model based on optimal feedback control (Todorov & Jordan, 2002) that could simulate the task under different conditions and with different objectives (Todorov & Jordan, 2002). The model involved a controller that generated optimal motor commands based on a given strategy to perform the CST via a simple effector model. The model also contained a state estimation block that estimated the states of the system based on the given feedback (Todorov, 2005). In this case, cursor position and cursor velocity were used as feedback to the controller at each time step. **Figure 3**A illustrates a block diagram of this model.

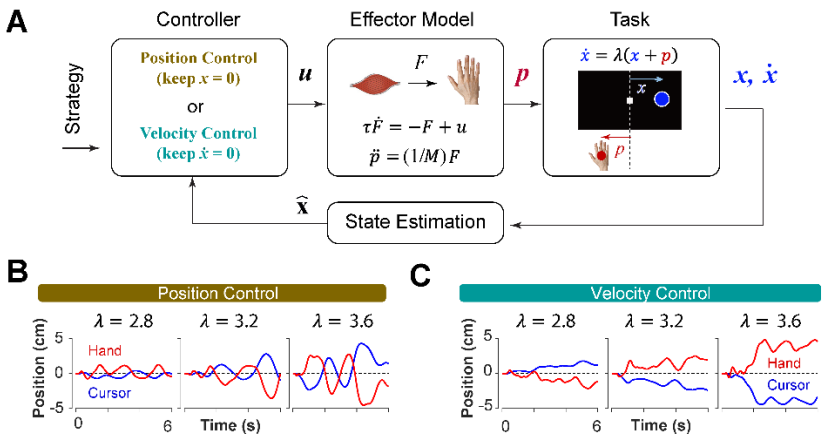

***Figure 3: A generative model to performs the CST. A**. An optimal feedback controller generates motor commands based on two control objectives, position and velocity control. The motor command leads the movement of the effector (hand), which performs the CST. The cursor position and velocity are provided as feedback from which all the states are estimated and fed back to the controller. **B** and **C**. Example trials simulated under the two control objectives for different difficulty levels: keeping the cursor at the center (**B**; position control) and keeping the cursor still (**C**; velocity control).*

The control gains used in the controller to generate the motor commands were optimally found by minimizing the sum of two cost functions: the cost of effort to reduce energy, as well as the cost of accuracy that prevented the states of the system from making large deviations (2):

$$J = \sum_{t=1}^{n} \left(\mathbf{x}_t^T \boldsymbol{Q} \mathbf{x}_t + u_t^T R u_t \right) \qquad (2)$$

where u and $\mathbf{x}$ represented the motor command and the state vector of the system, respectively. In this model, the state vector consisted of six states: the position, velocity and acceleration of the hand, as well as the position, velocity and acceleration of the cursor (see Methods). Variables t and n represent the time, and the total number of time steps, respectively, in a trial. The matrix $\boldsymbol{Q}$ and the scalar R determined the weight of accuracy and effort in the cost function, respectively. Importantly, the matrix $\boldsymbol{Q}$ allowed for determining which states of the system were of primary importance in the control

process. Therefore, the implementation of different control objectives in the controller was done through setting the $\boldsymbol{Q}$ matrix appropriately. As such, a Position Control strategy was implemented by setting the weight of cursor position in the $\boldsymbol{Q}$ matrix to a large value, emphasizing the primacy of cursor position as a control variable. Similarly, to implement the Velocity Control strategy, the weight of the cursor velocity in the $\boldsymbol{Q}$ matrix was set to a large value (see Methods). By simulating the task for each control strategy, we could generate synthetic behavior similar to that of humans and monkeys. **Figure 3**B and C illustrate a few example simulations of the task under different difficulty levels for the Position Control and Velocity Control, respectively. As exemplified, the simulated trials for Position Control show oscillatory movements of the cursor around the center, whereas the trials generated based on Velocity Control, exhibited slow drift of the cursor from the center with minimal attempt to correct for such drift. These characteristics were similar to the observed patterns of behavior in human and monkey data (Figure 1B and D).

To further identify the behavioral signatures associated with each control objective, beyond the apparent differences between single trials, we conducted a series of simulations in which the model performance was examined for a range of task difficulties, and novel predictions of the model for each control objective were assessed. For each control objective, the task was simulated for different difficulty levels, ranging from $\lambda = 1.5$ to $\lambda = 7$, with increments of $\Delta\lambda = 0.2$. For each difficulty level, 500 trials were simulated (see Methods for details). In the first step, we performed the same set of analyses as reported in **Figure 2** to evaluate how the model compared to human and monkey behavior at an aggregate level of CST performance. **Figure 4**A illustrates the overall performance of the model for both Position Control and Velocity Control strategies. As shown, for each metric, the model exhibited comparable behavior to experimental data with regard to the task difficulty: the success rate dropped in a sigmoidal fashion, the correlation between hand and cursor movements increased, and the response lag between hand and cursor as well as the hand/cursor gain decreased. These results showed that, overall, both simulated control strategies were capable of producing similar behavioral characteristics as humans and monkeys. But more interestingly, despite no apparent advantage of one strategy over the other in the task success (**Figure 4**A, top panel), they showed differences in the magnitude of hand-cursor correlation, lag and gain. Namely, Position Control consistently showed larger magnitudes for correlation, lag, and gain for any given task difficulty.

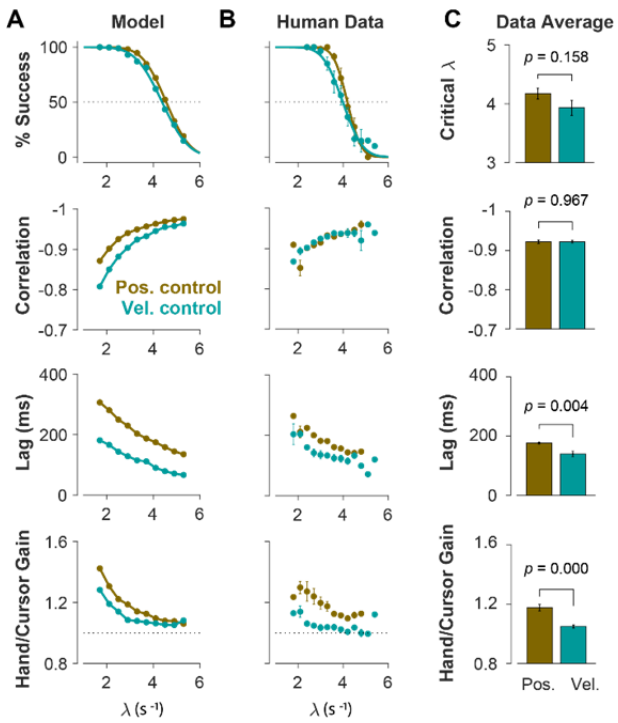

Figure 4: Different control objectives result in measurably different behavior. *Overall performance of the model (**A**) and human subjects (**B**) for two control objectives, Position Control and Velocity Control. The four rows show success rate (first row), correlation between hand and cursor movement (second row), sensorimotor lag between cursor and hand movements (third row), and the hand/cursor gain, defined as the RMS of hand movement over the RMS of cursor movement during each trial (last row). The error bars on the human average data indicate the standard error of the mean across subjects for each group. **C**. The average performance across difficulty levels and subjects within each group. The Critical λ (first row) indicates the difficulty level at which the success rate crosses 50%.*

Experiment 2: CST performance under explicit instructions

The model indicated that differences in behavioral metrics exist for Position vs Velocity Control. This led to a new experiment for which we recruited two new groups of human subjects (*n*=6 per group). Each group performed the CST under the same procedure as described in Experiment 1, except that this time each group was explicitly instructed to use a specific control strategy. One group was asked to perform the task with the objective of "keeping the cursor at the center of the screen at all times". This instruction was to induce a Position Control strategy. The second group was asked to "keep the cursor still anywhere within the boundaries of the screen". This instruction aimed to induce a Velocity Control strategy (see Methods for details). In each group, the kinematic behavior of hand and cursor was collected, and the control metrics were calculated. The goal was to elicit differences in performance between the two groups and, if such differences were found, determine whether they matched the behavior of the corresponding model.

The summary of performance for both human subject groups is shown in **Figure 4**B. The general trends of all four measures with respect to the task difficulty were consistent with the data generated by the model, as well as the human data from Experiment 1 (**Figure 2**). Importantly, the behavioral differences between the two control strategies in human data matched the predictions of the model relatively well (**Figure 4**A, B): the rate of success was similar, and with the exception of hand-cursor correlation, the group with Position Control instruction showed significantly larger hand-cursor lag (unpaired t-test: t_{10} =

3.79, p = 0.004) and hand-cursor gain (unpaired t-test: t_{10} = 5.27, $p < 10^{-3}$) compared to the group with Velocity Control instructions (**Figure 4**C).

These results showed that the model not only captured the overall performance features observed in the data, it also successfully demonstrated the redundancy of control strategies in CST performance, and qualitatively distinguished between such strategies at an aggregate level of performance. To ask further, can we identify, in a quantitative way, the control strategy employed by an individual, or even in a given trial, when no explicit information about their preferred strategy is available? To this end, we examined performance at single-trial level and introduced quantitative measures that evaluated the degree to which a particular control strategy was used in that trial, as described in the next section.

Behavioral traces of control strategy in an individual's overall performance

To further investigate what control strategy was preferred by an individual or in a given trial, we examined the predictions of the model about the cursor behavior in state space, and then tested these predictions using experimental data from Experiment 2. Two metrics were defined that captured the state-space behavior of the cursor in each trial. First, we examined the average cursor position and cursor velocity in each trial, represented in the state space of cursor movement. This provided a single data point for each trial in state space, indicating whether on average there was a drift in cursor position and its velocity away from zero ($x = \dot{x} = 0$). It was expected that for Position Control, all trials scattered around the origin of the state space, whereas for Velocity Control, they could deviate from the origin. We also examined whether the states of the cursor correlated. **Figure 5**A illustrates the state-space representation of cursor movement based on model simulations for both Position Control (top) and Velocity Control (bottom), where each data point represents one simulated trial. As shown, the distribution of trials in this space differed markedly between the two control objectives. The Position Control strategy resulted in a distribution with little correlation between cursor position and its velocity, and closely scattered around the center. In contrast, the Velocity Control strategy revealed an elongated distribution with a relatively strong correlation between the cursor position and its velocity. This allowed us to distinguish between different individuals' preferred control strategy.

To validate the model predictions, the same analysis was performed on the empirical data from Experiment 2. **Figure 5**B illustrates three example subjects from Position Control and Velocity Control groups, and **Figure 5**C shows a summary of how the correlation values differed across control strategies for the model and the empirical data. As shown, overall, subjects in the Velocity Control group showed significantly larger correlations than individuals in the Position Control group (unpaired t-test on the Pearson correlation coefficient: t_{10} =4.06, p=0.002). Based on the within-group variability, this allowed us to determine how pronounced a subject executed their respective strategy compared to other subjects in the same group. This metric, therefore, provided a quantitative way of estimating where on the spectrum of control strategy an individual's performance lies with respect to other performers.

