

The human motor system alters its reaching movement plan for task-irrelevant, positional forces

Joshua G. A. Cashaback,¹ Heather R. McGregor,^{1,2} and Paul L. Gribble^{1,3}

¹*Brain and Mind Institute, Department of Psychology, Western University Canada, London, Ontario, Canada;* ²*Graduate Program in Neuroscience, Western University Canada, London, Ontario, Canada;* ³*Department of Physiology and Pharmacology, Western University Canada, London, Ontario, Canada*

Submitted 11 November 2014; accepted in final form 10 January 2015

Cashaback JG, McGregor HR, Gribble PL. The human motor system alters its reaching movement plan for task-irrelevant, positional forces. *J Neurophysiol* 113: 2137–2149, 2015. First published January 14, 2015; doi:10.1152/jn.00901.2014.—The minimum intervention principle and the uncontrolled manifold hypothesis state that our nervous system only responds to force perturbations and sensorimotor noise if they affect task success. This idea has been tested in muscle and joint coordinate frames and more recently using workspace redundancy (e.g., reaching to large targets). However, reaching studies typically involve spatial and or temporal constraints. Constrained reaches represent a small proportion of movements we perform daily and may limit the emergence of natural behavior. Using more relaxed constraints, we conducted two reaching experiments to test the hypothesis that humans respond to task-relevant forces and ignore task-irrelevant forces. We found that participants responded to both task-relevant and -irrelevant forces. Interestingly, participants experiencing a task-irrelevant force, which simply pushed them into a different area of a large target and had no bearing on task success, changed their movement trajectory prior to being perturbed. These movement trajectory changes did not counteract the task-irrelevant perturbations, as shown in previous research, but rather were made into new areas of the workspace. A possible explanation for this behavior change is that participants were engaging in active exploration. Our data have implications for current models and theories on the control of biological motion.

minimum intervention principle; uncontrolled manifold hypothesis; reaching; positional forces; exploration

HUMANS ARE ABLE TO LEARN COMPLEX movements like kicking a soccer ball or hitting a tennis serve. This is quite remarkable when we consider that movement is performed in the presence of sensorimotor noise (Faisal et al. 2008; Harris and Wolpert 1998; Jones et al. 2002) and environmental forces (Shadmehr and Mussa-Ivaldi 1994, Franklin et al. 2007). To successfully perform goal-directed movements, our nervous system must select the appropriate motor commands from a set of infinite solutions. The vast number of solutions allows the nervous system to consider many criteria while performing an action (Rosenbaum 1991). It has been suggested that the relative importance of these criteria can be considered a free parameter that may change depending on task demands (Todorov 2004). In theory, this would allow the nervous system to flexibly respond to sensorimotor noise and forces and may also lead to improved performance (Liu and Todorov 2007; Todorov and Jordan 2002).

The minimum intervention principle (MIP; Todorov and Jordan 2002) and uncontrolled manifold hypothesis (UCMH; Latash 2012; Latash et al. 2007; Scholz and Schoner 1999), two current theories on the control of biological motion, are based upon the principle of flexible response (Liu and Todorov 2007). Both theories state that we only respond to sensorimotor noise and environmental forces if they affect task success. That is, our nervous system compensates for task-relevant noise and force, but ignores them if they are task-irrelevant. This flexible and robust response has been used to explain a wide range of human movement phenomena, such as individual trial-to-trial variability (Bernstein 1967; Scholz and Schoner 1999; Todorov and Jordan 2002) and synergistic behavior (Todorov 2004) and is theoretically possible by exploiting muscle, joint or workspace redundancy (Todorov and Jordan 2002).

Our ability to exploit redundancy has been primarily studied in joint and muscle coordinate frames. Many studies suggest that we take advantage of joint redundancy, but there is conflicting evidence about whether we also capitalize on muscle redundancy. In joint space, it has been shown that joint configurations can vary from trial-to-trial as long as they do not affect task success. Such a flexible response has been shown during throwing (Yang and Scholz 2005), postural (Scholz and Schoner 1999) and balancing tasks (Cluff et al. 2012), as well as unimanual (Yang et al. 2007) and bimanual (Domkin et al. 2002, 2005) reaching tasks. While there seems to be strong evidence of a flexible response in joint space, it is less clear in muscle space. In support of a flexible response, Valero-Cuevas and colleagues (2009) show that, during a finger press task where participants were required to match a desired force profile, the within-trial activation variability of seven index finger muscles follows the MIP. Specifically, they found the nervous system minimizes muscle activity variability if it affects task success, but was otherwise uncontrolled and allowed to vary. In opposition, it has been proposed that the muscle activity patterns are rigid and habitual (de Rugy et al. 2012). It would be interesting that the nervous system may exploit redundancy in one coordinate frame but not necessarily in another, particularly since they are closely coupled.

While muscle and joint redundancy have received a great deal of attention in the motor control and biomechanics literatures, much less research has focused on workspace redundancy. When performing goal directed movements, workspace redundancy describes how the endpoint (e.g., hand) can take multiple trajectories (Liu and Todorov 2007; Todorov and Jordan 2002), as well as different entry and terminal positions in large targets (Diedrichsen et al. 2010; Nashed et al. 2012,

Address for reprint requests and other correspondence: P. Gribble, Brain and Mind Institute, Dept. of Psychology, Western Univ. Canada, 1151 Richmond St., London, ON, Canada N6A 5B7 (e-mail: paul@gribblelab.org).

2014; Orban de Xivry 2013). For example, when reaching to push a swinging door, the hand can assume many locations on the door and take several paths to each location. Many studies examining workspace redundancy support a flexible response in-line with the MIP; however, other studies do not. The ability of the nervous system to capitalize on workspace redundancy has been shown behaviorally and modeled with optimal feedback controllers (Liu and Todorov 2007; Nashed et al. 2012, 2014; Todorov and Jordon 2002). In contradiction, both Schaefer et al. (2012) and Orban de Xivry (2013) reported that our nervous system does not follow the MIP when reaching to large targets. Schaefer and colleagues (2012) applied a rotational visuomotor perturbation that continuously increased the distance between cursor and hand position as participants reached from a common starting position to a large, curved rectangular target. This visuomotor perturbation did not affect task success, as it simply moved the cursor into a different position of the final target. Interestingly, they found that, when the visuomotor perturbation was removed, participants' early movements were in the opposite direction of the previously experienced visuomotor perturbation. These authors attribute this unnecessary compensation as an attempt of the nervous system to minimize sensory prediction errors. Similarly, Orban de Xivry (2013) found that participants resisted applied robot forces that caused their hand to travel along a path that was 2 cm off-center of an 8-cm-wide target, even though this hand trajectory did not affect task success. Compared with muscle and joint redundancy, workspace redundancy has received much less attention and deserves further investigation. Thoroughly examining a range of common coordinate frames and different motor control tasks allows for a more rigorous test of current theory axioms.

Human reaching experiments typically involve timed constrained, center-out reaches to small targets, with or without a viscous force field (Shadmehr and Mussa-Ivaldi 1994). These spatially and temporally constrained reaches may represent only a small portion of the movements we perform daily. When we reach for objects, there are often several locations to grasp and several hand paths to get to each location. Indeed, relaxing spatial constraints via workspace redundancy has been shown to produce rich and flexible behavior (Liu and Todorov 2007; Nashed et al. 2012, 2014; Todorov and Jordon 1998, c.f., Todorov and Jordon 2002). However, a potential limitation with several of these studies is the imposed temporal constraints (Diedrichsen et al. 2010; Liu and Todorov 2007; Nashed et al. 2012, 2014). It has recently been shown that constraining reach time can lead to hand paths resembling those that would be predicted by trajectory control models (Cluff and Scott 2014). That is, during fast-reaching movements, the hand appears to travel along an average trajectory, and even return to that average trajectory when perturbed by a force (Flash and Hogan 1985). However, when loosening these time constraints and applying the same perturbation, participants no longer return to an "average" trajectory but take a new hand path (Cluff and Scott 2014). Much of the research examining the MIP and UCMH has been done with spatially and or temporally constrained reaches. Relaxing constraints, whether through workspace redundancy (e.g., using larger targets) and or having participants choose their own pace, may allow for the emergence of a broader array of natural behavior. More realistic behavior can allow for a closer examination of

current theories and can be used to improve computational models.

To further test the MIP and UCMH while allowing for the emergence of natural behavior, we designed a task with relaxed spatial and temporal constraints. Specifically, we tested the hypothesis that humans flexibly respond to task-relevant forces and ignore task-irrelevant forces. Interestingly, we found that participants receiving a task-irrelevant force changed their movement trajectory before a perturbation, despite the force not affecting task success. For the large majority of participants, this behavior change did not counteract the task-irrelevant perturbation, as seen in previous experiments (Orban de Xivry 2013; Schaefer et al. 2012). Rather, we found participants moved in the opposite direction to unexplored areas of the workspace. To replicate the pattern of responses seen in *experiment 1* and to rule out potential biomechanical factors, we performed a second experiment. This second experiment involved a change in task geometry and direction of the applied task-irrelevant forces. Again, we found that participants receiving a task-irrelevant force changed their movement trajectory. These data do not agree with the minimization of sensory prediction errors, which has been useful to explain unnecessary compensations to task-irrelevant visuomotor perturbations (Schaefer et al. 2012). Furthermore, they are not fully consistent with the proposals of the MIP and UCMH on the control of human movement. We interpret our data from the framework of exploration during motor learning.

METHODS

Participants. One hundred and fifteen healthy participants (32 men and 83 women; age: 20.7 yr, 2.4 SD) participated in *experiments 1* and *2*. Participants reported they were right-handed, free of neuromuscular disease, and had normal or corrected vision. Participants provided informed consent to procedures approved by Western University's Ethics Board.

Apparatus. Participants performed horizontal reaching movements with their right arm while grasping a robotic handle (InMotion2, Interactive Motion Technologies, Cambridge, MA). Applied robot forces, as well as handle position and forces, were recorded at the handle at 600 Hz and stored for offline analysis. An airslid supported the participants' right arm while causing minimal friction with the desk surface. An LCD screen and a semi-silvered mirror provided participants with real-time hand position via a round yellow cursor (Fig. 1).

