Does the Sensorimotor System Minimize Prediction Error or Select the Most Likely Prediction During Object Lifting?

Joshua G. A. Cashaback[1], Heather R. McGregor[1,2], Henry C. H. Pun[3], Gavin Buckingham[4], Paul L. Gribble[1,3]

[1] Brain and Mind Institute, Department of Psychology, Western University
[2] Graduate Program in Neuroscience, Western University
[3] Department of Physiology and Pharmacology, Western University
[4] Department of Sport and Health Sciences, University of Exeter

Abbreviated Title: Making feedforward predictions in an uncertain environment.

Acknowledgements: Canadian Institutes of Health Research (CIHR) and the Natural Sciences and Engineering Council of Canada (NSERC).

The authors declare no competing financial interests.

Please Address Correspondence to:

Joshua G. A. Cashaback
Brain and Mind Institute, Dept. of Psychology
Western University
1151 Richmond St., London, ON, Canada
N6A 5B7
Email: cashabackjga@gmail.com

ABSTRACT

The human sensorimotor system is routinely capable of making accurate predictions about an object's weight, which allows for energetically efficient lifts and prevents objects from being dropped. Often however, poor predictions arise when the weight of an object can vary and sensory cues about object weight are sparse (e.g., picking up an opaque water bottle). The question arises, what strategies does the sensorimotor system use to make weight predictions when dealing with an object whose weight may vary? For example, does the sensorimotor system use a strategy that minimizes prediction error (minimal squared error) or one that selects the weight that is most likely to be correct (maximum a posteriori)? Here we dissociated the predictions of these two strategies by having participants lift an object whose weight varied according to a skewed probability distribution. We found, using a small range of weight uncertainty, that four indexes of sensorimotor prediction (grip force rate, grip force, load force rate, and load force) were consistent with a feedforward strategy that minimizes the square of prediction errors. These findings match research in the visuomotor system, suggesting parallels in underlying processes. We interpret our findings within a Bayesian framework and discuss the potential benefits of using a minimal squared error strategy.

KEYWORDS

Object lifting, Fingertip Force, Feedforward control, Prediction, Bayesian

NEW AND NOTEWORTHY

Using a novel experimental model of object lifting, we tested whether the sensorimotor system models the weight of objects by minimizing lifting errors, or by selecting the statistically most likely weight. We found that the sensorimotor system minimizes the square of prediction errors for object lifting. This parallels the results of studies that investigated visually guided reaching, suggesting an overlap in the underlying mechanisms between tasks that involve different sensory systems.

INTRODUCTION

Humans are remarkably adept at lifting and manipulating the hundreds of objects they interact with on a daily basis. To do so, we rely on relatively accurate predictions of an object's weight (Flanagan et al., 2006; Johansson and Flanagan, 2009; Johansson and Westling, 1988; Wolpert and Flanagan 2001). Prior knowledge from handling similar objects is integrated with sensory information about object size (Gordon et al, 1991a, b, c), material (Buckingham et al., 2009, 2010), shape (Jenmalm and Johansson, 1997) and density (Grandy and Westwood, 2006; Peters et al., 2016), to make a feedforward prediction of object weight (Buckingham and Goodale, 2010; Brayanov and Smith, 2010; Hermsdorfer et al., 2011). Often however, feedforward prediction errors can arise from having imperfect prior knowledge (e.g., environmental uncertainty), and also from misleading or sparse current information about an object's weight (Buckingham and Goodale, 2010; Brayanov and Smith, 2010; Buckingham et al., 2011).

When lifting an object of constant weight, humans can quickly reduce prediction errors within 2-3 lifts (Johansson and Westling, 1984). However, humans often operate in highly uncertain environments, making it impossible to make an accurate feedforward prediction on every lift. For example, a baggage handler at an airport must grasp and lift luggage for which the contents are not visible. If the baggage handler underestimates the true weight of the luggage it will not leave the ground or, if lifted, may slip from their grasp. Conversely, if weight is overestimated the luggage will accelerate at a much faster rate than predicted and will be gripped too tightly, both of which are energetically inefficient. Thus, given a lack of useful visual cues, the baggage handler must rely heavily on prior knowledge of the uncertainty associated with luggage weight. This will allow him or her to apply relatively appropriate lift and grip forces to efficiently move the luggage. In the presence of such environmental uncertainty, what strategy does the sensorimotor system employ to make a feedforward prediction? Two viable strategies to deal with environmental uncertainty are: 1) to minimize the squared error of potential feedforward predictions (Kording and Wolpert, 2004b), or 2) to select the feedforward prediction that is most likely to be correct (Peters et al., 2016).

Briefly, a minimal squared error strategy applies a quadratic penalization for linear increases in error magnitude. A feedforward prediction that minimizes squared error can be accomplished in many ways. For example, a minimal squared error strategy can be achieved by averaging somatosensory information from a single (Johansson and Westling, 1984) or several (Takahashi et al., 2001; Scheidt et al., 2001; Landy et al., 2012; Hadjiosif and Smith, 2015) previous lift(s) to predict the weight of a subsequent lift. A minimal squared error strategy can also be achieved using a Bayesian framework (Kording and Wolpert, 2004b; Zhang et al., 2015). Here the nervous system would have to build a representation of environmental uncertainty based on the somatosensory information gained from many previous lifts (Kording and Wolpert, 2004a). The attractiveness of the Bayesian framework is that it can account for many more behavioural features than a model based on simply averaging previous trials (Acerbi et al., 2014), such as reduced variability with practice (Kording and Wolpert, 2004a) and explaining perceptual illusions (Peters et al., 2016). Furthermore, in this framework environmental uncertainty can be integrated with available sensory information (e.g., object size, material, shape, density and other cues) to assign a probability to each possible weight that an object may have (Peters et al., 2016).

Ultimately however, the sensorimotor system must select a single weight, or ‘point estimate’, when forming a feedforward response to attempt to lift an object. One such point estimate corresponds to that generated by a minimal squared error strategy. While minimizing squared error does well to explain many patterns of behaviour (Scheidt et al., 2001; Kording and Wolpert, 2004b; Zhang et al., 2015), there are examples in the literature that suggest a departure from this strategy.

Instances in which the sensorimotor system departs from a minimal squared error strategy may occur when the controller attempts to predict the most likely occurrence. Again using a Bayesian framework, the point estimate that predicts the most likely occurrence is termed the maximum a posteriori estimate. As proposed by Wolpert (2007), there are likely many tasks in which the sensorimotor system may use a maximum a posteriori strategy, such as when maximizing externally provided reward (Trommerhausser et al., 2003), Mawase and Karniel (2010) provide evidence supporting the idea that the sensorimotor system may attempt to correctly predict the most likely weight of an object. The authors found that when participants experienced a sequential increase in object weight in a series of trials, they unconsciously and reliably predicted a heavier object weight on subsequent lifts. This predictive behaviour cannot be obtained using a model of object weight that relies on a minimal squared error estimate, but is consistent with a feedforward controller that predicts the weight of an object using a maximum a posteriori estimate (Mawase and Karniel, 2010; Karniel, 2011).

A challenge in attempting to determine whether a controller is using a minimal squared error or a maximum a posteriori strategy is that the optimal solutions of these two strategies often coincide. A feedforward controller using a minimal squared error strategy would, over many trials, converge on a prediction of object weight based on the statistical mean of the environment uncertainty. A controller that uses a maximum a posteriori strategy would base its prediction on the statistical mode of the environment uncertainty. In many experimental designs the stimuli, such as visual displacement or object weight, are held constant or they vary according to a symmetrical (e.g., Gaussian, bimodal or uniform) probability distribution. With constant (Gordan et al., 1993a) or Gaussian (Kording et al., 2004; Kording and Wolpert, 2004a; Hadjiosif and Smith, 2015) stimuli the mean and mode are identical, making it impossible to distinguish if the feedforward controller is using a minimal squared error or maximum a posteriori strategy. Further, another issue arises when stimuli are varied using uniform (Berg et al., 2016) or bimodal probability distributions (Scheidt et al., 2001; Kording and Wolpert, 2004a) that have an ill-defined mode. However, skewed probability distributions can be used to separate a well-defined mean and mode (Kording and Wolpert, 2004b).

To our knowledge, no one has varied object weight in a lifting task using a skewed distribution. By varying an object's weight according to a skewed probability distribution in which the mean and mode are distinct, we were able to dissociate the minimal squared error and maximum a posteriori point estimates. This dissociation allowed us to test whether the sensorimotor system uses a minimal squared error or maximum a posteriori strategy to make feedforward predictions of object weight.

METHODS

Participants.

90 healthy participants (age: 20.3 *yr*, 2.7 SD) participated in this experiment. Participants reported they were right-handed, free of neuromuscular disease, and had normal or corrected vision. Each participant was paid $10.00 CAN, and provided informed consent to procedures approved by Western University's Ethics Board.

Apparatus

A pair of six degree-of-freedom force transducers (ATI Industrial Automation, F/T model Nano17, North Carolina, United States) recorded forces and moments acting on three orthogonal axes. A digital computer with an A/D board (16-bit; National Instruments, model NI PCI-6033E, Texas, United States) sampled force transducer data at 770 *Hz*. The transducers were mounted to the top of a wooden platform that covered a hole in a table. **(Fig.1A,B)**. A metal cable attached to the bottom of the wooden platform was positioned under the centroid of the force transducer. This cable passed vertically (in line with the gravity vector) through a hole in the table, passed under the table through two pulleys, and was attached to a removable container that held lead shot. Thus, the additive weight of the force transducers, wooden platform, metal cable, container and lead shot determined the total weight of the object to be lifted. Different amounts of lead shot were placed in each container to produce 9 different object weights. The nine weights had an ordered, incremental difference of 0.1 kg and ranged from 0.4 *kg* to 1.2 *kg*. Participants were seated such that the object to lift was directly in front of them. A plastic block (height = 10 *cm*) was placed in front of participants, behind the object, and was used to specify the instructed lift height.

Protocol

Participants were pseudo-randomly assigned to one of six groups (n = 15 per group). Participants in all groups performed object lifting. The weight of the object was selected from a discrete probability distribution. Three of these probability distributions produced varying weights and the other three produced a constant weight **(Fig. 2)**. Each group of participants were assigned one of the following six probability distributions: 1) skewed heavy mode, 2) symmetrical, 3) skewed light mode, 4) constant heavy, 5) constant mean, and 6) constant light. See **Table 1** for complete statistics of these probability distributions.

Participants were instructed to use the beat of a metronome (40 *beats/min*) to time transitions between different phases of each lift. Pilot testing showed that this metronome frequency produced consistent and relatively quick lifts, allowing us to capture a feedforward response. Four successive metronome beats signified the following **(Fig 1C)**: Beat 1 - a warning noise that the trial was starting; Beat 2 - grip and lift the object in one motion; Beat 3 - the object should reach and then be held at the height of the plastic block (10 *cm*); and Beat 4 - lower and then release the object.

To practice lifting according to the beat of the metronome, participants performed ten training lifts with the weight of the object selected from their respective distribution (bin 1). Following practice, participants performed the main experiment. Participants made 21 lifts with object weight selected from their assigned probability distribution without replacement. That is, they lifted all of the weights in a given distribution until it was depleted. This process was performed nine times (bins 2 - 10) for a total of 189 lifts. By selecting object weight from a distribution without replacement, we were able to avoid random clustering of certain weights while ensuring that the statistical properties of any given probability distribution were preserved in each experimental bin.

We made sure that participants in the varying probability distribution groups (skewed heavy mode, skewed light mode, and symmetrical) had no knowledge of the weight they were about to lift by (1) hiding the attached and unattached containers from our participants' field of view and (2) when successive lifts had the same weight, we would remove the attached container, place it on the ground and then reattach the same container.

As mentioned above, participants in three of the groups repeatedly lifted an object with a constant weight of 0.6, 0.8 or 1.0 *kg*. These weights were chosen to match important statistics, the mean and mode, of the three skewed probability distributions. More specifically, the weight of the constant heavy probability distribution (1.0 *kg*) matched the modal weight of the skewed heavy mode probability distribution, the weight of the constant mean probability distribution (0.8 *kg*) matched the mean weight of the skewed heavy mode and the skewed light mode probability distributions, and the weight of the constant light probability distribution (0.6 *kg*) matched the modal weight of the skewed heavy mode probability distribution.

The inclusion of constant weight groups served two purposes. First, it allowed us to directly compare the sensorimotor system's feedforward response when participants lifted an object of varying weight relative to when they lifted an object of constant weight. That is, we were able to test whether feedforward responses in the context of skewed weight distributions would match those observed for constant weight distributions, where the constant weights were aligned with the mean or mode of the skewed probability distributions. Second, it allowed us to determine whether the dependent measures commonly used as indexes of a feedforward prediction during object lifting studies were sensitive enough to detect the weight difference between the mean and mode (Δ0.2 *kg*) of the skewed probability distributions weights. In this study, we used four dependent measures as indexes of the sensorimotor system's feedforward prediction. These dependent measures were grip force rate, grip force, load force rate, and load force, which were all taken at the time point that corresponded to the peak load force rate. This time point occurred several hundred milliseconds before object lift off.

The symmetrical group acted as a control to test whether load force variance alone influences the feedforward response of the sensorimotor system. Participants in this group lifted an object whose weight was selected from a symmetrical probability distribution (i.e., mean, median and mode were identical). This symmetrical probability distribution had very similar load force variance and identical complexity (discrete entropy) to the skewed light mode and skewed heavy mode probability distributions.

The skewed light mode and skewed heavy mode probability distributions had the same mean and variance, but opposite skew. As such, the mode of the skewed light mode distribution and skewed heavy mode distribution were on opposite sides of the mean at 0.6 *kg* and 1.0 *kg*, respectively. We designed these skewed distributions such that the mode had a much higher relative frequency (42.8%) than the other six weights

(9.5%). This difference in frequency increased the possibility that the sensorimotor system would be able to distinguish the modal weight from the other weights. Critically, the separation of the mean and mode in both of the skewed probability distributions allowed us to test whether the sensorimotor system uses a minimal squared error strategy (mean) or a maximum a posteriori strategy (mode).

In the context of a Bayesian framework, the predictions of minimal squared error and maximum a posteriori strategies are found by taking a point estimate (i.e., the mean and mode, respectively) from a posterior distribution. In this study, we have manipulated the prior probability distribution by imposing environmental uncertainty via the object weight distributions described above. During the time course of any given lift, participants obtain current somatosensory information of an object's weight. This current information (i.e., likelihood function) is then integrated with previously acquired somatosensory information (i.e., prior probability distribution) from past lifts. A point-wise multiplication of the prior probability distribution with the likelihood function results in a posterior probability distribution. Thus, at the start of a subsequent lift, a feedforward controller could draw upon this posterior (which is now the new prior) to select a set of motor commands. A minimal squared error feedforward strategy would select a set of motor commands that aligns with the mean of the posterior (i.e., the average weight of the imposed weight distribution). In contrast, a maximum a posteriori strategy would select a set of motor commands that aligns with the model weight (i.e., the most frequent) of the posterior.