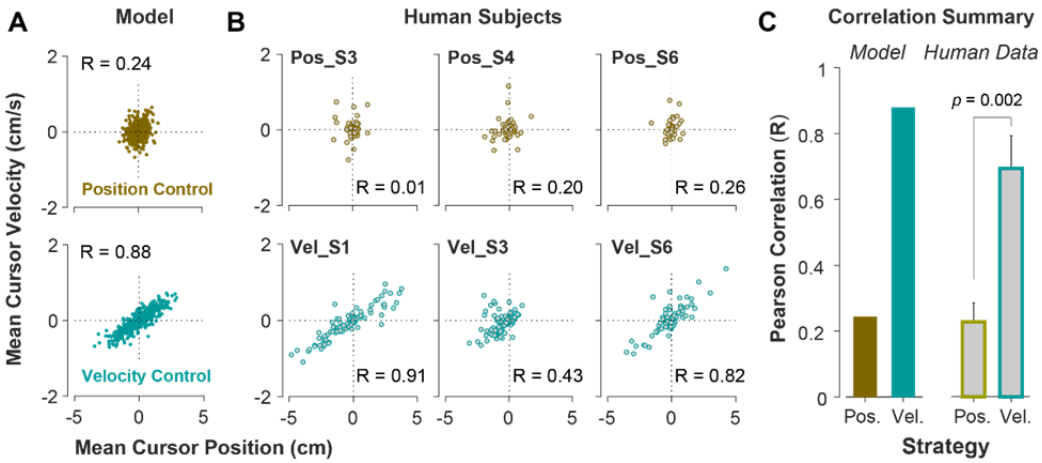

Figure 5: State-space distribution of trials reveals different control strategies. A*. Mean cursor velocity plotted against mean position for each trial, shown for the position control objective (top) and velocity control objective (bottom). Each data point represents one successful trial and was simulated for a range of difficulty levels up to the critical λ value (corresponding to 50% success rate).* ***B****. Three example human subjects from the position control group (top row) and velocity control group (bottom row). Each data point represents one successful trial. The data represents an ensemble of trials ranging in difficulty levels up to the critical λ value for each subject. R indicates the correlation between the trial position and velocities.* ***C****. Pearson correlation coefficient between cursor mean position and velocity for each control objective in the model (left) and human data (right). The human data shows the mean (±SE) across subjects for each control objective group.*

The effects of control strategy at a single-trial level of behavior

Due to the task's redundancy the choice of control strategy may not be fixed for an individual throughout their performance and might vary from one trial to the next. It is therefore of great interest to determine, in a given trial, to what extent the behavior is the outcome of Position versus Velocity Control strategies. To this end, we examined the magnitude of cursor movement calculated as the root mean squared (RMS) of its position and velocity in each trial. This was directly related to the objective functions used in the model (equation(2), which provided a more direct comparison regarding the primacy of position versus velocity in the control of the cursor: a Position Control strategy aimed to minimize the RMS of cursor position, while Velocity Control aimed to minimize the RMS of cursor velocity. This distinction could be well represented in the state-space of the cursor movement.

Figure 6A illustrates the model prediction for the RMS of cursor position and cursor velocity plotted against each other for the Position Control (top) and Velocity Control (bottom). For Position Control, the distribution of trials leans towards the vertical axis (restricting cursor position but allowing large cursor velocities), whereas for Velocity Control, it leans mainly towards the horizontal axis (a larger range of cursor positions but restricted velocities). This distinction could be quantified by the slope of a fitted regression line to the data, with relatively larger slopes indicating Position Control and smaller slopes signaling Velocity Control. Similar patterns of behavior could be observed in the human data from Experiment 2 as illustrated in **Figure 6**B and C, with the Position Control group showing significantly larger regression slope than the Velocity Control group (unpaired t-test, $t_{10} = 6.33$, $p<0.001$). The regression slope could more clearly distinguish between individual trials than could the correlation coefficient metric shown in **Figure 5**, regarding their corresponding control strategy: if a given trial in the RMS space of the cursor movement lay below/above a certain slope threshold, its performance could be considered the result of a Velocity/Position Control strategy. We could therefore use this behavioral

feature to develop a classifier that inferred, with a certain level of confidence, the underlying control strategy in the performance of an individual in any given trial.

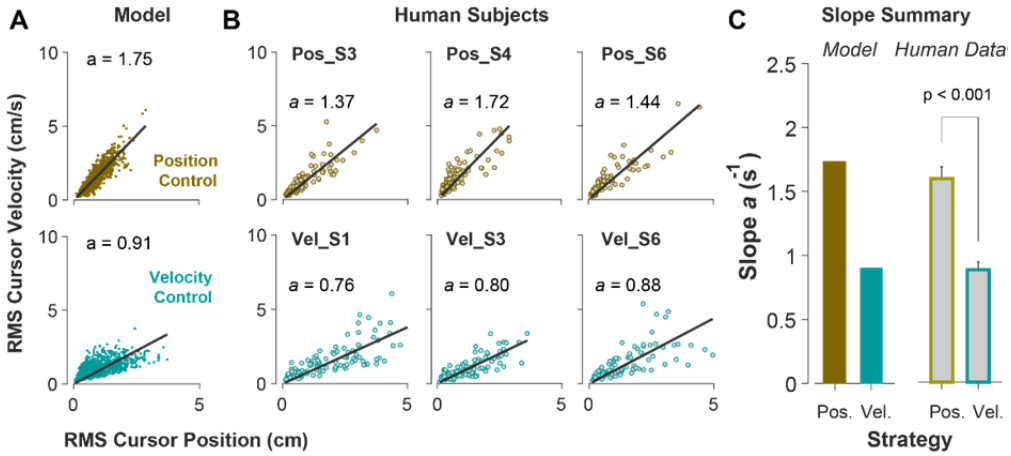

***Figure 6**: **Identifying control strategy based on magnitude of cursor movement in the state space. A**. Magnitude of cursor movements quantified by the RMS of position and cursor velocity for each trial, plotted against each other; position control objective (top) and velocity control objective (bottom). Each data point represents one successful trial and was generated based on the model simulations for a range of difficulty levels up to the critical λ value (corresponding to 50% success rate). **B**. Performance of three example subjects from the position control group (top row) and velocity control group (bottom row). Each data point represents one successful trial. The data represents an ensemble of trials ranging in difficulty level up to the critical λ value for each subject. The values of the regression slopes are also shown. **C**. Summary of the regression slopes for the RMS plots, shown for each control objective in the model (left) and human data (right). The human data shows the mean (±SE) across subjects for each control objective group.*

Inferring control strategies from behavior during CST performance

When monkeys performed the CST, we lacked explicit knowledge about which strategy they might have employed. This is similar to Experiment 1; when humans performed the CST with no specific instructions, their control objective was not explicitly available. To achieve the goal of inferring an individual's control objective based on their performance, we used the control characteristics that our computational approach introduced to distinguish between different control strategies. To this end, the simulation results based on the cursor movement in its RMS space (**Figure 6**A) were used to train a simple classifier, a support vector machine (see Methods). This classifier then determined, based on the learned regression slopes from the RMS distributions (**Figure 7**A), whether a given trial was likely performed under the Position Control, or Velocity Control strategy. We first tested the performance of the classifier on the empirical data from Experiment 2, where the intended control strategy used by each subject was known.

Figure 7B shows the cursor RMS data from three example subjects in each instructed group (similar to **Figure 6**); for each trial (data point) a probability was obtained from the classifier indicating to what extent a given trial was performed with the Position Control strategy (see Methods). A trial with the estimated probability of >70% was considered Position Control, while a probability of <30% for a trial signified Velocity Control. All other probabilities were considered as 'Uncertain' as to which of the two control objectives were used. As shown in **Figure 7**B, for the Position Control group, most of the trials were rightfully classified as Position Control trials, and similarly for the Velocity Control group, the majority of trials were classified under Velocity Control strategy. The average probability across all trials

for each individual was also obtained as an overall measure of the control objective for that subject. This average measure is shown in **Figure 7**B for the example subjects and summarized in **Figure 7**C for all subjects in each group. This showed that the classifier correctly determined the control strategy used by each individual without being trained on any experimental data.

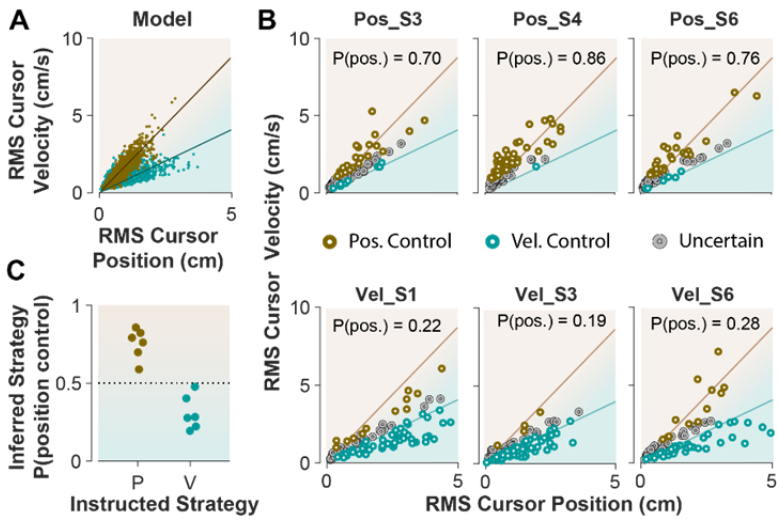

***Figure 7: Classifying control strategies in humans who received explicit instructions**. **A**. Simulated data in the RMS space of cursor movement used as training set for a classifier to determine the control objective of each trial. **B**. Data from three example subjects in each group, where each trial was classified as position control (brown), velocity control (cyan), or uncertain as to the control objective (grey). To obtain the control objective of each trial, the classifier (a support vector machine; see Methods) obtained the probability of that trial performed with position control objective, where P(pos)>70% was classified as position control, P(pos)<30% was classified as velocity control, and everything else was classified as uncertain. The average of P(pos) across all trials for each individual is shown inside the respective plot. **C**. Overall probability of Position Control summarized for all subjects instructed in the position and velocity control groups of Experiment 2.*

The ultimate test of our approach would be to infer the control strategy used by individuals whose strategy was unknown, that is the monkeys, and humans who received no instructions about control strategy in Experiment 1. After representing the performance of each subject in the RMS space, the classifier was used to determine what control strategy was used in each trial. **Figure 8** illustrates the classification results for human subjects who received no instructions (Experiment 1) as well as two monkeys (*Monkey I* and *J* from Quick et. al. 2018). The model simulations are also provided as reference in **Figure 8**A. **Figure 8**B and C show the data from three example human subjects, as well as two monkeys, in which each trial is either labelled as Position Control (brown), Velocity Control (cyan), or Uncertain (grey). Two example trials, one from each inferred control strategy are also singled out from each subject's performance in **Figure 8**B and C (bottom row) to show how the hand and cursor movement generally behaved under each control strategy. Calculating the average probability of control strategy for each individual, similar to **Figure 7**, we could infer which control strategy was of primary importance for each subject (**Figure 8**D). For example, human subject NI-S2 more likely adopted a Velocity control strategy, while human subject NI-S4 mainly performed the task with Position Control strategy (**Figure 8**B). Similarly, *Monkey I* seemed to prefer the Velocity Control strategy, while *Monkey J* most likely adopted a Position Control strategy (**Figure 8**C).

Ultimately, our procedure enabled us to not only infer the underlying control strategy at a single trial level, but also identify which control strategy was overall preferred by humans and monkeys when no explicit knowledge about their control strategy was available. These results are encouraging as they constitute an important step towards bridging our findings between human and monkey research, and ultimately guide neurophysiological analyses to identify the neural underpinnings of control strategy in the primates' brain.

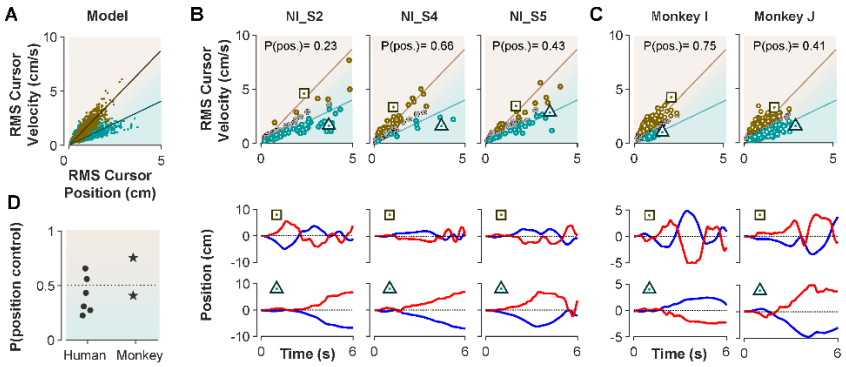

Figure 8: Inferring control strategies in monkeys and humans who received no instructions. A*. Simulated data in the RMS space of cursor movement was used as training set for a classifier to determine the control objective of a trial without explicit instructions.* ***B****. Data from three example human subjects with no instructions (NI) about the control objective. Each trial (data point) is classified based on the probability of position control, P(pos), obtained for each trial from the classifier. Trials with P(pos)>70% and P(pos)<30% were, respectively, labeled as position control (brown) and velocity control (cyan), while other probabilities were labeled as uncertain (grey). Two example trials, one from each control objective, are shown in the bottom row.* ***C****. The classifier was used on data from two monkeys (Monkey I and J) who performed the CST. Similarly, trials for each monkey were categorized as position control (brown), velocity control (cyan), or uncertain (grey).* ***D****. Overall probability of an individual preferring the position control strategy, shown for six humans and two monkeys. This measure was obtained for each individual as the average probability of position control across all trials.*

Discussion

As we seek to understand the neural basis of human motor control, it is important to build links between studies in humans, where behavior can be complex and naturalistic, and monkeys, where direct neural recordings are possible. Doing so requires close coordination between researchers who work with humans and animals(Badre et al., 2015). With the goal to advance insights into movement control, the current work developed a novel approach to parallel human-monkey behavior. In a matching task design humans and monkeys performed a virtual balancing task, where they controlled an unstable system using lateral movements of their right hand to keep a cursor on the screen. The task was challenging and, importantly, exhibited different ways to achieve task success. The task required skill, but that was conceptually simple enough for monkeys to learn the skill and ultimately achieve the same level of proficiency as humans.