Experiment 1. Participants were randomly assigned to the following force field groups: null (NULL₁), task-irrelevant (TI₁), and task relevant (TR₁). All trials started from a leftward home position, and participants were instructed to move the handle rightward through two via-point targets and to stop within the boundaries of the final, curved rectangular target (Fig. 1). The final target was part of the circumference of a circle centered about the second via-point target. This ensured that all participants traveled an equal distance during the reach, regardless of their preferred final location. When the cursor passed through the two via points and final target, they turned from white to blue. Movements were considered successful if no via-point targets were missed and the final target was not overshoot.

We provided all participants with the following standardized set of instructions: 1) to make smooth movements with no sudden changes in direction; 2) to move continuously through the via points; 3) the yellow cursor had to pass completely inside the via points; 4) that they could enter the final target anywhere along the final target surface that was closest to the second via point; 5) to come to a complete stop once inside the final target; and 6) to increase their speed to a comfortable pace as they improved, while emphasizing it was "not a race" (Yang

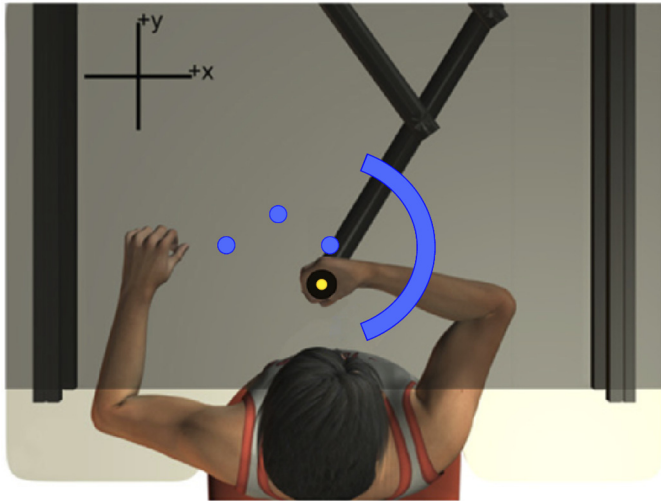


Fig. 1. Participants sat and grasped the robotic handle. A monitor projected real-time hand position (yellow cursor) and multiple blue targets onto a semi-silvered mirror. For both experiments, participants started at the leftmost blue circle, moved rightward through the two via point targets, and then stopped in the curved rectangular target.

et al. 2007). We did not inform participants about the perturbation forces.

All participants trained for 250 trials in the absence of applied force. Following a brief break (~1 min), the NULL group continued the task without any applied robot forces, while the TI₁ and TR₁ groups received a Gaussian, position-dependent force (see APPENDIX) for the remaining 250 trials. Robot forces were gradually introduced by linearly increasing force magnitude over 50 trials and held at a constant magnitude for the next 200 trials. The TI₁ group experienced task-irrelevant counterclockwise forces between the second via point and the final target. Crucially, these forces always acted parallel to the closest surface of the final target. Thus participants did not have to respond to this task-irrelevant force, as it would simply push the hand into another area of the final target and was irrelevant to task success. The TR group received posteriorly directed forces (anatomical reference frame) between the first and second via points. Conversely, these forces were relevant to task success because if participants did not compensate for the force, they would miss the second via point.

From pilot testing, we found the via-point positioning seen in Fig. 2A discouraged movements toward the uppermost area of the final target. Only in this particular region did the task-irrelevant forces become relevant, as it would push participants outside of the final target. However, despite our efforts to design a task where participants did not reach to this region, some still chose to. To ensure the counterclockwise force was task-irrelevant, participants were removed from further analysis if they reached to the uppermost region (50.0° to 67.5°) of the final target during the last 50 trials of baseline ($n = 2$). Another participant was removed from the TR₁ group because he was unable to compensate for the force and successfully complete the task. This resulted in 25 participants per group in *experiment 1*.

Experiment 2. We conducted a second experiment to replicate findings and be assured that potential biomechanical factors, such as joint comfort, interaction torques, control effort, or other possible influences, did not lead to the results of *experiment 1* (see DISCUSSION for further details). To do this, we inverted *experiment 1* about the x -axis (see Fig. 2B), and the TI group now received a task-irrelevant clockwise force. This configuration encouraged reaching toward the upper half of the final target. Participants were randomly assigned to each of the following force field groups: null (NULL₂) and task-irrelevant (TI₂). In addition to the 36 participants in *experiment 2*, 1 participant reached to the lower half of the final target. A Grubb's test confirmed that this participant's final target entry location was signif-

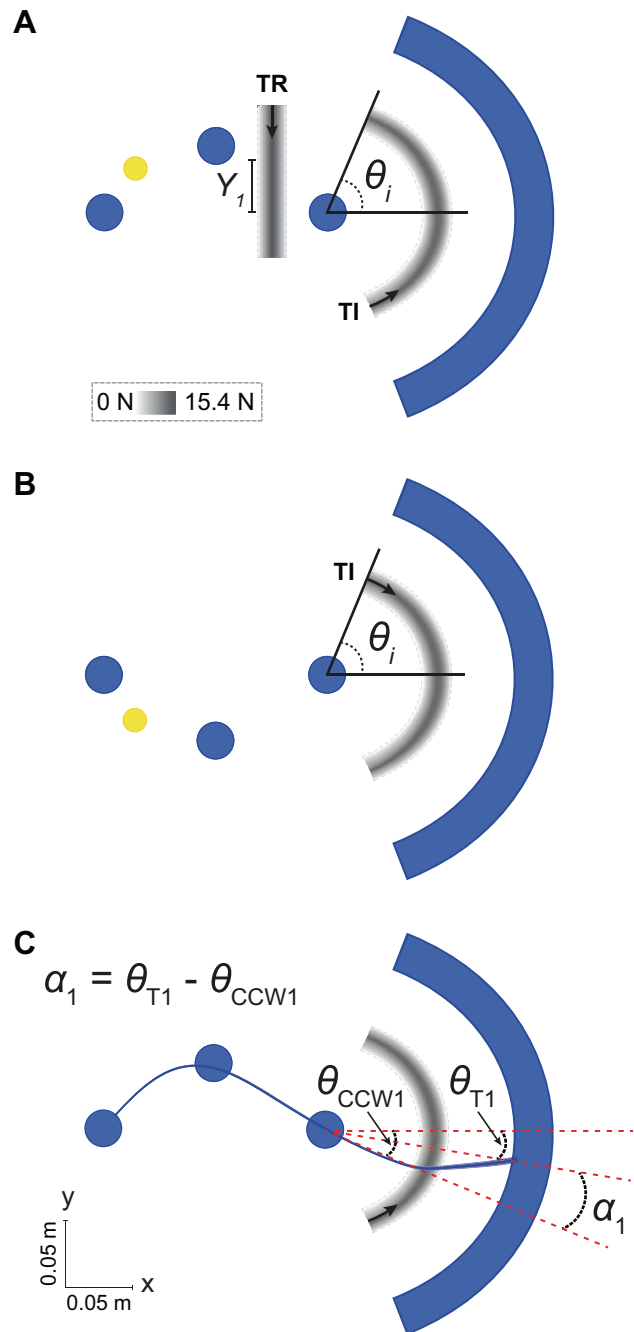


Fig. 2. *A* and *B* correspond to *experiments 1* and *2*, respectively. The robot did not perturb participants in the null force field groups (NULL₁ and NULL₂). In *experiment 1*, the task-relevant force field (TR₁) group had forces applied to their hands between the two via points (within the gray rectangle labeled TR). Entry position (Y_1) into the task-relevant force field was calculated as the y -coordinate relative to the last via point. In *experiments 1* and *2*, the task-irrelevant force field (TI₁ and TI₂) groups received a counterclockwise and clockwise force, respectively, between the second via point and the final, curved rectangular target (within the curved, gray rectangles labeled TI). These task-irrelevant forces, which acted parallel to the surface of the final target, simply pushed participants into a different area within this final target. *C*: a sample trajectory of a TI₁ participant is used to demonstrate how entry angle into the final target (θ_T) and task-irrelevant clockwise force field (θ_{CCW}) were subtracted from one another to calculate α_1 . Angles are relative to the x -axis, centered about the second via point target, and were calculated with the atan2(x,y) function. For *A*, *B*, and *C*, light to dark gray in the force fields represents an increase in force magnitude applied by the robot. The small black vector in each force field shows the direction of the applied, Gaussian robot forces. θ_i , Angle made with the x -axis.

icantly different from the other participants and was removed from further analysis. This resulted in 18 participants per group in *experiment 2*.

Data reduction and analysis. We performed offline data analyses in Python 2.7.3. All calculated measures below were averaged across blocks of 50 movements. From pilot testing, we found participants required ~ 200 trials before performance and movement duration became consistent. Thus we considered the last 50 movements before the brief break as trained, baseline motion. During the remaining 250 movements (*blocks 1–5*), the robot applied forces to the TR₁, TI₁, and TI₂ groups.

For both experiments, we examined several kinematic aspects of the participants' reaching movements. We measured the entry angle into the task-irrelevant counterclockwise (θ_{CCW1}) and task-irrelevant clockwise (θ_{CW2}) force fields, and the final target (θ_{T1} and θ_{T2}). These entry angles were quantified in polar coordinates, centered on the second via point, where θ_i is the angle made with the x-axis (see Fig. 2C). The MIP and UCMH would predict no change in θ_{CCW1} and θ_{CW2} since these forces were irrelevant to task success.

Both theories would also predict that participants would not resist this force, but rather let it push them downstream into another area of the final target. Here, we define "downstream" as the direction of the perturbing force. Thus the difference between entry angles into the final target and the task-irrelevant counterclockwise ($\alpha_1 = \theta_{T1} - \theta_{CCW1}$) or task-irrelevant clockwise ($\alpha_2 = \theta_{T2} - \theta_{CW2}$) force fields would increase and decrease for the TI₁ and TI₂ groups, respectively (see Fig. 2C).