There were a total of 8 a priori comparisons per dependent measure (32 comparisons in total) that could be made to assess whether the sensorimotor system uses a minimal squared error strategy or a maximum a posteriori strategy. For a visual representation of all predictions made by each strategy, please refer to Fig. 3. As an example, if the feedforward controller were using a minimal squared error strategy (Fig. 3A), we would expect grip force rate, grip force, load force rate, and load force to be the same between the skewed heavy mode group and the constant mean group. Contrastingly, if the feedforward controller were attempting to use a maximum a posteriori strategy (Fig. 3B), we would expect the skewed heavy mode and the constant mean groups to have a significantly different grip force rate, grip force, load force rate, and load force.

Data Reduction and Analysis

Raw force and moment signals were smoothed using a dual low-pass, 2nd order, 14 *Hz* cut-off (Flanagan et al., 2003; Buckingham and Goodale, 2010), critically damped filter (Dowling, Robertson, 2000). Grip force (*N*) was calculated by averaging the normal forces recorded from the two force transducers (Flanagan et al, 2003; Fig. 1A). Load force (*N*) was calculated by summing the vertical forces recorded from the two force transducers. Grip force rate (*N/s*) and load force rate (*N/s*) are the time derivatives of grip force and load force, respectively, and were calculated using a 4th order, central-difference method. Grip force rate, grip force, load force rate, and load force before object lift off often serve as an index of the sensorimotor system's feedforward prediction of object weight (Flanagan and Beltzner, 2000; Buckingham and Goodale, 2010).

To capture only a feedforward response, we analyzed grip force rate, grip force, load force rate, and load force at the time point that corresponded to peak load force rate (Johansson and Westling, 1988; Flanagan and Beltzner, 2000; Flanagan et al., 2008; Baugh et al., 2012). In the last bin of trials, for each participant and trial we estimated object lift off from the load force traces recorded by the force transducers. Specifically, for each trial we found the point in time where the load force magnitude had just exceeded the current weight of the object. Further, we inspected the data to be assured that the four dependent measures were representative of a feedforward response and were taken before any online feedback corrections.

Error analysis

An error analysis was performed to assess whether the behavioural data was better explained by a minimal squared error strategy or a maximum a posteriori strategy. The main advantage of this approach is that it considers all of the experimental data of a particular measure, allowing for a single comparison to be made between the two strategies. To do this, we used a bootstrap procedure that allowed us to simultaneously contrast several groups to one another.

Briefly, for each group, this bootstrap procedure involved the random resampling without replacement (*n* resamples = group size) of a recorded measure (i.e., grip force, grip force rate, load force, or load force rate), taking the average of each group's resampled data, and from these averages summing the absolute error (i.e., difference) between several key groups. The predictions of each strategy dictated which groups were contrasted to one another. This process was repeated a total of 10 000 times and performed for each strategy. If a particular strategy has significantly less absolute error than a competing strategy, this indicates it better explains the behavioural data.

Here we provide an example group contrast made during the bootstrap procedure. The maximum a posteriori strategy predicts that the skewed light mode group would have the same grip force, grip force rate, load force, and load force rate as the constant light group. Therefore, if a maximum a posteriori strategy were dictating the feedforward response, we would expect a small amount of absolute error between these groups. However, instead of considering just one individual prediction like the example above, this error analysis simultaneously considers several of the a priori predictions depicted in **Fig. 3**. For complete details of this error analysis, we refer the reader to the Appendix.

Statistical Analysis

Our research question was focused on the stable behaviour of the feedforward controller, after learning had occurred, during an object-lifting task. That is, we were interested in the state of the feedforward controller after it had reached some stable pattern of behaviour in response to the imposed environmental uncertainty. As such, we performed statistical analyses on bin 10 (the last bin of the main experimental trials). We performed four separate one-way Analyses of Variance (ANOVA) on the four dependent measures of grip force rate, grip force, load force rate and load force. In these four ANOVA the independent variable was group (skewed light mode, skewed heavy mode, symmetrical, constant light, constant mean, and constant heavy).

All post-hoc pairwise comparisons and error analysis comparisons (4 in total) were computed using a nonparametric bootstrap hypothesis test [*resamples* =

1,000,000] (Gribble and Scott, 2002; Good, 2005). This test provides a more reliable p-value estimate than traditional parametric tests (e.g., t-tests). Briefly, they make no parametric assumptions (e.g., Normality), are less biased by samples with unequal sample size or unequal variance, and are better suited to analyse heteroscedastic data that is present in several commonly recorded biological measures (e.g., neural activity, electromyography and force production) due to sensorimotor noise (Gribble and Scott, 2002; Faisal et al., 2008; Cashaback et al., 2014). Holm-Bonferroni corrections were used to correct for inflated Type-I error due to multiple comparisons (Holm, 1979). Reported p-values are Holm-Bonferroni adjusted. The effect size for each main effect was calculated using partial eta squared (η_P^2). Statistical significance was set to $p < 0.05$.

RESULTS

Individual Data

Fig. 4 shows the average traces of grip force rate, grip force, load force rate and load force trial traces, taken from the last bin of trials, of a participant from the constant light group and another participant from the skewed light mode group. For all dependent measures, both participants had similarly shaped force and force rate traces that differed only in magnitude before object lift off. Based on the load force traces, the average object lift off time across participants occurred at 0.134 (± 0.036 SD) seconds after peak load force rate (Figures 4D, 5D). After lift off, the displayed participant in the constant light group maintains relatively consistent traces for all dependent measures, indicating that their feedforward response was well aligned to the force requirements of the constant weight they repeatedly lifted during the experiment. In contrast, for all measures, the displayed skewed light mode participant had a large amount of variability beyond object lift off in response to experiencing weights that varied on a trial-to-trial basis. This reflects a shift from feedforward to feedback control that, importantly, occurred well after our recorded dependent measures of the feedforward response. These patterns of behavior were consistent across participants.

Group Data

Fig. 5 shows the average traces of each group, from their last bin of trials, of grip force rate, grip force, load force rate and load force. For all measures, these traces are similar in terms of shape, but not necessarily magnitude, for participants experiencing either a constant or varying object weight on a trial-to-trial basis.

Fig. 6 shows each group's average grip force rate, grip force, load force rate and load force, taken at the time point corresponding to peak load force rate, across the ten different bins of trials. Qualitatively, we found that both load force rate and load force reached a stable pattern of behaviour during bin 1 (practice), while grip force rate and grip force took longer (~ bin 5 or 6) to reach a stable pattern of behaviour.

In bin 10 (**Fig. 7**), we found that all four dependent measures were inline with the predictions of a feedforward controller that uses a minimal squared error strategy, rather than a maximum a posteriori strategy, to predict object weight. Compare **Fig. 7** to **Fig. 3** for a visualization of the data relative to each of the strategy predictions.

Grip Force Rate

We found a significant effect of group on grip force rate (**Fig. 7A**) in the final bin of trials [$F(5, 84) = 8.321$, $p < 0.001$, $\eta_P^2 = 0.331$]. For grip force rate, eight pairwise comparisons were made to determine how the sensorimotor system makes a feedforward prediction. We found that four of the comparisons matched the predictions of a minimal squared error strategy (**Table 2A**). The remaining four comparisons did not match the predictions of a maximum a posteriori strategy (**Table 2B**). Thus, taken together the eight pairwise comparisons support the idea that the sensorimotor system uses a minimal squared error strategy to make feedforward predictions about object weight.

Grip Force

For grip force (**Fig. 7B**), we found a significant effect of group in the final bin of trials [$F(5, 84) = 5.955$, $p < 0.001$, $\eta_P^2 = 0.262$]. Again, we made eight pairwise comparisons to test whether the sensorimotor system uses a minimal squared error or maximum a posteriori strategy. Three of four comparisons matched the predictions of a minimal squared error strategy (**Table 2A**). Of the remaining four comparisons, only one matched the maximum a posteriori prediction (**Table 2B**). In other words, six of the eight pairwise comparisons were consistent with the idea that the sensorimotor system uses a minimal squared error strategy to make feedforward predictions of object weight.

Pairwise comparisons that did not match with a minimal squared error strategy involved the skewed heavy mode and constant heavy groups. Consistent with the maximum a posteriori strategy predictions, the skewed heavy mode group did not have a significantly different grip force from the constant heavy group ($p = 0.466$, two-tailed).

Load Force Rate

We found a significant effect of group on load force rate (**Fig. 7C**) in bin 10 [F(5, 84) = 9.348, p < 0.001, $\eta_P^2 = 0.357$]. Six of the eight pairwise comparisons were consistent with the idea that the sensorimotor system uses a minimal squared error strategy (**see Tables 2A and 2B**). For load force rate, pairwise comparisons that did not support a minimal squared error strategy involved the skewed light mode and constant light groups. Consistent with the maximum a posteriori strategy, the load force rate was not significantly different between the skewed light mode group and constant light group (p = 0.075, two-tailed).

Load Force

For load force (**Fig. 7D**), we found a significant effect of group [F(5, 84) = 16.756, p < 0.001, $\eta_P^2 = 0.499$]. We found that four pairwise comparisons matched the predictions of a minimal squared error strategy (**Table 2A**). The remaining four tests did not follow the predictions of a maximum a posteriori strategy (**Table 2B**). Thus, for load force, all eight pairwise comparisons were consistent with the idea that the sensorimotor system uses a feedforward controller that minimizes squared error.

Error Analysis

For each dependent measure, the error analysis provided a single, comprehensive comparison between the two candidate strategies (minimize squared error versus maximum a posteriori). The results of the error analysis are shown in **Fig. 8**. For all four dependent measures, a model based on minimizing squared error explained significantly more of the behavioural data (i.e., had less error) compared to the maximum a posteriori model ($p < 0.001$ for all four comparisons). Across measures, the model based on minimizing squared error had 56.8% less absolute error relative to the model based on maximum a posteriori estimates of object weight.

Sensitivity of Dependent Measures to Different Weights

We found the four dependent measures were sensitive to object weight differences of 0.2 kg, which matched the weight difference between the mean and mode of the skewed probability distributions. We found that mean values of each dependent measure were significantly greater for the constant mean group compared to the constant light group (**Table 3**). Similarly, for three of the four dependent measures we found that mean values for the constant heavy group were significantly greater than those for the constant mean group (**Table 3**). The only non-significant comparison between these two groups was for load force rate ($p = 0.054$, one-tailed).

Influence of Load Force Variance

For all four dependent measures, we found that mean values for the symmetrical group were not significantly different from those of the constant mean group (**Table 4**). This was predicted by both the minimal squared error and maximal a posteriori strategies. More importantly, this shows that the load force variance alone, at least within the range as dictated by our probability distributions, did not significantly influence the sensorimotor system's feedforward controller for object lifting.

DISCUSSION

An important feature of our experimental task was the randomization of object weights from trial to trial using skewed probability distributions. This allowed us to dissociate the predictions of minimal squared error and maximum a posteriori strategies for predicting object weight. We found that for object lifting, the sensorimotor system minimizes the square of prediction errors in the presence of environmental uncertainty. This finding is consistent with results found in studies of visually guided reaching (Kording and Wolpert, 2004b). Below we discuss how minimizing the square of feedforward errors may be beneficial in terms of the interplay between feedback and feedforward systems for sensorimotor control.

The finding that the sensorimotor system uses a minimal squared error strategy was supported by all four dependent measures that we used as indexes of the feedforward response (grip force rate, grip force, load force rate and load force). The results of twenty-eight of the thirty-two pairwise comparisons made among these four measures were consistent with a minimal squared error strategy. Further, for each of the four dependent measures, our error analysis showed that a minimal squared error feedforward strategy explained significantly more behaviour than a maximum a posteriori feedforward strategy.

In our task, we found that the sensorimotor system used a minimal squared error strategy to make a feedforward prediction of object weight. This strategy could be accomplished by predicting the weight of a subsequent lift by using somatosensory information from a previous lift (Johansson and Westling, 1984), or by taking an unweighted (Takahashi et al., 2001; Scheidt et al., 2001) or weighted (e.g., exponential decay: Landy et al., 2012; Hadjiosif and Smith, 2015) moving average of somatosensory information over several previous lifts. The use of a single previous lift, or averaging several previous lifts to make weight predictions, is often termed 'sensorimotor memory' (Chouinard et al., 2005). However, the concept of sensorimotor memory in itself is unable to explain phenomena such as reduced variability with practice (Kording and Wolpert, 2004a; Acerbi et al., 2014), explaining perceptual illusions (Peters et al., 2016) or incorporating sensory cues (Trampenau et al., 2015). A Bayesian framework is able to account for all these phenomena.

If participants used a Bayesian-like process they would build a prior representation of the environmental uncertainty. Similar to the sensorimotor memory strategy, they would use somatosensory information from previous lifts to build up a prior. However, where the Bayesian framework and sensorimotor memory strategies differ relates to how the somatosensory information from previous lifts is weighted. The sensorimotor memory strategy would suggest a constant weighting scheme while the Bayesian approach uses an adaptive weighting process due to the evolving prior over the course of learning. For example, decreases in movement variability in the presence of environmental uncertainty noise can be explained by an adaptive (un)weighting process that places less emphasis on trial-by-trial perturbations as a prior representation of environmental uncertainty is built (Kording and Wolpert, 2004a).

In the context of our task it would be difficult to track the prior over time, since these weightings would be convoluted with the safety margin that took time to stabilize (see Figure 6). However, we were still able to answer our research question because we used a small range of object weight uncertainty and analyzed only the last bin of trials after the safety margin stabilized. While previous work has tracked the evolution of a prior with learning (Berniker et al., 2010), an interesting direction would be exploring how previously acquired sensorimotor information becomes adaptively (un)weighted in a Bayesian, statistically optimal way during the course of learning.

