Holding the arm still through subcortical mathematical integration of cortical commands

Scott T. Albert[1], Alkis M. Hadjiosif[2], Jihoon Jang[1], John W. Krakauer[2], and Reza Shadmehr[1]

1. Department of Biomedical Engineering
2. Department of Neurology
Johns Hopkins School of Medicine
Baltimore, Maryland 21205 USA

Correspondence: Scott Albert or Reza Shadmehr, 416 Traylor Building, Johns Hopkins School of Medicine, 720 Rutland Ave., Baltimore, MD 21205, USA. Email: salbert8@jhmi.edu or shadmehr@jhu.edu. Phone: 410-614-3424.

Acknowledgements: This work was supported by grants from the National Institutes of Health (1-R01-NS078311, NS095706), the National Science Foundation (1723967), and the Sheikh Khalifa Stroke Institute. We would like to thank Kahori Kita for help with performing and scoring the assessments for the stroke patients, and Jennifer Keller for help performing some of these assessments and allowing us to borrow dynamometry equipment. We would also like to thank all study participants, especially stroke survivors, for their time and contribution to this work.

A period of holding still follows every movement. It has been assumed that for the arm, moving and holding are functionally independent: movement is via an adaptive, feedback-dependent controller that generates commands to transport the arm, while holding is via setting of reflexes that produce a postural field at movement endpoint. This assumption predicts that commands that move the arm should not affect the postural field at movement termination. Surprisingly, we found that as the reach commands changed, so did the ensuing postural field. The postural commands depended on mathematical integration of the reach commands. Following damage to the corticospinal tract, despite severe reach impairments, the holding system faithfully integrated the imperfect reach commands on a trial-by-trial basis. Together, these findings suggest that holding the arm still is accomplished via a separate, likely subcortical structure that acts as a mathematical integrator of the commands generated by the cortical reach controller.

Posture accompanies [movement] "like a shadow".
Sir Charles Sherrington

During a reaching movement, motor commands not only transport the arm, but also instantiate a feedback controller that continuously stabilizes the arm's trajectory against potential perturbations [1-6]. When the reach ends, this stability persists during the ensuing period of stillness; arm position is maintained through the balanced activity of muscles that create a postural field with forces that converge to zero at the endpoint [7,8]. While the transport phase of reaching depends on sensory-to-motor pathways in the motor cortex [9,10], we know little about the postural controller that maintains arm position after the reach ends. Current models assume that control of reaching and holding are independent of each other [11-13], possibly mediated through distinct neurons in the motor cortex [14].

Notably, this independence is not present in control of the eyes. Like the arm, when a saccade ends, the eyes are held still via balanced activity of extraocular muscles [15]. However, unlike the arm, the neural system that holds the eyes is well understood; the "move" and "hold" controllers are anatomically distinct but causally linked [16]. The hold circuitry resides in a brainstem structure that mathematically integrates the motor commands that moved the eyes during the preceding saccade [17,18]. As a result, the move and hold controllers for the eyes are not independent, but serially linked; changes in the commands that moved the eyes produce distinct changes in the commands that hold the eyes. Could this design principle be shared with control of the arm?

To answer this question, we analyzed the natural variability in the reaching and holding commands of a large number of people (n=149) as they made simple point-to-point movements (Fig. 1A). On most trials the participants reached freely (Fig. 1B, null field trial) to the target while grasping a robotic arm. On other trials, the robot created a virtual channel along a straight line to the target (Fig. 1C, channel trial), allowing us to measure the forces that participants produced during both the reach and the hold periods. Did the natural variability in the generation of reach commands influence the commands during holding?

We separated movements into trials where participants happened to produce no lateral reaching forces (Fig. 1D, top inset, middle trace), lateral reaching forces to the left (Fig. 1D, top inset, bottom two traces), or lateral reaching forces to the right (Fig. 1D, top inset, top two traces). Notably, because these movements were performed in a channel, hand trajectory was virtually identical across

these trials, as was the final hand position (Fig. 1D, bottom). Despite this, changes in reach commands were associated with changes in hold commands; participants applied hold forces that were highly correlated to the preceding reach forces.

The hold forces could arise if at the end of the reach, the balance of muscle activities in the arm created a postural field that had a null point slightly different than the current hand position. To explore this idea, we set out to measure the postural field created by the hold system. On some trials, after the reach ended, the robot slowly displaced the arm while participants (n=27) were engaged with a working memory task (Fig. 1E, postural probe, 2-back task). In response to the displacement, participants produced restoring forces (Fig. 1F, top) that converged to zero at a null position near the endpoint of the preceding reach (Fig. 1G, top). Did the null position of this postural field depend on the forces that transported the arm during the preceding reach?

To answer this question, we imposed a velocity-dependent curl field that applied lateral forces to the arm during movement, but not during holding still. Now in order to reach to the target, participants needed to generate lateral forces, similar to, but larger than, those that naturally arose on null field trials (Fig. 1D). Indeed, participants modified their reaching forces over the series of many trials (Fig. 1H). Strikingly, even though the reach ended at the same target location as before, there was an unexpected change in the postural field; the null position of the postural field was no longer aligned with the target (Fig. 1F, bottom), shifting laterally by approximately 1 cm (Fig. 1G; paired t-test, $p<10^{-4}$), while the orientation (paired t-test, p=0.84) and stiffness (paired t-test, p=0.62) of the field remained unchanged. This shift in the null point altered the process of holding still; because the postural null position was no longer aligned with hand position at reach end, the arm now produced a static holding force (Fig. 1H, hold), like that observed in the null field trials (Fig. 1D). This holding force persisted as the hand was held at the target for the entire 6 second inter-trial interval (Fig. S1).

We reasoned that if the holding force was truly coupled to the null position of the postural field, we could eliminate it by displacing the hand toward the postural null position. Indeed, the holding force gradually approached zero as the arm was displaced, and then switched direction and grew larger as the hand was moved beyond the postural null position (Fig. 1I). Critically, holding force at reach end was linearly related to the location of the postural null position (Fig. 1J). This implied that by measuring the holding force at reach end, we could infer the postural null position; the larger the holding force, the farther the null position of the postural field.

How did the location of the postural null point depend on the preceding forces that transported the arm to the target? To answer this question, we first considered normal reaching movements of a large group of participants (n=220) in the channel trials within the null field period (a subset of these trials is shown in Fig. 1D). To mathematically construct the relationship between the postural null position and the reach forces, we hypothesized that like the oculomotor system, the change in the postural null position depended on the change in the time-integral of the reach force: the more force that accumulated during movement, the farther the postural null position. To test this hypothesis, we integrated, with respect to time, the reaching forces exerted laterally against the channel walls, and compared this force-time integral to the holding forces produced after the reach ended. We found that the holding forces were precisely predicted by the time-integral of the moving forces that had brought the hand to the target, both in an analysis of trials pooled across all participants (Fig. 2A, left), and from single trial analysis within each participant (Fig. 2A, right).

To extend the range of reaching forces beyond those that naturally arose during the execution of reaching, we again employed a force field paradigm. We instructed another set of participants (n=32) to reach between two locations and then gradually imposed various force fields during their reach (Fig. 2B), resulting in systematic changes in reach forces. We found that as the reach forces changed, so did the hold forces (Fig. 2C), with holding forces closely yoked to the time-integral of the reach forces (Fig. 2D). During reaching, the force-time integral reached a level three times larger than that which naturally arose in null field trials. Despite this, across both null field and force field trials, more than 95% of the variance in holding forces was predicted by the integral of the forces that brought the hand to the target (Fig. 2E).

To thoroughly test the robustness of this relationship, we performed 14 additional experiments involving a total of n=175 subjects. The time-integration hypothesis accounted for holding forces when reaches started at different locations, but ended at the same target (Figs. S2A-D), when reaches and targets were located in various parts of the workspace (Figs. S3), when reaches were made at angles oblique to the midline of the body (Fig. S4), for reaches of various speeds and durations (Fig. S5), for forces exerted during trials where learning was generalized from other locations in the workspace (Figs. S2A,E-G), and finally in cases where on a trial-by-trial basis, the moving forces alternated for reaches that started and ended at the constant locations (Fig. S6). Remarkably, across all of these experiments a single function accounted for 96% of the variance in the data (Fig. 2E); the postural controller produced holding forces at the endpoint that were proportional to the time-integral of the reach force that had brought the arm to that endpoint.

There are alternatives to our mathematical integration hypothesis; for example, holding forces may be a trivial continuation of the forces applied near the end of the preceding reach, not an integration of the entire history of the reaching forces. To test this possibility, we designed a force field that produced a large perturbation near the end of the reach, but had a time-integral for the entire movement that was equal to zero (Fig. 3B). The experiment began with a force perturbation during the second half of the reach (Fig. 3A, Phase 1, unidirectional). As expected, participants (n=14) produced holding forces that increased with the integral of the moving forces (Fig. 3C, unidirectional) and remained stable across hundreds of additional trials (Fig. S5, control participants, n=11). Next, we gradually added an opposing force field during the first half of the reach (Fig. 3A, Phase 2, bidirectional). In this way, we guided the participants to produce reach forces that had a time-integral of approximately zero. Remarkably, as the integral of the reach forces approached zero, the holding force gradually vanished (Figs. 3C and 3D, bidirectional field). A single linear function correctly predicted holding forces during both the unidirectional and bidirectional phases of the experiment (Fig. 3E). Together, these observations confirmed that the holding forces were consistent with mathematical integration of the entire time-history of the preceding reach, and not simply the forces exerted before movement termination.

If holding forces are causally determined by the time-integration of reaching forces, the postural controller should be situated downstream of the reach controller (Fig. 3F, left). In this arrangement, the moving forces change due to random fluctuations, or due to learning. In both cases, this change produces a change in the holding forces. Alternatively, moving and holding forces may be non-causally coupled, driven instead by a common input, the reach perturbations (Fig. 3F, right column). Critically, this alternate hypothesis cannot account for the coupling between moving and holding forces that we

observed in the null field (Figs. 1D and 2A). Still, to further test between these alternatives, we considered random fluctuations in the reach forces during periods when the perturbation was kept constant (Fig. 3G). Despite invariance in the perturbation, we again observed a strong relationship between the time-integral of reach forces in a given trial and the ensuing hold forces on the same trial (Figs. 3H, 3I). The co-variations in the reach and hold forces remained after we regressed out the effects of a changing reach perturbation across all trials (Fig. S7). Collectively, analysis of these residuals ruled out the possibility that moving and holding correlations were due to a confounding variable (external perturbations). Rather, the temporal integration of moving forces was present during holding still, whether or not the reach had experienced a perturbation (compare Figs. 1D, 3H, S7).

Our results describe a paradox. On the one hand, the reach controller adapts readily to changing dynamics, allowing the arm to arrive at the target despite changes in the environment. On the other hand, if the postural controller always integrates the reach forces, then the null position of the postural field will no longer be at the desired endpoint. As a result, postural stability would be compromised in the face of an adapting reach controller. A possible solution is that like the reach controller, the postural controller is also adaptive [19,20]. That is, following a change to the reach controller, the postural system should learn to reunite the null point of the postural field with the desired endpoint of the reach. A key feature of our design was that at reach end, we held the hand in an endpoint clamp (Fig. 1B). Without this endpoint clamp, non-zero holding forces would move the hand off the target. Under natural circumstances, these holding-related errors would provide a teaching signal for the postural controller, resulting in its adaptation and subsequent elimination of the holding forces.

To test this idea, we asked whether positional errors made while ending the reaching movements affected the stability of ensuing holding forces. In line with our hypothesis, we found that individuals that made errors in the direction of the holding forces while attempting to slow and stop the hand at the end of the reach, ultimately produced smaller holding forces (Fig. S8A-F). We next performed a control experiment in which reaching movements ended without an endpoint clamp (Fig. S8G). In this experiment, positional errors committed during reach termination were even larger (Fig. S8H and S8I, compare red and black), leading holding forces to gradually decrease over trials (Fig. S8J and S8K) in line with earlier reports [21]. Therefore, the endpoint clamp used in our primary experiments was critical for preventing adaption in the postural controller, thereby unmasking the link between the reach and the hold systems.

Descending commands in the corticospinal tract (CST) are critical for execution of voluntary movements. Are postural signals also present in these same descending pathways or does a separate, downstream postural controller receive and then integrate the reach commands? If reach and postural commands are conveyed in the CST, then damage to the CST should disrupt both the generation of forces during reaching, and its integration during holding still. However, if the postural controller is separate and downstream to the CST then damage to descending cortical tracts should selectively impair reaching, but spare the process of mathematical integration of those imperfect reach commands.

We recruited stroke patients (n=13) who had suffered lesions affecting the CST pathway from the cortex through the internal capsule. The patients exhibited profound impairments in motor control, manifesting clearly in an extreme difficulty with extension of the arm during unsupported reaching (Fig. 4A, patient S015, inability to extend at the elbow). To partially alleviate limitations in the arm's range of motion, we supported the weight of the arm with a frictionless air-sled. With arm support, patients

were better able to extend the arm at the elbow, but movements of the paretic arm continued to be plagued with erratic trajectories (Fig. 4B, patient S015; Fig. 4C, all patients), requiring more time for completion (Fig. 4D), and arriving at endpoints further away from the target location (Fig. 4E).

To measure the forces associated with the reach trajectories, the robot occasionally produced a channel (Fig. 1C). It was clear that the impairment in the control of reaching coincided with a marked increase in the variability of reach forces (Fig. 4F, example traces exhibiting different reach forces for the same movement trajectory). In fact, the time-integral of the reach forces in the paretic arm was nearly 3 times more variable than healthy controls (Fig. 4G, metric based on standard deviation). Despite this, impaired reach trajectories still terminated in sustained holding forces (Fig. 4F, dashed lines). Remarkably, the within-trial, within-subject coupling between time-integral of reach forces and hold forces was preserved in the paretic arm of the patients (Figs. 4H and 4I), even in their most impaired movements.

We next used adaptation to systematically manipulate reach forces of the patients. Because force field adaptation is largely a cerebellar dependent process [22], despite damage to the CST, the patients learned to alter their reach forces (Fig. 4J). Strikingly, we continued to observe holding forces in the paretic arm that were proportional to the time-integral of reach forces (Fig. 4K); the relationship between hold forces and the time-integral of reach forces was identical across the paretic arm, non-paretic arm, and dominant arm of age-matched controls (Fig. 4L).

These results suggest that whereas damage to the descending cortical tracts severely affected the reach controller, resulting in erratic reach forces, it spared the integration performed by the postural controller. It appears that in humans, the postural field for the arm is generated through integration of the preceding reach commands, likely via a subcortical structure.

Our findings suggest that a design principle may be shared between the oculomotor [23] and reach systems of the brain. Like the oculomotor system, control of reaching and holding may be performed by two distinct neural circuits (Fig. 4M). The reach controller is located in the cortex, and is impaired by damage to the CST. Perturbation of reaching causes adaptation of the reach controller, likely through a cerebellar error-based learning circuit [24]. The reaching forces are integrated by a subcortical neural system, resulting in a postural field with a null position at the reach endpoint. The postural controller is adaptive, itself learning from endpoint errors. Intriguingly, a similar process occurs in the oculomotor system; visual slip during gaze holding engages an error-based learning process that modulates the hold commands produced by the brainstem neural integrator following conclusion of a saccade [25].

Our model (Fig. 4M) provides a potential solution to a number of puzzles, including why decoding of neural activity from M1 can provide robust prediction of reach trajectories, but not arm posture during periods of holding still [26,27], and why reach adaptation accompanies paradoxical illusions in perception of arm position [28,29]. Future work will need to determine the relationship between muscle activation patterns that displace the arm to an endpoint, and activation patterns that are necessary to hold the arm at that endpoint. Experiments need to test whether like the oculomotor system there are distinct regions within the cerebellum that are responsible for adaptation of the reach and hold circuits. It remains to be asked whether postural control is solely influenced by a distinct subcortical neural integrator that accumulates descending commands, or if there are additional independent postural controllers; after all, while neural integration was normal despite CST lesions,

stroke patients still exhibited postural abnormalities (Fig. 4A). These abnormalities could arise from recurrent connections from the integrator back to the cortex, or perhaps more provocatively, from chronic but normal integration of abnormal reach commands.

Reference List

1. Burdet E., Osu R., Franklin D.W., Milner T.E., and Kawato M. (2001). The central nervous system stabilizes unstable dynamics by learning optimal impedance**.** Nature *414***:** 446-449.

2. Todorov E. and Jordan M.I. (2002). Optimal feedback control as a theory of motor coordination**.** Nat. Neurosci. *5***:** 1226-1235.

3. Franklin D.W., Osu R., Burdet E., Kawato M., and Milner T.E. (2003). Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model**.** J. Neurophysiol. *90***:** 3270-3282.

4. Franklin D.W., Liaw G., Milner T.E., Osu R., Burdet E., and Kawato M. (2007). Endpoint stiffness of the arm is directionally tuned to instability in the environment**.** J. Neurosci. *27***:** 7705-7716.