In addition to assessing the mean θ_i for each block, we assigned a 1 or 0 to *blocks 1–5*, if participants reached downstream or upstream of the task-irrelevant forces relative to baseline. Beyond a change in average behavior, this binary measure was useful in quantifying how participants changed their movement plan in a particular direction relative to the perturbing force.

For the TR₁ group, we also calculated the y-coordinate entry position (Y_1) into the task-relevant force field between the two via points. Since this force was task-relevant and the TR₁ participants would miss the second via point without a response, we expected them to compensate by having a greater Y_1 .

Within each force field, we examined the impulse applied by the robot to the hand. Robot impulse was calculated by multiplying the sum of the resultant applied robot forces by the time spent in the force field. To determine how much the hand resisted the applied robot forces, we used the dot product to calculate the length (projection) of the force vectors recorded at the handle when mapped onto the perturbing force vector. This calculation was also performed during baseline and for the NULL groups, where the applied robot force magnitude was zero, allowing us to compare hand forces between all groups. To calculate hand impulse, the projected lengths were summed and multiplied by the time spent in the force field. Here, a positive or negative impulse indicates the hand impulse acted with or against the applied robot force vector, respectively. For each trial, movement duration was calculated as the difference in time between leaving the home position and entering the final target.

As stated above, we found participants changed their entry angle into the task-irrelevant force field. A potential explanation for this behavior change is that participants may have been actively exploring the workspace. To provide a measure of exploration beyond just an average change in behavior, we calculated the spread [i.e., interquartile range (IQR)] of participants' change in entry angle, relative to baseline, into the task-irrelevant force field. IQR is a measure of distribution spread and is robust against outliers (Reimann et al. 2011). IQR calculations were done at the group level, across blocks and for individual blocks (1–5), for all NULL and TI groups. For the NULL₁ and TR₁ groups, we also calculated IQR for entry position changes, relative to baseline, into the task-relevant force field. If participants in the TI groups were exploring, we would expect a greater IQR relative to the NULL groups. Conversely, if participants

in the TI and TR groups were not exploring and simply changing their mean behavior, we would expect their IQR to remain relatively constant.

Statistical analysis. We performed five, two-way, mixed-factorial analyses of variance (ANOVA) for kinematic measures. For *experiment 1*, ANOVA-dependent variables were as follows: 1) entry angle

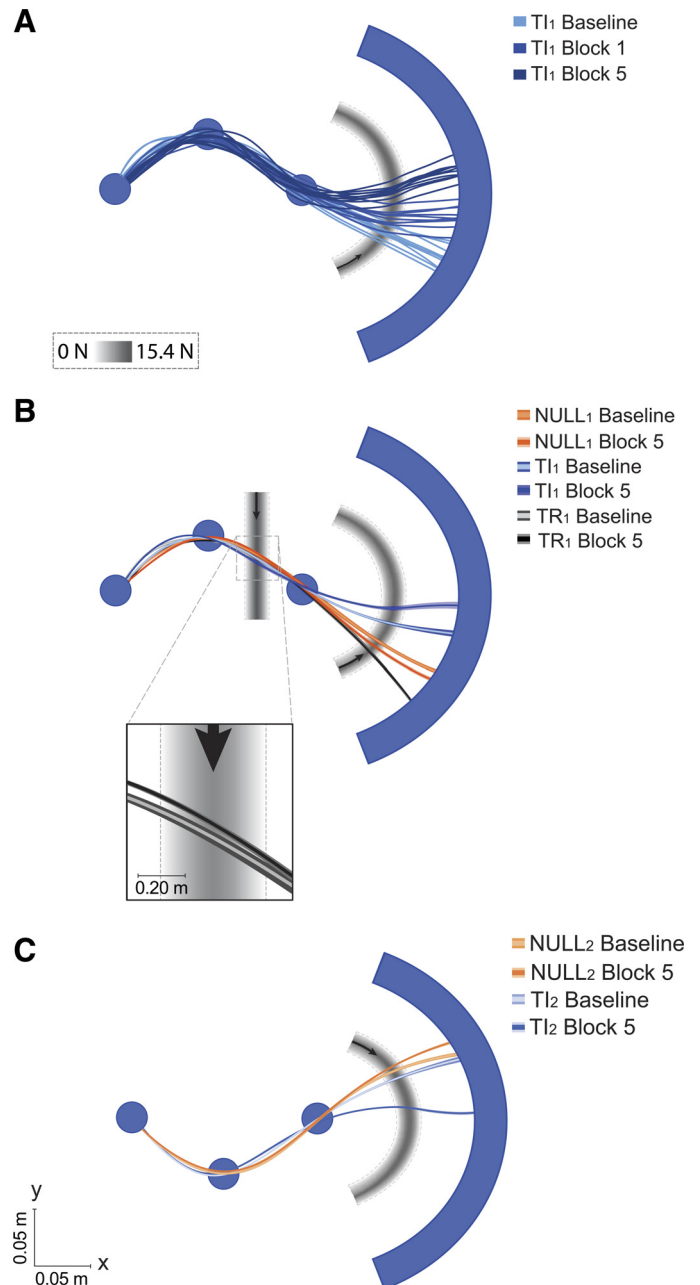


Fig. 3. A: representative sample trajectories of a TI₁ participant from baseline, *block 1*, and *block 5*. Average baseline and *block 5* hand trajectories (\pm SE, $n = 50$ trials per block) of five participants, each from one of the five groups within *experiment 1* (TI₁, TR₁, and NULL₁); B) and *experiment 2* (TI₂ and NULL₂); C) are shown. As seen in this figure, the two participants from the task-irrelevant groups changed their entry angle downstream of the upcoming force (higher and lower entry angles in *experiments 1* and *2*, respectively), despite this perturbation being unavoidable and irrelevant to task success. The square inset of B isolates and magnifies the TR₁ participant's mean trajectory over the task-relevant force field. This representative participant had greater Y_1 in *block 5* relative to baseline. Light to dark gray in the force fields represents an increase in force magnitude applied by the robot. The small black vector in each force field shows the direction of the applied, Gaussian robot forces.

into the task-irrelevant counterclockwise force field (θ_{CCW1}); 2) the difference between the final target and task-irrelevant counterclockwise force field entry angles (α_1); and 3) Y_1 into the task-relevant force field. Independent variables were group (NULL₁, TI₁, and TR₁) and block (6 levels: baseline to 5). For *experiment 2*, the kinematic-dependent variables were as follows: 1) entry angle into the task-irrelevant clockwise force field (θ_{CW2}); and 2) the difference between the final target and task-irrelevant clockwise force field entry angles (α_2). Independent variables were group (NULL₂ and TI₂) and block (baseline to 5). We also performed two, repeated-measures logistic regressions (RMLR) on the binary counts that quantified whether participants moved downstream or upstream of the task-irrelevant forces compared with baseline. Here, independent variables for the two separate regressions were group (*experiment 1*: NULL₁, TI₁, and TR₁; *experiment 2*: NULL₂ and TI₂) and block (1 to 5).

We performed four, two-way, mixed-factorial ANOVAs for kinetic measures. An ANOVA compared the robot impulses applied to the different force field groups (TR₁, TI₁, and TI₂) over 5 blocks (1 to 5). For the different force fields (task-relevant, task-irrelevant counterclockwise, and task-irrelevant clockwise), three separate ANOVAs were performed on the dependent variable of hand impulse. Independent variables were group (*experiment 1*: NULL₁, TI₁, and TR₁; *experiment 2*: NULL₂ and TI₂) and block (baseline to block 5). Two additional two-way repeated-measures ANOVAs were performed on movement duration, with independent variables of group (*experiment 1*: NULL₁, TI₁, and TR₁; *experiment 2*: NULL₂ and TI₂) and block (baseline to block 5).

For both TI groups, we calculated the coefficient of determination (R^2) between baseline and block (1 to 5) entry angle into the task-irrelevant force field. We expected nonsignificant correlations, which would indicate changes in entry angle were not related to initial entry position. This assured that entry angle changes were not due to

participants deciding to reach to the center of the final target or choosing to make more horizontal reaching movements upon experiencing the task-irrelevant force field, while also controlling for spurious biomechanical factors.

A bootstrap analysis was used to determine whether there was significant differences between the NULL and experimental TI and TR groups. This was done across blocks (overall) and for individual blocks (1–5). We performed 10,000 resamples for each comparison.

SPSS (version 21.0; IBM, Armond, NY) was used for all statistical analyses. We used Greenhouse-Geisser corrections to control for sphericity. Holm-Bonferroni corrections were used to make multiple comparisons for the RMLR and bootstrapping analyses, while Tukey's honestly significant difference was used for mean comparisons from the ANOVA. Significance was set to $P < 0.05$ for all statistical tests.