Our finding that the sensorimotor system uses a minimal squared error strategy during object lifting parallels research that examined visually guided reaching (Kording and Wolpert, 2004b; Zhang et al., 2015). We recently examined how the visuomotor system deals with environmental uncertainty during an implicit learning task (Cashaback et al., submitted). We found that the visuomotor system uses a minimal squared error strategy when updating where to aim reaches when using visual error feedback (i.e., the visual distance from a target), but can also switch to a maximum a posteriori strategy when using only binary reinforcement feedback (visual, auditory and monetary reward per target hit). Surprisingly, when both error and reinforcement feedback were made available the visuomotor system used a minimal squared error strategy, as opposed to a maximum a posteriori strategy that maximized both target hits and reward. This suggests during implicit learning that the visuomotor system heavily weights error feedback over reinforcement feedback when updating where to aim reaches. Likewise, it is possible that the sensorimotor system may be able to perform a maximum a posteriori feedforward prediction when using reinforcement feedback, but perhaps only in the absence of sensorimotor error feedback. Future research involving individuals with peripheral nerve deafferentation (Buckingham et al., 2016), or the blocking of ascending tactile (Johansson and Westling, 1984) and proprioceptive (Buffenoir et al., 2013) signals in healthy individuals would likely provide valuable insights into how the sensorimotor system uses error and reinforcement feedback to update feedforward predictions. Nevertheless, with error feedback available, the sensorimotor system appears to use a minimal squared error strategy when lifting objects and making visually guided reaches. This parallel in behaviour may be explained by the use of common brain areas to represent uncertainty or similar neuronal features, such as individual neuronal firing rates (Ma et al., 2006; Schultz, 2013) and neural population coding (Vilares et al., 2012; Pouget et al., 2013). Some reported brain areas that may represent environmental uncertainty include the putamen, amygdala, insula, orbitofrontal cortex, posterior parietal cortex, and the anterior cingulate cortex (Vilares et al., 2012; O'Reilly et al., 2013). However, theories and empirical studies on how the brain represents either sensorimotor noise or environmental uncertainty are currently sparse (Faisel et al., 2008; Kording, 2014).

Kording and Wolport (2004b) also examined the effects of environmental uncertainty in a visuomotor task. They had participants operate a virtual peashooter. When shot, the peas were visually displaced by an amount drawn from a skewed noise distribution. On separate trials, the authors also manipulated the amount of uncertainty (variance) of these skewed noise distribution. Participants were required to move a cursor to a location such that the shot peas were "on average as close to the target as possible". With low variance skewed noise, that is, when visual displacements were less than approximately ±1.5 *cm*, Kording and Wolpert found that the visuomotor system minimized approximately squared error. However, as visual displacement variance increased beyond this range, they found the visuomotor system shifted away from a minimal square error strategy and became less sensitive to larger errors (Kording and Wolpert, 2004b; Wolpert, 2007). In our task, both the skewed light mode and skewed heavy mode probability distributions that we used to determine object weight on a trial-to-trial basis each had a standard deviation of ±0.22 *kg*. With this relatively low level of uncertainty, participants used a feedforward response that was closely aligned with the mean (0.8 *kg*) of these skewed probability distributions. That is, with this amount of load force variance, the sensorimotor system used the same feedforward response as if it was lifting an object with a constant weight of 0.8 *kg*. This shows that the amount of load force variance associated with the two skewed distributions had little or no influence on the feedforward response. This was further supported by no behavioural differences between participants in the constant mean and symmetrical (no skew) groups. Thus, given that the variance of the probability distributions used to vary object weight did not influence behaviour, and that the sensorimotor system was sensitive to weight differences of 0.2 *kg*, we were able to directly assess whether the sensorimotor system was using a minimal squared error or maximum a posteriori strategy to deal with environmental uncertainty. With low amounts of load force variance, we found that the sensorimotor system used a minimal squared error strategy to make feedforward predictions of object weight.

Our finding that the sensorimotor system was not influenced by load force variance differs from research by Hadjiosif and Smith (2015). However, these differences are likely caused by difference in experimental design. We used a task where the load forces were acceleratory (gravitational) in nature and had relatively low amounts of load force variance relative to the mean (i.e., coefficient of variation = standard deviation / mean x 100.0 = 27.5%). In contrast, Hadjiosif and Smith (2015) had participants pinch grip a force transducer that was mounted on a robotic arm.

Participants then made reaching movements to a target in a velocity dependent (viscous) force field. The strength of this force field was either held constant or varied according to a Gaussian distribution. For the different blocks of trials where the force-field strength varied, the corresponding coefficient of variation ranged from 40% to 250%. Hadjiosif and Smith (2015) found that participants applied larger grip forces with greater variability in force field strength. The authors relate this finding to the idea of a ‘flexible safety margin’. Briefly, a safety margin refers to the finding that individuals grip with a higher force than is required to prevent an object from slipping, in the event of an inaccurate feedforward prediction. This safety margin is present when repeatedly lifting an object with a constant weight (Westling and Johansson, 1984), and is ‘flexible’ in the sense that it scales with environmental uncertainty (Hadjiosif and Smith, 2015). In our task, given the relatively low coefficient of variation (27%), the safety margin used for a constant weight of 0.8 kg may have been sufficient to absorb the majority of the load force variance. This load force variance was dictated by the spread of the three probability distributions (skewed heavy mode, skewed light mode, and symmetrical) used to vary object weight. However, with greater load force variance, as seen in Hadjiosif and Smith (2015), a feedforward response aligned with the mean of the environmental uncertainty may be unable to absorb the whole range of the load force variability. Taking into account both our current work and that of Hadjiosif and Smith (2015), it is possible that with larger amounts of load force variability that the sensorimotor system becomes sensitive to environmental uncertainty and places less emphasis on using a minimal squared error strategy.

A change in emphasis from using a minimal squared error strategy to becoming sensitive to environmental uncertainty may occur when the sensorimotor system is unable to fully compensate for high levels of load force variability. In other words, the feedback response may not have enough time to respond to the larger prediction errors, which in some instances could be detrimental to task success (e.g., dropping an object). An inability of the feedback system to respond quickly enough to the whole range of load force variability may explain the finding of Berg and colleagues (2016). They found in their ball catching experiment that the sensorimotor system seems to use a feedforward response aligned with the heaviest object. This may represent an upper bound of how the sensorimotor system deals with very high levels of weight uncertainty, where the feedforward response seems to scale its motor commands to the greatest weight that is lifted or caught. Nevertheless, in our experiment the safety factor seemed able to absorb the relatively small range of load force variability, providing the feedback system sufficient time to make small corrections in response to feedforward prediction errors.

Currently we do not know why the sensorimotor system uses a minimal squared error strategy, or how this strategy is implemented by the nervous system. Regardless, there are instances where a minimum squared error strategy is advantageous. As mentioned above, a minimum squared error strategy corresponds to the mean of the environmental uncertainty. From a computational point of view, the mean is always defined unlike other point estimate statistics. For example, unlike the mean, the mode and median become ill-defined when the environmental uncertainty follows a uniform (Berg et al., 2016) or certain bimodal probability distributions (Scheidt et al., 2001;

Kording and Wolpert, 2004a). Thus, using the mean may ensure an efficient updating of internal models when using noisy error-based feedback.

Another potential advantage of a minimal squared error strategy relates to how errors are penalized. This strategy considers all potential errors, but applies a greater penalization to large errors relative to smaller ones. As a result, a minimal squared error strategy will produce a feedforward response that protects against large feedforward prediction errors. By using a feedforward response that protects against large errors, this would allow the feedback system to respond more quickly to potentially detrimental feedforward prediction errors. For example, consider participants experiencing weights selected from the skewed light distribution. If the participants had used a maximum a posteriori strategy, they would have used a feedforward response corresponding to the lightest weight of 0.6 *kg*. However, this would place the feedforward grip and load forces far from the appropriate force magnitudes required to lift and grasp the maximum weight (1.2 *kg*) of the skewed light mode probability distribution. However, the minimum squared error strategy that participants used aligned them with the mean (0.8 kg) of the skewed light mode probability distribution, which was closer to the maximum weight of this distribution. As such, the feedback system would be able to respond more rapidly to the heaviest weight, since the required corrective adjustments would be smaller. Although a feedforward controller using maximum a posteriori strategy would predict the correct object weight at a higher frequency, this comes at the potential cost of having larger prediction errors with inaccurate feedforward responses. Conversely, the minimal squared error strategy would have a higher frequency of prediction errors, but these errors would be smaller and would subsequently allow for a more rapid feedback response. Thus, by using a minimum squared error strategy, it is possible that the feedforward system hedges against larger errors in order to setup the feedback system for success. To test this idea, future work should manipulate both the magnitude of feedforward prediction errors and the time the feedback system has to respond to such errors. Such work would improve our understanding on the interplay between the feedback and feedforward system.

It is noteworthy that many authors make the assumption of a maximum a posteriori strategy, often termed as maximum likelihood (equivalent to maximum a posteriori estimate when using a non-informative, flat prior). A convenient advantage of using maximum a posteriori estimates is that they are more readily calculated with explicit equations, making it easier to solve the optimal solution(s). Some examples of where authors have assumed a maximum a posteriori strategy include performing state estimation (Crevecoeur and Scott, 2013; Diedrichson, 2007), integrating information from multiple senses (Angelaki et al., 2009), making a choice in a forced decision-making task (Resulaj et al., 2009; Wolpert and Landy, 2012; Acuna et al., 2015), making feedforward predictions with the aid of visual cues (Trampenau et al., 2015), and predicting the weight of novel objects (Peters et al., 2016). While these studies have provided valuable information about how the sensorimotor system generates predictions in the presence of noise, the present study addresses a different question. Namely, how are humans able to generate feedforward predictions in the presence of asymmetrical noise? In the current study we separated the optimal solutions of a maximum a posteriori strategy and a minimum squared error strategy by using skewed probability distributions. We found that the sensorimotor system uses a minimal squared error strategy in the presence of a small range of environmental uncertainty, and that the maximum a posteriori estimate was inferior in predicting our behavioral measures. However, we do not argue that the sensorimotor system never uses a maximum a posteriori strategy (Mawase and Karniel, 2010). Rather, we propose that the chosen strategy is likely task and goal dependent. Nevertheless, our work highlights the importance of determining the underlying processes that drive the control of our movements.

In summary, in the presence of a relatively narrow range of object weight uncertainty we found that the sensorimotor system minimizes the square of potential prediction errors during object lifting. This finding parallels previous research that examined visually guided reaching. The apparent overlap in strategy when lifting objects and making visually guided reaches suggests common underlying mechanisms to deal with environmental uncertainty. These mechanisms may include an overlap in brain areas that integrate environmental uncertainty or similarities in neuronal features (e.g., firing rate properties and population coding). Finally, we propose that the sensorimotor system may use a minimum squared error strategy to hedge against potentially large prediction errors. Such error hedging may maximize the probability of a successful feedback response. Future work testing this hypothesis may provide important insights on the interplay between feedforward and feedback components of the sensorimotor system.

REFERENCES

Acerbi L, Vijayakumar S, Wolpert DM. On the origins of suboptimality in human probabilistic inference. *PLoS Comput Biol* 10(6): e1003661, 2014.

Acuna DE, Berniker M, Fernandes HL, Kording KP. Using psychophysics to ask if the brain samples or maximizes. *J Vision* 15(3)7: 1-16, 2015.

Angelaki DE, Gu Y, DeAngelis GC. Multisensory integration: psychophysics, neurophysiology, and computation. *Curr Opin Neurobiol* 19(4): 452-458, 2009.

Baugh LA, Kao M, Johansson RS, Flanagan JR. Material evidence: Interaction of well-learned priors and sensorimotor memory when lifting objects. *J Neurophysiol* 108(5): 1262-1269, 2012.

Berg WP, Richards BJ, Hannigan AM, Biller KL, Hughes MR. Does load uncertainty affect adaptation to catch training? *Exp Brain Res* :1-13, 2016.

Berniker M, Voss M, Kording K. Learning Priors for Bayesian Computations in the Nervous System. *PLoS ONE* 5(9), e12686, 2010.

Brayanov JB, Smith MA. Bayesian and “anti-Bayesian” biases in sensory integration for action and perception in the size-weight illusion. *J Neurophysiol* 103(3): 1518-1531, 2010.

Buckingham G, Cant JS, Goodale MA. Living in a material world: how visual cues to material properties affect the way that we lift objects and perceive their weight. *J Neurophysiol* 102: 3111-3118, 2009.

Buckingham G, Goodale MA. Lifting without seeing: the role of vision in perceiving and acting upon the size weight illusion. *PLoS ONE* 5: e9709, 2010.

Buckingham G, Michelakakis EE, Cole J. Perceiving and acting upon weight illusions in the absence of somatosensory information. *J Neurophysiol* 115(4): 1946-1953, 2016.

Buckingham G, Ranger NS, Goodale MA. The Role of Vision in Detecting and Correcting Fingertip Force Errors During Object Lifting. *J Vision* 11(1): 2011.

Buffenoir K, Decq P, Lambertz D, Perot C. Neuromechanical assessment of lidocaine test block in spastic lower limbs. *Appl Physiol Nutr Me* 38(11): 1120-1127, 2013.

Cashaback JGA, Fewster K, Potvin JR, Pierrynowski MR. Musculotendon translational stiffness and muscle activity are modified by shear forces. *Clin Biomech* 29: 494:499, 2014.

Cashaback JGA, McGregor HR, Mohatarem A, Gribble PL. Dissociating Error-Based and Reinforcement-Based Loss Functions During Implicit Learning, submitted.

Chouinard PA, Leonard G, Paus T. Role of the Primary Motor and Dorsal Premotor Cortices in the Anticipation of Forces During Object Lifting. *J Neurosci 25*(9): 2277–2284, 2005.

Crevecoeur F, Scott SH. Priors engaged in long-latency responses to mechanical perturbations suggest a rapid update in state estimation. *PLoS Comput Biol* 9(8): e1003177, 2013.

Diedrichsen J. Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr Biol* 17(19): 1675-1679, 2007.

Faisal AA, Selen LP, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci* 9(4): 292-303, 2008.

Flanagan JR, Beltzner MA. Independence of perceptual and sensorimotor predictions in the size-weight illusion. Nat Neurosci 3(7): 737-741, 2000.

Flanagan JR, Bittner JP, Johansson RS. Experience can change distinct size-weight priors engaged in lifting objects and judging their weights. *Curr Biol* 18(22): 1742-1747, 2008.

Flanagan JR, Bowman MC, Johansson RS. Control strategies in object manipulation tasks. *Curr Opin Neurobiol* 16: 650–659, 2006.

Flanagan JR, Vetter P, Johansson RS, Wolpert DM. Prediction precedes control in motor learning. *Curr Biol* 13:146-150, 2003.

Good PI. Permutation, Parametric and Bootstrap Tests of Hypotheses: A Practical Guide to Resampling Methods for Testing Hypotheses, 2005.

Gordon AM, Forssberg H, Johansson RS, Westling G. Integration of sensory information during the programming of precision grip: comments on the contributions of size cues. *Exp Brain Res* 85(1): 226–229, 1991a.

Gordon AM, Forssberg H, Johansson RS, Westling G. The integration of haptically acquired size information in the programming of precision grip. *Exp Brain Res* 83(3): 483–488, 1991b.

