

# Using Haptic fMRI to Enable Interactive Motor Neuroimaging Experiments

Samir Menon, Hari Ganti and Oussama Khatib

**Abstract** Combining haptics with functional magnetic resonance imaging (Haptic fMRI) has enabled complex motor neuroimaging experiments that non-invasively map real-world motor tasks on to the human brain. The technique's resolution, fidelity and susceptibility to scanning artifacts, however, have not yet been estimated in a quantitative manner. Here, we demonstrate that unconstrained three degree-of-freedom Haptic fMRI experiments can reliably activate brain regions involved in planning, motor control, haptic perception, and vision. We show that associated neural measurements are reliable, heterogeneous at the millimeter scale, and free from measurable artifacts, and that their anatomical localization is consistent with past neuroscience experiments. In addition, we demonstrate the feasibility of using electromagnetic actuation in Haptic fMRI interfaces to apply high fidelity open-loop three-axis haptic forces (0.5–2N; square or 0.1–65Hz sine waveforms) while maintaining negligible temporal noise in pre-motor, motor, somatosensory, and visual cortex (<1 % of signal). Our results show that Haptic fMRI is a robust and reliable technique for characterizing the human brain's motor controller.

**Keywords** Haptic fMRI · MRI-compatible robotics · Motor control · Neuroscience · Neuroimaging

## 1 Introduction

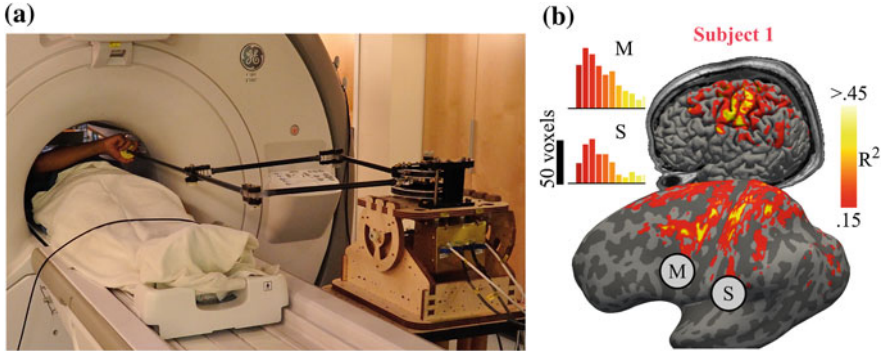
Human motor neuroimaging experiments presently involve either uncontrolled limb motions or simple visuo-motor tasks within a constrained workspace, which has limited our understanding of motion and force control in the brain. Haptic fMRI promises to overcome this limitation by imaging the brain with fMRI [1, 2] while

---

S. Menon (✉) · H. Ganti · O. Khatib  
A.I. Lab, Stanford University, Stanford, CA 94305, USA  
e-mail: smenon@stanford.edu

H. Ganti  
e-mail: hganti@stanford.edu

O. Khatib  
e-mail: ok@cs.stanford.edu



**Fig. 1 Mapping motor responses in cortex: a** HFI, a three degree-of-freedom fMRI-compatible haptic interface. **b** Reliable responses in motor (*M*) and somatosensory (*S*) cortex during a task that involves motor planning, reaching, and feedback trajectory control. Voxels with >15 % variance explained ( $R^2 > 0.15$ ; see Appendix) are shown on the brain’s surface and on an inflated mesh. Distribution histograms for motor and somatosensory cortex are inset

subjects simultaneously perform unconstrained three dimensional tasks in a virtual haptic simulation environment [3]. Engineering fMRI-compatible haptic interfaces for neuroscience experiments, however, requires achieving high-fidelity force control, backdrivability, natural operator-motion statistics [4], and uniform inertial properties [5] across a large three-dimensional workspace. Moreover, the devices must operate in a robust manner in large magnetic fields, elicit reliable subject motions and neural activation, and avoid injecting noise into fMRI measurements.

In this paper, we demonstrate that Haptic fMRI experiments elicit reliable neural activation for motor planning, unconstrained three-dimensional motions, and visual and force perception (Fig. 1). All neural activation patterns are localized to expected anatomical regions based on past experiments [6–10]. Our results use two experiments that independently contrast either planning and motion, or vision and force perception. To ensure accurate force rendering, we demonstrate that our electromagnetically actuated haptic interface can achieve high fidelity open-loop three-axis force control inside the MRI scanner while limiting performance decreases (<2–3 %) to levels below human perception [11]. Finally, we show that our experiments achieve temporal noise levels similar to the scanner’s baseline (median noise-to-signal = 0.85 %) in pre-motor, motor, somatosensory, parietal, and visual cortex. Combined with past results that demonstrate HFI’s haptic transparency, high fidelity motion monitoring, uniform inertial properties, reliable subject motions, and lack of confounding artifacts [12–14], this paper establishes Haptic fMRI as an effective motor neuroscience technique for three dimensional manipulation tasks.