RESULTS

Ten sample trajectories of a TI₁ participant are shown at the end of baseline, the middle of *block 1*, and the end of *block 5* (Fig. 3A). Average trajectories, across baseline and *block 5*, of a participant from each of the five groups in *experiments 1* and *2* are displayed in Fig. 3, B and C. Here, it can be seen between baseline and *block 5*, that the TI₁ participant increased his or her entry angle into the task-irrelevant force field (Fig. 3B), while the TI₂ participant decreased his or her entry angle into the task-irrelevant force fields (Fig. 3C). The NULL₁ and NULL₂ participants kept relatively consistent trajectories between baseline and *block 5* (Fig. 3, B and C). The TR₁ participant had a slight increase in entry position (4.0 mm)

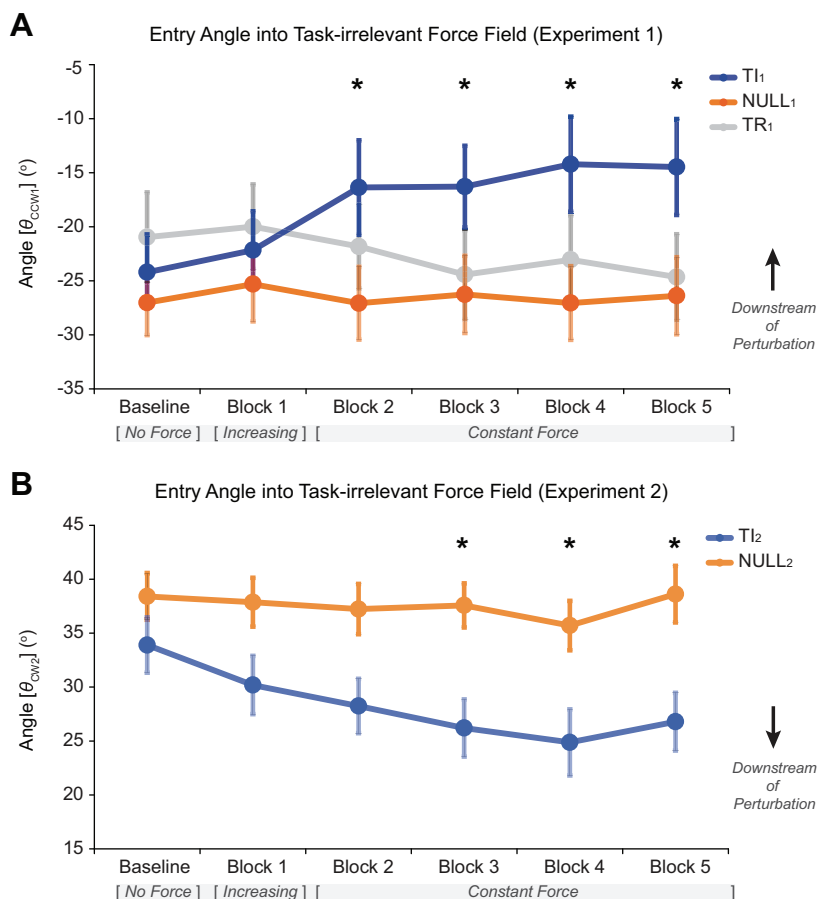


Fig. 4. Entry angle ($^{\circ}$) into the task-irrelevant counterclockwise (θ_{CCW1} ; A) and task-irrelevant clockwise (θ_{CW2} ; B) force fields of *experiments 1* and *2*, respectively. In both experiments, the task-irrelevant force field groups had significantly different entry angles compared with their corresponding NULL force field group. Relative to baseline, positive and negative changes in A and B represent an entry angle downstream of the upcoming force. Standard error bars are presented. * $P < 0.05$.

when experiencing the task-relevant force field (see Fig. 3*B*, *inset*).

Kinematics. We found that participants significantly changed their entry angle (θ_{CCW1} and θ_{CW2}) into the task-irrelevant force fields (see Fig. 4), representing a change in movement plan. For θ_{CCW1} , there was a significant interaction between group and block [$F(5.9,211.7) = 4.8, P < 0.001$]. Here, the TI_1 group had a significantly greater entry angle compared with the $NULL_1$ group from *blocks 2–5*. Despite the relaxed spatial constraints, at baseline there were no significant differences between the three groups. Likewise, we found a significant interaction of group and block for θ_{CW2} [$F(3.5,102.4) = 3.5, P = 0.018$]. At *blocks 3, 4, and 5*, the TI_2 group had a significantly lower entry angle than the $NULL_2$ group. Thus, in both experiments, the TI groups changed their average entry angle, aiming downstream of the task-irrelevant force. Again, downstream is a position change in the direction as the perturbing force. Figure 5 shows the percentage of participants per group who moved downstream of the task-irrelevant force. The RMLR showed that the TI groups had a significantly greater probability of moving downstream of the task-irrelevant force than the other groups (*experiment 1*: $\chi^2 = 17.1, P < 0.001$; *experiment 2*: $\chi^2 = 4.8, P < 0.028$). Across blocks and experiments, 88.4% of participant in the TI groups moved downstream of the upcoming task-irrelevant force field. This represents an unnecessary change in movement planning that had no bearing on task success.

As shown in Fig. 6, participants allowed the task-irrelevant force field to push them downstream into a new area of the final target (α_1 and α_2). For α_1 , there was a significant interaction of group and block [$F(5.9,211.6) = 12.4, P < 0.001$]. Compared with $NULL_1$, the TI_1 group had a significantly greater difference between final target and task-irrelevant counterclockwise entry angles from *blocks 2–5*. There were no significant differences between the three groups at baseline. Similarly, we found a significant group and block interaction for α_2 [$F(3.0,101.9) = 19.1, P < 0.001$], where TI_2 was significantly different from $NULL_2$ from *blocks 1–5*.

As expected, the TR_1 group increased their Y_1 into the task-relevant force field (Fig. 7). For Y_1 , there was a significant group and block interaction [$F(9.0,324.9) = 6.1, P < 0.001$]. From *blocks 1–5*, TR_1 had a significantly greater Y_1 compared with $NULL_1$. There were no significant differences between the three groups at baseline.

Kinetics. All groups received similar amounts of applied robot impulse (Fig. 8). The ANOVA examining the applied robot impulse found no significant interaction of group and block [$F(5.7,184.9) = 1.5, P = 0.197$] or main effect of group [$F(2.0,65.0) = 3.0, P = 0.058$]. To measure the participants' resistance to the applied robot impulses, we recorded the impulse at their hand. For both experiments, there was a significant group and block interaction for hand impulse in the different force fields [task-relevant: $F(4.9,175.8) = 19.5, P < 0.001$; task-irrelevant counterclockwise: $F(4.5,163.7) = 39.5, P < 0.001$; and task-irrelevant clockwise: $F(2.3,79.2) = 32.2, P < 0.001$]. In all cases, the force field groups (TR_1, TI_1 , and TI_2) had significantly more negative hand impulse than their respective $NULL$ group (see Fig. 9). Again, a negative hand impulse represents a resistance to the perturbing force vector. However, it is important to note that it is difficult to dissociate whether a more negative hand impulse was due to inertia when

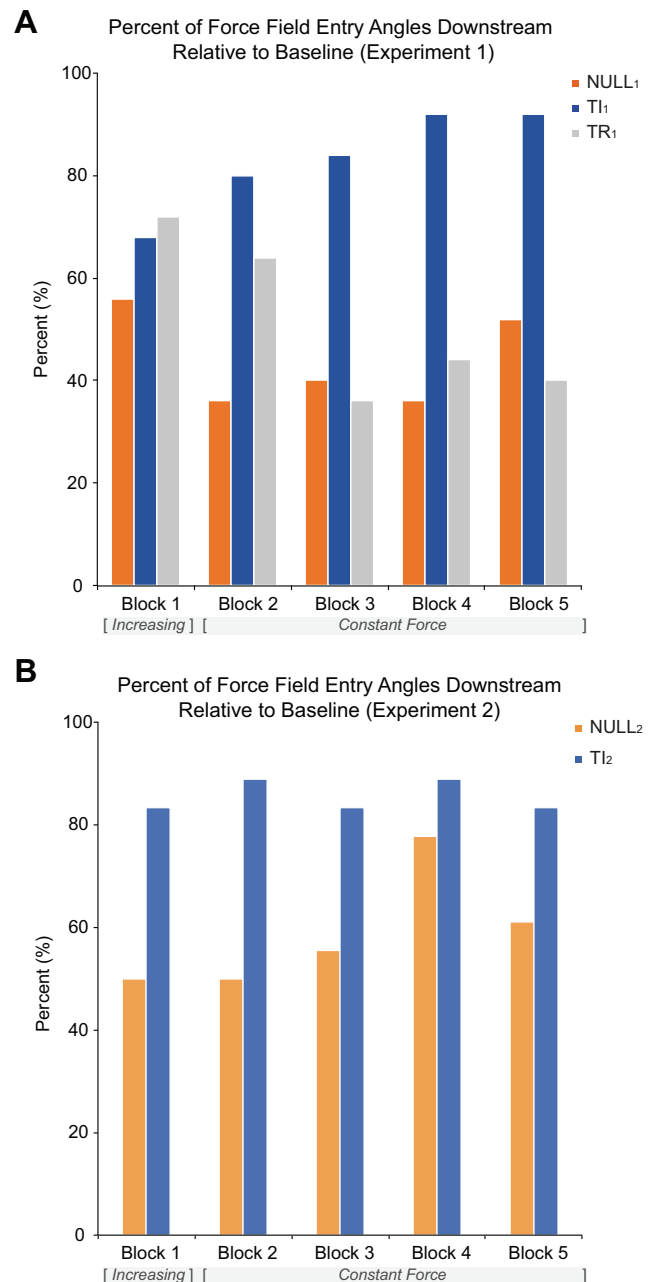


Fig. 5. Percentage (%) of participants who entered downstream of the upcoming task-irrelevant counterclockwise (A) and task-irrelevant counterclockwise force fields (B), relative to baseline.

perturbed, or participants actively resisting the applied robot forces.

Movement duration and trial successes. For movement duration (Fig. 10), the group and block interaction was significant in *experiment 1* [$F(6.7,242.6) = 3.7, P = 0.001$] and insignificant in *experiment 2* [$F(3.1,104.8) = 1.4, P = 0.254$]. However, no significant differences between groups were found in *experiment 1*. Furthermore, despite allowing participants to move at a comfortable pace, all group differences were less than 85 ms across both experiments. Overall, the average trial success rate for *experiments 1* and *2* was 90.2% and 94.1%, respectively, with differences between groups being less than 1.6%.