5. Selen L.P., Franklin D.W., and Wolpert D.M. (2009). Impedance control reduces instability that arises from motor noise**.** J. Neurosci. *29***:** 12606-12616.

6. Franklin S., Wolpert D.M., and Franklin D.W. (2012). Visuomotor feedback gains upregulate during the learning of novel dynamics**.** J. Neurophysiol. *108***:** 467-478.

7. Mussa-Ivaldi F.A., Hogan N., and Bizzi E. (1985). Neural, mechanical and geometric factors subserving arm posture in humans**.** J. Neurosci. *5***:** 2732-2743.

8. Shadmehr R., Mussa-Ivaldi F.A., and Bizzi E. (1993). Postural force fields of the human arm and their role in generating multi-joint movements**.** J. Neurosci. *13***:** 45-62.

9. Pruszynski J.A., Kurtzer I., Nashed J.Y., Omrani M., Brouwer B., and Scott S.H. (2011). Primary motor cortex underlies multi-joint integration for fast feedback control**.** Nature *478***:** 387-390.

10. Omrani M., Pruszynski J.A., Murnaghan C.D., and Scott S.H. (2014). Perturbation-evoked responses in primary motor cortex are modulated by behavioral context**.** J. Neurophysiol. *112***:** 2985-3000.

11. Lametti D.R., Houle G., and Ostry D.J. (2007). Control of movement variability and the regulation of limb impedance**.** J. Neurophysiol. *98***:** 3516-3524.

12. Yadav V. and Sainburg R.L. (2011). Motor lateralization is characterized by a serial hybrid control scheme**.** Neuroscience *196***:** 153-167.

13. Scheidt R.A. and Ghez C. (2007). Separate adaptive mechanisms for controlling trajectory and final position in reaching**.** J. Neurophysiol. *98***:** 3600-3613.

14. Kurtzer I., Herter T.M., and Scott S.H. (2005). Random change in cortical load representation suggests distinct control of posture and movement**.** Nat. Neurosci. *8***:** 498-504.

15. Cannon S.C. and Robinson D.A. (1987). Loss of the neural integrator of the oculomotor system from brain stem lesions in monkey. J. Neurophysiol. *57*: 1383-1409.

16. Shadmehr R. (2017). Distinct neural circuits for control of movement vs. holding still. J. Neurophysiol. *117*: 1431-1460.

17. Cohen B. and Komatsuzaki A. (1972). Eye movements induced by stimulation of the pontine reticular formation: evidence for integration in oculomotor pathways. Exp. Neurol. *36*: 101-117.

18. Miri A., Daie K., Arrenberg A.B., Baier H., Aksay E., and Tank D.W. (2011). Spatial gradients and multidimensional dynamics in a neural integrator circuit. Nat. Neurosci. *14*: 1150-1159.

19. Darainy M., Malfait N., Gribble P.L., Towhidkhah F., and Ostry D.J. (2004). Learning to control arm stiffness under static conditions. J. Neurophysiol. *92*: 3344-3350.

20. Darainy M., Towhidkhah F., and Ostry D.J. (2007). Control of hand impedance under static conditions and during reaching movement. J. Neurophysiol. *97*: 2676-2685.

21. Sing G.C., Joiner W.M., Nanayakkara T., Brayanov J.B., and Smith M.A. (2009). Primitives for motor adaptation reflect correlated neural tuning to position and velocity. Neuron *64*: 575-589.

22. Smith M.A. and Shadmehr R. (2005). Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. J. Neurophysiol. *93*: 2809-2821.

23. Crawford J.D., Cadera W., and Vilis T. (1991). Generation of torsional and vertical eye position signals by the interstitial nucleus of Cajal. Science *252*: 1551-1553.

24. Smith M.A. and Shadmehr R. (2005). Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. J. Neurophysiol. *93*: 2809-2821.

25. Optican L.M. and Miles F.A. (1985). Visually induced adaptive changes in primate saccadic oculomotor control signals. J. Neurophysiol. *54*: 940-958.

26. Sachs N.A., Ruiz-Torres R., Perreault E.J., and Miller L.E. (2016). Brain-state classification and a dual-state decoder dramatically improve the control of cursor movement through a brain-machine interface. J. Neural Eng *13*: 016009.

27. Shalit U., Zinger N., Joshua M., and Prut Y. (2012). Descending systems translate transient cortical commands into a sustained muscle activation signal. Cereb. Cortex *22*: 1904-1914.

28. Mattar A.A., Nasir S.M., Darainy M., and Ostry D.J. (2011). Sensory change following motor learning. Prog. Brain Res. *191*: 31-44.

29. Darainy M., Vahdat S., and Ostry D.J. (2013). Perceptual learning in sensorimotor adaptation. J. Neurophysiol. *110*: 2152-2162.

30. Mussa-Ivaldi F.A., Hogan N., and Bizzi E. (1985). Neural, mechanical and geometric factors subserving arm posture in humans. J. Neurosci. *5*: 2732-2743.

31. Hwang E.J. and Shadmehr R. (2005). Internal models of limb dynamics and the encoding of limb state**.** J. Neural Eng *2***:** S266-S278.

32. Pekny S.E., Izawa J., and Shadmehr R. (2015). Reward-dependent modulation of movement variability**.** J. Neurosci. *35***:** 4015-4024.

33. Heald J.B., Ingram J.N., Flanagan J.R., and Wolpert D.M. (2018). Multiple motor memories are learned to control different points on a tool**.** Nat. Hum. Behav. *2***:** 300-311.

Methods

Our results draw upon behavioral observations collected in 15 separate experiments. To organize our methodological description, here we first describe conceptual and experimental constructs that were common to the various experiments, and then conclude with a description of each individual experiment.

A total of n=223 healthy individuals participated in various reaching experiments. Healthy participants ranged from 18-61 years of age (mean ± SD, 25.2 ± 7.9) and included 128 males and 95 females. We also recruited a total of n=13 older adults that had suffered damage to the CST. The stroke patients ranged from 30-80 years of age (mean ± SD, 57.8 ± 14.8) and included 6 males and 7 females. For comparison, we recruited a cohort of healthy age-matched controls who ranged from 28-81 years of age (mean ± SD, 60.6 ± 16.3) and included 5 males and 5 females. There was no significant difference in age between the patient and older healthy control populations (2-sample t-test, p=0.68). A further description of lesion location and motor impairment in the stroke population is provided in our description of Experiment 15. All experiments were approved by the Institutional Review Board at Johns Hopkins School of Medicine.

Apparatus

In each experiment, participants held the handle of a planar robotic arm (Fig. 1A) and made point-to-point reaching movements between different target locations in the two-dimensional workspace. The forearm was obscured from view by an opaque horizontal screen. An overhead projector displayed a small white cursor (diameter = 3mm) on the screen that tracked the motion of the hand. Visual feedback of the cursor was provided continuously throughout the entirety of each testing period, except where otherwise noted. Throughout testing we recorded the position of the handle using a differential encoder with submillimeter precision, the forces applied at the handle by torque motors attached to the robotic arm, and the force produced on the handle by the subject using a 6-axis force transducer. Data were recorded at 200 Hz. Except where otherwise noted, kinematic timeseries were aligned to the onset of movement at the time point where hand velocity crossed a threshold of 1 cm/s.

Trial structure

At trial onset, a circular target (diameter = 1 cm) appeared in the workspace, coincident with a tone that cued subject movement. Participants then voluntarily reached from the starting position to the target. After stopping the hand within the target, a holding period of various durations (1.8 to 6.5 seconds) ensued where subjects were instructed to continue holding the handle within the target. After this holding period, a random inter-trial-interval sampled uniformly between 0.3 and 0.4 seconds elapsed prior to the start of the next trial.

At the end of each reach, coincident with the start of the holding period, movement timing feedback was provided. If the preceding reach was too slow, the target turned blue and a low tone was played. If the reach was too fast, the target turned red and a low tone was played. If the reach fell within the desired movement interval (450-550ms except where otherwise noted) the target “exploded” in rings of concentric circles, a pleasing tone was played, and a point was added to a score displayed in the upper-left-hand corner of the workspace. Participants were instructed to obtain as many points as

possible throughout the experimental session. In all experiments, each trial was a distinct point-to-point reaching movement, followed by a hold period.

The reaching period

Experiments began with a period of null field trials (no perturbations from the robot). After this period, participants were exposed to force field perturbations during reaching. For the majority of experiments, the force field was a velocity-dependent curl field in which the robot generated forces proportional and perpendicular to the velocity of the hand according to:

$$\begin{bmatrix} f_x \\ f_y \end{bmatrix} = f \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} \tag{1}$$

where v_x and v_y represent the x and y velocity of the hand, f_x and f_y represent the x and y force generated by the robot on the handle, and f represents the magnitude of the force field. We varied the force field magnitude across trials within a task. Participants were exposed to both clockwise (CW, $f > 0$) and counterclockwise (CCW, $f < 0$) curl fields. In most experiments, the perturbation was introduced in a gradual manner where the magnitude of the force field was increased very slowly from one trial to the next. In some experiments, participants were exposed to an abrupt force field where the force field magnitude immediately transitioned from 0 to the maximal field strength at the start of the perturbation block.

In Experiments 13 and 14, a position-dependent force field was applied to the hand, as opposed to the curl field described in Eq. (1). We constructed the position-dependent fields from two individual force fields that were applied separately during the first (called FF_1) and second (called FF_2) halves of the reaching movement, as a function of displacement. Each force field was programmed as a quadratic function of distance, and applied a force solely along the horizontal axis perpendicular to the vertical reaching movement. The combined output of FF_1 and FF_2 produced a force field with zero force at the start position, target position, and midpoint of the reach. FF_1 applied peak force at 25% of the displacement from start to target. FF_2 applied peak force at 75% of the displacement from start to target (see Figs. 3A&B for a visual depiction). Mathematically, the entire position-dependent field can be represented as:

$$f_x(y) = \begin{cases} 0, & y \le y_{start} \\ m_1(y - b_1)^2 + a_1, & y_{start} < y \le y_{midpoint} \\ m_2(y - b_2)^2 + a_2, & y_{midpoint} \le y \le y_{target} \\ 0, & y \ge y_{target} \end{cases} \tag{2}$$

where f_x represents the force applied by the robot along the axis perpendicular to the movement (the x-axis of the workspace), y represents the y-position of the hand, y_{start} represents the y-coordinate of the starting position, $y_{midpoint}$ represents the y-coordinate of the midpoint between the starting position and the target, and y_{target} represents the y-position of the target. The terms m_1, m_2, b_1, b_2, a_1, and a_2 are constants that provide the appropriate shape of the position-dependent field according to:

$$a_1 = F_1 \qquad b_1 = \frac{y_{start} + y_{midpoint}}{2} \qquad m_1 = \frac{4F_1}{\left(y_{start} - y_{midpoint}\right)^2}$$
$$a_2 = F_2 \qquad b_2 = \frac{y_{target} + y_{midpoint}}{2} \qquad m_2 = \frac{4F_2}{\left(y_{target} - y_{midpoint}\right)^2} \tag{3}$$

where F_1 and F_2 specify the maximum force produced during FF_1 and FF_2 respectively. The sign of each parameter also determines the direction of the applied force (if F_1 and F_2 are of opposite sign, the position-dependent force field applies a force to the left during one half of the reach and a force to the right during the other half of the reach). Trial-by-trial changes in F_1 and F_2 are discussed in our description of Experiments 13 and 14.

The holding period

At the end of each reaching movement, a holding period elapsed where the hand was maintained within the target location for some period of time. Participants were instructed to simply hold the handle of the robot and wait for the start of the next trial. In most experiments, each holding period lasted for a period of at least 1.8 seconds (a fixed 1.5 seconds with an additional inter-trial-interval of at least 0.3 seconds) and up to 6.5 seconds. In Experiments 1, 2, and 3, the holding period lasted only 0.5 seconds, but was following with a long cognitive task.

In contrast to typical force field paradigms, these long inter-trial holding periods were necessary in order for us to test whether the integration of moving forces predicted behavior during the maintenance of posture. To isolate this serial effect of moving on holding, we attempted to remove any disturbances that could lead to adaptation of the holding system. We hypothesized that spatial errors in the stabilization and maintenance of the final hand position might provide such a teaching signal. To minimize this potential source of holding error, we applied a two-dimensional endpoint clamp. This endpoint clamp prevented motion of the hand during the hold period, despite any forces the participant might have applied to the handle. The endpoint clamp was programmed as a “well” within the target location that attracted the hand in two dimensions, with stiff spring-like mechanics (stiffness = 4000 N/m, viscosity = 75 N-s/m). The endpoint clamp was applied when the hand entered the target location and the hand velocity fell below a threshold value of 3.5 cm/s, coincident with audiovisual feedback concerning the accuracy and speed of the terminated reaching movement.

Measuring the postural field

The first aim of this study was to determine how the holding system maintains a constant hand position at the end of a reaching movement. Traditionally, this holding controller is thought of as an impedance controller that is set to maintain an effector at a specific location [30]. We set out to characterize this impedance controller over small displacements of the hand, and understand how it might be impacted by the preceding reaching movement. Our question was, does the process of holding still at a given location depend on the motor commands that moved the arm to that location? To answer this question, we measured how hand position was controlled before and after adaptation to a force field. This force field provided a tool that we used to induce participants to systematically vary their reaching force between the same two points in space.

To interrogate the holding controller, in Experiments 1 and 3, we measured how the arm reacted to displacements in its final position. To do this, we used the robot to move the hand slowly and covertly in a random direction after the hand stopped within the target. As the hand was moved, visual feedback of hand position was prevented: the display cursor was frozen at the holding location throughout the robotic intervention. To quantify the output of the holding controller, we measured the forces the subject applied to the handle while the robot moved the hand. To prevent participants from voluntarily opposing the imposed hand displacement, we distracted each participant with a difficult working memory task. We did not inform participants as to the nature or presence of the postural perturbation. Instead, we instructed participants to solely concentrate on the working memory task and obtain as many points as possible by answering memory questions correctly. Points for correct responses were combined with the points awarded for successful reaching movements.

The postural perturbation consisted of a straight-line kinematic path designed to make the probe as imperceptible as possible. To move the hand along the path, we placed the hand in a two-dimensional clamp with stiff spring-like mechanics (stiffness = 4000 N/m, viscosity = 75 N-s/m) and moved the equilibrium position of the clamp through the workspace a total of either 2.5 cm, 4 cm or 5 cm, depending on the trial. The imposed motion consisted of three phases. In the first phase the hand was moved from its resting position to some desired position and velocity in accordance with a minimum jerk trajectory. In the second phase, the hand was maintained at its current velocity for some additional displacement. Finally, the hand was then brought to rest over some desired displacement in accordance with a minimum jerk trajectory.

At the start of each displacement, the clamp equilibrium position was first moved a short distance (0.15 cm for 2.5 cm probes, 0.15 cm for 4 cm probes, and 0.3 cm for 5 cm probe) along a minimum jerk trajectory, over a short duration (0.75 seconds for 2.5 cm probes, 0.75 seconds for 4 cm probes, and 1.5 seconds for 5 cm probes). At the end of this displacement the velocity of the hand was equal to 0.375 cm/s. The hand was then moved at this constant velocity for a specified displacement (2.2 cm for 2.5 cm probes, 3.7 cm for 4 cm probes, and 4.4 cm for 5 cm probes). This constant velocity displacement lasted for 5.87 seconds for 2.5 cm probes, 9.87 seconds for 4 cm probes, and 11.73 seconds for 5 cm probes. After this constant velocity period, the hand was slowed to rest over a short distance (0.15 cm for 2.5 cm probes, 0.15 cm for 4 cm probes, and 0.3 cm for 5 cm probe) along a minimum jerk trajectory, over a short duration (0.75 seconds for 2.5 cm probes, 0.75 seconds for 4 cm probes, and 1.5 seconds for 5 cm probes). Finally, an additional buffer period of 0.3 seconds was added after reaching the final displaced position, prior to the end of the probe trial. The total duration of the probe was therefore 7.67 seconds for 2.5 cm probes, 11.67 seconds for 4 cm probes, and 15.03 seconds for 5 cm probes.

Critically, as stated earlier, the participant was not provided position feedback during the postural probe. Instead, cursor feedback of hand position was frozen at its holding location. Therefore, at the end of the postural probe, there was a discrepancy between the location of the hand and the location of cursor feedback. To seamlessly reunite the hand with its cursor feedback without drawing the attention of the participant, we manipulated visual feedback during the following reaching trial: as the next reach was executed, we projected the cursor position onto the line connecting the frozen cursor position and the position of the next target. In this way, it appeared to the participant as if they were reaching perfectly straight between the start and target position. At the same time, we confined

the motion of the hand to a straight line connecting its displaced position with that of the next target. When the hand entered the target, a small and brief force pulse was applied to move the hand to the center of the target at which point *x* and *y* feedback was reunited with the hand position.