Gordon AM, Forssberg H, Johansson RS, Westling G. Visual size cues in the programming of manipulative forces during precision grip. *Exp Brain Res* 83(3): 477–482, 1991c.

Grandy MS, Westwood DA. Opposite Perceptual and Sensorimotor Responses to a Size-Weight Illusion. *J Neurophysiol* 95(6): 3887–3892, 2006.

Gribble PL, Scott SH. Overlap of internal models in motor cortex for mechanical loads during reaching. *Nature* 417: 938-941, 2002.

Hadjiosif AM, Smith MA. Flexible control of safety margins for actions based on environmental variability. *J Neurosci* 35(24): 9106-9121, 2015.

Hermdorfer J, Li Y, Randerath J, Goldenberg G, Eidenmuller S. Anticipatory scaling of grip forces when lifting objects of everyday life. *Exp Brain Res* 212(1): 19–31, 2011.

Holm S. A simple sequentially rejective multiple test procedure. *Scand J Stat*: 65-70, 1979.

Jenmalm P, Johansson RS. Visual and somatosensory-motor information about object shape control manipulative fingertip forces. *J Neurosci* 17(11): 4486-4499, 1997.

Johansson RS, Flanagan JR. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nat Rev Neurosci* 10: 345–359, 2009.

Johansson RS, Westling G. Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. *Exp Brain Res* 56(3): 550-564, 1984.

Johansson RS, Westling G. Coordinated isometric muscle commands adequately and erroneuously programmed for the weight during lifting task with precision grip. *Exp Brain Res* 71: 59–71, 1988.

Karniel A. Open questions in computational motor control. *J Integr Neurosci* 10(3): 391-417. 2011.

Körding KP. Bayesian statistics: relevant for the brain? *Curr Opin Neurobiol* 25: 130-133, 2014.

Körding KP, Ku SP, Wolpert DM. Bayesian integration in force estimation. *J Neurophysiol* 92(5): 3161-3165, 2004.

Körding KP, Wolpert DM. Bayesian integration in sensorimotor learning. *Nature* 427(6971): 244-247, 2004a.

Körding KP, Wolpert DM. The loss function of sensorimotor learning. *PNAS* 101(26): 9839-9842, 2004b.

Landy MS, Trommershäuser J, Daw ND. Dynamic estimation of task-relevant variance in movement under risk. *J Neurosci* 32(37): 12702-12711, 2012.

Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. *Nat Neurosci* 9(11): 1432-1438, 2006.

Mawase F, Karniel A. Evidence for predictive control in lifting series of virtual objects. *Exp Brain Res* 203(2): 447-452, 2010.

O'Reilly JX, Schuffelgen U, Cuell SF, Behrens TE, Mars RB, Rushworth MF.

Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *PNAS* 110(38): E3660-E3669, 2013.

Peters MAK, Wei JM, Shams L. The size-weight illusion in not anti-bayesian after all: a unifying Bayesian account. *Peer J* 4: e2124, 2016.

Pouget A, Beck JM, Ma WJ, Latham PE. Probabilistic brains: knowns and unknowns. *Nat Neurosci* 16(9): 1170-1178, 2013.

Resulaj A, Kiani R, Wolpert DM, Shadlen MN. Changes of mind in decision-making. *Nature* 461(7261): 263-266, 2009.

Robertson DGE, Dowling JJ. Design and responses of Butterworth and critically damped digital filters. *J Electromyogr Kines* 13(6): 569-573, 2003.

Scheidt RA, Dingwell JB, Mussa-Ivaldi FA. Learning to move amid uncertainty. *J Neurophysiol* 86(2): 971-985, 2001.

Schultz W. Updating dopamine reward signals. *Curr Opin Neurobiol* 23(2): 229-238, 2013.

Takahashi CD, Scheidt RA, Reinkensmeyer DJ. Impedance control and internal model formation when reaching in a randomly varying dynamical environment. *J Neurophysiol* 86(2): 1047-1051, 2001.

Trommerhausser J, Maloney LT, Landy MS. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A* 20(7): 1419-1433, 2003.

Trampenau L, Kuhtz-Buschbeck JP, van Eimeren T. Probabilistic information on object weight shapes force dynamics in a grip-lift task. *Exp Brain Res* 233(6): 1711-1720, 2015.

Vilares I, Howard JD, Fernandes HL, Gottfried JA, Körding KP. Differential representations of prior and likelihood uncertainty in the human brain. *Curr Biol* 22(18): 1641-1648, 2012.

Wolpert DM. Probabilistic models in human sensorimotor control. *Hum Movement Sci* 26(4), 511-524.

Wolpert DM, Flanagan JR. Motor prediction. *Curr Biol* 11: R729–R732, 2001.

Wolpert DM, Landy MS. Motor control is decision-making. *Curr Opin Neurobiol* 22(6): 996-1003, 2012.

Zhang H, Daw ND, Maloney LT. Human representation of visuo-motor uncertainty as mixtures of orthogonal basis distributions. *Nat Neurosci* 18(8): 1152-1158, 2015.

APPENDIX: ERROR ANALYSIS

Here we describe the error analysis we used to compare whether the experimental data were better explained by a minimize squared error strategy or a maximum a posteriori strategy. The main advantage of this error analysis is that it considers all of the experimental data to allow for a single comparison to be made between the two strategies. To do this, we bootstrap the experimental data and sum the absolute error between several key groups. The predictions of each strategy are used to select which groups are compared to one another. For example, the maximum a posteriori strategy predicts that the skewed heavy mode group would be no different from the constant heavy group. Thus, if the maximum a posteriori strategy was driving behaviour, we would expect a small amount of error between the groups. However, instead of just considering one individual prediction like the example directly above, the error analysis simultaneously considers all the predictions of a given strategy. Below, we describe this error analysis in detail.

First, let X represent all the data, from all groups, of one dependent measure (grip force rate, grip force, load force rate or load force) in the final BIN of trials. $\bar{X}$ represents the overall mean of a dependent measure, which we will use later to normalize the estimated absolute error. Further, let $X^j = x^j_1, x^j_2, ..., x^j_n$, where X^j represents a vector of the dependent measures for some group (j) and x^j_i represents some individual's (i) data point in that group. The six groups are the skewed heavy mode (shm), skewed light mode (slm), symmetrical (s), constant heavy (ch), constant mean (cm) and constant light (cl).

To perform bootstrapping, it is necessary to resample (with replacement) n times from a group of interest to generate a single bootstrap resample. This bootstrap resample is the same length as the original group (here, n = 15, matching the number of participants per group) and only contains individual data points from the original group it resampled from. This resampling procedure is performed N times to generate N bootstrap resamples (here, N = 10 000). We denote a bootstrap resample as X_k^{j*}, where * represents a resampled vector and k represents the bootstrap resample iteration. The average of a bootstrap resample is $\bar{X}_k^{j*}$.

As an example of some bootstrap resample, if we were resampling from the skewed heavy mode group and were on the 1054th iteration, it may look as follows: $X_{1054}^{shm*} = x_2^{shm}, x_3^{shm}, x_7^{shm}, x_{11}^{shm}, x_4^{shm}, x_5^{shm}, x_7^{shm}, x_{10}^{shm}, x_{14}^{shm}, x_4^{shm}, x_{14}^{shm}, x_8^{shm}, x_2^{shm}, x_{10}^{shm}, x_{12}^{shm}$. Notice that this bootstrap resample vector contains the same number of data point as their are participants in the group being sampled (n = 15). Also, due to the resampling with replacement, notice that that some data points are represented more than once (e.g., x_4^{shm}) while others are not present (e.g., x_1^{shm}). The data points in a bootstrap resample can vary on any given iteration. Further, each bootstrap resample is composed of individual data points from one group.

In the equations below (1 and 2), we describe how we use the experimental data and a bootstrap procedure to calculate the normalized, absolute error of a minimize squared error (mse) strategy and a maximum a posteriori (map) strategy. Briefly, each equation sums the absolute differences between each group lifting an object of varying weight to their corresponding group that lifts a constant weight. A particular strategy dictates the groups that are compared to one another (e.g., map strategy; shm = ch). The normalized absolute error of the mse strategy (ϵ_k^{mse*}) on any particular bootstrap iteration is

$$\epsilon_k^{mse*} = \frac{|\bar{X}_k^{shm*} - \bar{X}_k^{cm*}| + |\bar{X}_k^{s*} - \bar{X}_k^{cm*}| + |\bar{X}_k^{slm*} - \bar{X}_k^{cm*}|}{\bar{X}} \quad eq.(1).$$

Likewise, the normalized absolute error of the mse strategy (ϵ_k^{map*}) on any particular bootstrap iteration is

$$\epsilon_k^{map*} = \frac{|\bar{X}_k^{shm*} - \bar{X}_k^{ch*}| + |\bar{X}_k^{s*} - \bar{X}_k^{cm*}| + |\bar{X}_k^{slm*} - X_k^{cl*}|}{\bar{X}} \quad eq.(2).$$

Following the bootstrap procedure, we then compiled all iterations of ϵ_k^{mse*} and ϵ_k^{map*} to form a distribution of normalized absolute error for each strategy. $\hat{\epsilon}^{mse*}$ represents the distribution of normalized absolute error for the mse strategy, while $\hat{\epsilon}^{map*}$ represents the distribution of normalized absolute error for the map strategy.

We then compared whether $\hat{\epsilon}^{mse*}$ and $\hat{\epsilon}^{map*}$ were statistically different by using a two-tailed bootstrap hypothesis test. For graphical purposes (**Fig. 6**), we calculated the mean ($\bar{x}_{\hat{\epsilon}^{mse*}}$ and $\bar{x}_{\hat{\epsilon}^{map*}}$) and standard deviation ($\sigma_{\hat{\epsilon}^{mse*}}$ and $\sigma_{\hat{\epsilon}^{map*}}$) of $\hat{\epsilon}^{mse*}$ and $\hat{\epsilon}^{map*}$, respectively.

FIGURE CAPTIONS

Figure 1: Experimental Apparatus and Protocol. **A)** Participants used a pinch grip when grasping the transducers. Grip forces were perpendicular to the contact surfaces of the transducers. Load forces acted vertically and were parallel with the contact surfaces of the transducers. **B)** The force transducers were mounted to the top of a wood platform that covered a hole in the table. A cable was attached to the wood platform, passed through two pulleys and held up a container containing lead shot. There were a total of nine possible containers that participants could lift. Each container was filled with different amounts of lead shot (0.1 *kg* increments), such that the total object weight varied from 0.4–1.2 *kg*. **C)** The beginning of the trial was signaled by a warning noise timed to a metronome beat (40 *bpm*). On the second beat, participants were instructed to grip and lift the object in a single motion. At the time of the following beat, the participant was to lift the object to the height of a block (10 *cm*). They held the object there until the fourth and final beat, at which time they would lower and then release the object. For each new trial, the experimenter would attach a container that was selected according to the participant's assigned probability distribution.

Figure 2: Discrete probability distributions that describe the different object weights to be lifted (x-axis) and the frequency count of a particular weight (y-axis). Participants were assigned one of the displayed distributions. There were three probability distributions that resulted in a constant weight (**A = constant heavy, B = constant mean, C = constant light**) and three probability distributions that resulted in a varying weight (**D = skewed heavy mode, E = symmetrical, F = skewed light mode**). Each distribution had a total frequency count of 21 weights, matching the number of lifts per bin of trials. On each trial, object weight was randomly drawn from a distribution until its depletion. This process was performed 9 times (bins 2-10) for a total of 189 experimental lifts (bin 1 was a set of 10 practice trials). For each distribution, the thin solid line, thin dashed line, and thin dotted line correspond to its mean, median, and mode, respectively. The constant light distribution had a weight of 0.6 *kg* that was aligned to the mode of the skewed light mode. The constant mean had a weight of 0.8 kg that was aligned to the mean of the skewed light mode, symmetrical and skewed heavy mode probability distributions. The constant heavy had a weight of 1.0 *kg* that was aligned to the mode of the skewed heavy mode. The symmetrical distribution had variance, no skew (mean, median, and mode identical) and acted as a control to see if load force variance alone influenced feedforward predictions. Both the skewed light mode and skewed heavy mode had their mean and mode separated (by 0.2 *kg*), allowing us to investigate whether the sensorimotor feedforward system attempts to minimize the square of prediction errors (feedforward response aligned with the mean weight of a distribution) or attempts to select the most likely weight (feedforward response aligned with the modal weight of a distribution).

Figure 3: Predictions of feedforward controller that uses a: **A**) minimal squared error strategy or **B**) maximum a posteriori strategy. These predictions apply to the four dependent measures, grip force rate, grip force, load force rate and load force, which we used to characterize the feedforward response of the sensorimotor system. Under the heading, 'Predictions', we summarize the expected outcome of group mean comparisons for a minimal squared error strategy (**3A**: light blue text) and a maximum a posteriori strategy (**3B**: dark blue text). Black text (i.e., S = CM) indicates an identical prediction between the two strategies. Here, =, <, and > indicate whether we expect the dependent measures of a group to be equal to, significantly less than, or significantly greater than another group, respectively.

Figure 4: Individual participant traces, averaged across the last bin of trials, of A) grip force rate (*N/s*), B) grip force (*N*), C) load force rate (*N/s*), and D) load force (*N*) from a participant in the constant light group and a participant in the skewed light mode group. For all measures, individual trial traces were aligned to peak load force rate. Dashed vertical lines represent the time of peak load force rate, which intercepts the x-axis at 0.0 *s*. Both participants had consistently shaped force and force rate traces before object lift off, which on average occurred at 0.134 ± 0.036 *s*, differing only in magnitude. By recording all four measures at the peak load force rate (0.0 *s*), before object lift off, we were able to capture each participant's feedforward response. Beyond object lift off, the increased trace variability of the skewed light mode participant reflects feedback modulation in response to lifting weights that varied on a trial-to-trial basis. Contrastingly, the constant light participant showed more consistent traces throughout the entire trial, indicating that their feedforward response was well matched to the force requirements of the constant weight they repeatedly lifted throughout the experiment. Shaded regions represent ±1 standard deviation.

Figure 5: Average group traces, using the last bin of trials, of A) grip force rate (*N/s*), B) grip force (*N*), C) load force rate (*N/s*), and D) load force (*N*). For all measures, individual trial traces were aligned to peak load force rate. Dashed vertical lines represent the time of peak load force rate, which intercepts the x-axis at 0.0 *s*. The shape, but not necessarily the magnitude, of all four measures was quite consistent across groups. For all four measures that were recorded at the dashed line, representing an index of the feedforward response, there were no significant differences between the groups whose participants lifted varying weights (skewed heavy mode, symmetrical, skewed light mode) and the constant mean group. This finding aligns with the prediction of a feedforward response using a minimal squared error strategy. Beyond the time of object lift off, which on average occurred at 0.134 ± 0.036 *s*, there appears to be slight separation of grip force between the constant mean group compared to the skewed heavy mode, symmetrical, skewed light mode groups. This separation likely represents feedback modulation in response to lifting weights that varied on a trial-to trial basis (see Figure 4B). Shaded regions represent ±1 standard error.