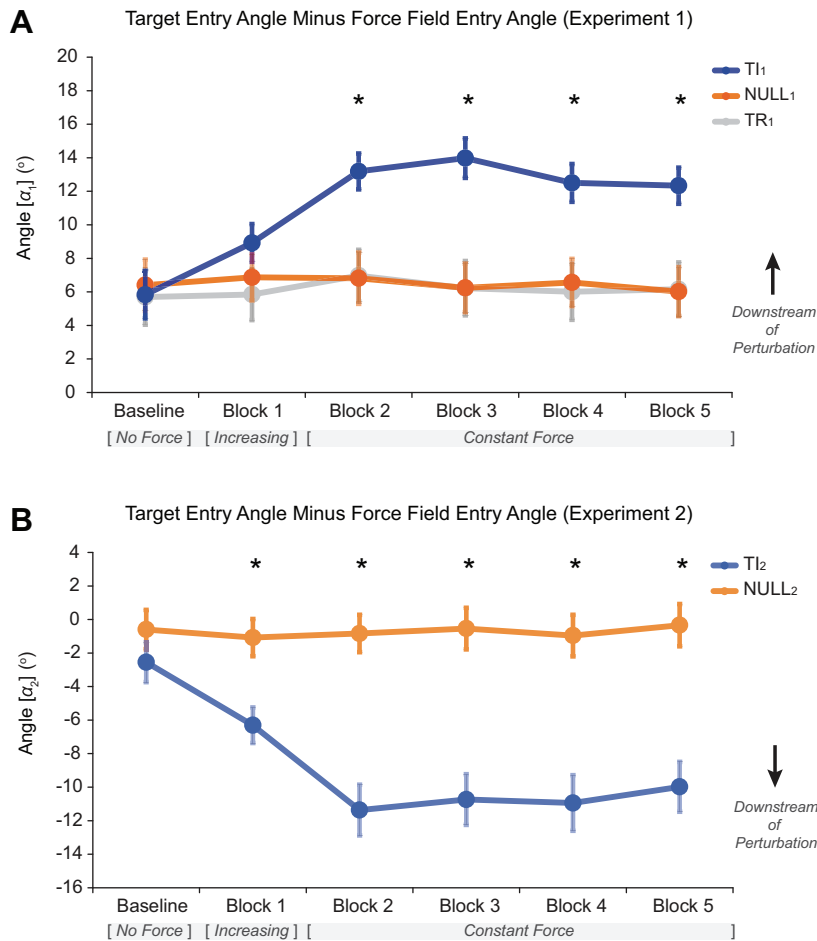


Fig. 6. Final target minus task-irrelevant counterclockwise (α_1) (A) and task-irrelevant clockwise (α_2) force field entry angles ($^\circ$) (B). Positive and negative changes from baseline for α_1 and α_2 , respectively, indicate that the task-irrelevant groups (TI₁ and TI₂) allowed the task-irrelevant force to push them into a new area of the final target. Standard error bars are presented. * $P < 0.05$.

Correlation between baseline and changes in task-irrelevant entry angle. As expected, there were no significant correlations between baseline entry angle and subsequent changes in entry angle from blocks 1–5 (Table 1). For the TI₁ group, the R^2 and P values, respectively, across blocks ranged from 0.00 to 0.04 and 0.35 to 0.91. Likewise for the TI₂ group, the R^2 and P values across blocks ranged from 0.00 to 0.07 and 0.29 to 0.94, respectively. These nonsignificant findings indicate that any entry angle changes from blocks 1–5 were not related to participants choosing to reach to the middle of the target,

deciding to make more horizontal movements, or spurious biomechanical factors.

Spread of entry angle and entry position. To provide a metric of exploration beyond mean changes in entry angle and entry position into the task-irrelevant and task-relevant force fields, we also looked at the spread of these changes using IQR. Here, a larger spread represents greater exploration. Figure 11, A–D, shows the overall IQR and histogram, collapsed across blocks, of entry angle changes into the task-irrelevant force field for the TI_{1&2} and NULL_{1&2} groups. Figure 11, E and F,

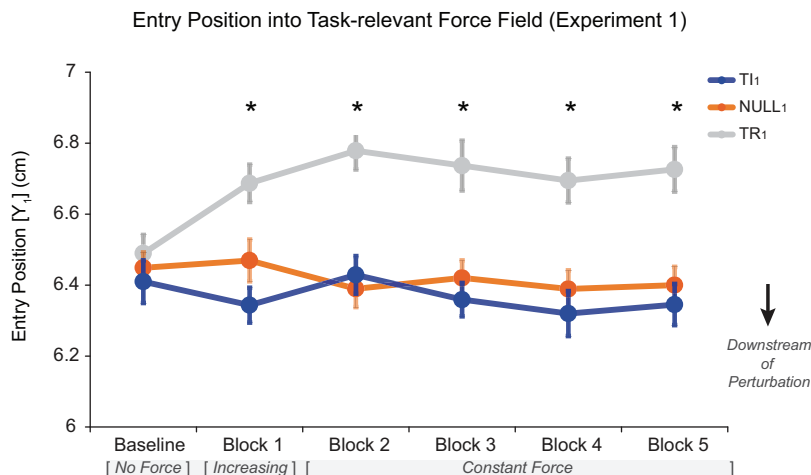


Fig. 7. Y_1 (cm) into the task-relevant force field. The TR₁ group had a significantly greater Y_1 to compensate for the task-relevant force field. Standard error bars are presented. * $P < 0.05$.

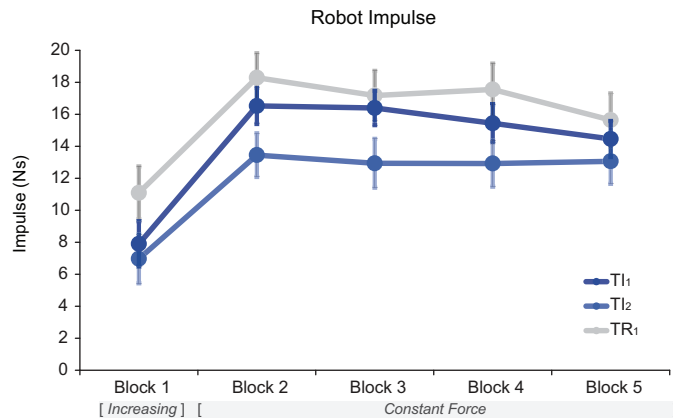


Fig. 8. Applied robot impulse (Ns) for the TI_1 , TI_2 , and TR_1 groups. All groups received similar amounts of applied forces as they passed through the different force fields. No significant differences were found between groups. Standard error bars are presented.

displays the overall IQR and histogram, collapsed across blocks, of entry position changes into the task-relevant force field for the TR_1 and $NULL_1$ groups. Table 2 lists each group's overall and individual block (1–5) IQRs.

For both experiments, the overall IQR was greater in the TI groups relative to their corresponding NULL group. This finding was statistically significant for *experiment 1* ($P < 0.001$). In *experiment 1*, the IQRs of the TI_1 group were greater than the $NULL_1$ group at *blocks 2 and 3* by 115.7% ($P = 0.11$) and 360.2% ($P < 0.01$), respectively. This increase in IQR at *block 2* also aligned with initial and significant changes of entry angle (θ_{CCW1}). In *experiment 2*, the TI_2 group IQRs were 40.9% ($P = 0.18$) and 135.9% ($P = 0.01$) larger than those of the NULL group at *blocks 3 and 4*. These IQR increases also aligned with initial and significant mean changes in entry angle (θ_{CW2}). Thus the TI groups may have explored different entry angles into the task-irrelevant force field while simultaneously changing their average behavior. However, after these initial and mean changes, the IQR differences decreased between the TI and NULL groups in later blocks. This may signify that, after an increase in exploration and behavior change, there is a subsequent decrease in exploration with continued training.

Conversely, for the entry position changes into the task-relevant force field, we found that the overall IQR of the TR_1 and $NULL_1$ group were not statistically different. Despite significant mean changes of entry position into the task-relevant force field, the IQRs of the TR_1 group were typically smaller and not statistically different than those of the $NULL_1$ group. Furthermore, the TR_1 group's IQRs from *blocks 2–5* were qualitatively less than that of *block 1*. This may indicate that the TR_1 participants did not explore, despite significantly changing their behavior in response to the task-relevant force field.

DISCUSSION

The main finding of this study is that participants responded to task-irrelevant, positional forces by changing their movement plan. Using a reaching task with loose spatial and temporal constraints, we found that participants changed their kinematics before being pushed by a force that had no bearing on task success. Interestingly, ~90% of the participants reached downstream of the upcoming force. This behavior

change may have been caused by active exploration. Taken together, our experimental results are not consistent with current models and theories on human movement.

We found that the nervous system responded to both task-irrelevant and -relevant forces. As expected, aspects of the task seemingly followed the MIP. When introduced to a task-relevant force, TR_1 participants significantly changed their movement path (Y_1) to complete the task successfully. Furthermore, similar to

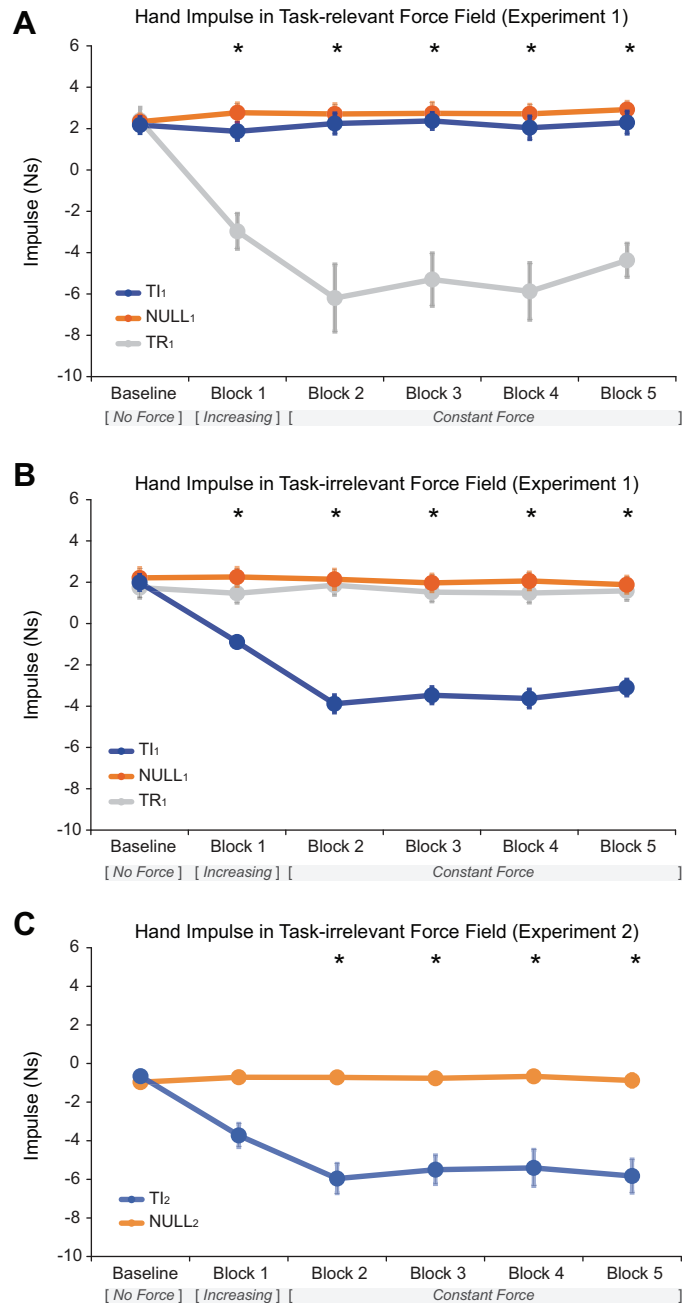


Fig. 9. Hand impulses (Ns) projected onto the applied robot force vectors for the different groups as they passed through the task-relevant (A), task-irrelevant counterclockwise (B) and task-irrelevant clockwise force fields (C). Hand impulse was also calculated during baseline and for each group whenever the applied robot force magnitude was zero (e.g., task-relevant group passing through the task-irrelevant force field), allowing for hand impulse comparisons between all groups. More negative values represent a greater resistance to the applied robot forces. Standard error bars are presented. * $P < 0.05$.