The results showed that both humans and monkeys exhibited the same behavioral characteristics as the task was made progressively more difficult: success rates dropped in a sigmoidal fashion, the correlation magnitude between hand and cursor increased, and the response lag from cursor movement to hand response decreased. Further observations based on single trials showed that the task was possibly achieved with different control strategies, both across subjects and across trials. Our goal was to identify the underlying control objectives that led to different behavior, a model based on optimal feedback control was developed that identified two different control objectives that successfully captured the

average performance features of humans and monkeys: Position Control and Velocity Control. Both strategies produced behavior that was consistent with observations even at the single trial level. Additional experiments revealed that humans who followed specific instructions as to performing the task with Position Control ("keep the cursor at the center") or Velocity Control ("keep the cursor still") matched the behavior predicted by the two simulated control policies. Model simulations exhibited features that served to identify control strategies of humans and monkeys who received no specific instructions at a single trial level.

Studies in motor neurophysiology have largely relied on simple paradigms such as center-out movements (Batista et al., 1999; Cisek et al., 2003; Georgopoulos et al., 1986; Pruszynski et al., 2011; Scott & Kalaska, 1997), which were brief in duration, highly stereotypical across repetitions, and could be performed to a reasonable degree of success with limited sensory feedback. Such characteristics were needed to make sense of noisy neural data through averaging trials over repeats of highly similar behaviors. However, such tasks are not common in natural settings, where we continually utilize sensory feedback to respond to our environment, interact with objects around us, and never do the same action the exact same way. Indeed, such fluid, prolonged and feedback-driven interactions are what we seek to understand both at the behavioral and neural levels. To this end, we need to investigate more complex tasks that involve sensory-driven control and allow for different control strategies while still within a sufficiently controlled scope. The task employed here, the Critical Stability Task (CST) offers advantages for the study of sensorimotor control that complements previously used tasks. The task continuously engages feedback-driven control mechanisms for a prolonged period of time and is rich in its trial-to-trial and subject-to-subject variability. As we can titrate the difficulty of the task, both monkeys and humans can learn it and we can study and model their behavior. This opens the gate towards understanding the neural principles of skill learning beyond simple reaching tasks. This study showed that CST afforded the examination of control strategies through a computational approach that modelled monkey and human behavior in comparable fashion.

A critical step for bridging insights between human and monkey behavior is through the computational approach that could explain behavior equally well in both human and monkey performance (Badre et al., 2015; Rajalingham et al., 2022). In an earlier attempt of modeling CST, a simple PD controller with delay in sensory feedback was proposed to explain the recorded behavior (Quick et al., 2018). However, the model was limited in its ability to capture most features observed in the data, such as success rate, or correlation between hand and cursor movements. In the past years, Optimal Feedback Control (OFC) has been introduced as an effective approach to understanding the control mechanisms of reaching movements at the level of behavior (Diedrichsen et al., 2010; McNamee & Wolpert, 2019; Pruszynski & Scott, 2012; Scott, 2004; Todorov, 2004), separately in human research (Liu & Todorov, 2007; Nagengast et al., 2010; Nashed et al., 2014; Razavian et al., 2023; Ronsse et al., 2010; Todorov, 2005; Todorov & Jordan, 2002; Yeo et al., 2016) and monkey research (Benyamini & Zacksenhouse, 2015; Cross et al., 2023; Kalidindi et al., 2021; Kao et al., 2021; Takei et al., 2021). Here, the OFC framework was used to account for and make novel predictions about behavioral features in CST. Note that there are fundamental differences between reaching and CST movements, which needed to be accounted for in the modeling process. Unlike center-out reaching, the CST did not have a stationary target toward which the hand needed to move; rather, it required the hand/cursor to remain anywhere within a predefined area for a prolonged period of time. Also, the behavior was not tracking a point on the screen, but rather moving in opposite direction of the cursor, a behavior that probably requires more cognitive

resources. Despite these advanced task features, OFC as a feedback control framework proved an appropriate approach to examine this demanding interactive and sensory-driven task.

Two aspects in our computational approach are worth discussing. First, we examined control strategies that only involved two main kinematic quantities of movement: cursor position and cursor velocity. One might argue that other kinematic features could be explored, such as acceleration or other higher derivatives of the cursor and/or the hand. However, it is important to note that, given the task of keeping the cursor within a specified area for a period of time, cursor position and velocity are the most directly related quantities to the goal of the task. These quantities were also less demanding to predict from sensory feedback, compared to, for example, acceleration (Hwang et al., 2006; Sing et al., 2009). Also note that the kinematics of the hand were not the variables of interest in the task, as the goal was to control the cursor, and not the hand.

Second, we mainly explored Position and Velocity Control strategies separately to identify distinctive behavioral features associated with each one. Experimental data, however, shows that a large number of trials fall somewhere between the Position and Velocity Control boundaries (Figure 7 and 8). This could be due to a mixed control strategy, where both Position and Velocity control strategies contribute simultaneously to achieving the task goal, or where subjects switch strategies of their own accord. Here, we aimed to determine the behavioral signatures of the extreme cases, either predominantly based on position, or velocity of the cursor movement. This may increase the chance to detect differences more clearly in neural activity associated with each control objective in further analysis of monkeys' neurophysiological data. Even though in this experiment only a subset of trials is amenable to a clear identification as using one control strategy or another, with monkeys it is possible to collect tens of thousands of trials over many days accumulating enough trials for analysis.

Despite potential limitations, our approach was successful in two main ways. First, it provided a normative explanation for the macro-level characteristics of behavior observed in human and monkey data. Second, due to its generative nature, model simulations provided for not yet seen conditions and made predictions about the behavior under new control objectives. In the future, our behavioral analysis can serve as a foundation to classify or parse neural activity in monkeys performing complex actions where trial averaging is no longer possible. This behavioral analysis holds promise to generate crucial insights into neural principles of skillful manipulation, not only in monkeys but also, by induction, in humans.

Methods

Participants and Ethics Statement

18 healthy, right-handed university students (age: 18—25 years; 8 females) with no self-reported neuromuscular pathology volunteered to take part in the experiments. All participants were naïve to the purpose of the experiment and provided informed written consent prior to participation. The experimental paradigm and procedure were approved by the Northeastern University Institutional Review Board (IRB# 22-02-15).

The data from two adult male Rhesus monkeys (Macaca mulatta) used in this study was taken from a previously published work (Quick et al., 2018). All animal procedures were approved by the University of Pittsburgh Institutional Animal Care and Use Committee, in accordance with the guidelines of the US

Department of Agriculture, the International Association for the Assessment and Accreditation of Laboratory Animal Care, and the National Institutes of Health. For details of experimental rig and procedure see the Methods in (Quick et al., 2018).

Critical Stability Task (CST)

The CST involved balancing an unstable cursor displayed on the screen using the movement of the hand (Jex et al., 1966; Quick et al., 2014, 2018). The CST dynamics was governed by a first-order differential equation as shown in equation (1). The difficulty of the task was manipulated by changing the parameter λ : by increasing λ the task became more unstable, hence more difficult to accomplish. To perform the task, subjects sat on a sturdy chair behind a small table, with their right hand free to move above the table (**Figure 1**). A reflective marker was attached to the subject's back of the hand on the third metacarpal, and the hand position was recorded using a 12-camera motion capture system at a sampling rate of 250Hz (Qualisys, 5+, Goetheburg, SE). The mediolateral component of the hand position was used to solve the CST dynamics with the initial condition of $x(t = 0) = 0$ (Quick et al., 2018). The calculated cursor position was real-time projected as a small blue disk (diameter: 4mm, approximately 0.8deg in visual angle) on a large vertical screen in front of the subject at a 150cm distance. The processing delay of the visual rendering was roughly 50ms.

Experimental Design

Task

At the beginning of the experiment, human subjects held their right hand comfortably above the table and in front of their right shoulder as shown in **Figure 1**, where the hand position was mapped to the center of the screen. The visual display of the cursor and hand position was scaled such that the lateral hand movements of ±10cm corresponded to ±20deg of visual angle from the screen center and served as the boundaries of the workspace. Each trial started with the hand position displayed on the screen as a red cursor (diameter: 4mm, or approximately 0.8deg in visual angle). Subjects were asked to bring the red cursor to the center of the screen depicted by a small grey box (**Figure 1**). Once the red cursor was at the center, and after a delay of 500ms, the trial started. The red cursor disappeared and a blue cursor representing the x position in equation (1) appeared at the center. Subjects were instructed to keep (or 'balance') the blue cursor within the boundaries of the workspace for 6s for the trial to be considered successful. If the cursor escaped the workspace at any time, the trial would abort and considered as failed. Subjects were informed of the outcome of the trial by a message on the screen, reading "Well Done!" for success, and "Failed!" for failure. The next trial started after an intertrial interval of 1000ms. This feedback matched the binary reward that monkeys were given in the experiment by Quick and colleagues.

Experimental Paradigm and Conditions

Each human subject participated in the experiment for three consecutive days. At the beginning of the first day, subjects were familiarized with the experimental setup and the objectives of the task. Familiarization consisted of five CST trials with moderate difficulty level. These trials were later excluded from the analyses. The main experiment consisted of three main phases that were repeated on each day. The first and second phases of the experiment involved 15 reaction time trials and 10 tracking trials, respectively (data for reaction time and tracking trials are not reported in this study). Phase three involved the CST trials, which were performed in three blocks. In Block 1, subjects performed 30 CST trials, where the difficulty level was determined in each trial using an up-down method: starting from

$\lambda = 2.5$ in the first trial, if subjects succeeded/failed on the current trial, λ was increased/decreased by $\Delta\lambda = 0.2$ in the next trial. By the end of Block 1, subjects had gradually converged to λ values in which the success rate was approximately 50%. This value was considered as the critical instability value (Quick et al., 2018), denoted by λ_c, and was obtained by averaging the λ's of the last 5 trials of block 1.

In Block 2, a stepwise increase in λ was adopted: subjects started with a difficulty level of $\lambda = 70\%\ \lambda_c$ (using λ_c from the previous block). They continued until they completed 10 successful trials, or 20 trials in total (whichever occurred first). The difficulty level was then increased by $\Delta\lambda = 0.2$, and the procedure repeated. This incremental increase of λ continued until the subjects' success rate for the ongoing λ dropped below 10% (i.e., less than 2 successful trials out of 20). This marked the end of the second block. In total, subjects performed approximately 120-200 trials in Block 2, depending on the individual's performance.

In Block 3, subjects performed the CST under three selected difficulty levels of easy, medium, and hard, with 20 trials for each difficulty level. These levels corresponded to λ values that led to 75% success rate (easy), 50% success rate (medium) and 25% success rate (hard) obtained from each individual's performance in Block 2. The exact values of $\lambda_{75\%}$, $\lambda_{50\%}$, and $\lambda_{25\%}$ were calculated by fitting a psychometric curve to the success rate data from Block 2 as a function of λ (see **Figure 2**). The order of difficulty was pseudo-randomly selected for each subject. For this study, we only analyzed the CST data from Block 2 (stepwise increase in λ) as it matched the procedure used in the monkey experiment (Quick et al., 2018). Subjects repeated the same experimental procedure on Day 2 and 3.