Quantifying the null point and shape of the postural field

As the hand of the participant was moved by the robot during postural probe trials, the displacement of the hand was opposed by postural restoring forces. We used these forces to quantify how postural control of arm position responded to changes in the movement forces that transported the arm. To mathematically characterize the two-dimensional field of restoring forces, we fit a simple mathematical model [30] that treated the arm as a linear two-dimensional spring with a single equilibrium point:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = K \begin{bmatrix} x - x_{null} \\ y - y_{null} \end{bmatrix} \tag{4}$$

where F_x and F_y are the forces applied to the handle due to displacement of the hand from the null point of the system (x_{null},y_{null}) to some position (x,y). The stiffness matrix K,

$$K = \begin{bmatrix} k_{xx} & k_{xy} \\ k_{yx} & k_{yy} \end{bmatrix} \tag{5}$$

describes the magnitude and orientation of the stiffness field. We constrained K to be a symmetric matrix (*i.e.*, $k_{xy} = k_{yx}$). We fit this linear spring model to the postural restoring forces by identifying the parameter set (5 free parameters, x_{null}, y_{null}, k_{xx}, k_{xy}, k_{yy}) that minimized the sum of squared error between the hand forces (collapsed across the x and y axes) predicted by Eq. (4) and the hand forces measured during all of the postural probe trials. For this fit, we used the forces measured within the ellipse bounded by -2.25 to 2.25 cm along the x-axis and -1.5 to 1.5 cm along the y-axis, relative to the final hand position. To locate the optimal parameter set, we used the genetic algorithm in MATLAB R2018a. We repeated the genetic algorithm search 8 times and selected the one that minimized the squared error cost function. The optimal parameter set provided a good fit to the data, accounting for approximately 70% of the variance in the observed postural field (R^2 during baseline period, mean ± SEM: 0.70±0.03; R^2 after adaptation: 0.69±0.02).

After fitting this linear model to the postural restoring field, we evaluated if the location and shape of the field changed after the reaching forces changed due to exposure to a force field. First, we looked for within-subject changes to the location of the null point (Fig. 1C, x_{null} and y_{null}) of the field. Second, we looked for within-subject changes to the orientation (Fig. 1C, orientation) of the field, and the overall stiffness magnitude of the field (Fig. 1C, stiffness). To calculate the orientation of the field, we considered the eigenvector of the stiffness matrix K corresponding to the largest eigenvalue of K. We calculated the angle of this eigenvector in the *x*-*y* plane. To compute the stiffness of the field, we calculated the Frobenius norm of the stiffness matrix K.

Working memory task

In Experiments 1, 2, and 3 we employed a cognitive task to distract participants during measurement of holding forces. The working memory task consisted of a modified 2-back task where subjects were randomly shown an integer between 1 and 4. The integers appeared one at a time so that the next

integer replaced the previous integer on the screen (Fig. 1A). Participants were told to determine if the integer on the display matched the integer shown two numbers in the past. If the integers matched, participants verbally responded with the keyword “same”. If the integers did not match, participants verbally responded with the keyword “different”. If the response was correct a pleasant tone was played and a point was added to the experiment score. If the response was incorrect a low pitch tone was played and no point was awarded. To confirm that participants were engaged in the cognitive task, we recorded each correct and incorrect response. Participants were clearly engaged in the cognitive task and responded correctly to 91.8 ± 0.6% of items correctly, at rates of approximately 0.77 ± 0.6 items per second.

Force measurement during reaching and hold periods

Measurements of the postural restoring field in Experiments 1 and 3 indicated that the desired hold position of the arm was dependent on the forces that moved the arm to that position. This raised the possibility that the control of reaching and the control of holding may be coordinated in a serial manner, similar to the oculomotor system. We wondered if a process of neural integration could also explain the link between the forces that moved the arm and the postural field used to hold the arm at the end of movement. To test this idea, we measured moving and holding forces throughout the adaptation process. For each movement, we calculated the time-integral of the forces that moved the arm during the reach, and compared this integral to the holding forces the participant applied to the handle during the holding period.

To calculate the reaching force integral, we numerically integrated the forces applied to the handle during movement. Movement onset and offset were detected by identifying the time-points at which the hand crossed a velocity threshold of 1 cm/s. To calculate the holding forces at the end of the reach, we computed the average force applied to the handle over a 900 ms period after the termination of the reaching movement (100-1000 ms after reach termination). To study the relationship between holding forces and the reaching force time-integral, we linearly regressed holding forces onto the moving force integral. At times, we performed this regression within single subject datasets, and at others, on quantities that were averaged across subjects. To measure the strength of reaching and holding correlations within each subject, we also calculated correlation coefficients.

Reaching and holding forces were measured on designated probe trials. The perturbing force field only intervened in the axis orthogonal to the primary motion of the hand. In this way, changes to the moving forces were always oriented along the axis perpendicular to the primary movement direction. The forces we report here, along with any analysis of the time-course of force production, refer to the component of force perpendicular to the axis of the primary reaching movement. The only exception to this is the forces we report in the measurement of the postural fields of Experiments 1 and 3.

Reaching forces were measured during “channel” trials where the motion of the handle was restricted to a linear path connecting the start and target locations (Fig. 1A). To restrict hand motion to the straight-line channel trajectory, the robot applied perpendicular stiff spring-like forces with damping (stiffness = 6000 N/m, viscosity = 250 N-s/m). This channel condition was also used to measure holding forces. That is, on channel trials, the robot did not apply a two-dimensional endpoint clamp as we

described earlier (see *Holding period*). Rather, the channel condition remained unchanged during holding so that the robot only maintained equilibrium of the hand along the perpendicular axis.

Reaching and holding forces were measured at regular intervals throughout the experiment. In general, every 5th outwards reaching trial (the reaching direction along which the perturbation was applied) was performed in channel trials to probe moving and holding forces. All backwards reaching movements (the reaching direction that never experienced a perturbation) were performed in either a channel trial or a partial channel trial where the channel was inactivated partway through the reach, after the hand had traveled 40% of the desired movement amplitude. By starting backwards reaching movements in the channel condition, we prevented the possibility that holding forces might move the hand off the target if the channel was disengaged.

Offline we isolated forces along the perpendicular direction that were related to the process of adaptation. To do this, we calculated the average force during channel trials interspersed throughout the baseline period, prior to the introduction of force perturbations. We then subtracted this baseline force time-series from all of the force time-series recorded during channel trials within the adaptation period. All force profiles and measurements (*e.g.,* the integral of moving forces, the static holding force) are composed from these force time-series where baseline reaching force is removed (the only exception to this are the forces reported during the measurement of the postural restoring field in Experiments 1 and 3).

Testing the causal effect of reaching on holding

We found that holding forces were tightly coupled to the integral of the reaching forces that preceded them. We noted that this coupling could be explained by two different models of moving and holding circuits. In one model, the postural controller sits downstream from the reach controller, integrating the moving forces to generate the holding forces (Fig. 3F, left, serial model). In the other model, the two controllers operate in parallel, but are both adapted by a common perturbation (Fig. 3F, right, parallel model). Critically, in the serial model, reaching forces causally determine the holding forces, but in the parallel model, a spurious non-causal correlation between reaching and holding arises based on a common perturbation input.

To test the hypotheses, we first isolated the initial trials of each experiment before participants had ever experienced a force field (Figs. 3G-I). For each subject we calculated the mean force profile during this period, and then subtracted this profile from each individual trial. What remained were residual forces relative to the mean. We calculated the integral of these residual forces during reaching, and the averaged residual force during holding as described earlier (*Force measurement during reaching and holding*). We then sorted these residuals based on the time-integral of reaching forces. For each subject, we placed reaching force time-integrals and holding forces into one of twenty bins spanning reaching force time-integrals of -0.5 to 0.5 N-s. We then calculated the average reaching force time-integral and holding force in each bin and performed a linear regression (Fig. 3H). For visualization of these residuals, we performed the same process, but using only 5 bins (Fig. 3G). To report the strength of the moving and holding force relationship within single subjects, we calculated the correlation coefficient between these quantities for each subject (Fig. 3I).

We next repeated a very similar analysis but for a larger subset of trials where the perturbation was constant over many consecutive trials (Fig. 3J). Here, we subtracted baseline reach forces from each

trial as in our primary analysis (see *Force measurement during reaching and holding*). Then, for each separate period of constant perturbation trials, we calculated the average reaching force time-integral and holding force, and subtracted these quantities from each trial measurement. We then repeated the process described for baseline trials, where trials were sorted based on the residual reaching force time-integral into one of twenty bins, averages were calculated across trials within each subject, and then the trial means were averaged across subjects to perform a linear regression (Fig. 3K). Again, as for the analysis of the baseline trials, we also computed the correlation coefficient between the residual moving and holding forces for each individual subject (Fig. 3L).

Finally, we casted the widest net around trials that could be used to analyze the causal effect of reaching on holding, namely, all trials. In order to perform the same analysis as before, we needed to de-trend the reaching and holding forces with respect to the perturbation. To do this, for each subject, we linearly regressed the reaching force time-integrals and holding forces for each trial onto the perturbation sequence according to a linear model:

$$\begin{aligned} \hat{F}_m^{(t)} &= a_m p^{(t)} + b_m \\ \hat{F}_h^{(t)} &= a_h p^{(t)} + b_h \end{aligned} \tag{6}$$

where $\hat{F}_m^{(t)}$ and $\hat{F}_h^{(t)}$ are the model predictions for the reaching force time-integral and holding force on trial t, $p^{(t)}$ is the perturbation on trial t, and a_m, a_h, b_m, and b_h are the gains and biases for the linear model. Trial-by-trial, we then subtracted the linear model predictions for the reaching and holding forces, from the actual reaching and holding forces, to construct a distribution of residuals with zero mean. We analyzed these residuals (Fig. S7) in exactly the same manner as the constant perturbation periods described above.

Measuring errors in the process of holding still

Our primary analysis demonstrated how changes to reaching forces led to changes in the process of holding still. In most cases, we forced the reaching forces to change by adapting the move controller. We attempted to eliminate adaptation in the holding controller by applying an endpoint clamp during the holding period. These endpoint clamps, however, could not eliminate errors individuals made as they attempted to stop the hand within the target. We considered the possibility that the holding controller could be modified by these stabilization errors.

To isolate errors near the end of the reaching movement, we spatially aligned each trajectory by subtracting the final x and y coordinates of the hand from each point in the trajectory. We quantified the endpoint error in a manner agnostic to the exact source of the error. We looked for any correction that occurred as the hand was decelerated to rest at the end of the reaching movement. Specifically, we calculated the magnitude and sign of the largest lateral deviation from the final hand position, after the hand had exceeded 80% of the target-to-target displacement. This endpoint error could either be positive or negative; our convention is that positive errors indicate errors in the direction of the holding force (should reduce holding force magnitude) and negative errors indicate errors opposite the direction of the holding force (should stabilize and increase the holding force magnitude).

Familiarization

All participants were given familiarization trials at the start of each experiment. During this familiarization period, participants learned the desired movement speed through audiovisual feedback and scoring, as described in *Trial Structure*. The experimenter clarified the meaning of all visual objects in the workspace and instructed participants to obtain as many points as possible. After participants were able to reach between targets comfortably and accurately, the experiment score was reset and the experimental paradigm began.

After practicing reaching movements, further familiarization was provided to participants who were to perform the working memory task during holding (Experiments 1, 2, and 3). First, these individuals were exposed to the working memory task in isolation (*i.e.*, without having to perform reaching movements). Then, a practice period was given where postural probe trials were interspersed within normal reaching trials. For the first few postural probe trials in this practice period, the experimenter warned the participant that the reaching movement would end with the working memory task.

Experiment 1 – Altering the reach forces and measuring its effects on the postural field

To determine if the postural field changed after changes to the reaching forces, we adapted a set of participants (n=27) to a CCW velocity-dependent force field. To measure the postural field in two-dimensional space, postural probes moved the hand in 1 of 12 directions (0°, 15°, 30°, 45°, 90°, 135°, 180°, 225°, 270°, 315°, 330°, and 345° with respect to the x-axis, while participants were distracted with a working memory task (see *Working Memory Task*). We measured the postural field before and after adaptation. Before adaptation, participants completed 3 blocks of trials, each separated by a short break. In each block, all 12 postural probe directions were visited a single time. The probe displacement was 2.5 cm for all probe directions. Postural probes were given on every 4th outward reaching movement. Therefore, participants completed a total of 288 baseline trials (3 blocks x 12 probes/block x 4 reaching movement pairs/probe x 2 movements/reaching movement pair). Outwards reaching movements of 10 cm were all performed directly away from the body.

Participants were then gradually adapted to a velocity-dependent force field. The field magnitude (f in Eq. 1.1) was decreased from 0 to -10 N-s/m in constant increments over the course of 65 outwards reaching trials (130 trials total). After this adaptation period, the postural field was re-measured. As before, all 12 probe directions were probed in a random order, 3 times. No breaks were provided in between blocks. We anticipated that the postural field would shift after adaptation to the force field. Therefore, we extended the probe displacement to 4 cm for probe angles of 0°, 15°, 30°, 45°, 315°, 330°, and 345°. Postural probes occurred at the same frequency as before adaptation for a total of 288 trials. To maintain participants in an adapted state, on all outwards reaching trials other than postural probe trials, a velocity-dependent perturbation was maintained at -10 N-s/m. Endpoint clamps were applied to the final hand position during all outwards trials (with the exception of postural probe trials) during the holding period (1.5 second duration, with an addition 0.3-0.4 inter-trial-interval).

To quantify the postural field before and after adaptation (Figs. 1B and 1C), we fit a two-dimensional linear spring (Eq. 4) to x and y forces measured on postural probe trials for each subject individually. The model and fitting process are described in *Quantifying the null point and shape of the postural field*. For visualization purposes (Fig. 1B), we constructed a two-dimensional postural field from

the forces measured during probe trials using linear interpolation. To do this, along each probe direction we resampled forces in x and y spatially in intervals of approximately 0.1 cm. For each of the resampled restoring forces we calculated the corresponding polar coordinates (*i.e.*, the radius and angle). In polar coordinates, all x and y forces lied along a rectangular grid. We used bilinear interpolation along these polar coordinates to estimate the postural field within the space between the 12 probe angles.

Experiment 2 – Measuring the duration of holding forces

Misalignment between the null point of the postural field and the hand position led to a static holding force (Fig. 1D). We recruited a set of subjects (n=7) to determine if this holding force was stable in time (*i.e.*, that the set point of the postural controller was stable in time). We gradually adapted participants to a CCW velocity-dependent force field. At various points during the experiment, we measured holding forces for a long period of time (6.5 seconds in total). To keep participants engaged in the holding process, during the final 6 seconds of the holding period, subjects were distracted with the working memory task (same task used in measuring the postural field). During this period of time, the hand was not displaced from the target, rather an endpoint clamp was engaged to hold the hand at its final position. Endpoint clamps were applied to the hand on all outwards reaching movements, but only lasted 1.8-1.9 seconds on normal reaching trials with no working memory task. The reaching movement preceding the measurement of holding force was performed in a channel. Outwards reaching movements of 10 cm were all performed directly away from the body.

Before adaptation, holding forces were measured in a block consisting of 100 trials (50 outwards reaches and 50 backwards reaches). On every 5th outwards reaching trial, the holding period was extended and the working memory task was presented. After this period, participants took a short break. After the break, the velocity-dependent perturbation was introduced. The field magnitude was decreased from 0 to -7.5 N-s/m in constant increments over the course of 50 outwards reaching trials (100 trials total). Over the next 50 outwards reaches (100 trials total) holding forces were again measured during an extended period on every 5th reach. On all other reaches, the force field was maintained at a constant level (-7.5 N-s/m). Next the force field was gradually deceased further over 50 outwards reaches (100 trials total) to -15 N-s/m. During the final 50 outwards reaches (100 trials total), holding forces were again measured during the working memory task on every 5th reach. On all other reaches, the force field was maintained at a constant level (-15 N-s/m).

Experiment 3 – Measuring the relationship between holding forces and the null point of the postural field

We reasoned that if holding forces truly reflected misalignment between the position of the hand and the null point of the postural field, we could gradually eliminate these forces if we displaced the hand towards its null point. We recruited a set of subjects (n=19) to test these predictions throughout the process of gradual adaptation to a velocity-dependent force field. Given our findings in Experiment 1 we expected that the null point of the postural system would be located laterally to the hand in the direction of the compensatory hand force. For this reason, to measure forces in the direction of the null point, we exposed participants to probes along this single direction (0° with respect to the x axis). As always, during the postural probe, participants were distracted with the working memory task. For the first ten participants, we used 5 cm postural probe trials. For the last nine participants, we shortened the

probe length to 4 cm. Here we analyze only the first 4 cm of displacement to combine both versions of the experiment.

Before adaptation, we measured the postural forces a total of 10 times. Postural probes were inserted regularly on every 5th outward reach, for a total of 100 trials in this baseline period (10 probes x 5 outwards reaches/probe x 2 reaches/outward reach). Outwards reaching movements of 10 cm were all performed directly away from the body. Next, we adapted participants gradually to a CCW velocity-dependent force field, where we decreased the field magnitude (*f* in Eq. 1.1) from 0 to -10.5 N-s/m in constant increments over the course of 175 outwards reaching trials (350 trials total). We measured hand forces evoked by the postural probe on every 5th outwards reaching movement. On postural probe trials, reaches were performed in a channel.