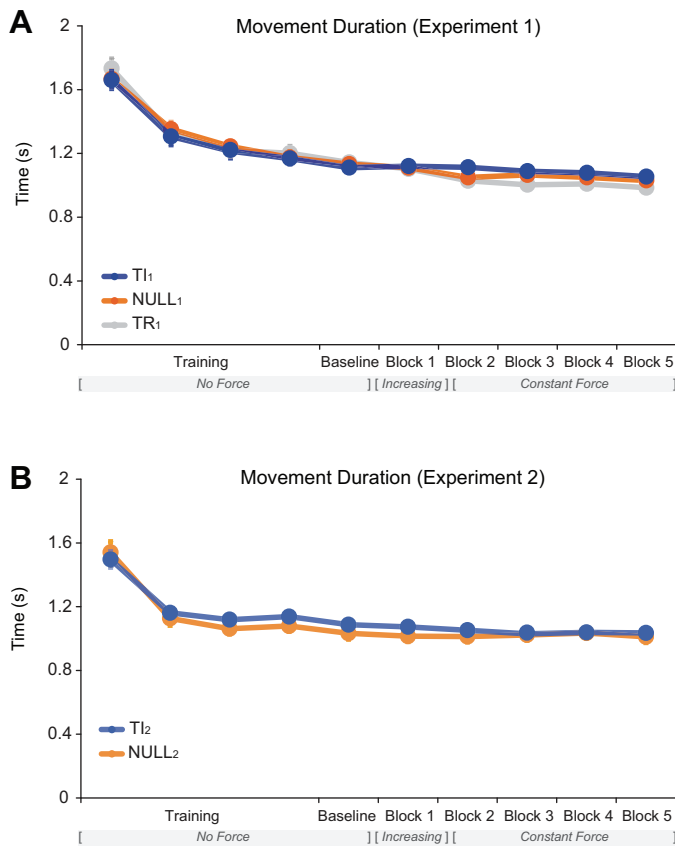


Fig. 10. Movement duration (s) for *experiment 1* (A) and *experiment 2* (B). No significant differences were found between groups. After training, participants had consistent movement times from baseline to *block 5*. Standard error bars are presented.

other experiments (Nashed et al. 2012, 2014), participants allowed the task-irrelevant force to push them into a different region of the large, final target (as shown by α_1 and α_2).

However, contradicting the MIP and UCMH, we also found that participants significantly changed their movement plan (θ_{CCW1} and θ_{CW2}) before being pushed by a task-irrelevant perturbation. Again, there was no need to respond before or after being perturbed by this force, as it had no bearing on task success and simply pushed the participants into a different area of the final target. Interestingly, in both experiments, almost 90% of these entry angle changes, relative to baseline, were made to an entry angle position that was downstream of the upcoming force. These entry angle changes became significant for *experiments 1* and *2* during the second and third blocks, respectively. Slight differences between *experiments 1* and *2* may have occurred because they, respectively, required more shoulder and elbow exertion near the end of the reaches when

entering the task-irrelevant force fields. Furthermore, relative to baseline, there was a larger spread of movement changes for both groups experiencing a task-irrelevant force field. This larger spread of movement changes was most prominent in both experiments during the initial two blocks, where we found significant mean changes in entry angle, which may suggest an active exploration strategy during these early changes in average behavior. Conversely, the TR group, which also had a significant movement plan change, had a relatively constant spread of movement changes and likely did not engage in exploration. It has been previously suggested that the nervous system explores more in the task-irrelevant dimensions (Latash et al. 2007, 2012; van Beers et al. 2013). However, we found not only exploration but also an unnecessary change in average behavior.

It has been previously found during reaching to large targets that humans will produce an unnecessary response to compensate for task-irrelevant perturbations. These have been found in experiments involving visuomotor rotation (Schaefer et al. 2012) and where the hand is constrained along a predefined path (Orban de Xivry 2013). In both of these studies, the unnecessary response acted in the opposite direction of the task-irrelevant perturbation. Schaefer and colleagues (2012) proposed these changes were a consequence of the nervous system minimizing sensory prediction errors. While this hypothesis does well to explain the observed task-irrelevant behavior changes in their visuomotor task, it does not match our data. In our experiments, participants did not appear to counteract the task-irrelevant force field as they did for the task-relevant force field. Otherwise, they would have moved “upstream” of the task-irrelevant force field. Rather, a large proportion of our participants moved “downstream” of the upcoming task-irrelevant force field. However, similar to Orban de Xivry (2012), it is possible that participants may have actively resisted the task-irrelevant force field. We did find a greater hand impulse magnitude, but, as previously mentioned, for our task it is difficult to dissociate how much of this impulse change was due to inertia or participants actively resisting the task-irrelevant force field. A major experimental difference between our experiments and that of Schaefer et al. (2012) and Orban de Xivry (2013) is our task was both spatially and temporally unconstrained, and involved positional force fields.

There were no correlations between baseline and entry angle changes for the TI groups. This confirms that movement plan changes were not caused by participants choosing to make more horizontal movements, participants deciding to reach to the target center, or biomechanical factors, such as joint comfort, interaction torques, and control effort. It was important to control for these factors, to be assured the applied task-irrelevant forces were

Table 1. Correlation between baseline and change in task-irrelevant force field entry angle

	Block 1		Block 2		Block 3		Block 4		Block 5	
	R^2	P	R^2	P	R^2	P	R^2	P	R^2	P
TI ₁	0.01	0.66	0.03	0.45	0.00	0.91	0.04	0.35	0.03	0.41
TI ₂	0.00	0.94	0.05	0.36	0.07	0.29	0.00	0.81	0.04	0.40

TI₁ and TI₂, task-irrelevant *experiments 1* and *2* groups, respectively. A lack of correlation and significance, as reported here, indicates behavior changes were not due to participants choosing to reach to certain areas of the target or make more horizontal movements upon experiencing the task-irrelevant force field. Furthermore, this also controls for potential biomechanical factors.