Three groups of human subjects participated in the experiment, where each group received different instructions about the task goal. The first group was instructed to perform the CST "without failing to the best of their ability" (no-instruction group); the second group was instructed to "keep the cursor at the center of the screen at all times" (Position Control group); and the third group was instructed to "keep the cursor still anywhere within the bounds of the screen" (velocity control group).

Analysis

To evaluate the overall performance of humans and monkeys during the CST, four quantities were calculated: success rate, hand-cursor correlation, hand-cursor lag, and hand/cursor gain. For each individual, the quantities were calculated as the average across trials for each bin of λ values (bin size: 0.3, starting from $\lambda = 1.5$).

The success rate was obtained as the percentage of successful trials within each λ bin. A psychometric curve (a Gaussian cumulative distribution function) was then fitted to the success rate data as a function of λ to estimate λ_c (critical stability, where success rate was 50%):

$$\%\ \text{Success} = 50\left[1 - \text{erf}\left(\frac{\lambda - \lambda_c}{\sqrt{2}\sigma}\right)\right] \qquad (3)$$

where, 'erf' indicates the error function, and σ denotes the standard deviation of the Gaussian cumulative. The correlation and lag quantities (**Figure 2**, B and C) were obtained by first cross-correlating the hand and cursor trajectories in each trial, and then finding the peak correlation, and the corresponding lag (**Figure 2**, see also (Quick et al., 2018)). The hand/cursor gain (**Figure 2**, D) was defined

as the ratio of the root mean squared (RMS) value of hand position over the RMS value of the cursor position in each trial.

Finally, to perform the classification analysis used in **Figure 7** and **Figure 8**, a Support Vector Machine method was applied to learn the two-class control objective labels. In order to build and train a classifier, we used 'fitcsvm.m' function in MATLAB, where synthetic data (RMS of cursor position and cursor velocity) was used as training set. To classify experimental data using the trained classifier, the MATLAB function 'predict.m' was used. Finally, the posterior probabilities over each classification (i.e., the confidence on classification) was calculated using the 'fitPosterior.m' function in MATLAB.

Optimal Feedback Control Model

A generative model approach was used to build control agents that performed the CST with different control strategies. The model involved an optimal feedback controller that moved the hand, a point mass of *m*=1kg, through a simple muscle-like actuator (Todorov, 2005; Todorov & Jordan, 2002).Click or tap here to enter text. The muscle model was approximated by a first-order low-pass filter that generated forces on the hand in the lateral direction as in equations (4) and (5):

$$\tau \dot{F} = -F + u \tag{4}$$

$$\ddot{p} = \frac{1}{m} F \tag{5}$$

where F is the actuator force acting on the hand, τ is the time constant of the low-pass filter, u is the control input to the muscle, and $\ddot{p}$ is the second derivative of the hand position. These equations consist of three states: hand position p, hand velocity $\dot{p}$ and muscle force F. Similarly, by taking first and second derivatives of equation (1), three more states were added to the dynamics of the system:

$$\ddot{x} = \lambda(\dot{x} + \dot{p}) \tag{6}$$

$$\dddot{x} = \lambda(\ddot{x} + \ddot{p}) \tag{7}$$

By combining equations (1) and (6) and then (6) and (7), the CST equation was expanded as follows:

$$\dddot{x} = \lambda^3 x + \lambda^3 p + \lambda^2 \dot{p} + \lambda \ddot{p} \tag{8}$$

The advantage of the higher derivatives of CST dynamics was that it made the cursor position x, cursor velocity $\dot{x}$ and cursor acceleration $\ddot{x}$ available to the controller. Hence, different control strategies that directly involved these states could be explored. Note that equation (8) required that the initial conditions of both hand and cursor position, velocity and acceleration all satisfied equations (6) and (7). The dynamics of the system could then be captured by equations (4), (5) and (8), and represented by the state vector: $\mathbf{x} = [x, \dot{x}, \ddot{x}, p, \dot{p}, F]$. By adding additive signal-independent noise $\boldsymbol{\xi}_t$, as well as multiplicative signal-dependent noise $\varepsilon_t C$, the full dynamics of the system could be presented in state-space format:

$$\mathbf{x}_{t+1} = A\mathbf{x}_t + B(1 + \varepsilon_t C)u_t + \boldsymbol{\xi}_t \tag{9}$$

where ε_t and $\boldsymbol{\xi}_t$ were zero-mean Gaussian noise terms, C was the signal-dependent noise scalar, and *A* and *B* represent the dynamics of the system:

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ \lambda^3 & 0 & 0 & \lambda^3 & \lambda^2 & \lambda \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & -1/\tau \end{bmatrix} \qquad (10)$$

$$B = [0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 1/(\tau M)]^T$$

Noisy sensory feedback $\boldsymbol{y}_t$ was given as:

$$\boldsymbol{y}_t = H\boldsymbol{x}_t + \boldsymbol{\omega}_t \qquad (11)$$

where $\boldsymbol{\omega}_t$ was a zero-mean additive Gaussian noise, and matrix *H* determined the available sensory feedback from the vector of states. For our simulations, the feedback included the cursor position x and velocity $\dot{x}$, therefore, *H* was defined as: $H = [1,1,0,0,0,0]$. An optimal controller determined the motor command $u_t$to minimize the cost function J as follows (Todorov, 2005):

$$u_t = argmin\ (J)$$
$$J = \sum_{t=1}^{n} \left(\mathbf{x}_t^T \boldsymbol{Q} \mathbf{x}_t \ + u_t^T R u_t\right) \qquad (12)$$

where n was the number of time samples throughout the movement, and $\boldsymbol{Q}$ and R determined the contribution of accuracy and effort cost, respectively. In all simulations, $R = 1$. The matrix $\boldsymbol{Q}$, however, was appropriately manipulated to implement different state-dependent control strategies as discussed below.

Position Control

The aim of the Position Control strategy was to maintain the cursor at the center of the screen throughout the trial. This was implemented by penalizing the deviation of the cursor position x from the center. In this case, the matrix $\boldsymbol{Q}$ was set to $\boldsymbol{Q} = \mathrm{diag}([q, 1, 1,1,1,1])$, where $q \gg 1$ was a constant. As such, the cost of deviation from the center for the cursor position was dominant represented in the value J of the cost function, making the regulation of cursor position at the center the primary goal of control.

Velocity Control

The Velocity Control strategy aimed to keep the cursor still at any point within the boundaries of the workspace. In this case, upon deviation of the cursor from the center, the main goal was to bring the cursor to a stop regardless of the location. This was implemented through penalizing the cursor velocity $\dot{x}$ by setting the matrix $\boldsymbol{Q} = \mathrm{diag}([1, v, 1,1,1,1])$, where $v \gg 1$ was a constant.

Simulations

Given a control strategy, the model was used to generate 500 trials of CST for each level of task difficulty from $\lambda = 1.5$ to $\lambda = 7$, with increments of $\Delta\lambda = 0.2$. The parameters of the hand and the muscle model (4)(5) were fixed to $m = 1$kg and $\tau = 0.06$s. A sensory delay of 50ms was considered when simulating the task with the optimal feedback controller (Todorov, 2005). The signal-dependent noise terms were set to $\varepsilon_t \sim N(0,1)$, and $C = 1.5$. The motor noise was $\boldsymbol{\xi}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma} = 0.4\, BB^T$. For each trial, the simulation started from the initial condition of $\mathbf{x} = \mathbf{0}$, and ran for 8s. Only the first 6s of each simulation were considered in the analysis for consistency with the experimental paradigm. The success or failure in each simulated trial was decided post hoc, by determining whether the cursor position x exceeded the limits of the workspace (±10cm from the center) within the 6s duration of the trial.

Acknowledgement

This research was funded by the National Institute of Health R01-CRCNS-NS120579, awarded to Dagmar Sternad and Aaron Batista. Dagmar Sternad was also supported by NIH-R37-HD087089 and NSF-M3X-1825942. Aaron Batista and Patrick Loughlin were also supported by NIH-R01-HD0909125.

References

Badre, D., Frank, M. J., & Moore, C. I. (2015). Interactionist Neuroscience. *Neuron*, *88*(5), 855–860. https://doi.org/10.1016/J.NEURON.2015.10.021

Batista, A. P., Buneo, C. A., Snyder, L. H., & Andersen, R. A. (1999). Reach Plans in Eye-Centered Coordinates. *Science*, *285*(5425), 257–260. https://doi.org/10.1126/science.285.5425.257

Benyamini, M., & Zacksenhouse, M. (2015). Optimal feedback control successfully explains changes in neural modulations during experiments with brain-machine interfaces. *Frontiers in Systems Neuroscience*, *9*. https://doi.org/10.3389/fnsys.2015.00071

Braun, D. A., Aertsen, A., Wolpert, D. M., & Mehring, C. (2009). Learning optimal adaptation strategies in unpredictable motor tasks. *Journal of Neuroscience*, *29*(20), 6472–6478. https://doi.org/10.1523/JNEUROSCI.3075-08.2009

Cisek, P., Crammond, D. J., & Kalaska, J. F. (2003). Neural Activity in Primary Motor and Dorsal Premotor Cortex In Reaching Tasks With the Contralateral Versus Ipsilateral Arm. *Journal of Neurophysiology*, *89*(2), 922–942. https://doi.org/10.1152/jn.00607.2002

Crevecoeur, F., Scott, S. H., & Cluff, T. (2019). Robust Control in Human Reaching Movements: A Model-Free Strategy to Compensate for Unpredictable Disturbances. *Journal of Neuroscience*, *39*(41), 8135–8148. https://doi.org/10.1523/JNEUROSCI.0770-19.2019

Cross, K. P., Guang, H., & Scott, S. H. (2023). Proprioceptive and Visual Feedback Responses in Macaques Exploit Goal Redundancy. *Journal of Neuroscience*, *43*(5), 787–802. https://doi.org/10.1523/JNEUROSCI.1332-22.2022

Croxson, P. L., Walton, M. E., O'Reilly, J. X., Behrens, T. E. J., & Rushworth, M. F. S. (2009). Effort-Based Cost-Benefit Valuation and the Human Brain. *Journal of Neuroscience*, *29*(14), 4531–4541. https://doi.org/10.1523/JNEUROSCI.4515-08.2009

Desrochers, T. M., Burk, D. C., Badre, D., & Sheinberg, D. L. (2016). The Monitoring and Control of Task Sequences in Human and Non-Human Primates. *Frontiers in Systems Neuroscience*, *9*. https://doi.org/10.3389/fnsys.2015.00185

Diedrichsen, J. (2007). Optimal Task-Dependent Changes of Bimanual Feedback Control and Adaptation. *Current Biology*, *17*(19), 1675–1679. https://doi.org/10.1016/j.cub.2007.08.051

Diedrichsen, J., Shadmehr, R., & Ivry, R. B. (2010). The coordination of movement: optimal feedback control and beyond. *Trends in Cognitive Sciences*, *14*(1), 31–39. https://doi.org/10.1016/j.tics.2009.11.004

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal Population Coding of Movement Direction. *Science*, *233*(4771), 1416–1419. https://doi.org/10.1126/science.3749885

Hwang, E. J., Smith, M. A., & Shadmehr, R. (2006). Adaptation and generalization in acceleration-dependent force fields. *Experimental Brain Research*, *169*(4), 496–506. https://doi.org/10.1007/s00221-005-0163-2

Izawa, J., Rane, T., Donchin, O., & Shadmehr, R. (2008). Motor adaptation as a process of reoptimization. *Journal of Neuroscience*, *28*(11), 2883–2891. https://doi.org/10.1523/JNEUROSCI.5359-07.2008

Jex, H. R., McDonnell, J. D., & Phatak, A. v. (1966). A ``Critical'' Tracking Task for Manual Control Research. *IEEE Transactions on Human Factors in Electronics*, *HFE-7*(4), 138–145. https://doi.org/10.1109/THFE.1966.232660

Kalidindi, H. T., Cross, K. P., Lillicrap, T. P., Omrani, M., Falotico, E., Sabes, P. N., & Scott, S. H. (2021). Rotational dynamics in motor cortex are consistent with a feedback controller. *ELife*, *10*. https://doi.org/10.7554/eLife.67256

Kao, T.-C., Sadabadi, M. S., & Hennequin, G. (2021). Optimal anticipatory control as a theory of motor preparation: A thalamo-cortical circuit model. *Neuron*, *109*(9), 1567-1581.e12. https://doi.org/10.1016/j.neuron.2021.03.009