We found that hand forces during the postural probe resembled a linear spring throughout the process of adaptation. To determine the location of the null point on a trial-to-trial basis we fit a line to the hand forces as a function of hand displacement in the probe, and recorded the x-intercept of the line. To do this, we first resampled subject forces spatially in increments of 0.05 cm. Next, to reduce noise inherent in the single trial force measurements, we used a bootstrapping approach. On each trial, we randomly sampled subjects with replacement, calculated the mean postural force as a function of distance across the sample, and fit a line to this mean behavior. We repeated this process 2000 times, and used this distribution to estimate 95% confidence intervals around the mean (Fig. 1F).

Experiments 4, 5, 6, and 7 – Measuring reaching and holding forces during adaptation and de-adaptation

After adaptation to a movement-related perturbation, postural control at the end of movement was changed. To test if, like the oculomotor system, changes in holding were related to the integration of movement-related commands, we compared holding forces with moving force integrals across a wide variety of tasks. In these tasks we systematically changed moving forces using different force field adaptation paradigms. We conducted different experiments to test how the relationship between holding force and the moving force integral varied as a function of the position of the arm in the workspace, the direction of the reaching movement, and the direction of the force perturbation.

In Experiment 4, we tested participants (n=15) in a paradigm consisting of 10 cm reaching movements, where the hand was centered in front of the body and outwards reaching movements were performed directly away from and the body. In Experiment 5, we tested participants (n=11) in a paradigm consisting of 10 cm reaching movements, where the starting position was placed in front of the body but outwards reaching movements were performed at an oblique angle of 135° (away from and to the left of the midline of the body). In Experiments 6 and 7, we tested participants (n=9 for Experiment 6, n=8 for Experiment 7) in a paradigm consisting of 10 cm reaching movements, where the hand was placed to the left and the right of the midline of the body by approximately 15 cm and 5 cm respectively, and outwards reaches were performed only along the y-axis of the workspace away from the body (essentially these experiments were analogues of Experiment 4, translated to the left and right). Experiments 6 and 7 were identical, except that in Experiment 6 participants were adapted to CCW force fields, and in Experiment 7 participants were adapted to CW force fields.

Experiments 4-7 all followed the same general trial structure (Fig. 2A). Before exposure to the perturbation, participants reached for 40 trials (20 outwards trials and 20 backwards trials) in a baseline

period. Every 2nd outwards reach was performed in a channel where moving and holding forces were measured. Next, the adaptation period started, which consisted of two blocks with 3 phases each. In the first phase, a CW/CCW velocity-dependent perturbation was introduced, and gradually increased/decreased from 0 to 15/-15 N-s/m over the course of 100 outwards reaching trials (200 trials in total). The perturbation magnitude was then maintained at a constant level of 15/-15 N-s/m over the course of 50 outwards reaches (100 trials total) and then brought back to zero gradually in a de-adaptation period of 100 outwards reaching trials (200 trials total). After this de-adaptation period, participants continued to reach in a washout period of 20 outwards reaches (40 trials total) where no force field was applied. Participants were then given a short break and this structure was repeated. In Experiments 4 and 5, participants were exposed to either a CCW or CW perturbation in the 1st block, and the opposite perturbation in the 2nd block (order counterbalanced across subjects). In Experiments 6 and 7, participants were exposed to either the left-most target or the right-most target in the 1st block, and the opposite target in the 2nd block (order counterbalanced across subjects). During the adaptation and de-adaptation periods, we measured subject moving and holding forces in channel trials on every 5th outwards reaching movement.

Experiments 8 and *9 – Reaching during a gradual, then sustained force field perturbation*

Our experiments utilized force field perturbations as a means of adapting the moving controller in order to change reach forces. To maintain postural stability, we expected that the integrator should also be able to adapt. That is, if the null point of the postural system is not aligned with the hand at the end of a reaching moving, the ensuing postural error should provide a teaching signal for the integrator. In search of evidence for this process of adaptation, we recruited two sets of subjects (n=15 for Experiment 8 and n=17 for Experiment 9). We gradually adapted participants to a velocity-dependent force field (CCW field for Experiment 8, CW field for Experiment 9). Then, we maintained the perturbation at a constant level for a long period of time. Reaches of 10 cm were performed at oblique angles of 135° as in Experiment 5. Along this reaching direction, CCW and CW perturbations led to large differences in endpoint errors (Fig. S8A-C). We wanted to know if these different endpoint errors resulted in changes in the size of holding forces, reflecting the process of adaptation. We calculated stabilization errors on all perturbation trials (see *Quantifying holding errors*) and compared these to the holding forces (Fig. S8).

Experiments 8 and 9 both followed the same trial structure. Before exposure to the perturbation, participants reached for 40 trials (20 outwards trials and 20 backwards trials) in a baseline period. Every 2nd outwards reach was performed in a channel where moving and holding forces were measured. Next, the adaptation period started, which consisted of 2 phases. In the first phase, a CW/CCW velocity-dependent perturbation was introduced, and gradually increased/decreased from 0 to 15/-15 N-s/m over the course of 100 outwards reaching trials (200 trials in total). In the second phase, the perturbation magnitude was then maintained at a constant level of 15/-15 N-s/m over the course of 200 outwards reaches (400 trials total). We measured subject moving and holding forces in channel trials on every 5th outwards reaching movement.

Experiment 10 – Holding forces disappear without endpoint clamps

Our findings in Experiment 9 suggested that the holding system is adapted by the experience of stabilization errors. In Experiment 10, we recruited a set of participants (n=20) to increase the size of these errors to determine if these errors could drive adaptation that would gradually eliminate holding forces over trials. To increase endpoint errors, we changed the task in two ways: (1) we perturbed individuals abruptly (as opposed to gradually) to increase the size of errors during the movement and (2) we never applied a two-dimensional endpoint clamp to stabilize the final reaching position (in other words, large holding forces could move the hand off the target as in typical force field adaptation paradigms). As in Experiment 9, we calculated stabilization errors on each perturbation trial.

The experiment paradigm consisted of two reaching blocks with a short break in the middle. In the first block, before exposure to the perturbation, participants reached for 40 trials (20 outwards trials and 20 backwards trials). Every 2nd outwards reach was performed in a channel where moving and holding forces were measured. Next, the abrupt adaptation period started and a 15 or -15 N-s/m velocity-dependent perturbation was immediately introduced. Participants adapted to this constant perturbation for 200 outwards reaching trials (400 trials total). At the end of the block, the perturbation was abruptly removed, and subjects reached in the absence of the force field for an additional 45 outwards trials (90 trials total). The second block followed the same structure as the first block. The blocks differed in the direction of the force field (CW or CCW) which was counterbalanced across participants.

Experiment 11 – Reaching from different start locations to the same endpoint

We found that holding forces were closely predicted by the time-integral of moving forces for reaching movement starting at one target position and ending at another. We next wondered, if different holding forces could co-exist at the same location in space, if the moving forces that brought the hand to that location started from different points. We first tested this idea by exploiting the phenomenon of spatial generalization [31,32] where adaptation in one location in space, can generalize to other locations in space where perturbations were never experienced. We designed an experiment, where two starting locations would lead to the same target location, but each reach would generalize moving forces with opposing directions.

The generalization paradigm had one start position, but two target positions. Target 1 was located 10 cm directly away from the body, relative to the start position. Target 2 was located 10 cm directly towards the body, relative to the start position. Participants (n=13) reached in a repetitive 4 trial structure (Fig. S2A): start position to Target 1, Target 1 back to start position, start position to Target 2, and finally Target 2 back to start position. Because of this geometry, generalization of learning could occur because reaches from the start position to Target 1 were made in the same direction as reaches from Target 2 to the start position, and reaches from the start position to Target 2 were made in the same direction as reaches from Target 1 to the start position. We perturbed reaching movements from the start position to Target 1 and Target 2, but never perturbed reaching movements made towards the start position. Instead, these reaches were always made in channel trials.

The experiment paradigm consisted of gradual adaptation to a CW velocity-dependent force field. In other words, reaches to Target 1 were pushed to the right, and reaches to Target 2 were pushed to the left. This presented an interesting scenario (Fig. S2D and S2E), where the forces generalized when

reaching to the start position would sometimes be oriented to the right (Target 2 to start), and other times be oriented to the left (Target 1 to start). Participants were adapted to the force field over the course of 100 cycles. In each cycle, all 4 possible reaching movements were completed one time. On every 5th cycle, reaches to Target 1 and Target 2 were made in a channel in order to measure moving and holding forces. The CW perturbation gradually increased from 0 to 15 N-s/m in even increments across trials.

Experiment 12 – Reaching to a single endpoint with variable reaching forces

In Experiment 11, when reaching towards the start position, we found that participants produced two different holding forces, depending on the moving forces that preceded them. In other words, holding forces were coupled to moving forces, even though they varied in direction from one trial to the next. The caveat here, is that these opposing moving and holding forces were produced on trials that started from different positions in space. In Experiment 12, We recruited a set of subjects (n=14) to determine if this phenomenon would hold when reaching movements began from the same location in space.

To do this we adapted participants to a modified version of a dual adaptation paradigm described recently [33]. In this paradigm, rather than controlling a single cursor, participants controlled the motion of a rectangular box centered on the hand (Fig. S6A). On some trials, a control point and virtual target were placed on the left side of the tool. On other trials, a control and virtual target were placed on the right side of the tool. Participants were instructed to move the displayed control point to the displayed target. Critically, even though a different control point was used on alternating trials, the reaching movement started and began at the same locations in space. The control points were instead used as an implicit cue for the direction of an upcoming perturbation. During the adaptation period of the experiment, on trials where Control point 1 was displayed, a CW perturbation was applied to the reaching movement. On trials where Control point 2 was displayed, a CCW perturbation was applied to the reaching movement.

At the start of the experiment, participants reached with each control point a total of 20 times (40 reaches total). Then the CW or CCW perturbations were abruptly increased to 15 or -15 N-s/m, respectively. Participants adapted to the perturbation over the course of 100 Control point 1 reaches and 100 Control point 2 reaches (400 trials total). On every 5th outwards reaching movement, the reach was performed in a channel, and moving and holding forces for either Control point 1 or Control point 2 were measured. Reaching movements were always 10 cm in magnitude, and directed away from or towards the body.

Experiments 13 and 14 – Reaching in a zero integral force field

We considered a simpler alternative hypothesis to neural integration that could explain the one-to-one correlations we observed between moving forces and holding forces: holding forces could be a trivial continuation of the forces produced at the end of movement, just prior to the start of holding. To test this hypothesis, we designed a position-dependent force field, composed of two components (see *The moving period*, Eqs. 1.2 and 1.3): one that acted on the first half of the reaching movement ($FF_{1,\text{ first 10 cm of the reach}}$) and one that acted on the second (FF_2, last 10 cm of the reach). By adapting individuals to a field combined from FF_1 and FF_2, we created a scenario where large forces were produced at the end of the reaching movement, but the reaching forces integrated to zero.

Throughout Experiment 13 (n=14) we measured moving and holding forces in channel trials on every fifth outwards reaching movement. The experiment started with 25 unperturbed outwards reaching trials (50 trials total). After this, we gradually adapted participants to a force field that only perturbed the hand during the second half of the reaching movement (Fig. 3A, left). To do this, we gradually increased the magnitude of FF_2 from 0 N to 3.5 N in even increments, over the course of 100 outwards reaching movements (200 trials total). During this time, the magnitude of FF_1 was maintained at 0 N. We reinforced this learning period with an additional 10 outwards reaches (20 trials total) where FF_2 was maintained at 3.5 N and FF_1 at 0 N. In the final part of the experiment (Fig. 3A, right), we maintained FF_2 at 3.5 N, but gradually decreased FF_1 from 0 N to -3.5 N, to teach individuals a new set of forces during the first half of the reaching movement. We adapted individuals to this bidirectional force field slowly by decreasing FF_1 to its terminal magnitude over the course of 200 outwards reaching trials (400 trials total). In this way, at the end of the experiment, participants were exposed to two force fields within the same reaching movement that perturbed the hand in opposite directions but with equal magnitude (Fig. 3B).

We found that holding forces gradually decreased in the final phase of the experiment as individuals adapted to FF_1, consistent with our hypothesis of integration. If adaptation to FF_1 was the sole driver of the reduction of holding force, if FF_1 was omitted entirely, the holding forces should not decrease. To test this idea (and also confirm that the reduction in holding force was not caused by others factors like fatigue or accumulation of endpoint errors), we recruited a control group (n=11). The control experiment (Experiment 14) also started with 25 unperturbed outwards reaching trials (50 trial total). As in Experiment 13, control participants were then adapted to FF_2 only. The magnitude of FF_2 was gradually increased from 0 to 3.5 N over the course of 100 outwards reaches (200 trials total), while FF_1 was maintained at 0 N. During the final phase of the experiment, FF_2 was maintained at 3.5 N and FF_1 was maintained at 0 N for an additional 210 outwards reaches (420 trials in total, to match the trial count of Experiment 13).

Experiment 15 – Measuring cortical involvement in the control of reaching and holding

Altogether, our findings indicated that the control of reaching was dependent on the process of moving; forces that moved the arm were integrated into a null point that was then applied to postural control at the end of movement. Like the eye, our experiments suggested that the moving controller and the holding controller were connected in serial (Fig. 3F, left). For the oculomotor system, move and hold commands are generated by separate regions in the brainstem. We wondered if this is also true in the control of reaching and holding.

We recruited patients (n=13) who had survived a stroke affecting cortical or subcortical white matter associated with the corticospinal tract (CST). We measured the degree of motor impairment in the patients using the Fugl-Meyer Assessment (FMA) and the Action Research Arm Task (ARAT). Two separate raters scored each assessment and scores were averaged across raters. In each limb, we measured the strength of elbow flexion/extension and shoulder horizontal adduction/abduction using a dynamometer (microFet 2). During measurements, participants rested their arm on a side table which supported their arm so it rested slightly below shoulder level. Strength measurements were repeated twice, the maximal force was recorded on each effort, and forces were averaged over repetitions. During each measurement, patients were verbally encouraged by the experimenter to produce as much

force as possible. FMA scores, ARAT scores, strength, and other patient characteristics are reported in Table S1. Missing entries in table indicate that the patient was unable to perform the desired action. Patients were selected based on MRI or CT scans, and/or available radiologic reports. Scans and/or reports were corroborated to determine the level at which the white matter of the corticospinal tract (CST) was lesioned. Table S1 provides the level of the brain at which the white matter was lesioned.

If moving used cortical control, and the holding integrator used subcortical control, we reasoned that stroke patients may show impairment in the generation of reaching movements, but the process of integrating moving forces into holding commands would be spared. We tested the stroke patients and a set of healthy age-matched controls (n=10) on a paradigm similar to Experiment 4. To account for motor impairment, we relaxed the timing requirements of our standard reaching experiment. Participants were awarded a point for reaches (10 cm in magnitude) that terminated within 600-800 ms (as opposed to 450-550 ms). Additionally, the arm of the patient was supported against gravity by an air sled that rested on a smooth table located underneath the visual display. Age-matched controls were tested in an identical protocol.

Each experiment block started with 20 unperturbed outwards reaches (40 trials total). After this initial period, we gradually introduced a CW or a CCW velocity-dependent force field. Over the course of 80 outwards reaches (160 trials total) we gradually increased (CW) or decreased (CCW) the force field magnitude in even increments from 0 N-s/m to 18 or -18 N-s/m. The force field magnitude was maintained for an additional 40 outwards reaches (80 trials total) and then brought gradually back to 0 N-s/m over the course of 80 outwards reaches (160 trials total). The experiment block ended with 20 unperturbed outwards reaching movements (40 trials total). We measured moving and holding forces on channel trials on every 5th outwards reaching movement.

We measured behavior in both the paretic (contralateral to the lesion) and non-paretic arms. In this way, we could also determine the effect of paresis on reaching and holding forces within each patient. The paretic and non-paretic arms were tested in different blocks, with 2 blocks for each arm. The order of the blocks was always paretic, then non-paretic, then paretic, then non-paretic. In each block either a CCW or CW force field was applied. Across blocks, the perturbation order was always A-B-B-A where A and B refer to either a CW or CCW force field. The force field orientation on the first block was counterbalanced across participants. Healthy age-matched controls were tested in a similar manner. In the absence of a paretic/non-paretic distinction, we tested either the left or the right arm in the first block (counterbalanced across participants).

To identify any impairment in reaching we focused on various kinematic measures. In all of these measures, we only considered trials during the null field period prior to the introduction of the force field. First, we considered time-courses for how far the arm deviated from a straight-line trajectory. For this, we calculated the unsigned lateral deviation from the straight-line trajectory for each trial in the null field period, averaged these time-courses across trials within each subject, and then finally averaged these mean traces across subjects (Fig. 4C). In addition to this, we also calculated several kinematic measures of performance: the duration of the reaching movement (Fig. 4D) and the displacement from the target at reach endpoint (Fig. 4E).