Figure 6: Average **A)** grip force rate (*N/s*), **B)** grip force (*N*), **C)** load force rate (*N/s*), and **D)** load force (*N*) of each group across separate bins of trials. Error bars represent ±1 standard error.

Figure 7: Average **A)** grip force rate (*N/s*), **B)** grip force (*N*), **C)** load force rate (*N/s*), and **D)** load force (*N*) of each group in the final, 10^{th} bin of trials. Under the heading 'Comparisons', we summarize key group mean comparisons that relate to how the sensorimotor system makes a feedforward prediction (for an exhaustive list, see **Table 1, 2 and 3**). For any dependent measure, =, <, and > indicate whether one group was equal to, less than, or greater than another group, respectively. Dark blue lettering indicates the comparison is aligned with a maximum a posteriori strategy, while light blue lettering indicates a comparison that supports a minimize squared error strategy. Black lettering indicates an identical prediction between the two strategies. As can be seen across dependent measures, the vast majority of comparisons support a minimal squared error strategy. Error bars represent ±1 standard error. $p < 0.05$.

Figure 8: For each dependent measure (x-axis), the resulting magnitude of error (y-axis) when predicting the data with a minimal squared error strategy (light blue) or maximum a posteriori strategy (dark blue). Error bars represent ±1 standard deviation. $p < 0.05$.

TABLE CAPTIONS

Table 1: Descriptive statistics of the six probability distributions that dictated the trial-by-trial weight of the object to be lifted. Participants were pseudorandomly assigned to one of the six probability distributions.

Table 2: For each measure, the adjusted p-values of each group mean comparison. The second row, leftmost four entries show the groups being compared and indicates the predicted results (i.e., equal to, greater than, less than) of a **A)** minimal squared error strategy and **B)** maximum a posteriori strategy. These predictions match those visually seen in **Fig. 3**. We have bolded comparisons where the p-value supports a specific prediction (corresponding to the cell above in the second row). When a strategy predicts two groups to be equal to one another (e.g., skewed heavy mode equal to constant mean), for the prediction to be true then the p-value would have to be greater than or equal to 0.05 (i.e., no difference between groups). In contrast, if the prediction expects one group to be significantly different from another group (e.g., skewed heavy mode less than constant heavy mode), then p-value has to be less than 0.05 for the prediction to be true. As can be seen in **1A**, 14 out of 16 comparisons are aligned with a minimal squared error strategy. Conversely, only 2 of 16 comparisons in **1B** are aligned with a maximum a posteriori strategy. Taken together, 28 of the 32 total comparisons support the idea of a sensorimotor system that minimizes the square of prediction errors.

Table 3: For each dependent measure, the corresponding adjusted p-values when comparing whether the constant light was significantly less than the constant mean group, and whether the constant mean group was less than the constant heavy group. Bold indicates significant differences between the specified group mean comparisons. All but one of the comparisons was insignificant, albeit trending towards a difference ($p = 0.054$). The results of these comparisons suggest that the dependent measures were sensitive to weight changes of 0.2 *kg*, which is the difference between the mean and mode in both the skewed light mode and skewed heavy mode probability distributions.

Table 4: For each dependent measure, the corresponding adjusted p-value when comparing the symmetrical and constant mean groups. Bold indicates significant differences between groups. As expected, all comparisons were insignificant, indicating that the dependent measures were not sensitive to the low range of load force variance used in this study.

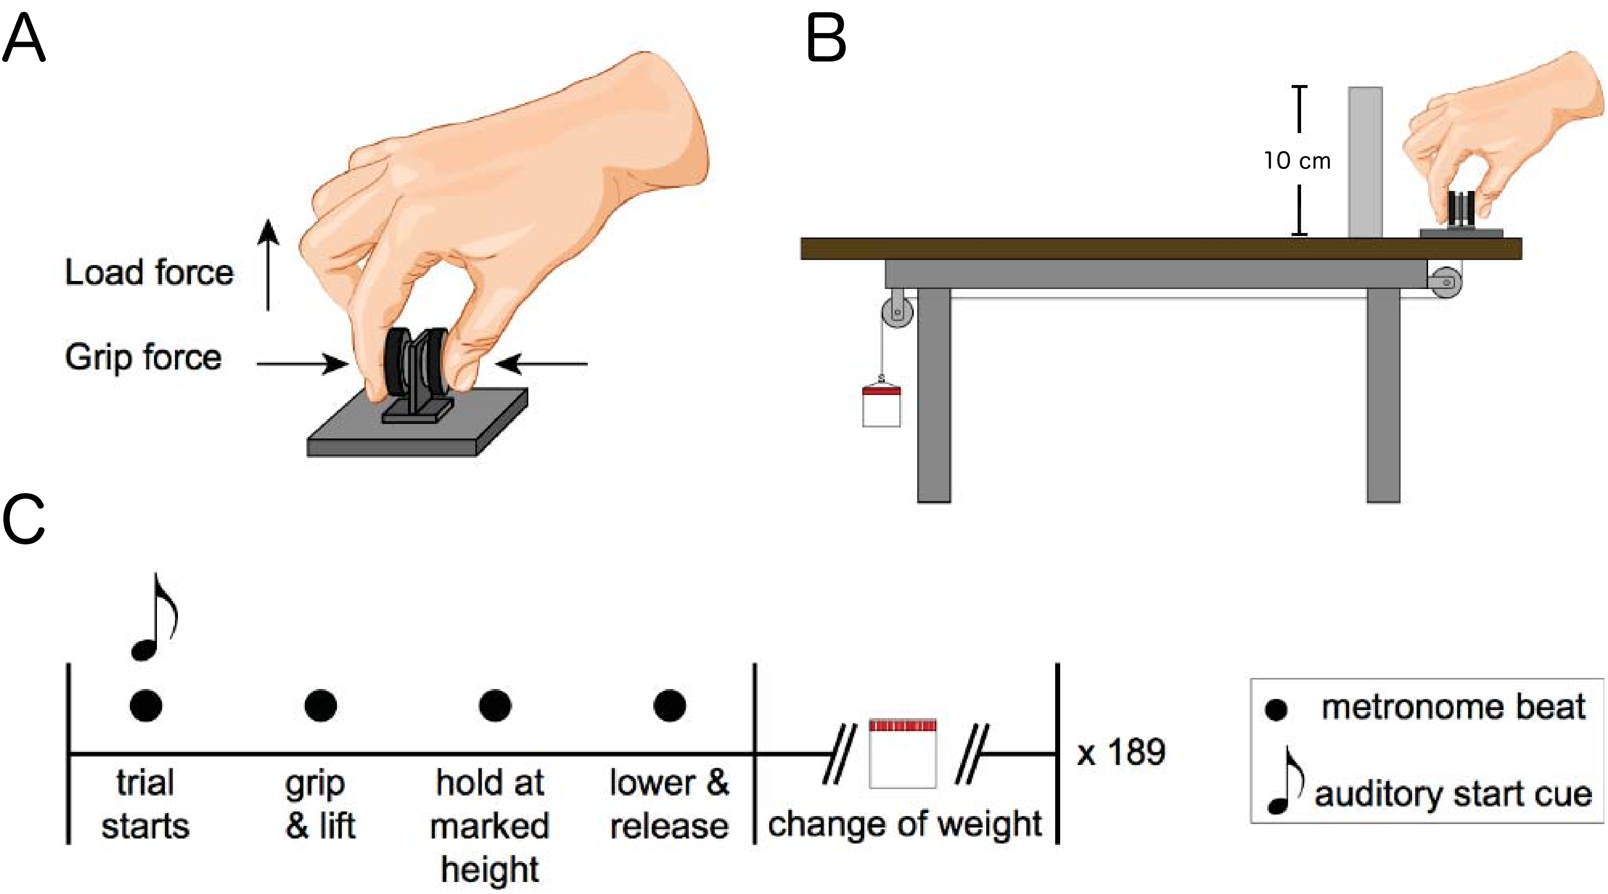
A
Load force
Grip force
B
10 cm
C
trial starts
grip & lift
hold at marked height
lower & release
change of weight
x 189
metronome beat
auditory start cue

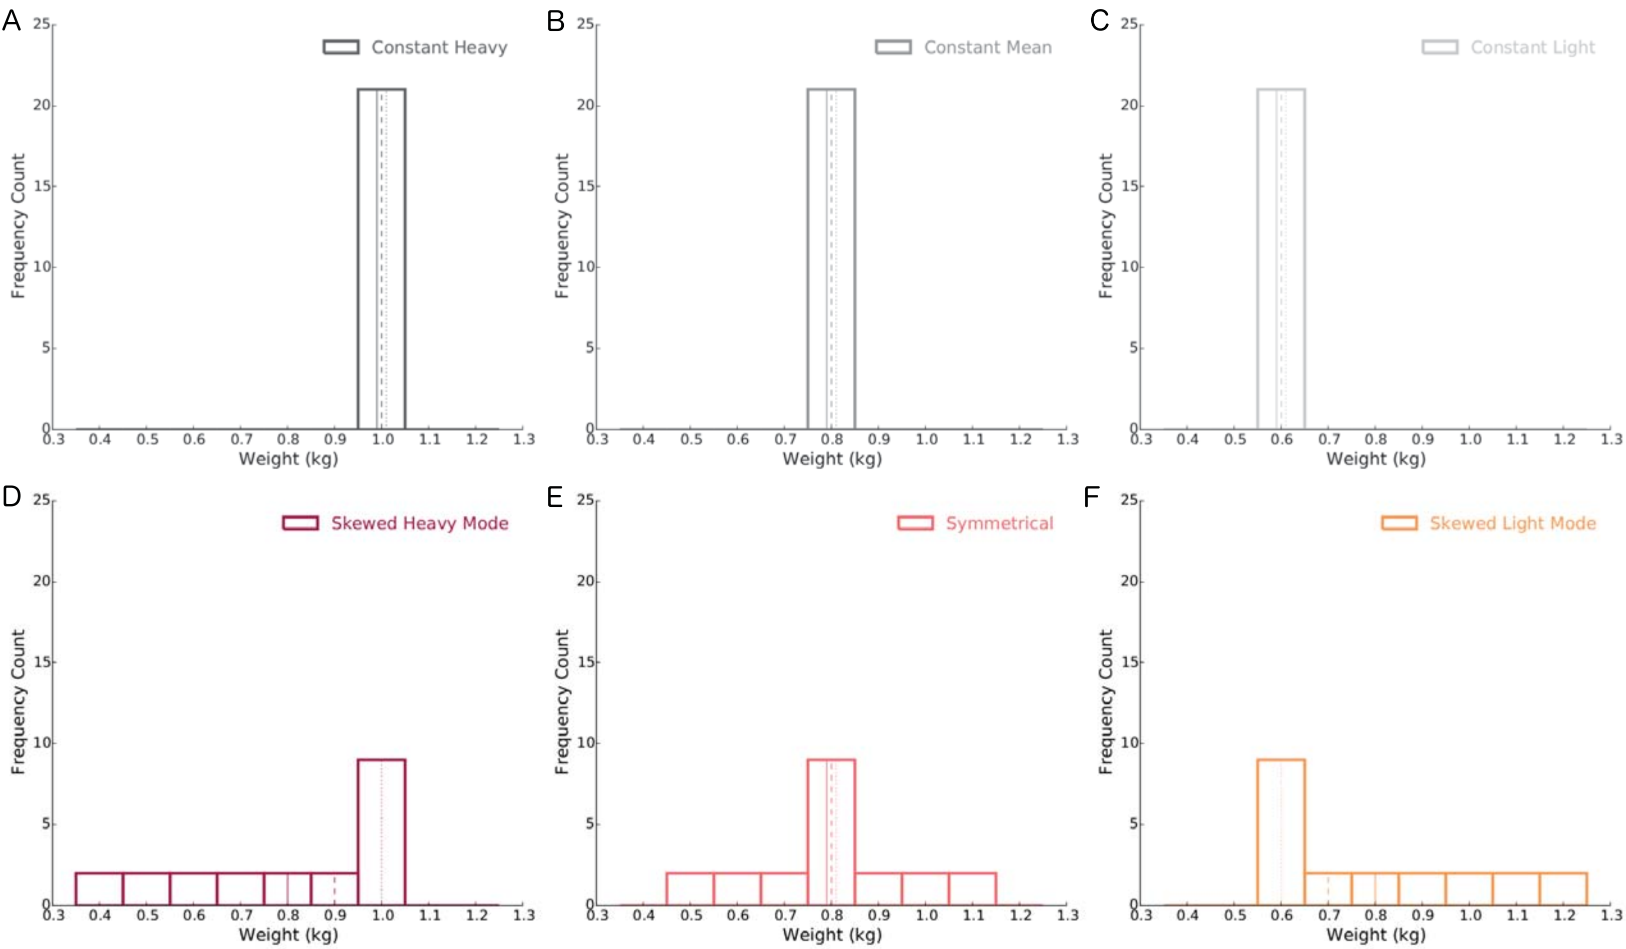
A
B
C
D
E
F
Constant Heavy
Constant Mean
Constant Light
Skewed Heavy Mode
Symmetrical
Skewed Light Mode
Frequency Count
Weight (kg)
0
5
10
15
20
25
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1.0
1.1
1.2
1.3