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, *9*(6), 718–727. https://doi.org/10.1016/S0959-4388(99)00028-8

Liu, D., & Todorov, E. (2007). Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *Journal of Neuroscience*, *27*(35), 9354–9368. https://doi.org/10.1523/JNEUROSCI.1110-06.2007

McNamee, D., & Wolpert, D. M. (2019). Internal Models in Biological Control. *Annual Review of Control, Robotics, and Autonomous Systems*, *2*(1), 339–364. https://doi.org/10.1146/annurev-control-060117-105206

Miall, R. C., Christensen, L. O. D., Cain, O., & Stanley, J. (2007). Disruption of State Estimation in the Human Lateral Cerebellum. *PLoS Biology*, *5*(11), e316. https://doi.org/10.1371/journal.pbio.0050316

Nagengast, A. J., Braun, D. A., & Wolpert, D. M. (2009). Optimal control predicts human performance on objects with internal degrees of freedom. *PLoS Computational Biology*, *5*(6), e1000419. https://doi.org/10.1371/journal.pcbi.1000419

Nagengast, A. J., Braun, D. A., & Wolpert, D. M. (2010). Risk-Sensitive Optimal Feedback Control Accounts for Sensorimotor Behavior under Uncertainty. *PLoS Computational Biology*, *6*(7), e1000857. https://doi.org/10.1371/journal.pcbi.1000857

Nashed, J. Y., Crevecoeur, F., & Scott, S. H. (2014). Rapid Online Selection between Multiple Motor Plans. *Journal of Neuroscience*, *34*(5), 1769–1780. https://doi.org/10.1523/JNEUROSCI.3063-13.2014

Nayeem, R., Bazzi, S., Sadeghi, M., Hogan, N., & Sternad, D. (2021). Preparing to move: Setting initial conditions to simplify interactions with complex objects. *PLoS Computational Biology*, *17*(12), e1009597. https://doi.org/10.1371/journal.pcbi.1009597

Omrani, M., Murnaghan, C. D., Pruszynski, J. A., & Scott, S. H. (2016). Distributed task-specific processing of somatosensory feedback for voluntary motor control. *ELife*, *5*. https://doi.org/10.7554/eLife.13141

Pruszynski, J. A., Kurtzer, I., Nashed, J. Y., Omrani, M., Brouwer, B., & Scott, S. H. (2011). Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature*, *478*(7369), 387–390. https://doi.org/10.1038/nature10436

Pruszynski, J. A., & Scott, S. H. (2012). Optimal feedback control and the long-latency stretch response. *Experimental Brain Research*, *218*(3), 341–359. https://doi.org/10.1007/s00221-012-3041-8

Quick, K. M., Card, N. S., Whaite, S. M., Mischel, J., Loughlin, P., & Batista, A. P. (2014). Assessing vibrotactile feedback strategies by controlling a cursor with unstable dynamics. *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2589–2592. https://doi.org/10.1109/EMBC.2014.6944152

Quick, K. M., Mischel, J. L., Loughlin, P. J., & Batista, A. P. (2018). The critical stability task: quantifying sensory-motor control during ongoing movement in nonhuman primates. *Journal of Neurophysiology*, *120*(5), 2164–2181. https://doi.org/10.1152/jn.00300.2017

Rajalingham, R., Piccato, A., & Jazayeri, M. (2022). Recurrent neural networks with explicit representation of dynamic latent variables can mimic behavioral patterns in a physical inference task. *Nature Communications*, *13*(1), 5865. https://doi.org/10.1038/s41467-022-33581-6

Razavian, R. S., Sadeghi, M., Bazzi, S., Nayeem, R., & Sternad, D. (2023). Body Mechanics, Optimality, and Sensory Feedback in the Human Control of Complex Objects. *Neural Computation*, *in press*, 1–43. https://doi.org/10.1162/neco_a_01576

Ronsse, R., Wei, K., & Sternad, D. (2010). Optimal Control of a Hybrid Rhythmic-Discrete Task: The Bouncing Ball Revisited. *Journal of Neurophysiology*, *103*(5), 2482–2493. https://doi.org/10.1152/jn.00600.2009

Scott, S. H. (2004). Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews. Neuroscience*, *5*(7), 532–546. https://doi.org/10.1038/nrn1427

Scott, S. H., & Kalaska, J. F. (1997). Reaching Movements With Similar Hand Paths But Different Arm Orientations. I. Activity of Individual Cells in Motor Cortex. *Journal of Neurophysiology*, *77*(2), 826–852. https://doi.org/10.1152/jn.1997.77.2.826

Sing, G. C., Joiner, W. M., Nanayakkara, T., Brayanov, J. B., & Smith, M. A. (2009). Primitives for motor adaptation reflect correlated neural tuning to position and velocity. *Neuron*, *64*(4), 575–589. https://doi.org/10.1016/j.neuron.2009.10.001

Takei, T., Lomber, S. G., Cook, D. J., & Scott, S. H. (2021). Transient deactivation of dorsal premotor cortex or parietal area 5 impairs feedback control of the limb in macaques. *Current Biology*, *31*(7), 1476-1487.e5. https://doi.org/10.1016/j.cub.2021.01.049

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, *7*(9), 907–915. https://doi.org/10.1038/nn1309

Todorov, E. (2005). Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation*, *17*(5), 1084–1108. https://doi.org/10.1162/0899766053491887

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

Uno, Y., Kawato, M., & Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement. *Biological Cybernetics*, *61*(2), 89–101. https://doi.org/10.1007/BF00204593

Wong, J. D., Cluff, T., & Kuo, A. D. (2021). The energetic basis for smooth human arm movements. *ELife*, *10*. https://doi.org/10.7554/eLife.68013

Yeo, S.-H., Franklin, D. W., & Wolpert, D. M. (2016). When Optimal Feedback Control Is Not Enough: Feedforward Strategies Are Required for Optimal Control with Active Sensing. *PLoS Computational Biology*, *12*(12), e1005190. https://doi.org/10.1371/journal.pcbi.1005190

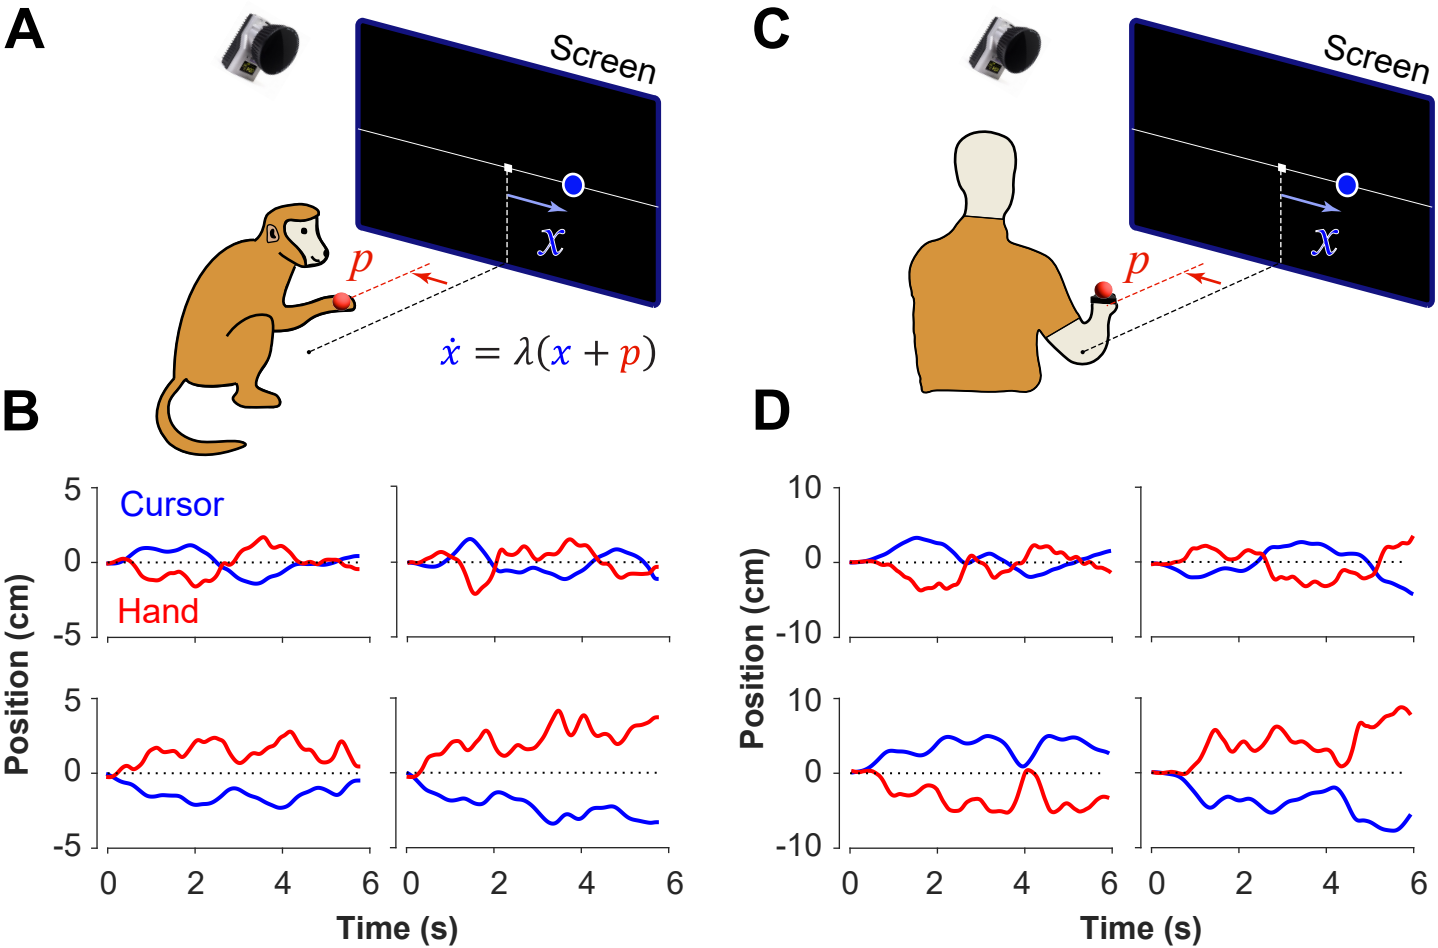
A
Screen
x
p
$\dot{x} = \lambda(x + p)$
C
Screen
x
p
B
Cursor
Hand
Position (cm)
5
0
-5
Time (s)
0
2
4
6
D
Position (cm)
10
0
-10
Time (s)