We found that erratic reaching forces on null field trials led to sustained holding forces (Fig. 4H). To determine if coupling between moving and holding forces is affected by stroke, we selected the most impaired reaching movements. To do this, we identified all reaches where the integral of moving forces

fell outside of 2 standard deviations of the moving force time-integral distribution of the healthy control population (Fig. 4G). We analyzed if these highly impaired reaching forces were integrated similarly to normal reaching movements (Fig. 4H). To do this, we separated trials into bins according to the integral of the moving force.

To further search for an impairment in the process of neural integration, we looked at the gain of the linear relationship between reaching and holding, after exposing participants to a velocity-dependent curl field. We fit a linear model to the trial-by-trial measurements of holding forces and moving force integrals for each individual participant (Fig. 4J) and tested if this gain differed across the paretic arm of patients, non-paretic arm of patients, left arm of healthy control participants, and right arm of healthy control participants (Fig. 4L). When collapsing moving and holding forces across participants (*e.g.*, Fig. 4K) we accounted for the increased movement variability observed in the paretic arm by first sorting the trial-by-trial behavior with respect to the magnitude of the moving force integral, before averaging moving and holding forces across participants (within 1 trial bins).

Additional experiment for analysis of null field forces

In our analysis of reaching and holding force residuals during null field periods (Figs. 3G-I), we included null period trials from an additional experiment not described or included in our primary analysis. In this additional experiment (n=23), we considered only the relevant initial period which consisted of null field trials. This period consisted of 40 total reaching movements, 20 outwards movements and 20 backwards movements. Every 2nd outwards reach was performed in a channel trial where forces were measured. All backwards trials were performed in channels.

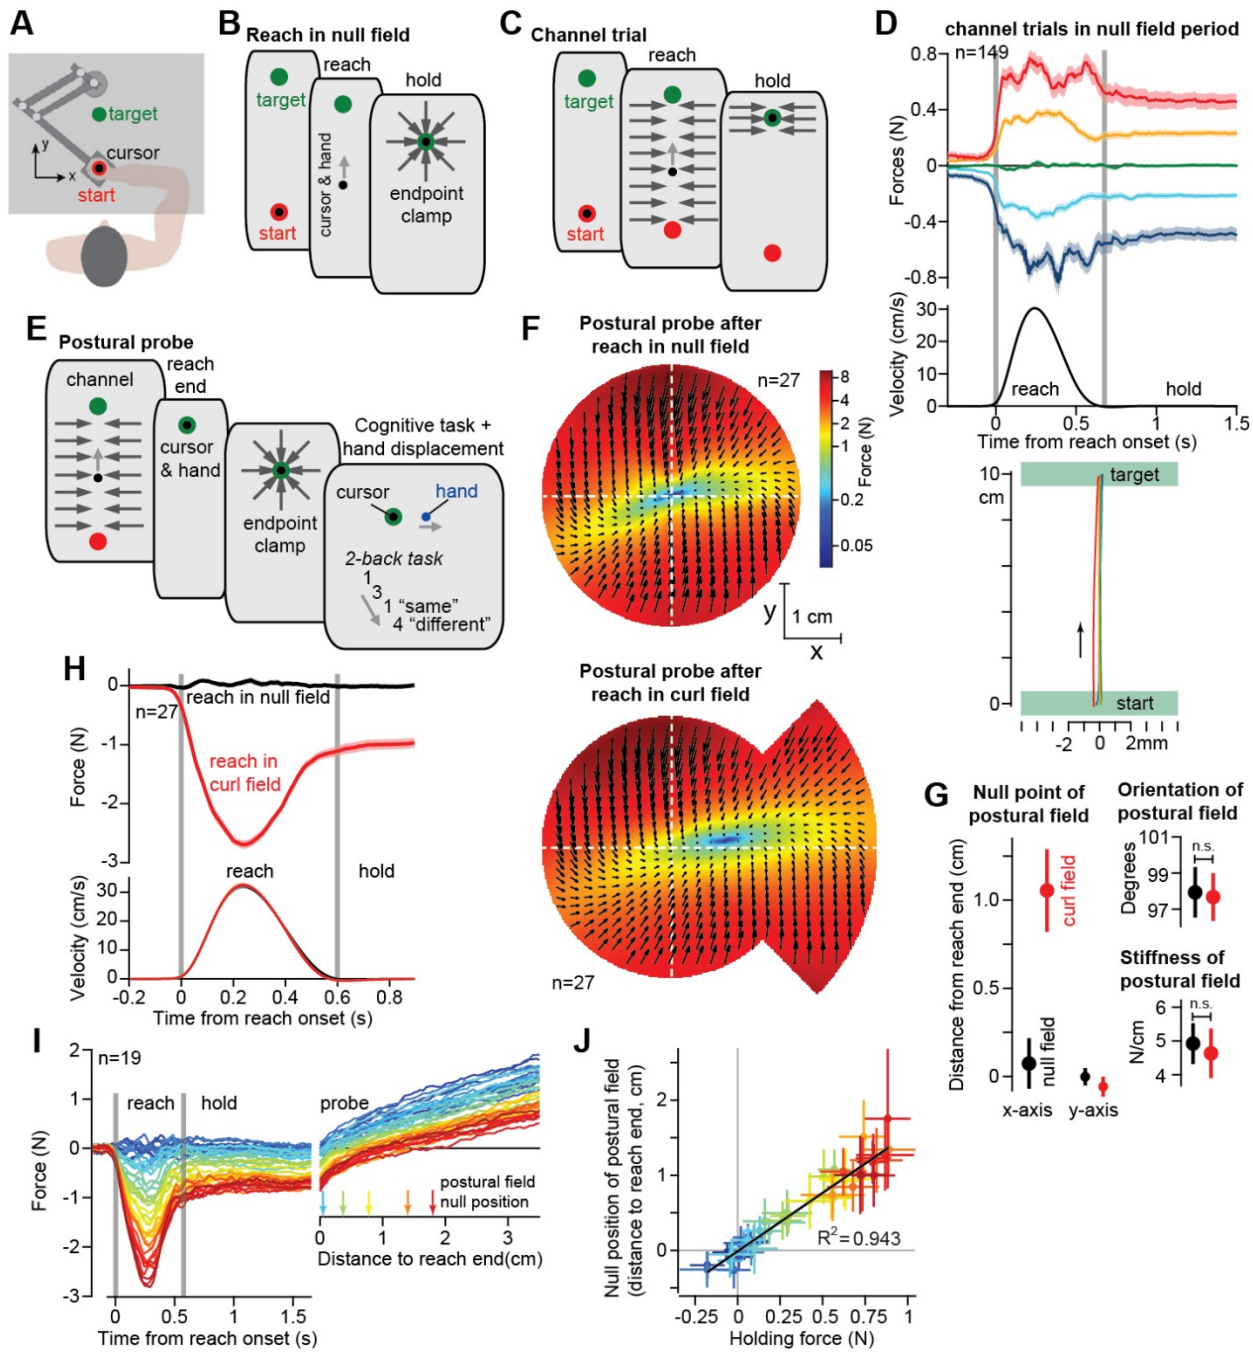

Figure 1. Change in reach forces affects the null point of the postural field. **A.** Participants held the handle of a robotic arm and made point-to-point reaching movements. **B.** First, participants reached free of any robotic intervention (null field). After the reach ended, an endpoint clamp held the hand in place. **C.** On some trials the robot created a channel that guided the hand straight to the target. Reaching and holding forces were measured on these channel trials. **D.** In the channel, natural variability in the execution of reaching forces could be measured. We separated trials into five groups depending on the magnitude and sign of lateral reaching forces. The mean velocity trace for all reaching movements is shown. The channel ensured that reaching trajectories and holding positions were virtually identical (bottom inset). Trajectory colors at bottom match the forces shown at top. Here we show data for 10 cm reaching movements that had holding periods of at least 1.5 seconds (n=149 subjects). Values are mean ± SEM across participants. **E.** We measured the postural field on postural probe trials by slowly

displacing the hand during the holding period, while participants were distracted with a working memory task. **F.** First, we measured the postural field during the null field trials. Next, we re-measured the postural field after participants were exposed to a velocity-dependent curl force field. The holding position at reach end is shown by intersection of two dashed white lines. Color reiterates the restoring force magnitude. Forces were measured in the null field (left) and curl field (right) conditions by displacing the hand outwards along 12 different lines. Interior estimates for the force were made using two-dimensional linear interpolation. **G.** We used a two-dimensional spring model to quantify postural field properties: null point, orientation, and stiffness (null field and curl field in black and red). Values are mean ± SEM across participants. **H.** We measured lateral forces (top) applied to the channel walls during reaching movements (null field period, black, and curl field period, red). Values are mean ± SEM across participants. **I.** Participants (n=19) were exposed to a curl field that gradually increased over trials. Gray vertical bars are movement start and stop times. During holding, we recorded hand forces (right) as the arm was displaced in the direction of holding forces. Arrows show the location of the null point (zero-crossing) on selected trials. **J.** We calculated the holding force before displacement of the hand, and the corresponding postural null point on each trial. Values are trial means and 95% CIs for distributions bootstrapped across participants. Linear regression was performed on the bootstrapped estimates (black line).

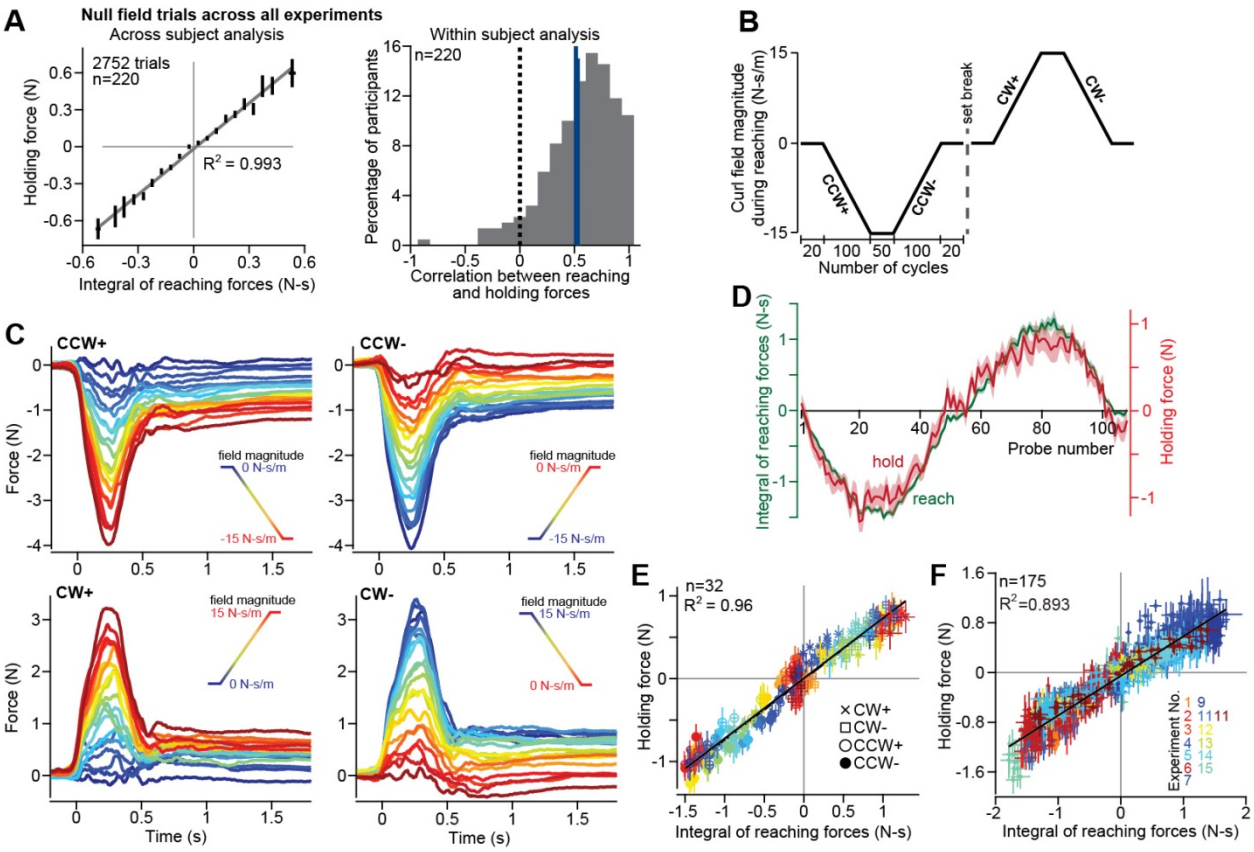

Figure 2. Holding forces are predicted by the time-integral of moving forces. **A.** We considered reaching behavior during the initial null field period for all healthy participants (n=220). For all channel trials, we calculated the integral of reaching forces, and the mean value of the holding force over a 900 ms period. We sorted trial pairs of holding forces and moving force time-integrals into twenty bins based on the moving forces. Shown at left, are the holding forces for each bin. Values are mean ± SEM across participants. Trials with moving force-time integrals between -0.75 N-s and 0.75 N-s were selected (>99% of all trials) for this analysis. At right, we calculated the correlation coefficient between moving force time-integrals and holding forces across all channel trials in the null field. **B.** To extend the range of reach forces, we exposed participants in Experiments 4, 6, and 7 (n=32) to gradually increasing (+) and decreasing (-) clockwise (CW) and counterclockwise (CCW) velocity-dependent curl force fields. A short break separated two testing periods. Reaching cycles consisted of an outwards and backwards movement. Only outwards movements were perturbed. **C.** On every fifth reaching cycle, moving and holding periods were executed in a channel in which lateral forces were measured. Each trace represents one trial, averaged across participants. Color indicates force field magnitude at each point in the experiment (blue is early in each period and red is late). **D.** We calculated the time-integral of forces during moving (green) and the mean holding force (red), as in our analysis of null field trials in **A**. Values are mean ± SEM across participants. **E.** We observed a tight coupling of holding forces to the moving force time-integral. Each point represents one trial. Colors represent trial ordering as in **C**. Values are mean ± SEM across participants. The solid line depicts a linear regression across all trials. **F.** The time-integral of moving force predicted holding forces across all experiments (total n=175). Each point represents one trial. Colors are provided to visually distinguish between experiments. The solid line shows a linear regression across all trials and all experiments. Experiment numbers are shown at bottom right.