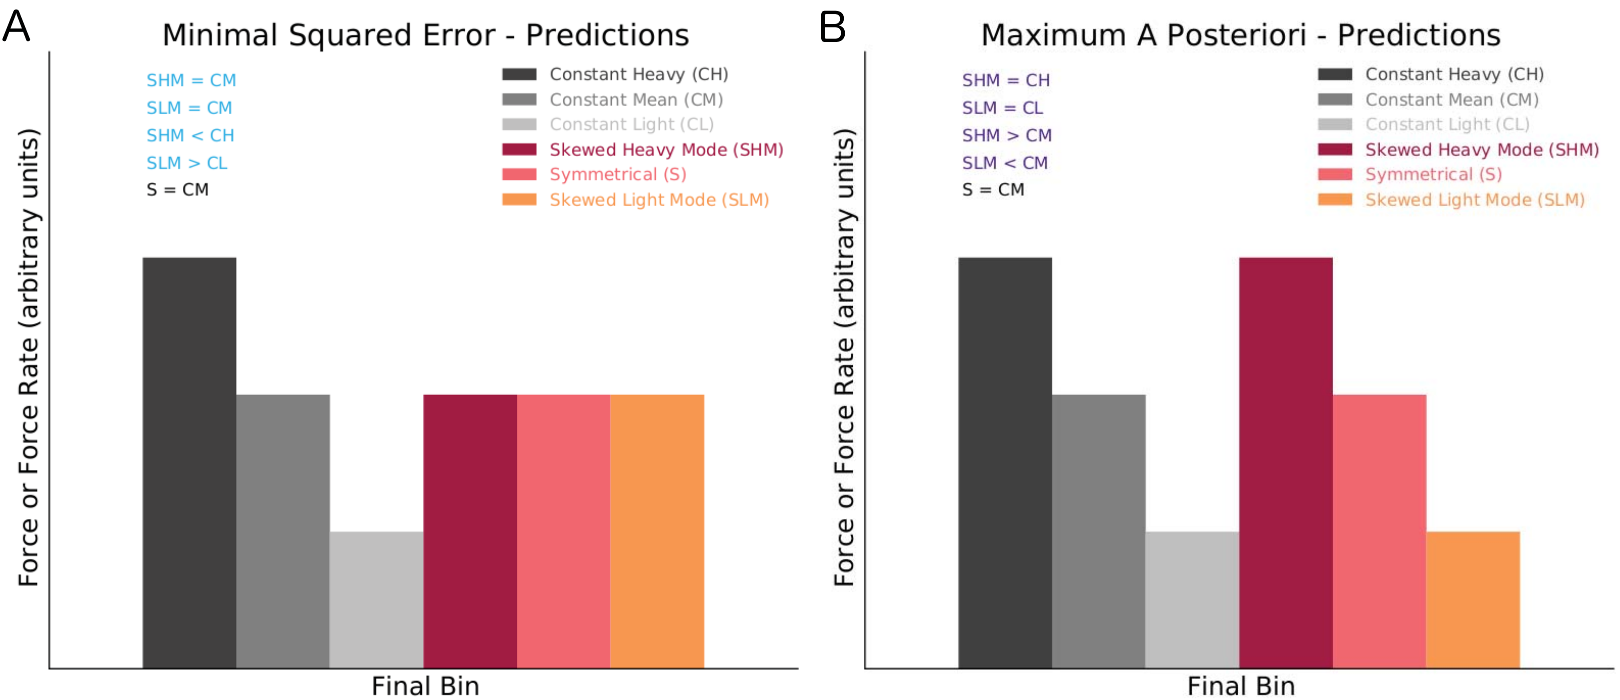
A
Minimal Squared Error - Predictions
SHM = CM
SLM = CM
SHM < CH
SLM > CL
S = CM
Constant Heavy (CH)
Constant Mean (CM)
Constant Light (CL)
Skewed Heavy Mode (SHM)
Symmetrical (S)
Skewed Light Mode (SLM)
Force or Force Rate (arbitrary units)
Final Bin
B
Maximum A Posteriori - Predictions
SHM = CH
SLM = CL
SHM > CM
SLM < CM
S = CM
Constant Heavy (CH)
Constant Mean (CM)
Constant Light (CL)
Skewed Heavy Mode (SHM)
Symmetrical (S)
Skewed Light Mode (SLM)
Force or Force Rate (arbitrary units)
Final Bin

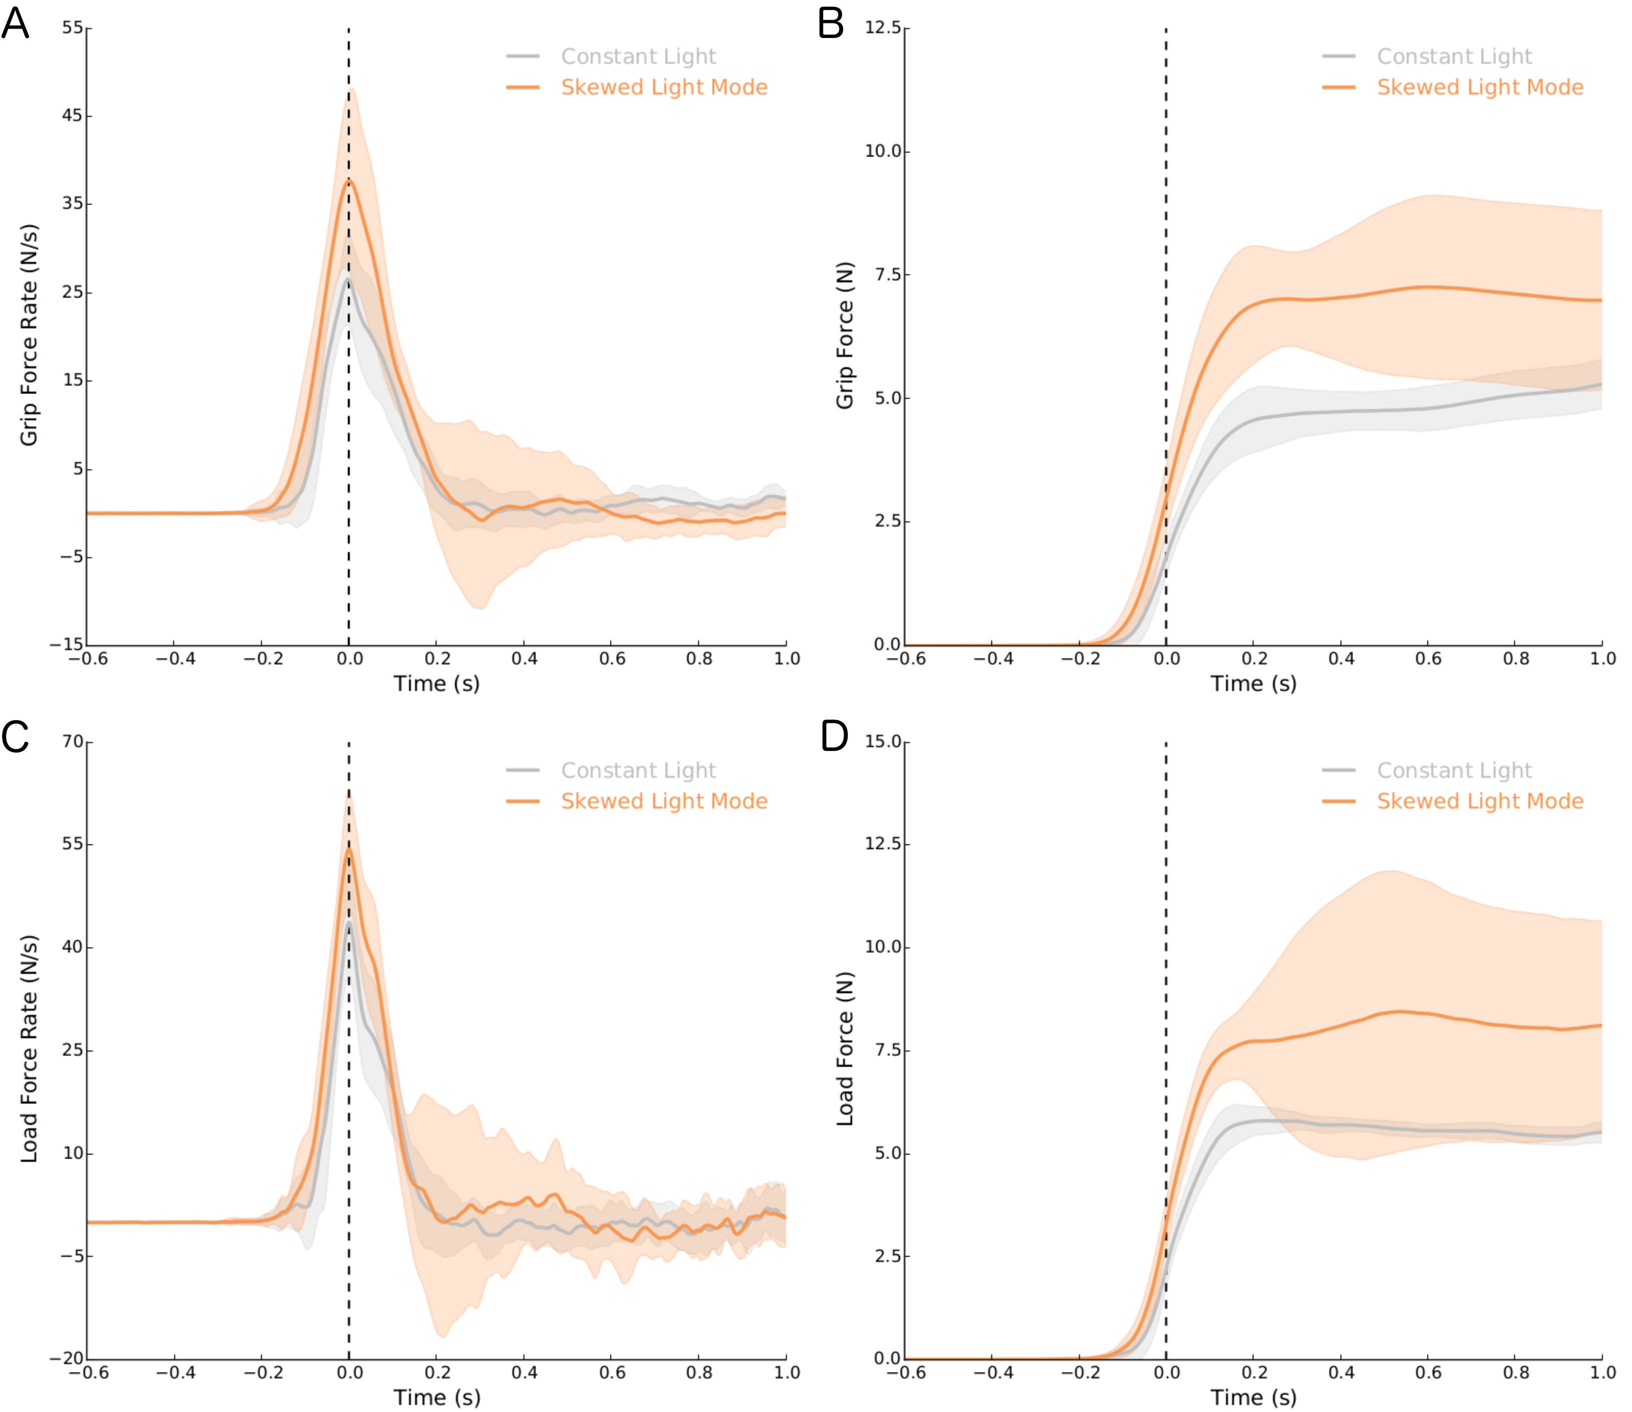
A
Constant Light
Skewed Light Mode
Grip Force Rate (N/s)
Time (s)
B
Constant Light
Skewed Light Mode
Grip Force (N)
Time (s)
C
Constant Light
Skewed Light Mode
Load Force Rate (N/s)
Time (s)
D
Constant Light
Skewed Light Mode
Load Force (N)
Time (s)

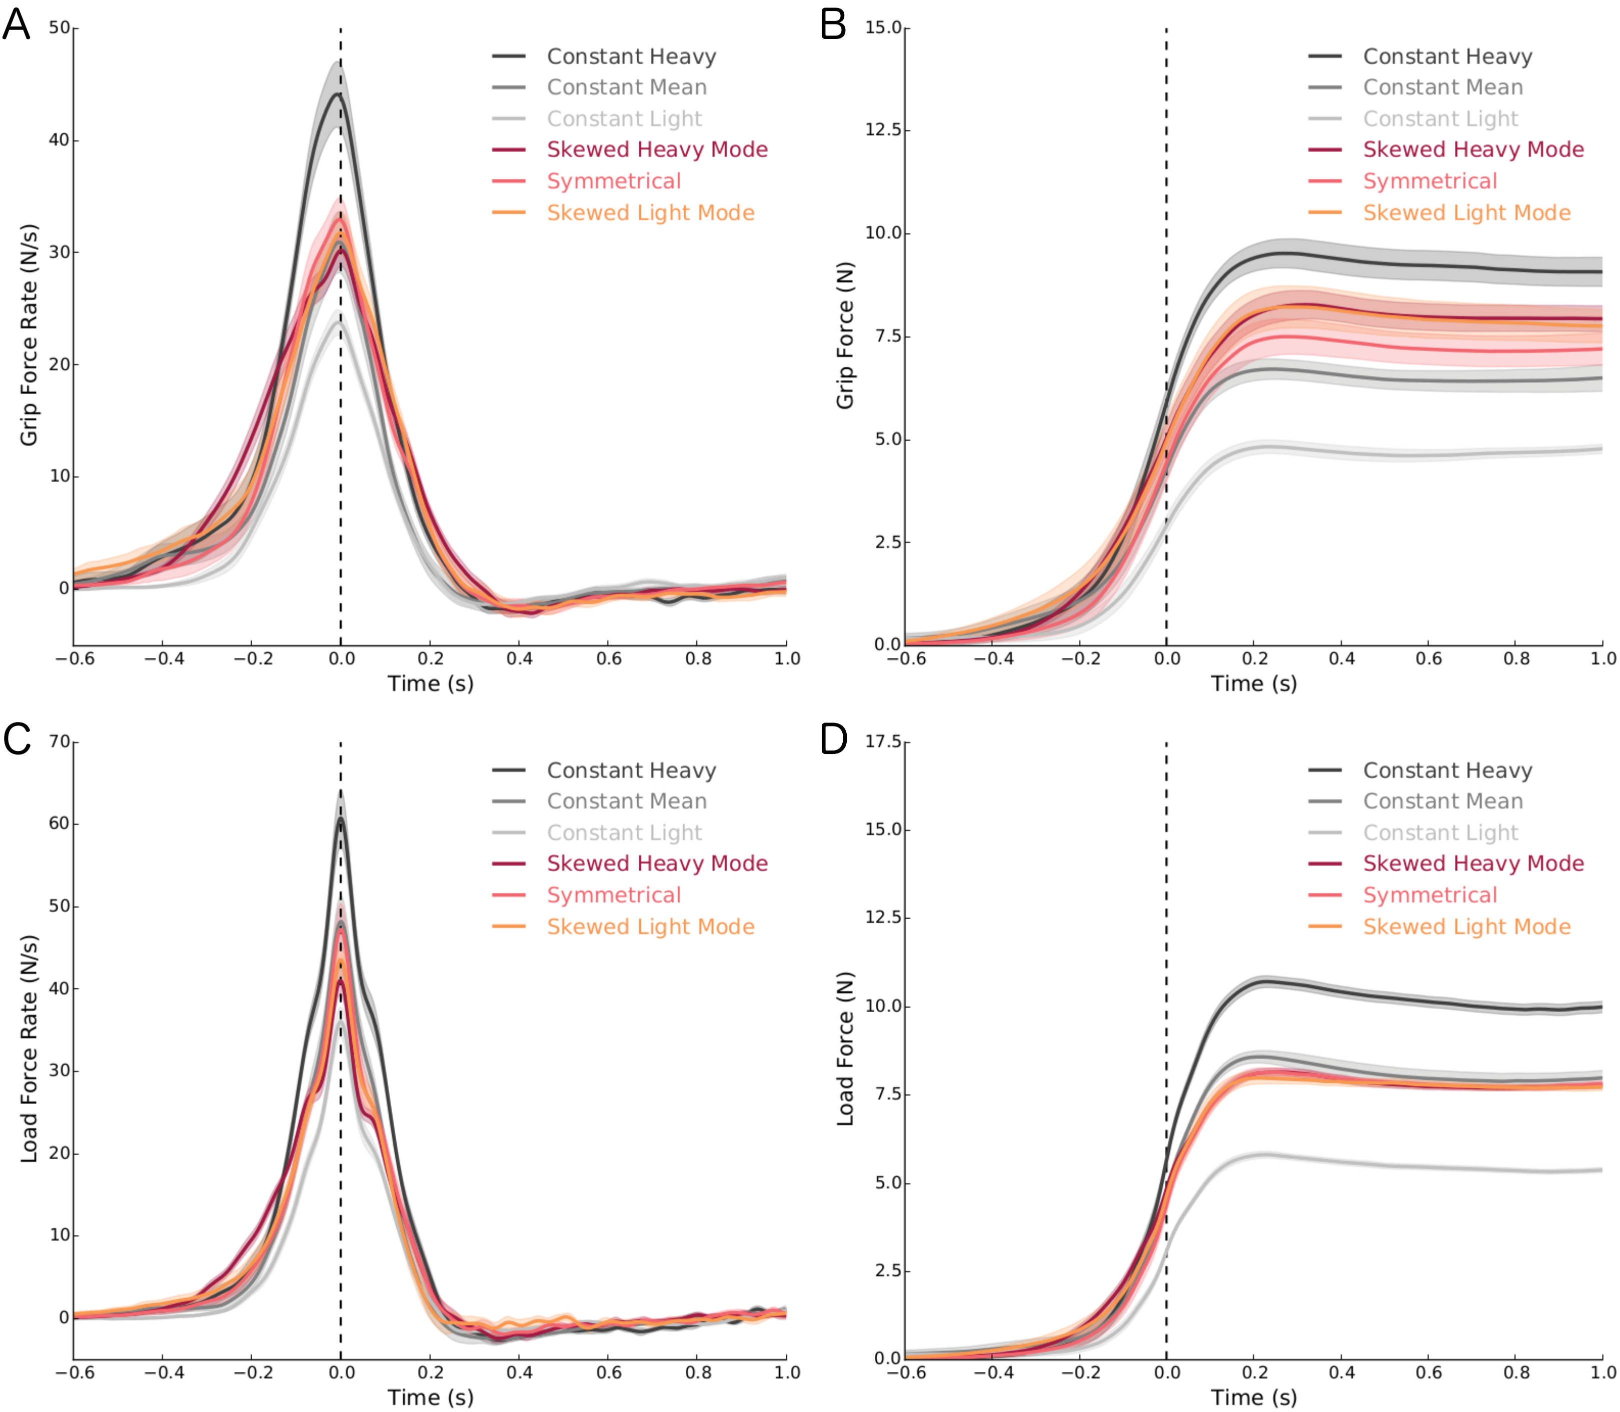
A
B
C
D
Constant Heavy
Constant Mean
Constant Light
Skewed Heavy Mode
Symmetrical
Skewed Light Mode
Grip Force Rate (N/s)
Grip Force (N)
Load Force Rate (N/s)
Load Force (N)
Time (s)