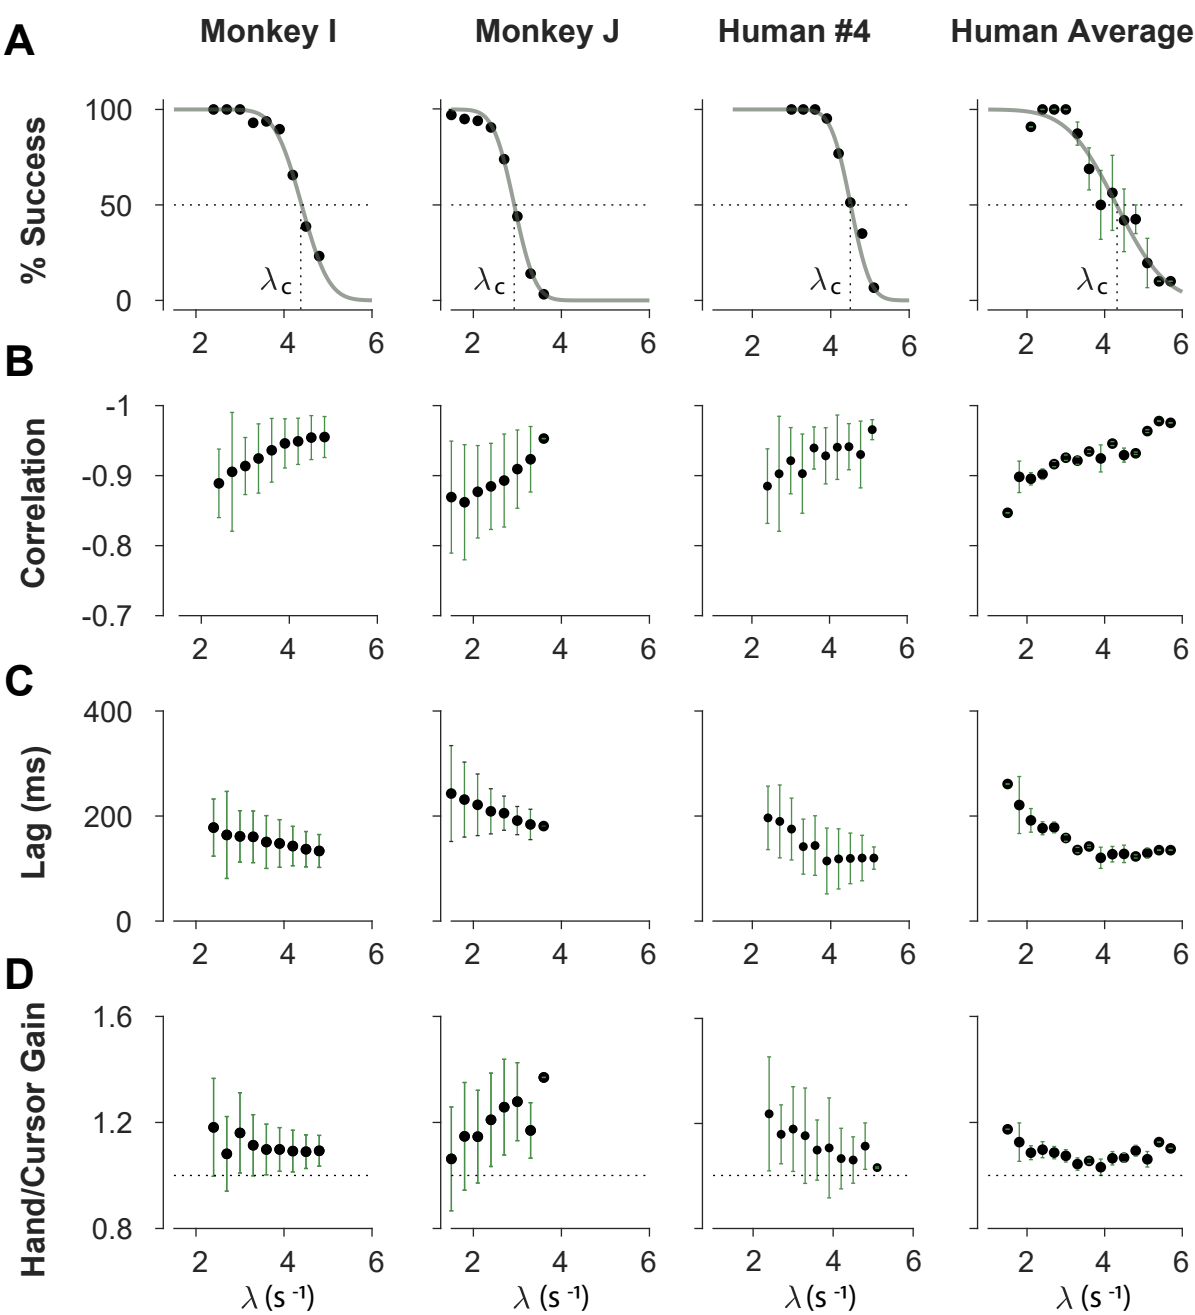
A
Monkey I
Monkey J
Human #4
Human Average
% Success
100
50
0
λc
B
Correlation
-1
-0.9
-0.8
-0.7
C
Lag (ms)
400
200
0
D
Hand/Cursor Gain
1.6
1.2
0.8
2
4
6
λ (s -1)

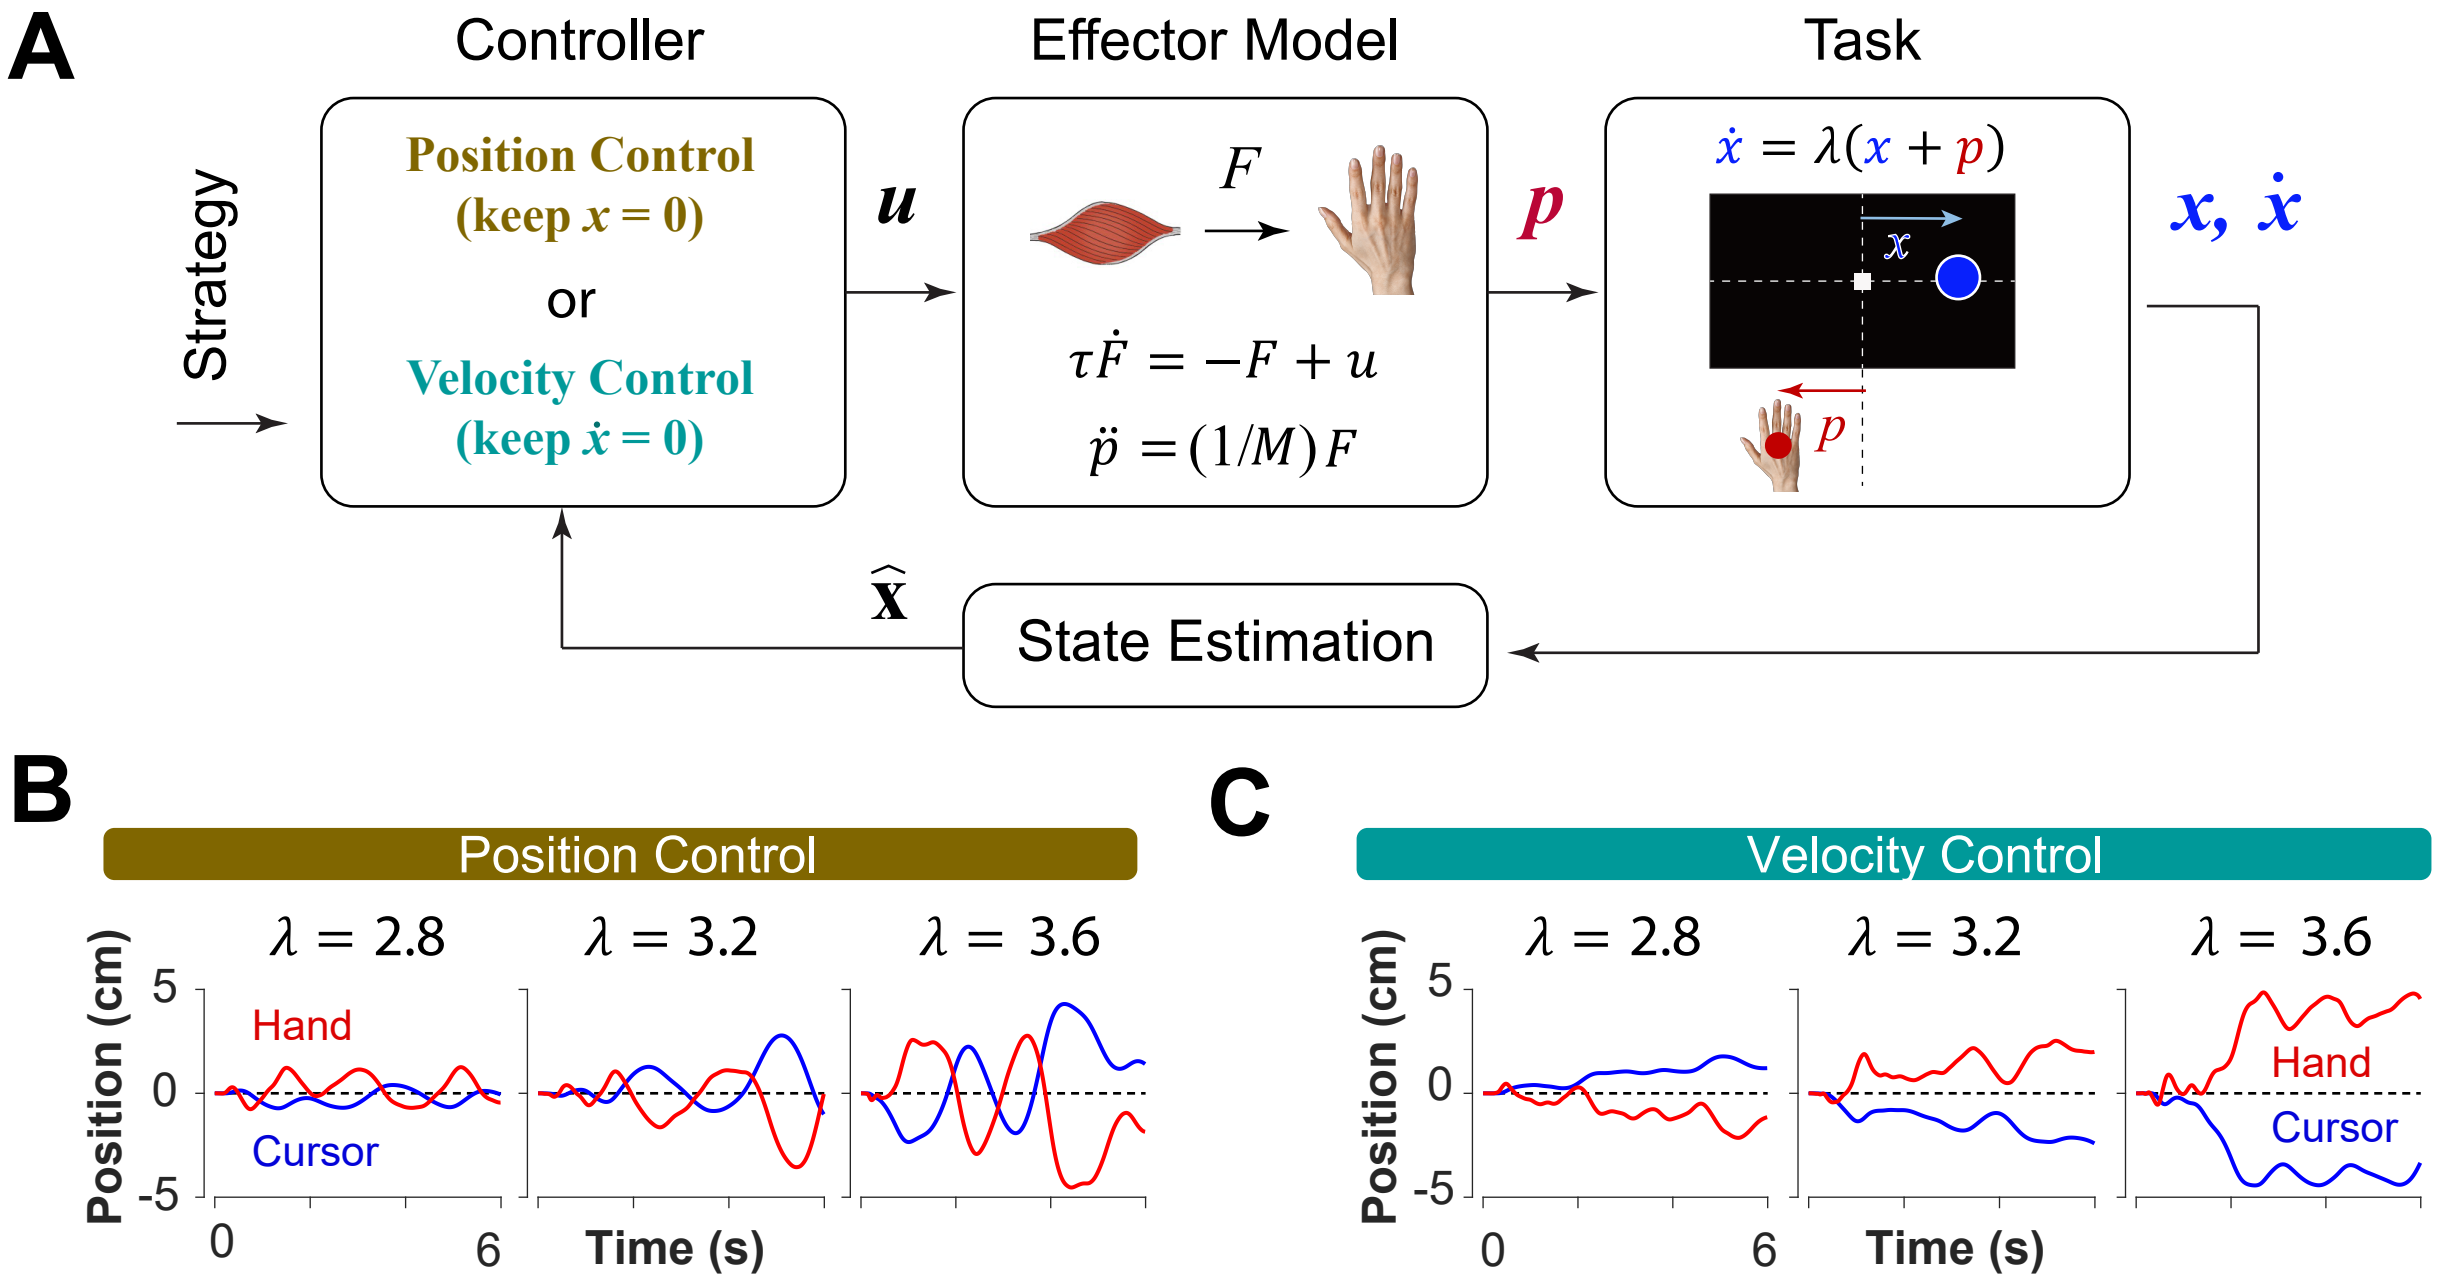
A
Controller
Effector Model
Task
Strategy
Position Control
(keep $x = 0$)
or
Velocity Control
(keep $\dot{x} = 0$)
u
F
$\tau\dot{F} = -F + u$
$\ddot{p} = (1/M)F$
p
$\dot{x} = \lambda(x + p)$
x
p
$x, \dot{x}$
$\hat{\mathbf{x}}$
State Estimation
B
Position Control
$\lambda = 2.8$
$\lambda = 3.2$
$\lambda = 3.6$
Position (cm)
5
0
-5
Hand
Cursor
0
6
Time (s)
C
Velocity Control
$\lambda = 2.8$
$\lambda = 3.2$
$\lambda = 3.6$
Position (cm)
5
0
-5
Hand
Cursor
0
6
Time (s)