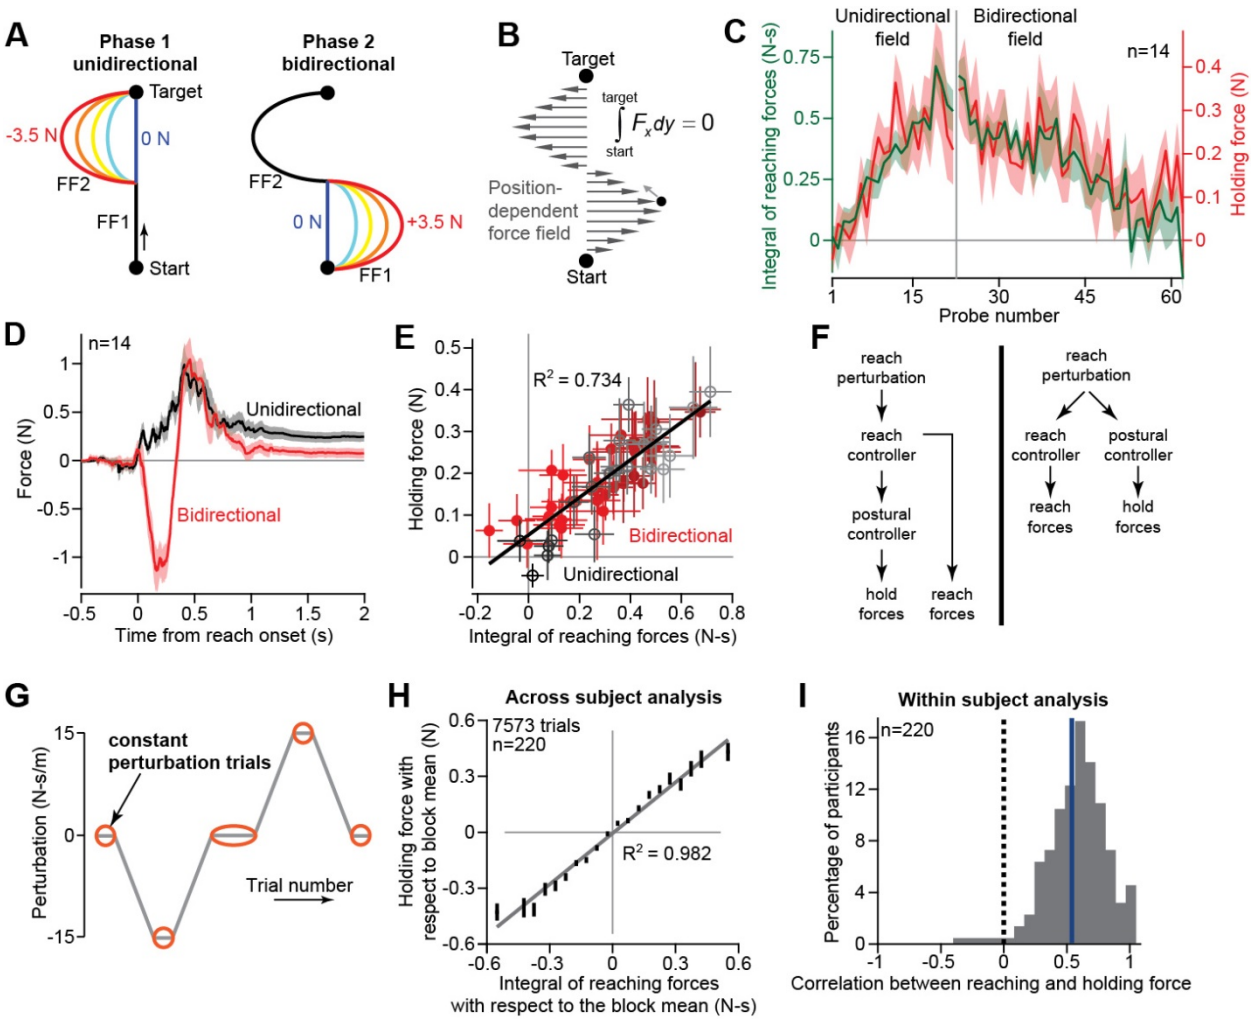

Figure 3. Holding forces are causally determined by the time-integral of the history of reach forces. **A.** Participants made 20 cm reaches between two targets. In the first phase of the experiment, a parabolic position-dependent force field was gradually applied to the second half of the reaching movement. In the second phase, an opposite force field was gradually applied to the first half of the movement. **B.** At the end of the experiment, successful compensation for the force field yields a force time-integral of zero. **C.** To compare moving and holding, we overlaid the time-integral of moving forces (green) and holding forces (red) on each trial. A gap and vertical line are provided to show the two separate phases of the experiment. Values are mean ± SEM across participants. **D.** We compared the time-course of force production during the unidirectional (black) and bidirectional (red) phases of the experiment. Vertical line shows movement end. Values are mean ± SEM across participants. **E.** We report the regression (black line) of holding forces onto the moving force time-integral for the unidirectional phase (black) and bidirectional phase (red). Each point represents one trial. Value are mean ± SEM across participants. **F.** Holding forces could be causally predicted by moving forces in the postural controller is in serial with the moving controller (left). Holding forces could be noncausally predicted by the moving controller if both systems are parallel and simultaneously adapted (right). **G.** We isolated all trials where the force field was constant (depicted by red circles). **H.** We subtracted off the mean holding force and moving force time-integral from each constant perturbation period. We next sorted these trial distances from the block mean into twenty bins based on the moving forces. Shown are the holding forces for each bin. Values are mean ± SEM across participants. **I.** For each subject, we calculated the correlation coefficient between moving force time-integrals and holding forces across constant perturbation trials, after mean subtraction as in **H**. Trials with moving force-time integrals between -0.75 N-s and 0.75 N-s were selected (>99% of all trials) across 220 participants for **G-I**.

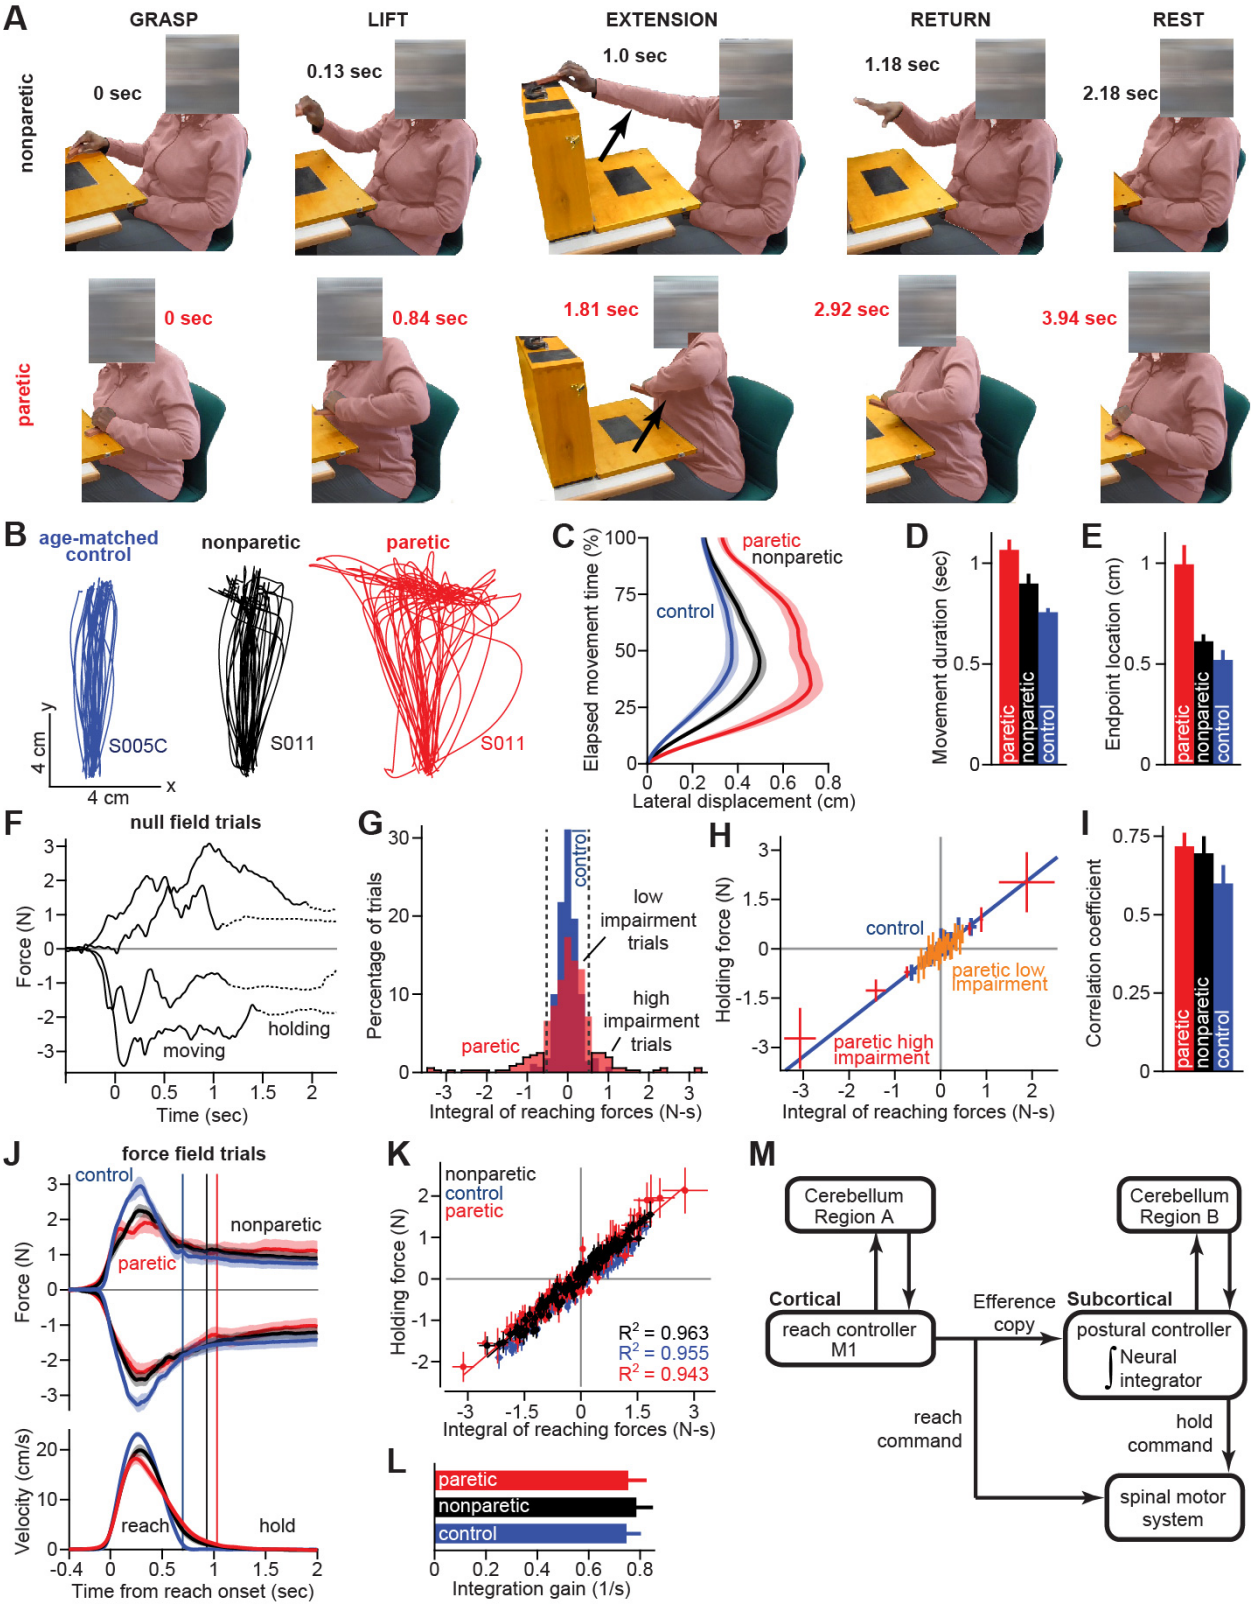

Figure 4. The cortical reach controller is separate from a subcortical holding integrator. **A.** Stroke survivors (n=13) participated in a set of clinical exams to measure functional impairment. Shown are isolated images for an

extension-based task, for the non-paretic (top) and paretic (bottom) arms of an example participant (S015). From left to right, the evolution of an outwards reach is shown. The instruction is to place a rectangular block on the elevated surface. The task could not be completed with the paretic limb. Time stamps provided near the top of each image. **B.** To improve the range of motion of the arm, patients and healthy controls performed reaching movements holding the robotic handle, with the arm supported by an air sled. Shown are example trajectories during an initial null field period for the representative patient (black is nonparetic, red is paretic) in **A**, and a control participant (blue). **C.** We measured the distance of the hand from the desired straight-line trajectory throughout the evolution of the reach. Values are mean ± SEM across participants. **D** and **E.** We calculated several kinematic measures that captured basic features of the motor impairment in stroke: **D** shows the required duration for the reach and **E** shows the distance of the hand from the intended target at reach end. **F.** Occasionally, the arm was restricted to a straight-line path and force was measured during reaching movements in the initial null field period. Shown are example trajectories. The solid line denotes forces during moving. The dashed line denotes forces during holding still. **G.** We calculated the time-integral of moving forces during null field period trials. Shown is a histogram of moving force integrals, collapsed across all participants. To identify the most erratic reaching movements in the paretic arm (high impairment, outlined with black border), we selected reaching movements of the paretic limb that fell outside 2 standard deviations of healthy control reaches (dashed lines). **H.** We calculated the holding forces on low (orange) and high (red) impairment trials in the paretic arm, and compared these to healthy control reaches (blue). The solid blue line is a linear regression for the control participants. Lines provide estimates for the 95% confidence intervals on the mean holding force and moving force time-integral across trials binned according to the integral of moving forces. **I.** We calculated the correlation coefficient between moving force integrals and holding forces during the baseline period. Values are mean ± SEM across participants. **J.** Patients and controls were exposed to gradually increasing force fields (top is CW, bottom is CCW). Shown are the forces at the peak of the force field. Velocity traces for the hand are provided at bottom. Vertical lines show the time of movement end. Values are mean ± SEM across participants. **K.** We interrogated the relationship between the holding force and time-integral of moving force. Each point represents one trial, sorted within each subject by the time-integral of the moving force. Values are mean ± SEM across participants. **L.** We fit a linear model to capture the relationship between the moving force time-integral and holding force for each participant separately. Here we report the slope of this line, *i.e.*, the gain of integration. Values are mean ± SEM across participants. **M.** A schematic of the proposed architecture for moving, holding, and neural integration.

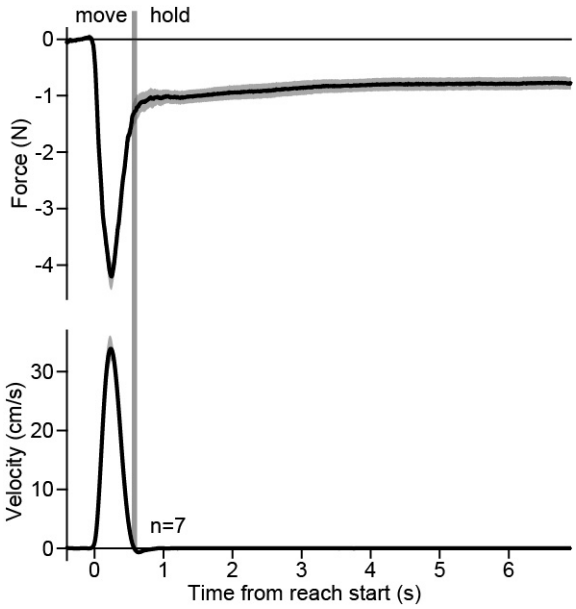

Figure S1. Holding forces are sustained across long time intervals. Participants (n=7) were exposed to a CCW velocity-dependent curl force that increased in magnitude across many trials. Here we show the lateral force (top) and hand velocity (bottom) after exposure to maximal field strength. The vertical line denotes the end of the movement, hence, the start of the holding period. On the trials shown, participants held the hand at the target and waited for a long inter-trial-interval to start the next reaching movement. During this period of time, participants engaged in a working memory task.

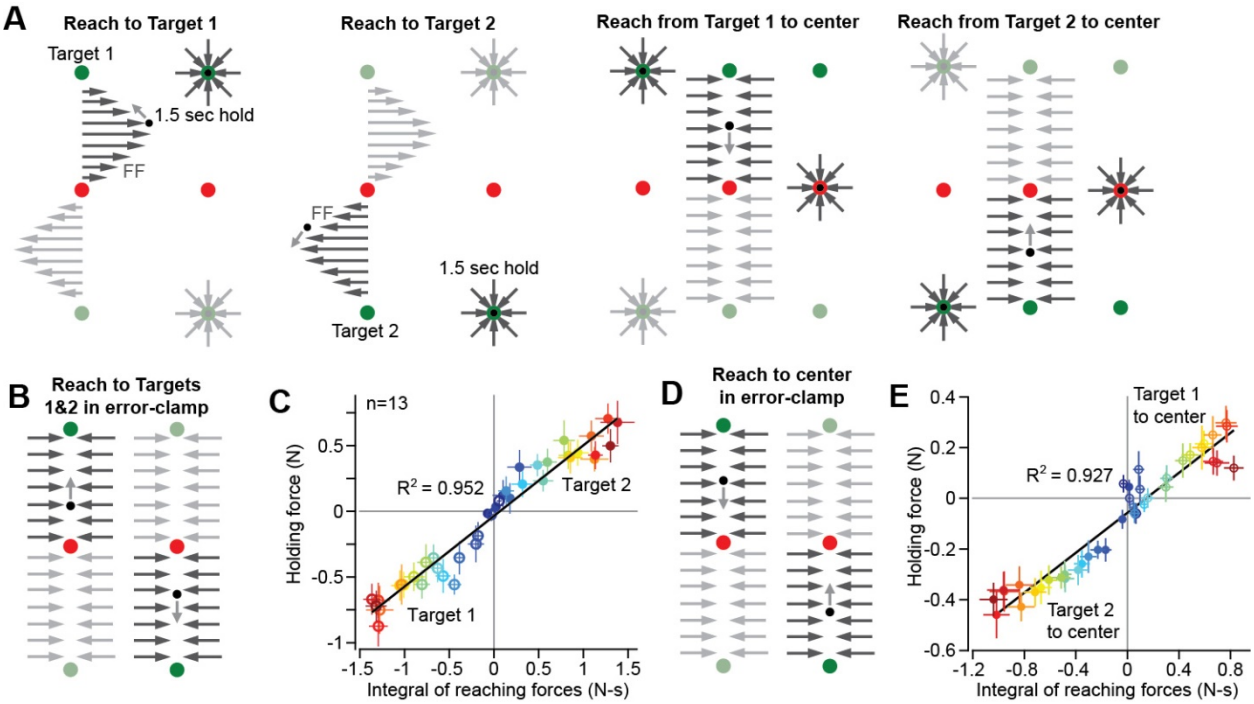

Figure S2. Holding forces at the same target position are specific to reaches initiated from different start positions. **A.** Participants (n=13) reached in cycles of four trials. On the first trial (first column), participants reached away from the body from a center position (red) to Target 1 (green). At the end of the movement, the hand was held in place for a long holding period. On the next trial (third column), the participant returned the hand from Target 1 to center. On the third trial of the cycle (second column), participants reached from the center position to Target 2. And finally, on the fourth trial (fourth column), participants returned the hand back to center again. After an initial familiarization period, participants were exposed to force fields when reaching from the center to Target 1 (rightwards field) and from the center to Target 2 (leftwards field). Reaching movements that terminated at the center position were always executed in a channel that prevented lateral errors. **B.** Occasionally, channel probe trials were inserted to measure hand forces when reaching to Targets 1 and 2. **C.** Shown is a scatterplot of the holding force as a function of the moving force time-integral for the probe trials depicted in **B**. Each point represents one trial. Color indicates the strength of the force field at that point in the experiment (blue means small, and red means large). Unfilled circles show reaches to Target 1 and filled circles show reaches to Target 2. The solid line shows the linear regression across trials. Values are mean ± SEM across participants. **D.** All reaching movements made towards the center position were executed in a channel. **E.** Same as in **C** but for probe trials depicted in **D**. Color indicates the force field magnitude at a specific point in the experiment. Filled circles depict reaching movements from Target 2 to center and unfilled circles depict reaching movements from Target 1 to center. Note that this implies that different holding forces coexisted at the same location in space (the center) on different trials, but at the same point in the experiment. The solid line shows the linear regression across trials. Values are mean ± SEM across participants.