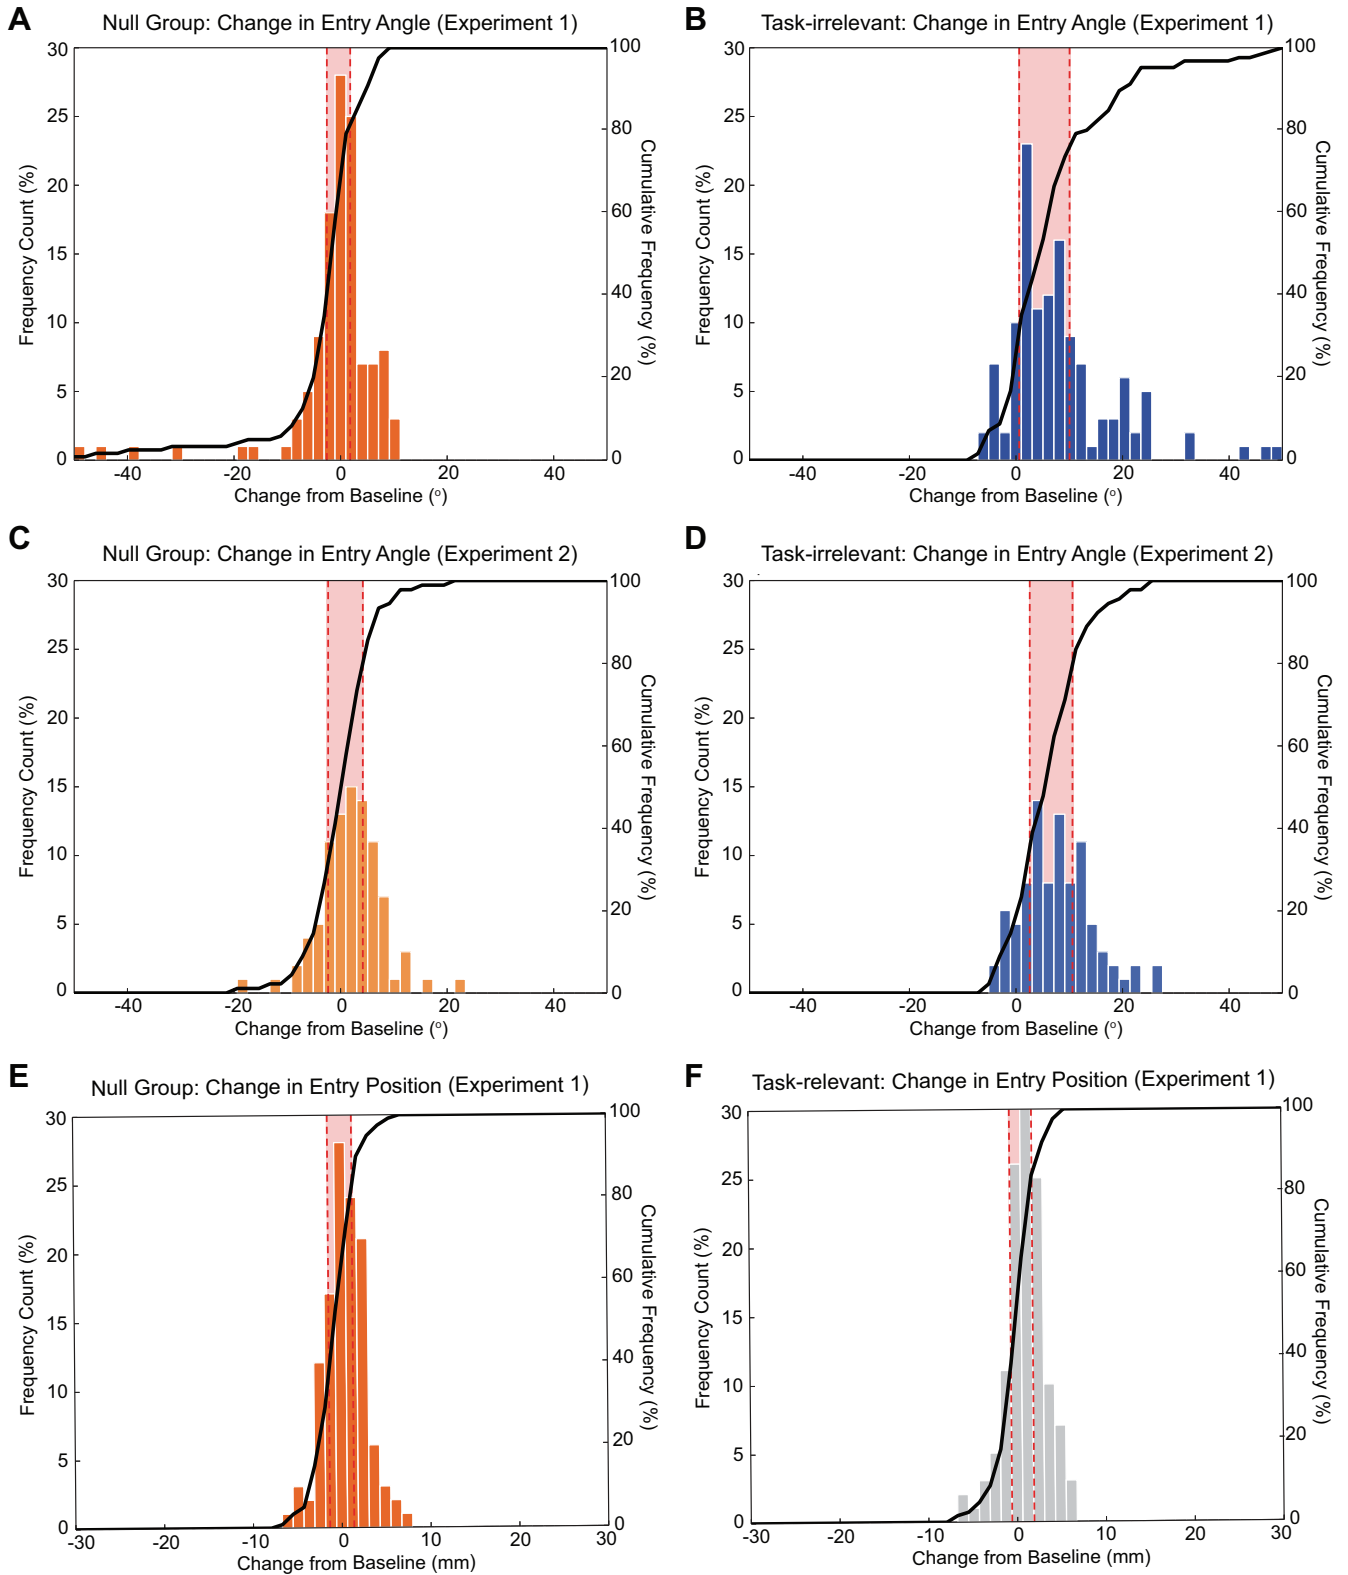


Fig. 11. Overall interquartile range (IQR: distance between the red dashed lines within the light red shaded area), histogram (bars correspond to left y-axis), and cumulative frequency (black line corresponds to right y-axis) of entry angle changes, relative to baseline, into the task-irrelevant force field for *experiment 1* participants in the NULL₁ (A) and TI₁ (B) force field groups, as well as *experiment 2* participants in the NULL₂ (C) and TI₂ (D) force field group. Subplots E and F represent the same measure for Y₁ changes, relative to baseline, into the task-relevant force field for *experiment 1* participants in the NULL₁ and TR₁ force field groups, respectively. IQR was used because it is robust to outliers, such as the few large entry angle changes seen in A. The evolution of IQR over blocks is listed in Table 2.

Table 2. IQR of entry angle and position into the task-irrelevant and task-relevant force fields

	Block 1	Block 2	Block 3	Block 4	Block 5	Overall
Task-irrelevant force field entry angle IQR, °						
NULL ₁	4.0	5.7	3.2	6.2	4.8	4.4
TI ₁	5.1	12.4	14.6	5.1	8.5	9.5
NULL ₂	5.6	5.0	6.7	4.6	11.0	6.5
TI ₂	4.5	5.4	9.4	10.9	9.6	8.0
Task-irrelevant force field position IQR, mm						
NULL ₁	2.4	2.6	1.7	2.2	2.8	2.7
TR ₁	2.8	2.2	2.4	1.2	2.1	2.5

NULL₁ and NULL₂, null groups for *experiments 1* and *2*, respectively; TR₁, task-relevant *experiment 1* group. Bolded numbers indicate a significantly different interquartile range (IQR) relative to the corresponding NULL group ($P < 0.05$). These large IQR differences correspond with initial and significant entry angle changes into the task-irrelevant force fields. This finding suggests that the TI groups were exploring the task-irrelevant force field during these early changes in behavior. Conversely, the TR₁ group had a relative consistent IQR that did not differ from the NULL₁ group. This suggests that they did not explore, despite having a mean change in entry position into the task-relevant force field.

indeed irrelevant. Furthermore, both experimental design and several other features of the data further assured us that these factors did not influence the findings. For *experiment 1*, it was possible that such biomechanical factors could have caused increases in force field entry angle. For example, in this experiment, reaching toward upper regions of the final target may have been more comfortable for participants. To control for this, we changed the direction of the applied task-irrelevant forces in *experiment 2*. Here, if increased entry angles were the result of biomechanical factors, we would expect to see similar, systematic increases in entry angle over similar regions of the workspace. As shown however, the TI₂ group decreased their entry angle. Furthermore, the TR₁ group changed their average entry position from *blocks 2–5* by only 3.3 mm. The TI₁ and TI₂ groups had larger average entry angle changes of 11.4° and 10.8°, which, respectively, correspond to 11.9 mm and 11.3 mm when converting from polar to linear coordinates. Thus the TI₁ and TI₂ groups had approximately a three and one-half fold greater change in entry position compared with the TR₁ group, despite the former not having to respond to the task-irrelevant force fields. This indicates it is very unlikely that movement plan changes were due to biomechanics. Finally, such factors cannot explain the larger spread of TI₁ and TI₂ movement plan changes when experiencing the task-irrelevant force fields. Therefore, given the statistical findings, experimental design, and supporting data, movement plan changes were not due to spurious biomechanical factors.

When examining the movement plan changes, it is important to consider the task from a participant's perspective. For example, after experiencing a force field, it would be unknown to a participant if another trajectory would have less applied force or better accomplish the goal of hitting the final target. That is, could they further descend some cost-function gradient? For the participant to determine this in a large workspace, they would likely have to actively explore. The term "active exploration" refers to exploration beyond what may naturally occur from sensorimotor noise (Izawa and Shadmehr 2011). In

other words, participants likely had to actively explore because sensorimotor noise alone would likely be insufficient to explore such a large workspace. Previous work has shown that sensorimotor noise accumulates in the task-irrelevant dimension(s), leading to greater variability (van Beers 2012). However, we observed additional variability along the task-irrelevant dimension beyond what was found during baseline. Thus it is unlikely that accumulated sensorimotor noise was the cause of the movement plan changes found in our task. Recently, it has been shown during reward and error-based tasks that greater, initial active exploration leads to faster rates of learning (Wu et al. 2014). These authors propose that early active exploration allows for more information to be gathered on the task, which can later be exploited to improve performance. In our task, active exploration was unnecessary, given the force field design. However, participants would not be aware of this and may have actively explored to see if there was another trajectory with less applied forces.

We found participants preferred to actively move downstream of the task-irrelevant force, despite most having enough workspace to explore in either direction. Here we offer some potential explanations. First, it is possible that this finding may have been due to force avoidance. Many common positional loads, like those produced by elastics and springs, dissipate as the point of application moves in the same direction of the acting force. Thus, based on previous experience, participants may have been actively exploring downstream for a trajectory with less applied positional forces.

Another possibility is that participants may have been influenced by the potential effects of what has been called "use-dependent learning." Diedrichsen and colleagues (2010) define this form of learning as changing movements to become more similar to the previous movement. In their experiment, participants reached in either a clockwise or counterclockwise viscous force field toward a large, rectangular target. They found that participants' movements shifted in the direction of the perturbing clockwise or counterclockwise viscous curl fields. That is, on any given trial where the movement was perturbed, on the following trial the movement plan would shift toward the previously perturbed movement.

There are a number of differences between the Diedrichsen et al. (2010) study and the current one. In the Diedrichsen et al. (2010) study, subjects were instructed to obey a relatively strict constraint on movement timing, unlike in the present study in which there was no constraint on movement timing or movement speed. In addition, the perturbation used in the Diedrichsen et al. (2010) study was a viscous curl field. These curl fields were not entirely task-irrelevant, as they would provide resistance with any deviation from a perfectly straight trajectory as participants moved toward the final target. Finally, subjects in the Diedrichsen et al. (2010) study tended to shift their movements over the course of the experiment such that they ended up reaching to the far edges of the movement target. In our task, if participants were only affected by this kind of use-dependent learning, one would expect that they might also continually shift their movement plan until they were reaching toward the edge of the final, curved rectangular target. Instead, we found that the TI group changed their movements in a way that asymptoted at a point far before the edge of the curved rectangular target. In addition, subjects in the present study showed an increase in movement variability shortly after being exposed to trials in which perturbations were present,

and this was followed by a later decrease in variability toward the end of the experiment. Taken together, the present results may be more consistent with an exploration-exploitation strategy than with an account solely based on changes due to use-dependent learning.