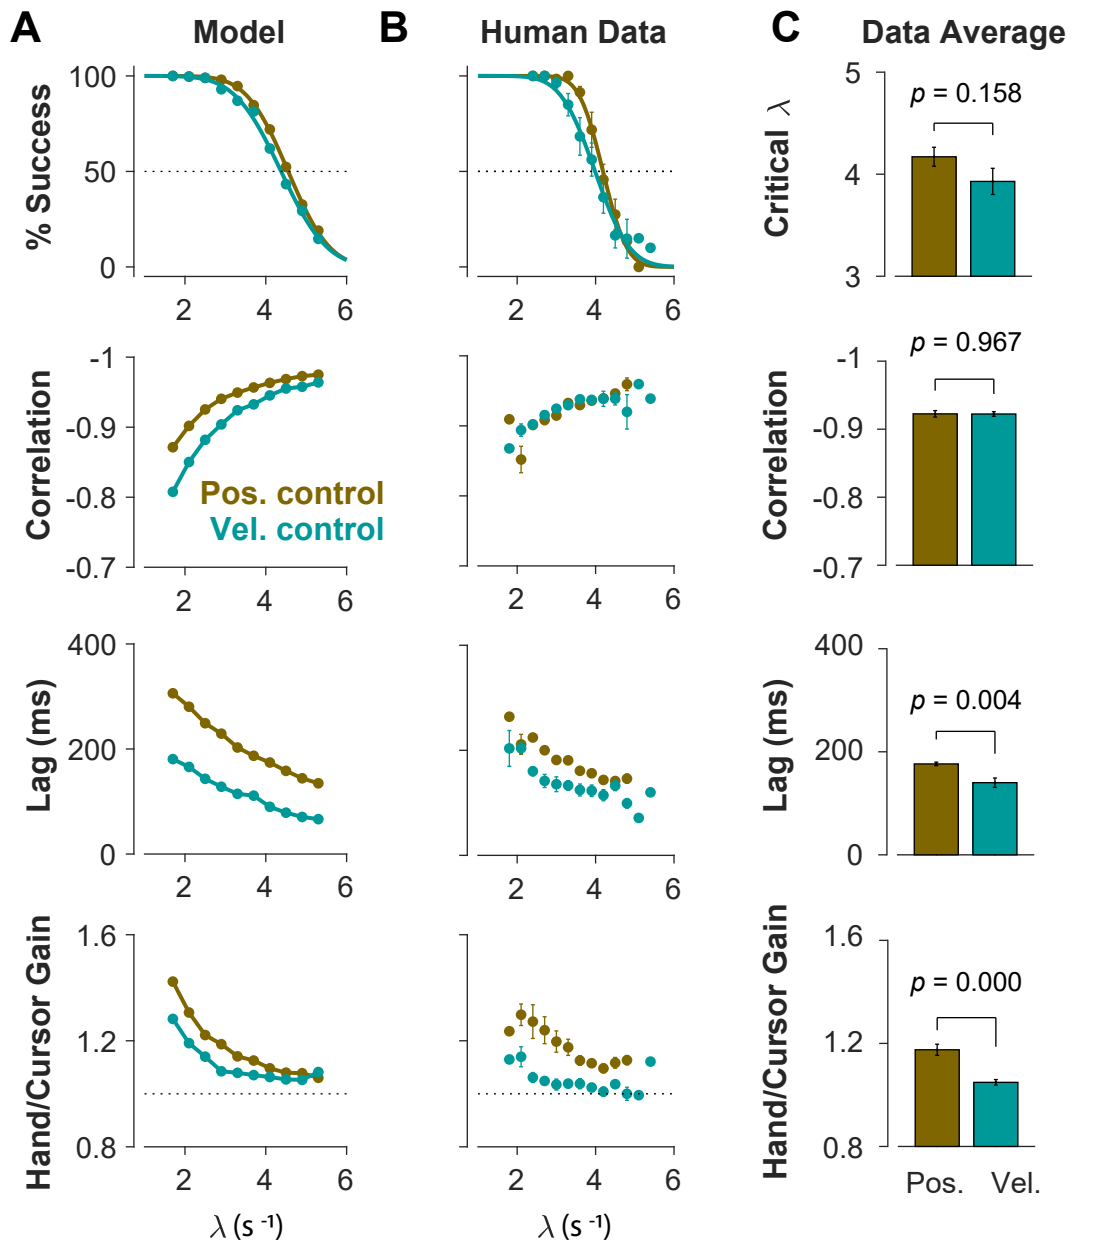
A
Model
B
Human Data
C
Data Average
% Success
100
50
0
Correlation
-1
-0.9
-0.8
-0.7
Pos. control
Vel. control
Lag (ms)
400
200
0
Hand/Cursor Gain
1.6
1.2
0.8
2
4
6
λ (s -1)
Critical λ
5
4
3
p = 0.158
p = 0.967
p = 0.004
p = 0.000
Pos.
Vel.

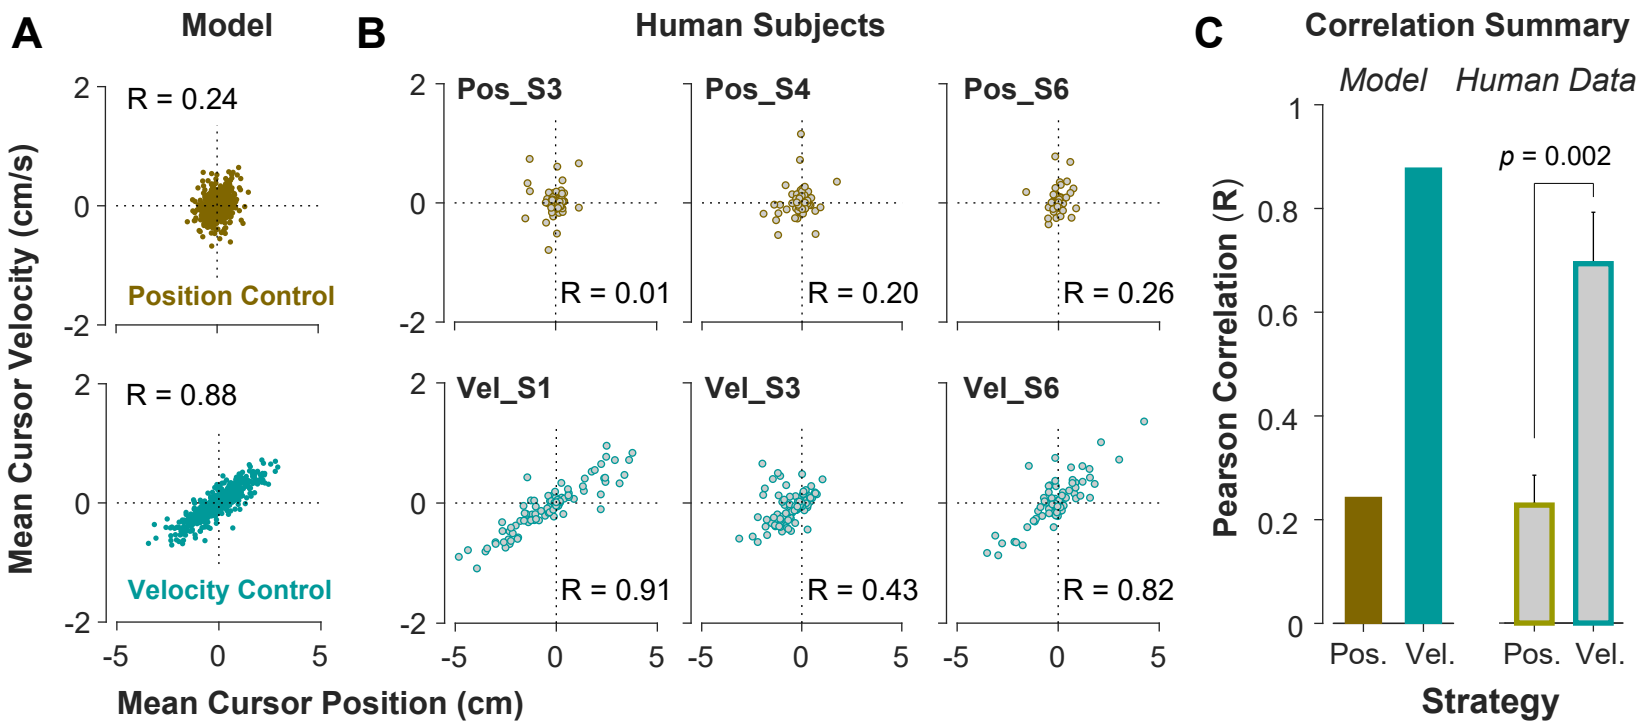
A
Model
R = 0.24
Position Control
R = 0.88
Velocity Control
Mean Cursor Velocity (cm/s)
Mean Cursor Position (cm)
B
Human Subjects
Pos_S3
R = 0.01
Pos_S4
R = 0.20
Pos_S6
R = 0.26
Vel_S1
R = 0.91
Vel_S3
R = 0.43
Vel_S6
R = 0.82
C
Correlation Summary
Model
Human Data
p = 0.002
Pearson Correlation (R)
Pos.
Vel.
Pos.
Vel.
Strategy