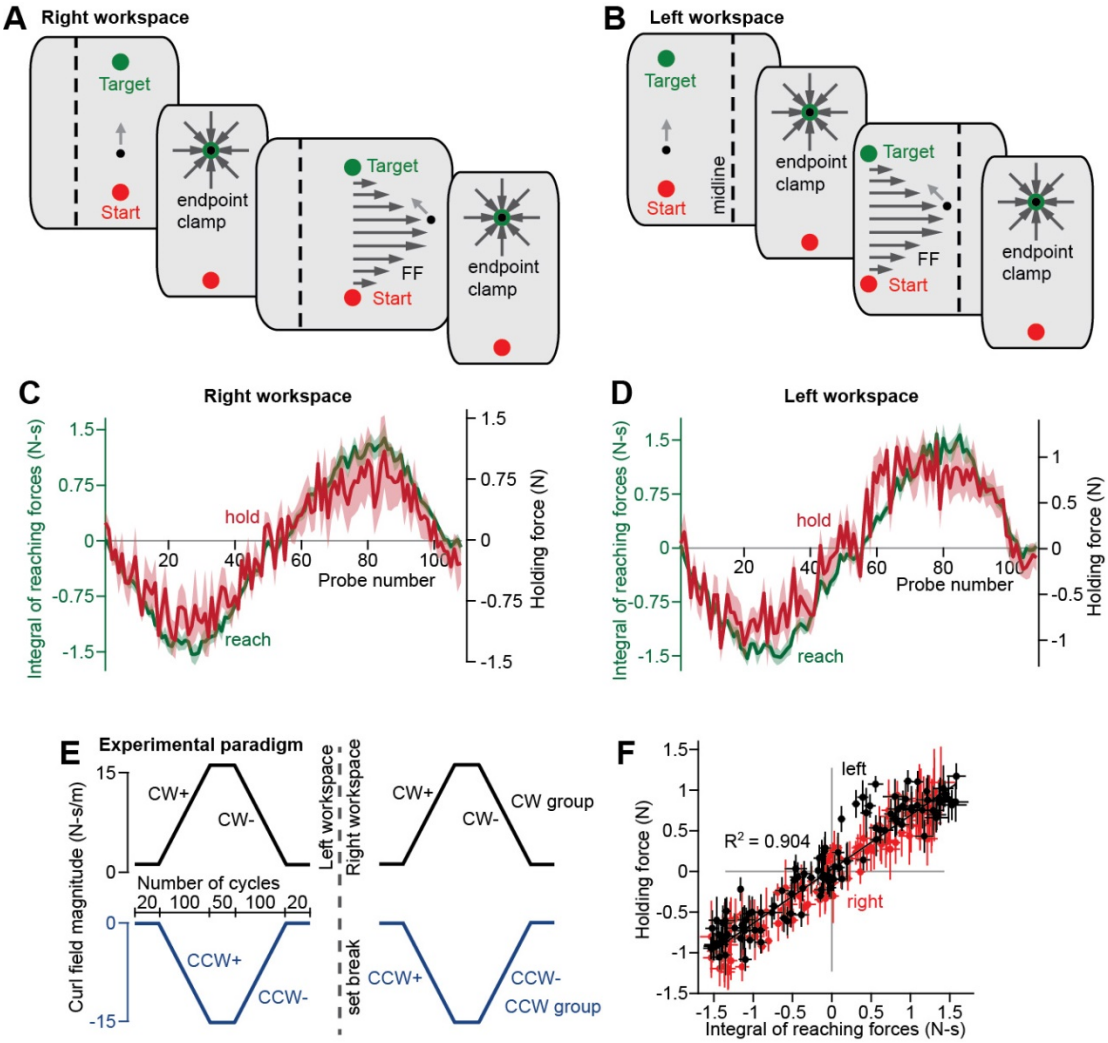

Figure S3. The relationship between reaching and holding forces is consistent across different arm postures. **A.** Participants reached from a starting position and target with their right arm. The movement was centered about 15 cm to the left of the midline of the body (dashed line). First, participants reached in a null field (top) but were then exposed to a CCW (Experiment 6, n=9) or CW (Experiment 7, n=8) force field that gradually increased in strength over many trials. At the end of each null field or perturbation trial, the hand was held in place during an extended holding period within an endpoint clamp. **B.** Participants also reached in a separate block in the same manner as in **A**, but with the arm centered 5 cm to the right of the midline of the body. **C.** Occasional probe trials were inserted to measure hand forces in the left workspace. The time-integral of the movement force is shown on each trial in green. The holding force is shown in red. The + and - indicators represent periods of time in which the force field was increasing and decreasing, respectively. Values are mean ± SEM across participants. **D.** Same as in C but for the right workspace. **E.** The experimental paradigm. Participants were gradually adapted and then deadapted to either a CCW (blue group, bottom) or CW (black group, top) force field. Participants were tested in two blocks separated by a short break. In one block, participants reached with the arm centered to the left of midline (shown to the left of dashed line). In the other block, participants reached with the arm centered to the right of midline (shown to the right of dashed line). The order was counterbalanced across participants. **F.** Shown is a scatterplot plot depicting the holding forces as a function of the moving force time-integrals. Points represent individual trials. The black line shows the linear regression. Values are mean ± SEM across participants.

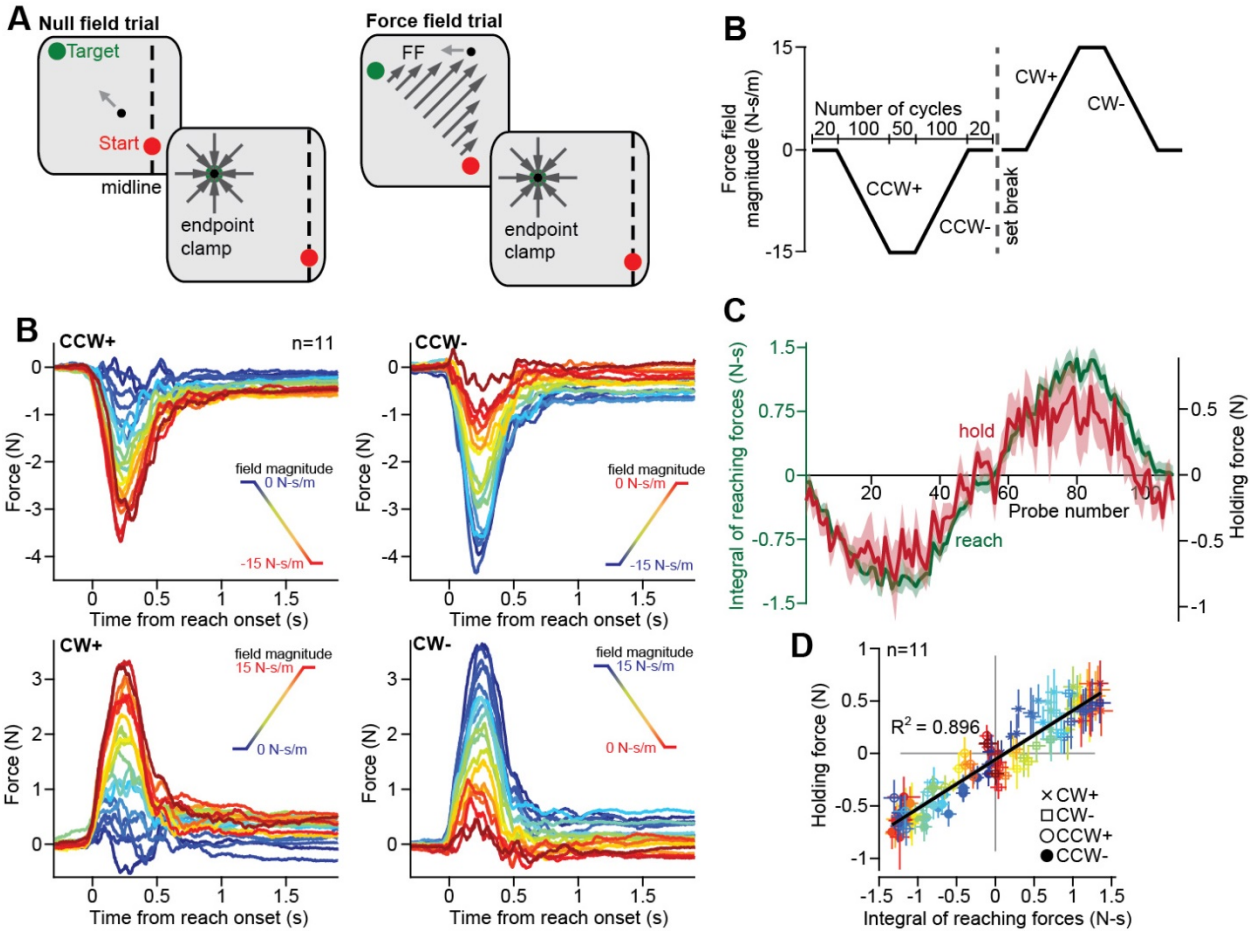

Figure S4. Moving and holding forces for reach movements of different angles. **A.** Participants (n=11, Experiment 5) reached between a start location and a target located at an oblique angle (135° with respect to the positive x-axis) from the midline of the body (represented by the dashed line). First, participants reached in a null field (left), but then were exposed to CW or CCW velocity-dependent force fields (right). At the end of null field and force field trials, the hand was held in place with an endpoint clamp. **B.** The force field schedule. Shown is the force field magnitude on each trial. The experiment was composed of cycles consisting of one outwards reach and backwards reach. On outwards trials were perturbed. The force field gradually increased to a constant value and then decreased. The + and - indicators represent periods of increasing and decreasing force field strength, respectively. CCW and CW force fields were applied in different blocks, separated by a short break. **C.** Occasionally, probe trials were inserted to measure forces during moving and holding. Shown are forces during different periods of the experiment (upper-left is increasing CCW field, upper-right is decreasing CCW field, lower-left is increasing CW field, and lower-right is decreasing CW field). Each trace is one trial. The color indicates the field magnitude in each respective period (blue is early and red is late). **D.** We calculated the time-integral of the moving forces (green) and the static holding force (zred) on each trial. Values are mean ± SEM across participants. **E.** Shown is a scatterplot of the holding forces as a function of the moving force time-integrals. Colors are consistent with trials in **C**. Marker shape depicts the period of the experiment (legend bottom-right). The solid line depicts the linear regression. Values are mean ± SEM across participants.

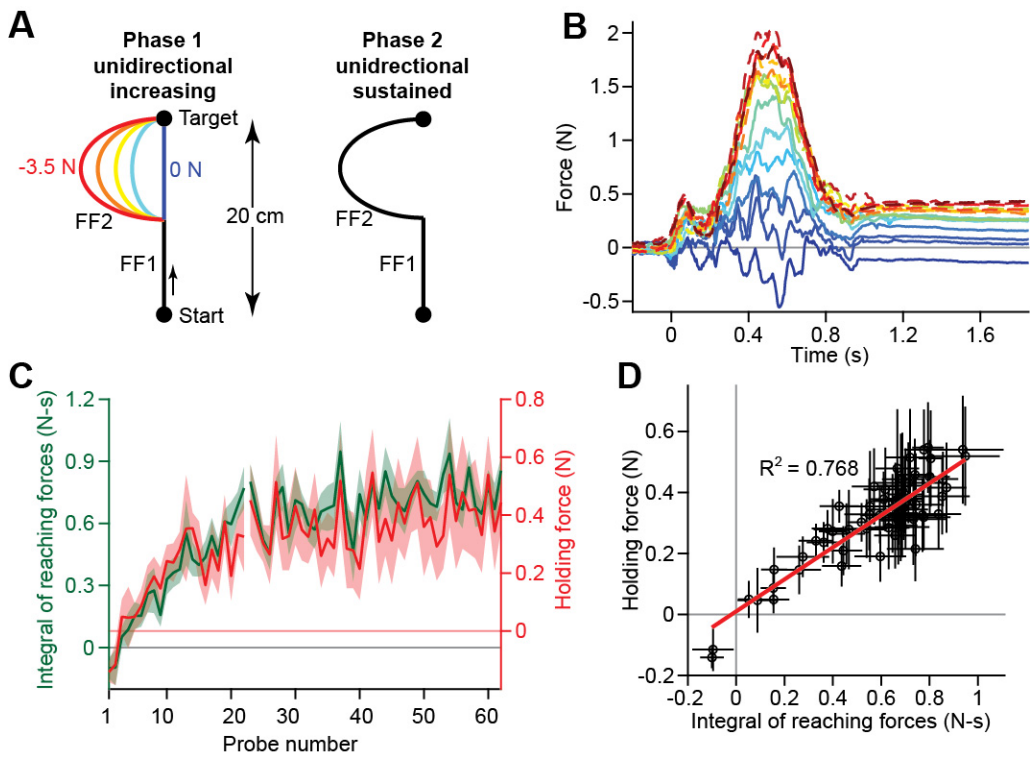

Figure S5. Holding forces in a sustained position-dependent force field. **A.** Participants (n=11, Experiment 14) reached between two locations separated by 20 cm. In the first phase of the experiment, participants were exposed to a position-dependent force field that was only active during the second half of the movement. The magnitude of the force field was increased over several trials (indicated by color, blue to red). In the second phase of the experiment, the perturbation was maintained for many trials. **B.** Occasionally, probe channel trials were inserted to measure hand forces. Shown are hand forces in time over sixteen different trial bins (roughly four probe trials in each bin). Blue to red denotes trials from early to late in the experiment. Solid traces correspond to Phase 1 and dashed traces correspond to Phase 2. **C.** We calculated the time-integral of moving force (green) and the final holding force (red) on each trial. Values are mean ± SEM across participants. **D.** Shown is a scatterplot of the holding forces as a function of the moving force time-integral. Each point is a trial. The solid red line depicts the linear regression. Values are mean ± SEM across participants.