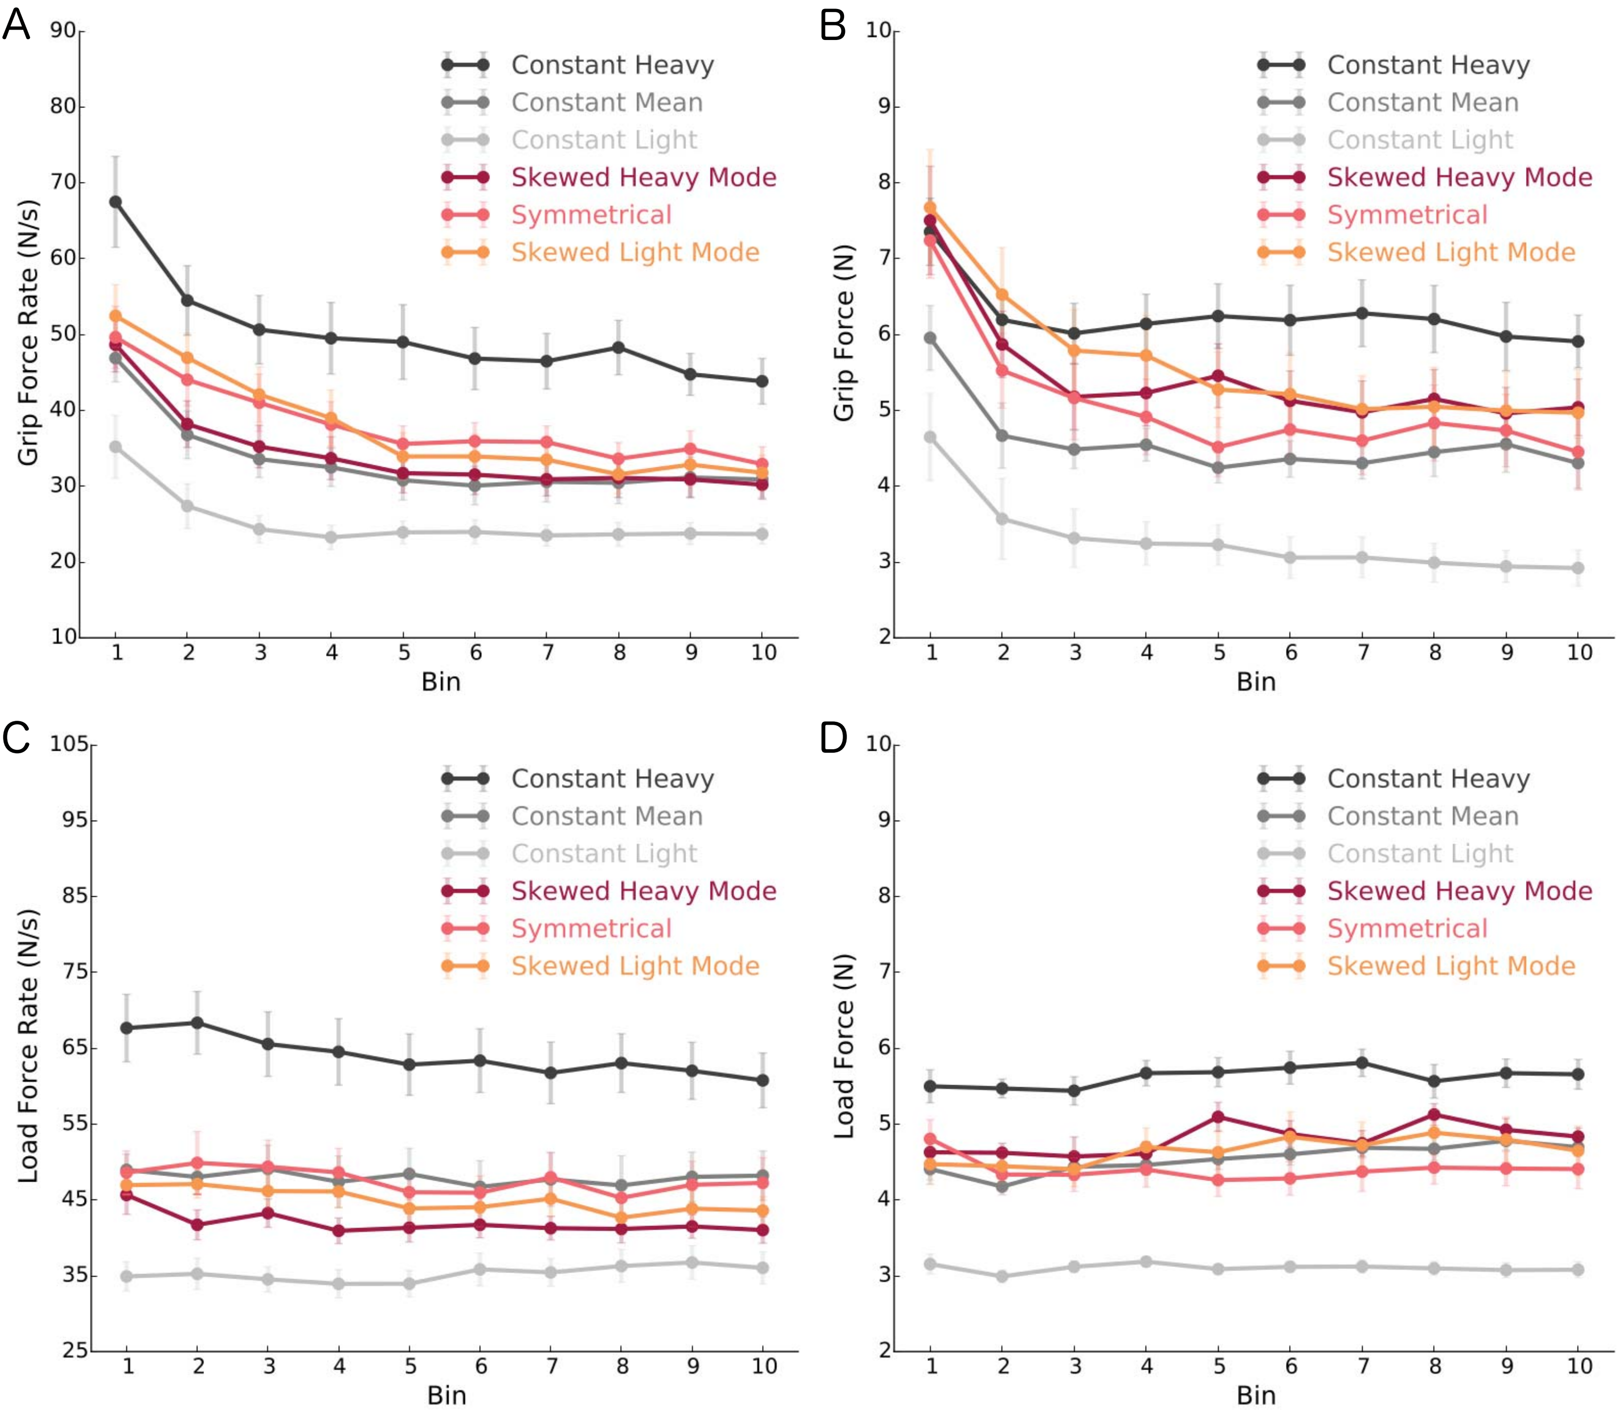
A
B
C
D
Grip Force Rate (N/s)
Grip Force (N)
Load Force Rate (N/s)
Load Force (N)
Bin
Constant Heavy
Constant Mean
Constant Light
Skewed Heavy Mode
Symmetrical
Skewed Light Mode