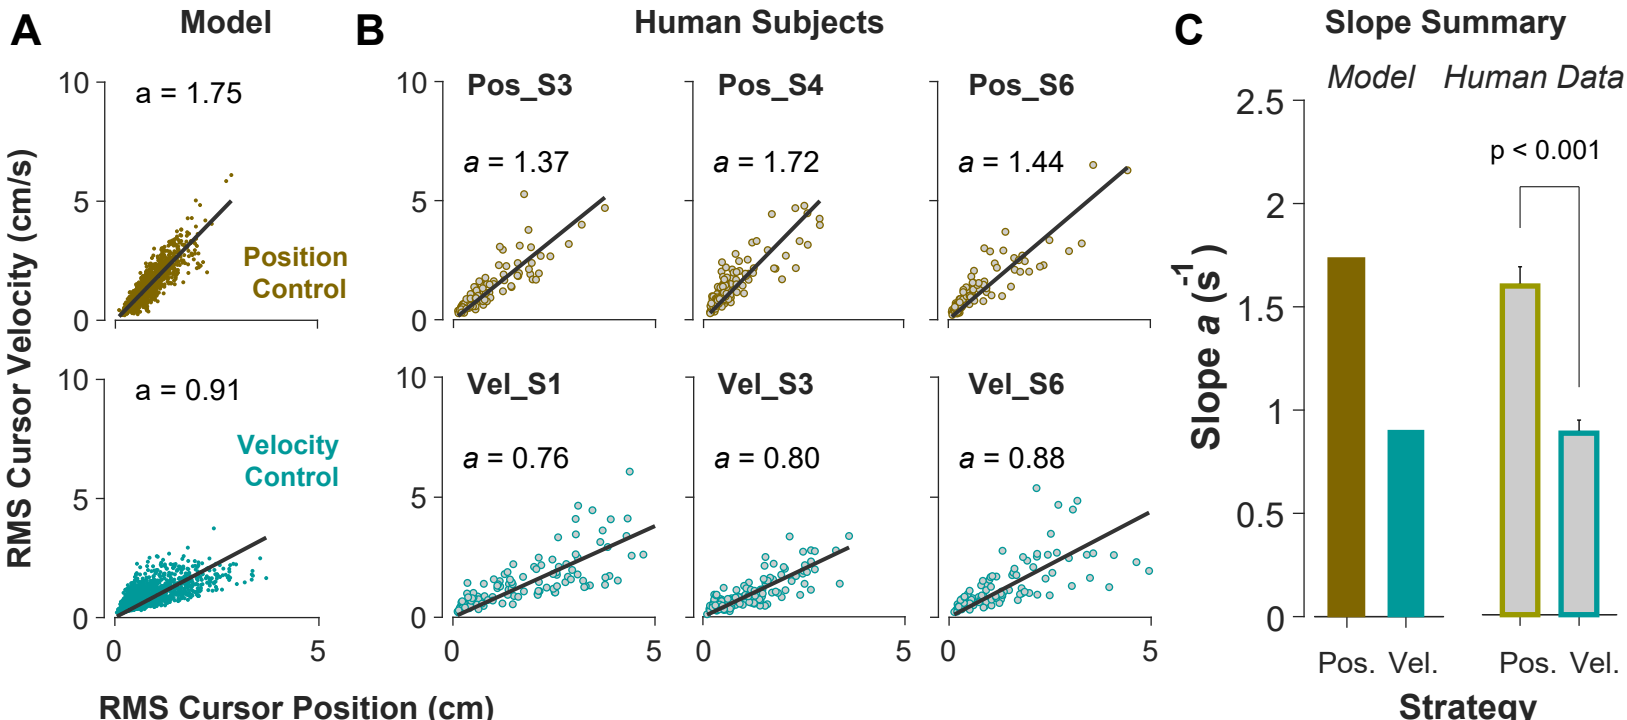
A
Model
a = 1.75
Position Control
a = 0.91
Velocity Control
RMS Cursor Velocity (cm/s)
RMS Cursor Position (cm)
B
Human Subjects
Pos_S3
a = 1.37
Pos_S4
a = 1.72
Pos_S6
a = 1.44
Vel_S1
a = 0.76
Vel_S3
a = 0.80
Vel_S6
a = 0.88
C
Slope Summary
Model
Human Data
p < 0.001
Slope a (s-1)
Pos.
Vel.
Pos.
Vel.
Strategy

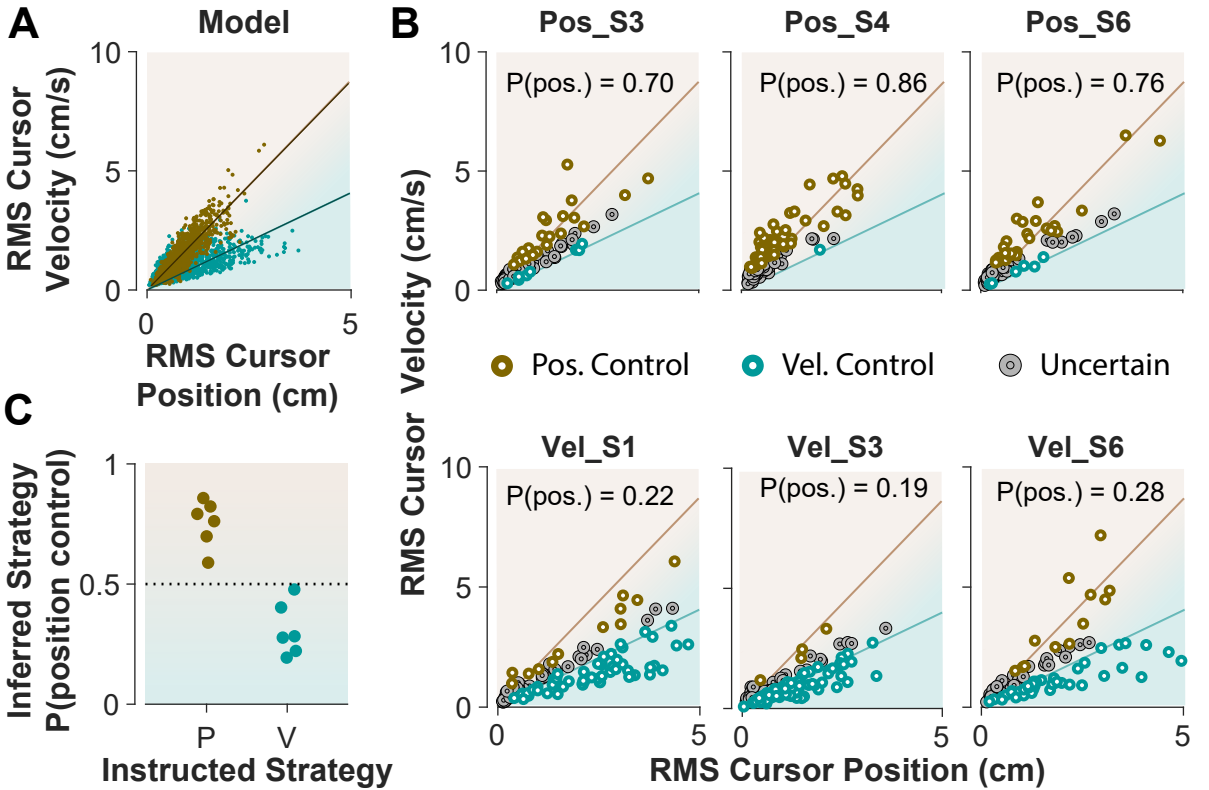
A
Model
RMS Cursor Velocity (cm/s)
RMS Cursor Position (cm)
B
Pos_S3
Pos_S4
Pos_S6
P(pos.) = 0.70
P(pos.) = 0.86
P(pos.) = 0.76
Pos. Control
Vel. Control
Uncertain
Vel_S1
Vel_S3
Vel_S6
P(pos.) = 0.22
P(pos.) = 0.19
P(pos.) = 0.28
RMS Cursor Velocity (cm/s)
RMS Cursor Position (cm)
C
Inferred Strategy P(position control)
Instructed Strategy
P
V

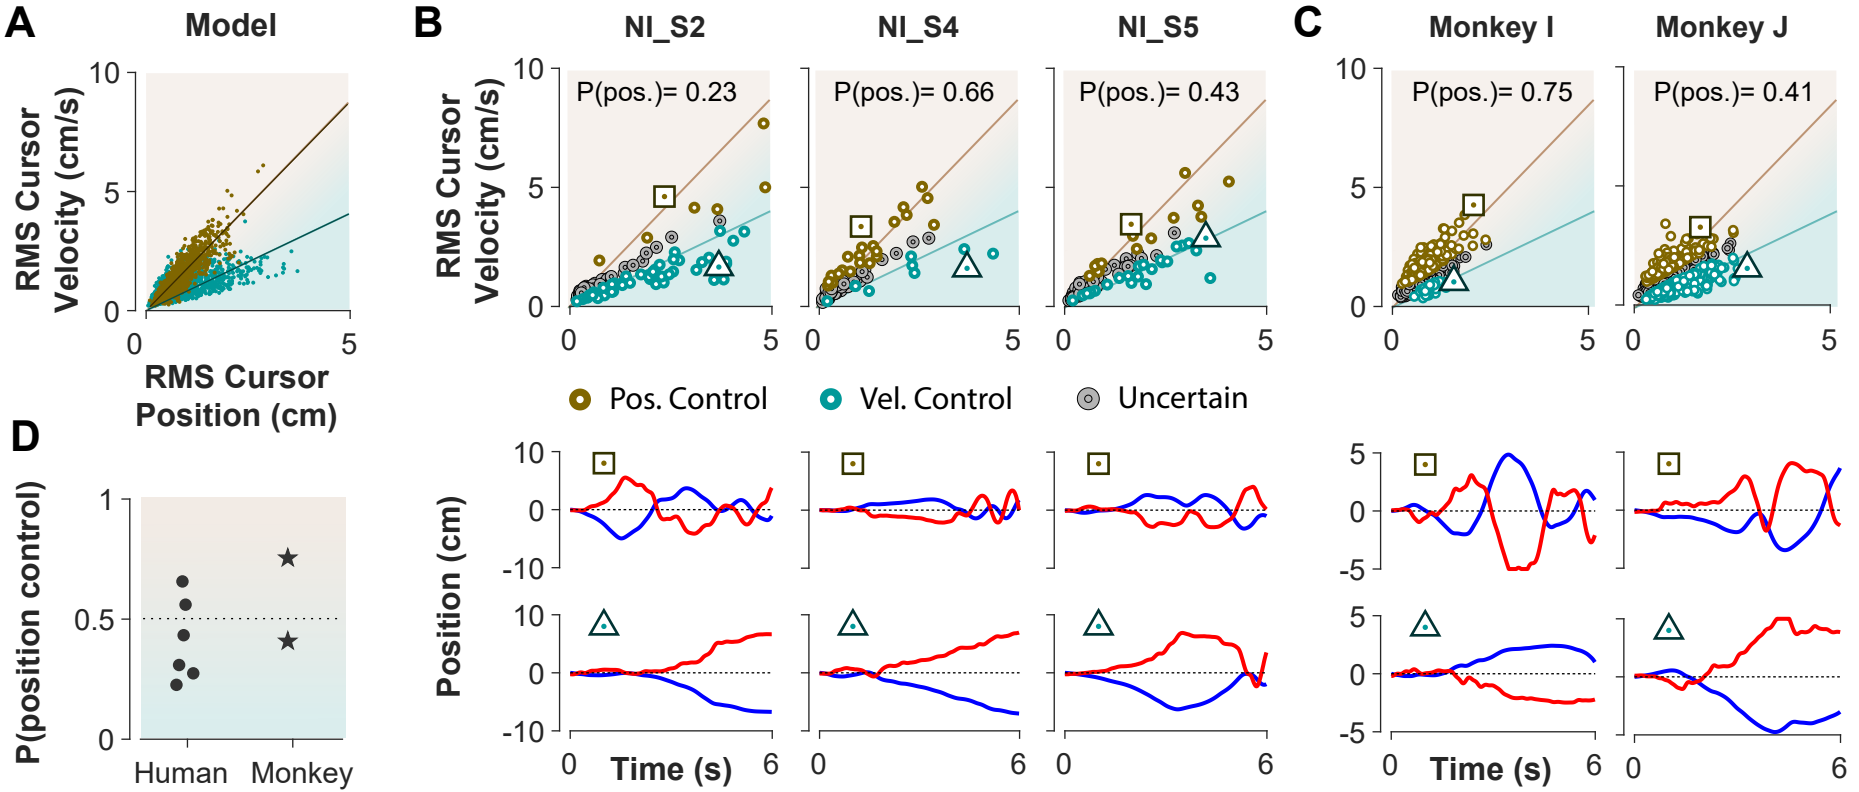
A
Model
RMS Cursor Velocity (cm/s)
RMS Cursor Position (cm)
B
NI_S2
NI_S4
NI_S5
P(pos.)= 0.23
P(pos.)= 0.66
P(pos.)= 0.43
C
Monkey I
Monkey J
P(pos.)= 0.75
P(pos.)= 0.41
Pos. Control
Vel. Control
Uncertain
Position (cm)
Time (s)
D
P(position control)
Human
Monkey