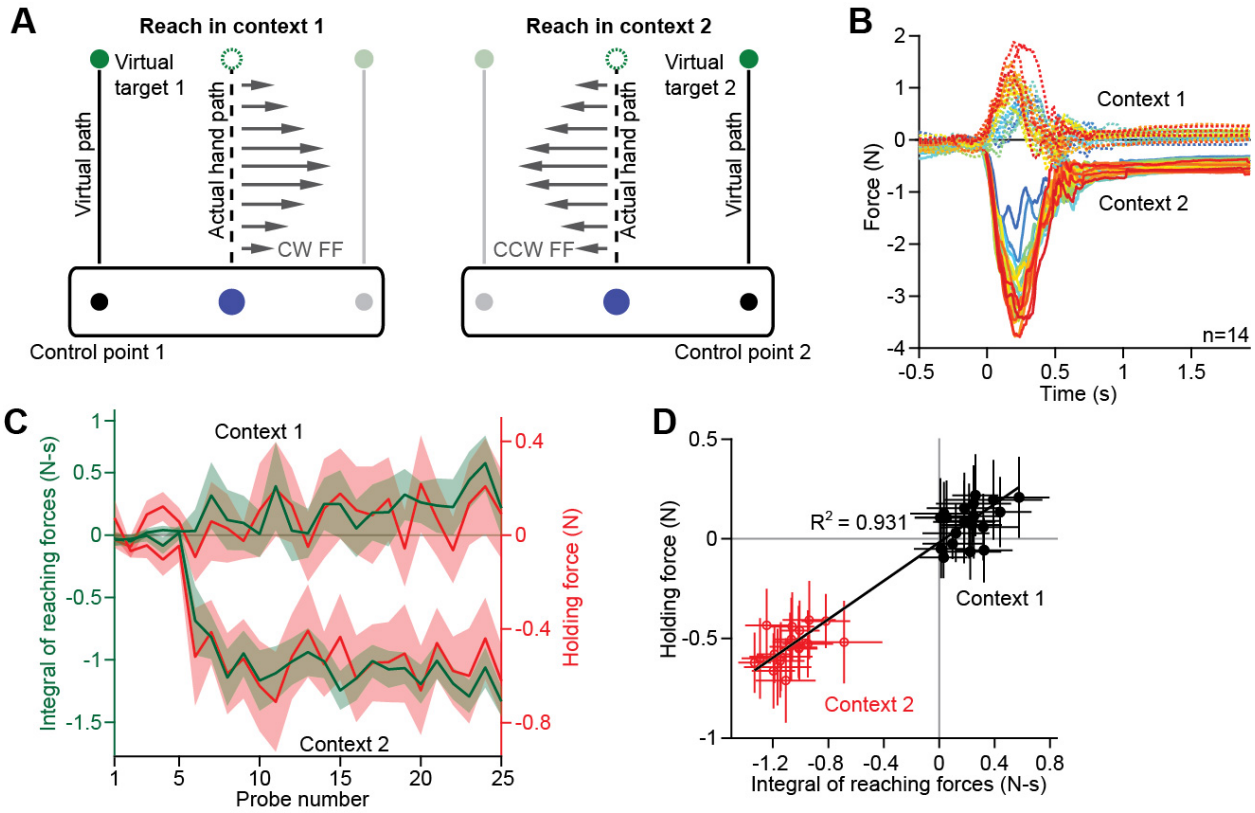

Figure S6. Moving and holding forces during exposure to dual force fields. **A.** Participants (n=14, Experiment 12) reached between two locations in space (the filled blue circle and the dashed central green circle). Rather than controlling a single cursor, participants were shown a rectangular tool that was centered on the hand. On some trials (context 1, left schematic), a control point appeared on the left of the tool and participants were told to move that point into a target that appeared to the left in the workspace. On other trials (context 2, right schematic), a control point appeared on the right of the tool and participants were told to move that point into a target that appeared to the right in the workspace. Critically, the physical starting position and target position were the same for all of the trials, simply the context changed. Participants were exposed to CW and CCW force fields. The orientation of the field was paired to the context (context 1 paired to CW field and context 2 paired to CCW field). **B.** Forces were measured on occasional channel probe trials. Here we show the forces produced on probes in context 1 (dashed lines) and context 2 (solid lines). Color indicates the trial number (blue is early in the experiment, red is late). **C.** We calculate the moving force time-integral (green) and holding force (red) for context 1 trials (top) and context 2 trials (bottom). Values are mean ± SEM across participants. **D.** Next we examined the scatterplot showing the holding force one each trial as a function of the moving force time-integral on that trial. Black points correspond to context 1 trials and red points correspond to context 2 trials. To reiterate, all of the trials start and end at the same location in space. The solid black line shows the linear regression. Values are mean ± SEM across participants.

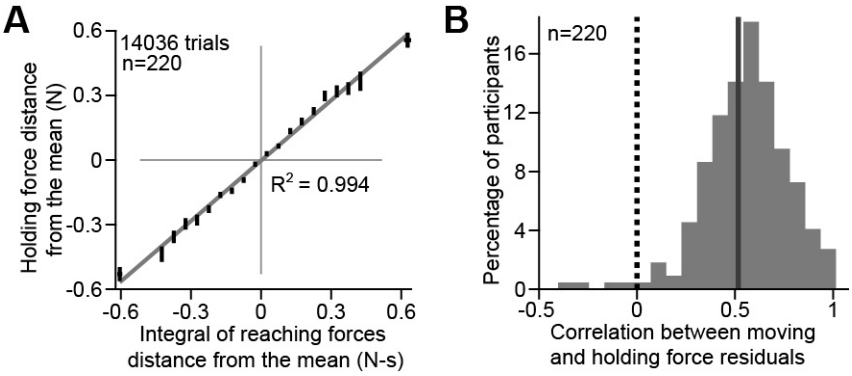

Figure S7. Moving and holding forces after removing perturbation-related changes. **A.** We considered all participants in experiments that used gradual force perturbations (n=220). In these paradigms, moving and holding forces trend linearly with the perturbation magnitude. To confirm that it is not the perturbation that drives the measured relationship between holding forces and moving force time-integrals, we detrended moving and holding forces from the perturbation. To do this, we linearly regressed the moving force time-integral and the holding force, separately, onto the trial-by-trial schedule of perturbations. Then on every trial, we subtracted the linear prediction for the moving force time-integral and holding force from the true measured values. This yields a moving force distribution and holding force distribution that have zero mean. We binned these perturbation-detrended residuals according to the moving force time-integral into 20 bins spaced evenly between moving force time-integrals of -0.5 N-s and 0.5 N-s. Trials that fell outside of this region were added to the first and last bins as appropriate. We placed moving and holding forces into these bins according, and then calculated the average moving force time-integral and holding force for each subject within each bin. Values are mean ± SEM across participants in each bin. The solid black line denotes a linear regression. **B.** We calculated the correlation coefficient for the perturbation-detrended moving force time-integrals and holding forces for each subject in the analysis in **A**. Shown is the histogram and mean (blue) of the within-subject moving and holding correlations.

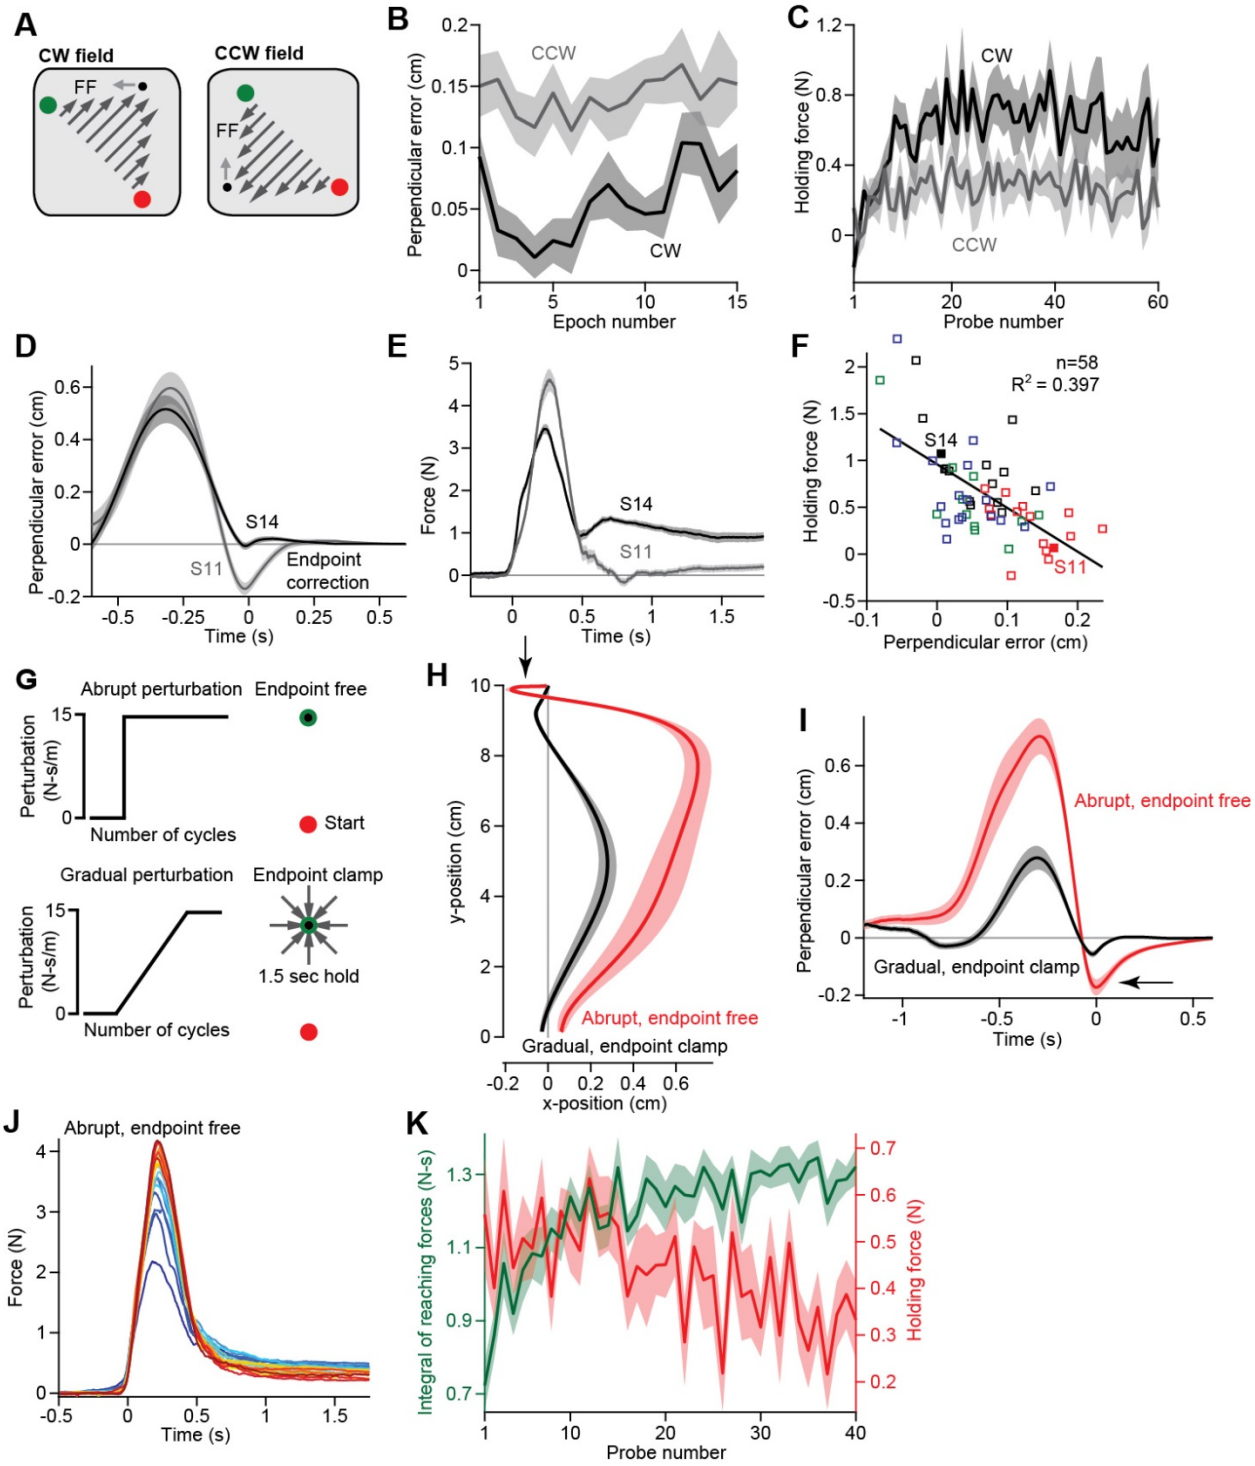

Figure S8. The postural controller adapts to positional errors. **A.** Participants reached between two oblique targets and were gradually adapted to either a CCW (n=15, Experiment 8) or CW (n=17, Experiment 9) velocity-dependent curl field. **B.** Compensatory forces led participants to make positional errors while attempting to decelerate and stop the hand on the target. To quantify these errors, we measured the maximum deviation of the hand from the straight-line path, in the late part of the movement after 80% of the reaching trajectory had elapsed. Positive values indicate errors in the direction of compensatory forces (opposite the robot forces). Due to the inertial properties of the robot and the arm, late reach termination errors were much larger in the CCW field than CW field. Values are mean ± SEM across participants. Each point represents a bin consisting of 16 outwards

movements. **C.** We measured holding forces on channel trials interspersed across the experiment. Larger reach termination errors in the CCW field led to smaller holding forces, in a process resembling error-based learning in the postural controller. Values are mean ± SEM across participants. **D.** Here we show the lateral deviation as for two example participants. Traces were aligned to the maximal holding-related error. Participant S11 experienced a much large error in position while attempting to stop the hand at the target than Participant S14. Values are mean ± SEM across trials. **E.** Same as in **D** but for moving and holding forces. Participant S14 had larger holding forces despite smaller moving forces than S11, likely because the postural system learned to eliminate holding forces due to holding-related errors. Values are mean ± SEM across trials. **F.** We collected participants across four experiments (colors denote different experiments) where perturbations were introduced gradually and at equal rates. The y-axis shows the holding force during peak force perturbations. The x-axis shows the perpendicular error on reaching trials that preceded the peak force perturbations. Positive values indicate errors in the direction of holding forces. Each point represents a different participant. The solid black line shows the linear regression across participants. **G.** Next, we designed a control experiment where reaching movements never terminated in an endpoint clamp (top, Endpoint free). To increase the size of late positional errors near reach termination, we abruptly perturbed participants (n=20, Experiment 10). **H.** Shown is the reaching trajectory for movements in the abrupt endpoint free experiment (red), and the standard gradual endpoint clamp experiment (black). Traces were constructed by interpolating along the x-axis based on even increments along the y-axis. Values are mean ± SEM across participants. The black arrow denotes the large late deviation of the hand during endpoint stabilization. **I.** The time-course of the lateral errors depicted in **H.** Values are mean ± SEM across participants. **J.** Moving and holding forces were measured regularly on channel trials. Here color denotes trial order (blue is early in the perturbation, red is late). Note the gradual reduction in holding force despite the increase in moving force. **K.** Shown are the moving force time-integral (black) and holding force (red) as a function of trial. Values are mean ± SEM across participants. The large positional error near reach termination shown in **H** clearly lead to a reduction in holding force, implying the existence of a postural mechanism for error-based learning.

Supplementary Table 1. Measures of impairment in stroke patients. Patients completed the Fugl-Meyer Assessment (FMA) and the Action Research Arm Test (ARAT), as well as strength testing in the shoulder and elbow. Strength measurements were repeated twice and the maximal force was recorded on each effort and then averaged across repetitions. Two separate raters scored the FMA and ARAT assessments, and scores were averaged across raters. Missing entries in table indicate that the patient was unable to perform the desired action. Patients were selected based on MRI or CT scans, and/or available radiologic reports. Scans and/or reports were corroborated to determine the level at which the white matter of the corticospinal tract (CST) was lesioned. Here we provide the level of the brain at which the white matter was lesioned.

ID	Age	Sex	Time since stroke	Handed-ness	Paretic arm	Elbow strength (N)		Shoulder strength (N)		FMA (/66)	ARAT (/57)	CST lesion location
						flexion	extension	adduction	abduction			
S001	80	M	2y	Right	Left	155.2 (P) 168.1 (NP)	82.3 (P) 73.0 (NP)	87.6 (P) 103.2 (NP)	115.7 (P) 108.5 (NP)	57.5	57	Right internal capsule
S002	51	M	6y	Right	Left	177.5 (P) 323.4 (NP)	114.5 (P) 204.6 (NP)	229.5 (P) 345.6 (NP)	152.6 (P) 209.1 (NP)	40	47.5	Right fronto-parietal white matter
S003	68	F	7y	Right	Right	142.8 (P) 225.5 (NP)	99.4 (P) 153.0 (NP)	128.1 (P) 120.5 (NP)	107.6 (P) 144.1 (NP)	34.5	19	Left corona radiata
S004	30	F	5y	Right	Left	116.1 (P) 157 (NP)	83.6 (P) 94.3 (NP)	95.2 (P) 119.7 (NP)	84.5 (P) 111.2 (NP)	55.5	43.5	Right precentral and postcentral gyri
S005	78	M	13mo	Right	Right	n/a	n/a	n/a	n/a	43.5	34	Left corona radiata
S007	54	F	2mo	Left	Right	140.6 (P) 144.6 (NP)	101.9 (P) 120.1 (NP)	107.6 (P) 109.4 (NP)	97 (P) 85.4 (NP)	63	57	Left corona radiata
S008	53	F	14mo	Right	Left	104.5 (P) 136.6 (NP)	65.4 (P) 148.1 (NP)	88.1 (P) 133 (NP)	91.6 (P) 125.9 (NP)	41	25	Right centrum semiovale
S010	70	M	5y	Right	Left	90.7 (P) 197.1 (NP)	68.1 (P) 161.9 (NP)	177.5 (P) 194.4 (NP)	82.7 (P) 144.1 (NP)	20	12	Posterior limb of right internal capsule
S011	43	F	20mo	Right	Right	116.5 (P) 130.8 (NP)	97 (P) 109 (NP)	86.7 (P) 91.6 (NP)	98.3 (P) 113.9 (NP)	64	57	Left corona radiata
S012	48	M	6y	Right	Left	145.9 (P) 358.1 (NP)	91.2 (P) 207.3 (NP)	178.8 (P) 292.2 (NP)	97.4 (P) 226 (NP)	18.5	6.5	Right corona radiata
S013	68	M	9y	Right	Left	217.1 (P) 355 (NP)	112.5 (P) 236.2 (NP)	220.2 (P) 282 (NP)	130.3 (P) 174.4 (NP)	14	8	Right internal capsule
S014	45	F	16mo	Right	Left	29.8 (P) 114.8 (NP)	33.4 (P) 113 (NP)	48.9 (P) 120.5 (NP)	46.7 (P) 87.2 (NP)	40	39.5	Right precentral and postcentral gyri
S015	64	F	10y	Right	Left	149.9 (P) 138.78(NP)	44 (P) 140.1 (NP)	73.4 (P) 100.5 (NP)	48.9 (P) 84.1 (NP)	22	4.5	Right corona radiata