Although our data are more suggestive of active exploration, an interesting possibility is that force avoidance and or use-dependent learning may have initially guided the direction of the exploration downstream of the perturbing force. Nevertheless, our task represents a scenario where active exploration was unnecessary, but may have been used by participants in an attempt to find a better solution that did not exist. It is important to note that the current experiments and hypotheses were not focused on examining exploration, but were designed to test a fundamental aspect of the MIP. Future research would benefit from experiments specifically designed to examine how the nervous system actively explores its environment and the potential mechanisms that might promote searching in specific regions of the workspace.

Active exploration has implications on sensorimotor computational models. As mentioned, active exploration has only recently been applied to human sensorimotor learning (Wu et al. 2014). We found that participants actively explored beyond signal-dependent noise and consequently changed their mean trajectory, with the large range of movement plan changes supporting the former statement. As Wu and colleagues state, current models are formulated on the assumption that human movement variability originates only from sensorimotor signal-dependent noise (Harris and Wolpert 1998; Todorov and Jordan 2002; van Beers et al. 2013). Present formulations of optimal feedback controllers do not allow for exploration beyond the normally distributed, sensorimotor noise added to the equations (Todorov and Jordan 2002). Furthermore, UCM mathematics is based on the assumption of an unchanging, optimal trajectory (Scholz and Schoner 1999). Our data do not support these assumptions. To improve our knowledge and computational models of biological motion, future research should further investigate the exploration-exploitation relationship. Insights into the amount of active exploration needed to optimally exploit an environment can likely be gained from foraging (Biesmeijer and de Vries 2001; Eliassen et al. 2007; Krebs et al. 1978) and reinforcement learning (Barto 1998; Poupart et al. 2006; Thrun 1992) literature.

Using a redundant workspace with loose spatial and temporal constraints, we found that certain features of reaching did not follow the MIP and UCMH. Specifically, participants changed their movement plan prior to being perturbed by a task-irrelevant position force. We suggest this behavior change may have been a result of active exploration. The exploration-exploitation relationship has only recently been examined in motor learning (Wu et al. 2014) and deserves further attention, given its potential relevance to our understanding and ability to model biological motion.

APPENDIX

For all groups, robot forces (F) in the x - and y -directions are described by

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = m \left\{ \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} - c \right\} \begin{bmatrix} a \\ b \end{bmatrix}$$

where m (39) is the force magnitude coefficient, c (~ 0.0044) is a constant, z (range = -3 to $+3$) is some position coordinate converted to a z -score, and $[a \ b]^T$ is a vector: $\text{NULL}_1 = \text{NULL}_2 = [0 \ 0]^T$; $\text{TR}_1 = [0 \ -1]^T$; $\text{TI}_1 = [-\sin(\theta) \ \cos(\theta)]^T$, $\text{TI}_2 = [\sin(\theta) \ -\cos(\theta)]^T$, where θ is the angle between current position and the x -axis, as shown in Fig. 2. This resulted in a peak applied robot force of ~ 15.4 N. An example of position z -scores, for the TR_1 and $\text{TI}_{1\&2}$ groups, x -coordinates (e.g., -0.0525 m, -0.0375 m, and -0.0225 m) and distance from the second via point (e.g., 0.06 m, 0.075 m, and 0.09 m) were respectively converted to corresponding z -scores (e.g., -3.0 , 0.0 , and $+3.0$).

GRANTS

This research was supported by research grants from the Canadian Institutes of Health Research and the Natural Sciences and Engineering Research Council of Canada.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

AUTHOR CONTRIBUTIONS

Author contributions: J.G.A.C., H.R.M., and P.L.G. conception and design of research; J.G.A.C. performed experiments; J.G.A.C. analyzed data; J.G.A.C., H.R.M., and P.L.G. interpreted results of experiments; J.G.A.C. and H.R.M. prepared figures; J.G.A.C. drafted manuscript; J.G.A.C., H.R.M., and P.L.G. edited and revised manuscript; J.G.A.C., H.R.M., and P.L.G. approved final version of manuscript.

REFERENCES

- Barto AG.** *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT, 1998.
- Bernstein NA.** *The Coordination and Regulation of Movements*. Oxford, UK: Pergamon, 1967.
- Biesmeijer JC, de Vries H.** Exploration and exploitation of food sources by social insect colonies: a revision of the scout-recruit concept. *Behav Ecol Sociobiol* 49: 89–99, 2001.
- Cluff T, Manos A, Lee TD, Balasubramaniam R.** Multijoint error compensation mediates unstable object control. *J Neurophysiol* 108: 1167–1175, 2012.
- Cluff T, Scott S.** *Feedback Responses Can Resemble Trajectory Control When Required by the Behavioral Constraints of the Ongoing Task*. Amsterdam: NCM, 2014.
- de Rugy A, Loeb GE, Carroll TJ.** Muscle coordination is habitual rather than optimal. *J Neurosci* 32: 7384–7391, 2012.
- Diedrichsen J, White O, Newman D, Lally N.** Use-dependent and error-based learning of motor behaviors. *J Neurosci* 30: 5159–5166, 2010.
- Domkin D, Laczko J, Djupsjöbacka M, Jaric S, Latash ML.** Joint angle variability in 3D bimanual pointing: uncontrolled manifold analysis. *Exp Brain Res* 163: 44–57, 2005.
- Domkin D, Laczko J, Jaric S, Johansson H, Latash ML.** Structure of joint variability in bimanual pointing tasks. *Exp Brain Res* 143: 11–23, 2002.
- Eliassen S, Jørgensen C, Mangel M, Giske J.** Exploration or exploitation: life expectancy changes the value of learning in foraging strategies. *Oikos* 116: 513–523, 2007.
- Faisal AA, Selen LP, Wolpert DM.** Noise in the nervous system. *Nat Rev Neurosci* 9: 292–303, 2008.
- Flash T, Hogan N.** The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci* 5: 1688–1703, 1985.
- Franklin DW, Liaw G, Milner TE, Osu R, Burdet E, Kawato M.** Endpoint stiffness of the arm is directionally tuned to instability in the environment. *J Neurosci* 27: 7705–7716, 2007.
- Harris CM, Wolpert DM.** Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.
- Izawa J, Shadmehr R.** Learning from sensory and reward prediction errors during motor adaptation. *PLoS Comput Biol* 7: e1002012, 2011.

- Jones KE, Hamilton AF, Wolpert DM.** Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88: 1533–1544, 2002.
- Krebs JR, Kaelinik A, Taylor P.** Test of optimal sampling by foraging great tits. *Nature* 275: 27–31, 1978.
- Latash ML.** The bliss (not the problem) of motor abundance (not redundancy). *Exp Brain Res* 217: 1–5, 2012.
- Latash ML, Scholz JP, Schöner G.** Toward a new theory of motor synergies. *Motor Control* 11: 276308, 2007.
- Liu D, Todorov E.** Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J Neurosci* 27: 9354–9368, 2007.
- Nashed JY, Crevecoeur F, Scott SH.** Influence of the behavioral goal and environmental obstacles on rapid feedback responses. *J Neurophysiol* 108: 999–1009, 2012.
- Nashed JY, Crevecoeur F, Scott SH.** Rapid online selection between multiple motor plans. *J Neurosci* 34: 1769–1780, 2014.
- Orban de Xivry JJ.** Trial-to-trial reoptimization of motor behavior due to changes in task demands is limited. *PLoS One* 8: e66013, 2013.
- Poupart P, Vlassis N, Hoey J, Regan K.** An analytic solution to discrete Bayesian reinforcement learning. In: *Proceedings of the 23rd International Conference on Machine Learning*. New York: ACM, 2006, p. 697–704.
- Reimann C, Filzmoser P, Garrett R, Dutter R.** *Statistical Data Analysis Explained: Applied Environmental Statistics with R*. Chichester, UK: Wiley, 2011.
- Rosenbaum DA.** *Human Motor Control*. San Diego, CA: Academic, 1991.
- Schaefer SY, Shelly IL, Thoroughman KA.** Beside the point: motor adaptation without feedback-based error correction in task irrelevant conditions. *J Neurophysiol* 107: 1247–1256, 2012.
- Scholz JP, Schöner G.** The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 126: 289–306, 1999.
- Shadmehr R, Mussa-Ivaldi FA.** Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14: 3208–3224, 1994.
- Thrun SB.** *Efficient Exploration in Reinforcement Learning*. Pittsburgh, PA: Carnegie Mellon University, 1992.
- Todorov E.** Optimality principles in sensorimotor control. *Nat Neurosci* 7: 907–915, 2004.
- Todorov E, Jordan MI.** Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002.
- Todorov E, Jordan MI.** Smoothness maximization along a predefined path accurately predicts the speed profiles of complex arm movements. *J Neurophysiol* 80: 696–714, 1998.
- Valero-Cuevas FJ, Venkadesan M, Todorov E.** Structured variability of muscle activations supports the minimal intervention principle of motor control. *J Neurophysiol* 102: 59–68, 2009.
- van Beers RJ.** How does our motor system determine its learning rate? *PLoS One* 7: e49373, 2012.
- van Beers RJ, Brenner E, Smeets JBJ.** Random walk of motor planning in task irrelevant dimensions. *J Neurophysiol* 109: 969–977, 2013.
- Wu HG, Miyamoto YR, Castro LNG, Ölveczky BP, Smith MA.** Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat Neurosci* 17: 312–321, 2014.
- Yang JF, Scholz JP.** Learning a throwing task is associated with differential changes in the use of motor abundance. *Exp Brain Res* 163: 137–158, 2005.
- Yang JF, Scholz JP, Latash ML.** The role of kinematic redundancy in adaptation of reaching. *Exp Brain Res* 176: 54–69, 2007.