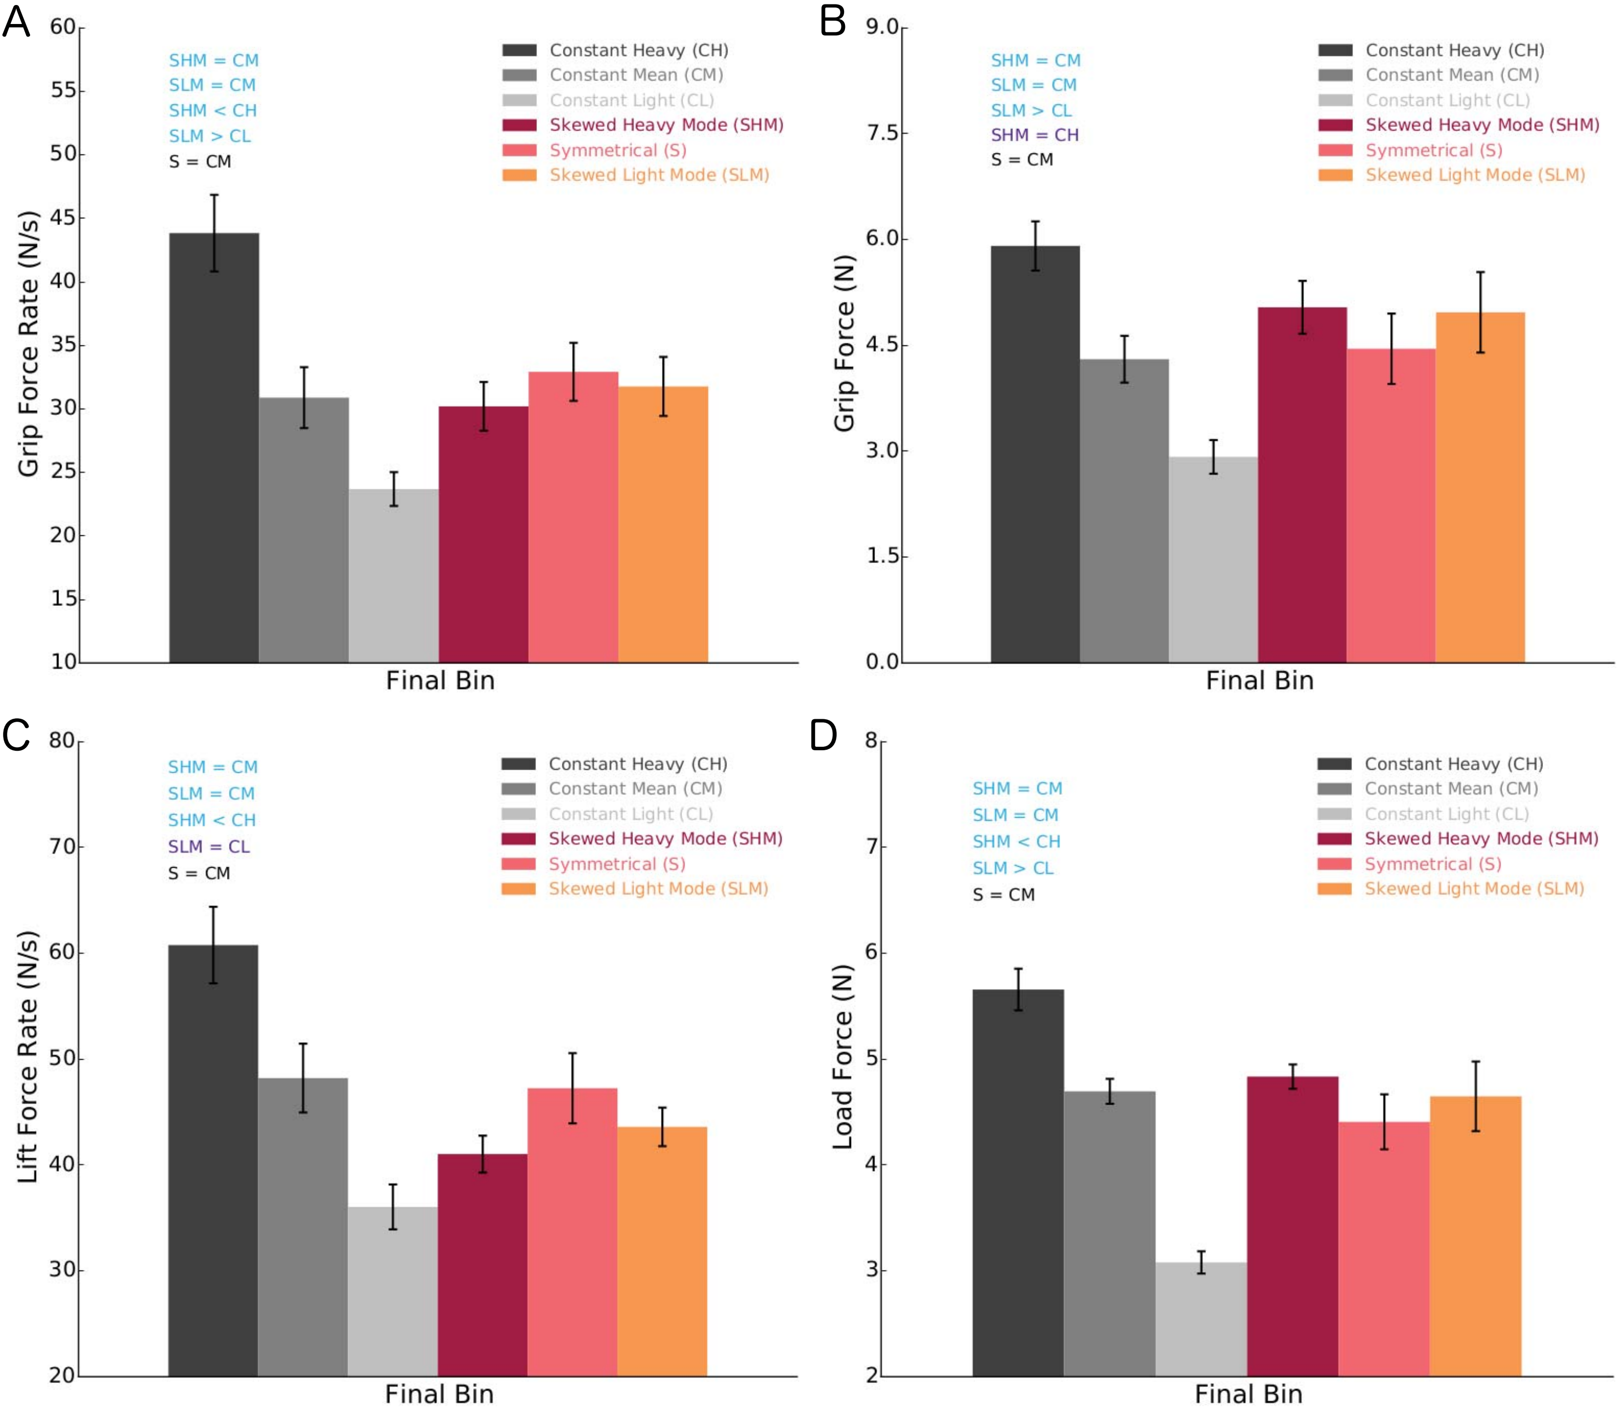
A
SHM = CM
SLM = CM
SHM < CH
SLM > CL
S = CM
Constant Heavy (CH)
Constant Mean (CM)
Constant Light (CL)
Skewed Heavy Mode (SHM)
Symmetrical (S)
Skewed Light Mode (SLM)
Grip Force Rate (N/s)
Final Bin
B
SHM = CM
SLM = CM
SLM > CL
SHM = CH
S = CM
Grip Force (N)
Final Bin
C
SHM = CM
SLM = CM
SHM < CH
SLM = CL
S = CM
Lift Force Rate (N/s)
Final Bin
D
SHM = CM
SLM = CM
SHM < CH
SLM > CL
S = CM
Load Force (N)
Final Bin

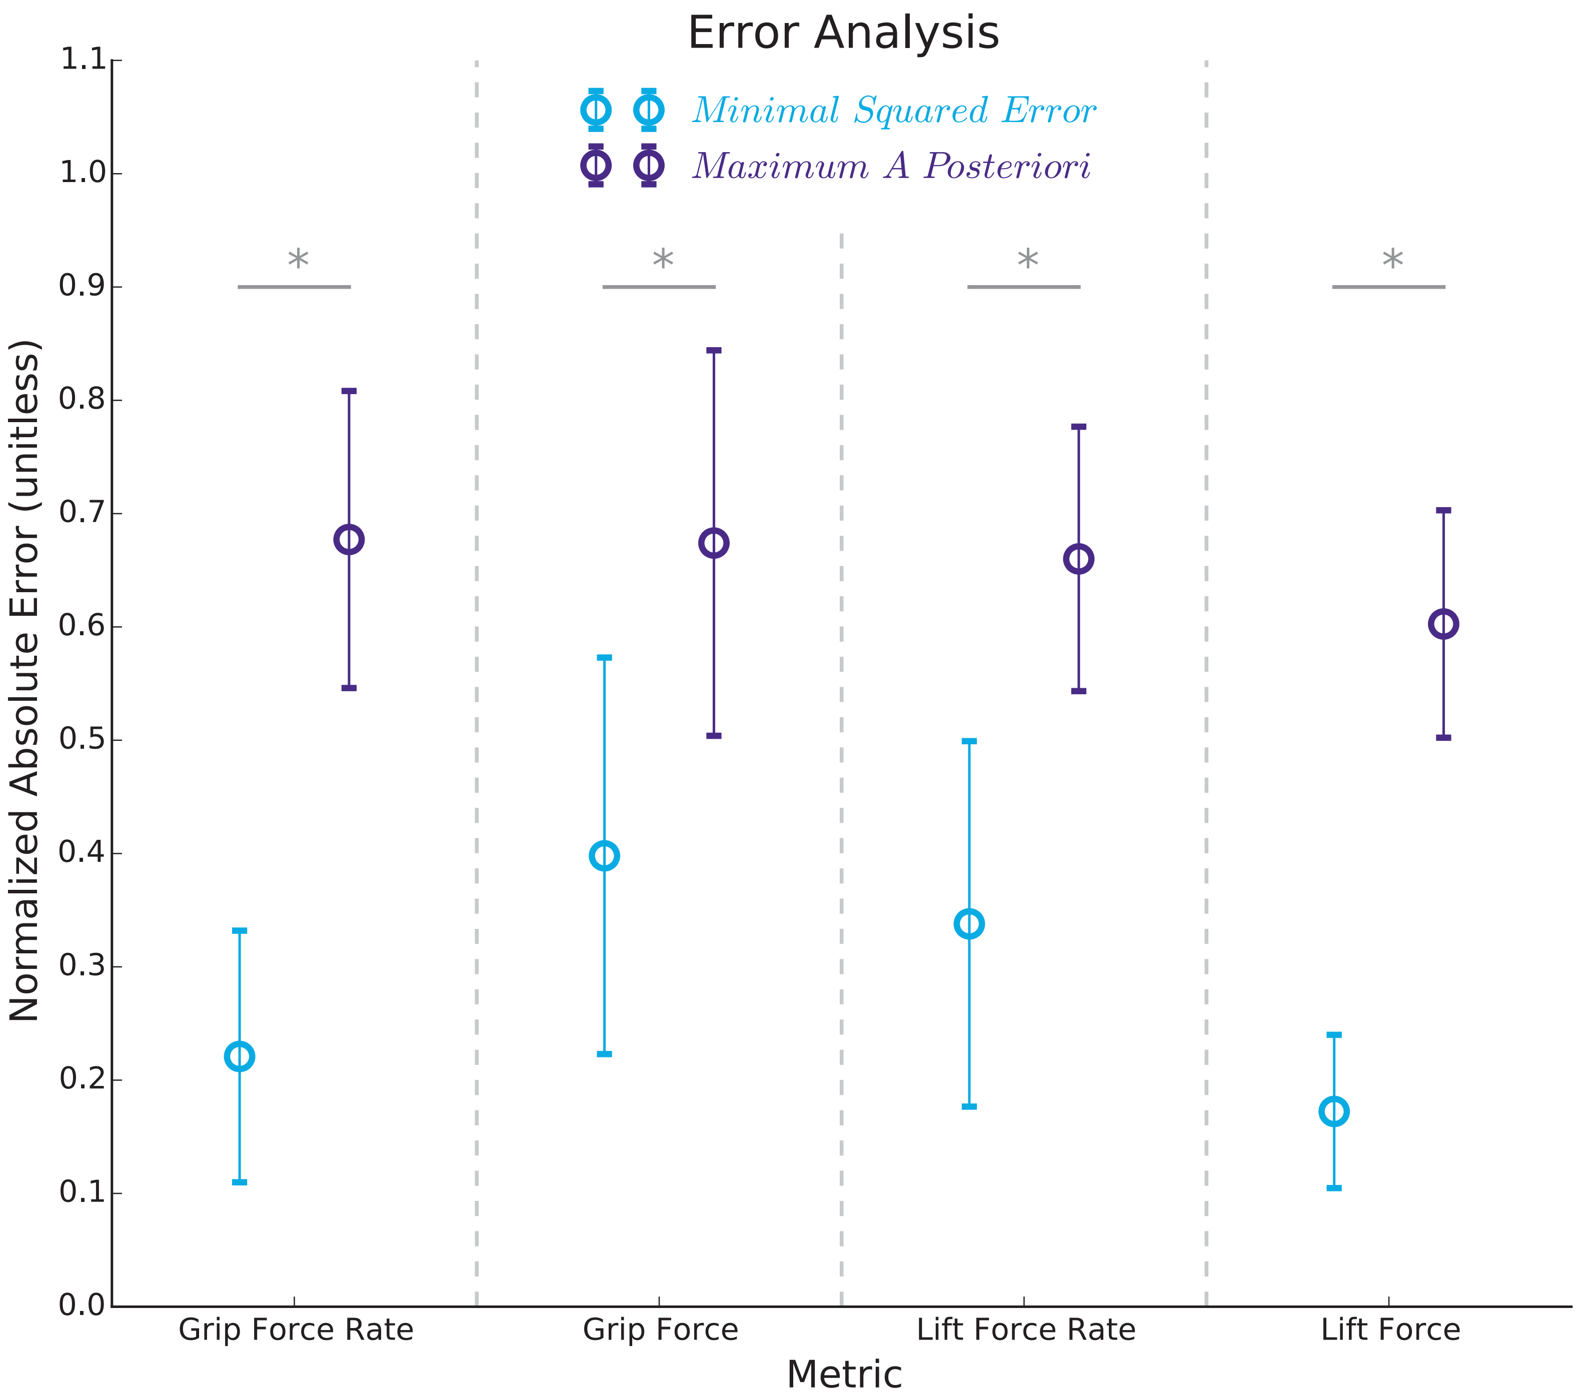
Error Analysis
Minimal Squared Error
Maximum A Posteriori
Normalized Absolute Error (unitless)
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1.0
1.1
*
*
*
*
Grip Force Rate
Grip Force
Lift Force Rate
Lift Force
Metric

Table 1

	Probability Distribution Statistics						
Probability Distribution	mean (kg)	mode (kg)	median (kg)	range (kg)	standard deviation (kg)	skew (kg)	discrete entropy (bits)
Constant Heavy	1.0	1.0	1.0	[1.0]	0.0	0.0	0.0
Constant Mean	0.8	0.8	0.8	[0.8]	0.0	0.0	0.0
Constant Light	0.6	0.6	0.6	[0.6]	0.0	0.0	0.0
Skewed Heavy Mode	0.8	1.0	0.9	[0.4, 1.0]	0.22	-0.6	1.7
Symmetrical	0.8	0.8	0.8	[0.5, 1.1]	0.16	0.0	1.7
Skewed Light Mode	0.8	0.6	0.7	[0.6, 1.2]	0.22	0.6	1.7

Table 2A

	Minimal Squared Error - Predicted Comparisons			
Measure	skewed heavy mode equal to constant mean	skewed light mode equal to constant mean	skewed heavy mode less than constant heavy	skewed light mode greater than constant light
Grip Force Rate (N/s)	**$p > 0.999$**	**$p = 0.490$**	**$p = 0.002$**	**$p = 0.003$**
Grip Force (N)	**$p = 0.565$**	**$p = 0.598$**	$p = 0.330$	**$p = 0.022$**
Lift Force Rate (N/s)	**$p = 0.294$**	**$p = 0.633$**	**$p = 0.001$**	$p = 0.051$
Lift Force (N)	**$p > 0.999$**	**$p > 0.999$**	**$p = 0.007$**	**$p = 0.002$**

Table 2B

	Maximum a Posteriori - Predicted Comparisons			
Measure	skewed heavy mode equal to constant heavy	skewed light mode equal to constant light	skewed heavy mode greater than constant mean	skewed light mode less than constant light
Grip Force Rate (N/s)	$p = 0.003$	$p = 0.008$	$p > 0.999$	$p > 0.999$
Grip Force (N)	**$p = 0.466$**	$p = 0.040$	$p = 0.424$	$p = 0.565$
Lift Force Rate (N/s)	$p = 0.002$	**$p = 0.075$**	$p > 0.999$	$p = 0.424$
Lift Force (N)	$p = 0.012$	$p = 0.004$	$p = 0.967$	$p > 0.999$

Table 3

Measure	Sensitivity to Weight	
	constant light less than constant mean	constant mean less than constant heavy
Grip Force Rate (N/s)	**$p = 0.009$**	**$p < 0.001$**
Grip Force (N)	**$p = 0.019$**	**$p = 0.020$**
Lift Force Rate (N/s)	**$p = 0.024$**	$p = 0.054$
Lift Force (N)	**$p < 0.001$**	**$p = 0.004$**

Table 4

	Sensitivity to Load Force Variance
Measure	symmetrical equal to constant mean
Grip Force Rate (N/s)	$p = 0.249$
Grip Force (N)	$p = 0.796$
Lift Force Rate (N/s)	$p > 0.999$
Lift Force (N)	$p > 0.999$