## 2 Related Work and Technical Approach

Past research on fMRI-compatible haptic interfaces for motor control experiments has primarily focused on avoiding electromagnetic actuation, which simplified fMRI compatibility but limited device transparency, backdrivability, or degrees-of-freedom. MRI-compatible actuation mechanisms developed in the past include electro-active polymers [15], pneumatics [16, 17], hydraulics [18, 19], and cables driven by remote actuators [20]. An alternative approach modified an electromagnetically actuated PHANTOM device [21, 22] for MRI by using RF shields and a long carbon fiber extension rod attached to a linear rail (to improve stiffness); the device, however, can not span the entire MRI workspace.

To overcome the limitations of earlier designs, we designed and built an fMRI-compatible haptic interface, HFI (see Fig. 1a), with three degrees-of-freedom that span the entire MRI workspace (also see [12]). We used electromagnetic actuators for high frequency force control, low gear ratios (20 in-plane, 30 vertical axis) to ensure backdrivability, lightweight composites to ensure haptic transparency, and a mechanical structure that places actuators beyond the 400 Gauss static-field to limit magnetic interference in the motors. Furthermore, we mitigated eddy currents in motors induced by time-varying magnetic fields during fMRI scans by limiting motor displacement to  $<2$  cm. The device's low and uniform inertia ( $x = 0.81$ ,  $y = 0.22$ ,  $z = 0.27$  kg; condition number 3.81) help minimize friction (0.01–0.3 N), make it transparent, and elicit stereotypical motion patterns across subjects [13].

HFI's ability to support Haptic fMRI experiments that involve unconstrained arm motions promises to enable complex motor studies capable of testing how generative motor coordination models [23–25] map on to the human brain. Moreover, allowing non-invasive human motor neuroscience studies promises to help cross-validate decades of past animal neuroscience experiments, which have studied how the brain represents forces [26], movement direction [27–30], spatial gradients [31, 32], muscular activation [33–35], arm orientation [34, 36], limb joints [37], reach distance [38], motor planning [39], and hand position, velocity and acceleration [40].


## 3 Experimental Results

We focused on testing whether Haptic fMRI experiments conducted with our haptic interface, HFI, could enable controlled neuroscience experiments that localize neural activation patterns for planning, motion, vision, and force perception to expected brain regions. Our goal was to demonstrate that neural activation patterns obtained with Haptic fMRI match past research. This is important because the technique's novelty and reliance on indirect fMRI neural measurements (blood oxygenation response; see Appendix for details) have raised questions about the limits of spatio-temporal resolution and the potential for inducing complex noise patterns. Moreover, past attempts to use electromagnetic actuation with fMRI have resulted in unac-

ceptable noise [18]—potentially because of inadequate radiofrequency shielding or unsuitable scanning hardware and protocols. In contrast, our earlier experiments found that HFI’s shields limited noise due to electromagnetic actuation during uncontrolled haptic experiments [12]. Our past work, however, did not test the shields’ effectiveness while applying arbitrary force magnitudes, directions, or frequencies. It also did not test HFI’s ability to generate forces at the end-effector while operating in the MRI scanner room’s large magnetic fields.

### 3.1 Neural Activation for Planning and Movement

We demonstrated HFI’s ability to elicit reliable neural activation in the brain using a motor task that involves planning and reaching to different locations of space, followed by visuo-motor trajectory tracking in two directions centered around each reach destination (Figs. 1b, and 2; see video for protocol). The experiment’s goal was to simultaneously activate as many motor-related regions as possible and test whether the resulting anatomical localization matched past research findings (see [6] for a review). As expected, the experiment evoked large reliable neural activation in



**Fig. 2 Neural activation for planning and movement:** Motion planning and reaching elicit reliable responses in pre-motor and motor cortex respectively. An impulse response model helped segregate neural activation during planning and movement into separate time series components (see Appendix for details). Time series fits are shown for two exemplar voxels for the left reaching condition; responses to the *center* and *right* conditions were similar. Heterogeneous time series in voxels that are a few centimeters apart indicate that neural activation signals dominate any potential motion-induced low spatial-frequency artifacts

primary motor and somatosensory cortex, which have been anatomically connected to low-level motor control and sensory perception. We also found activation in pre-motor, supplementary motor, and parietal cortex, which are involved in higher-level motor and visual processing.

In addition to mapping neural activation reliability across the brain's regions, we used a finite impulse response model (see Appendix for details) to estimate time series responses in individual  $2.5 \times 2.5 \times 2.5 \text{ mm}^3$  voxels for each task condition (see Fig. 2). As expected, neural activation in the pre-motor region increased during the motor planning period before actual movement, and decreased after movement onset. In contrast, primary motor cortex voxels typically showed a decrease in activation during planning and an increase during movement. Movement-related activation for visually-guided trajectory tracking (not shown) were also localized to the motor regions and were associated with a longer time series activation than reaching, as expected (trajectory tracking time was 8 s; reaching was 5 s).

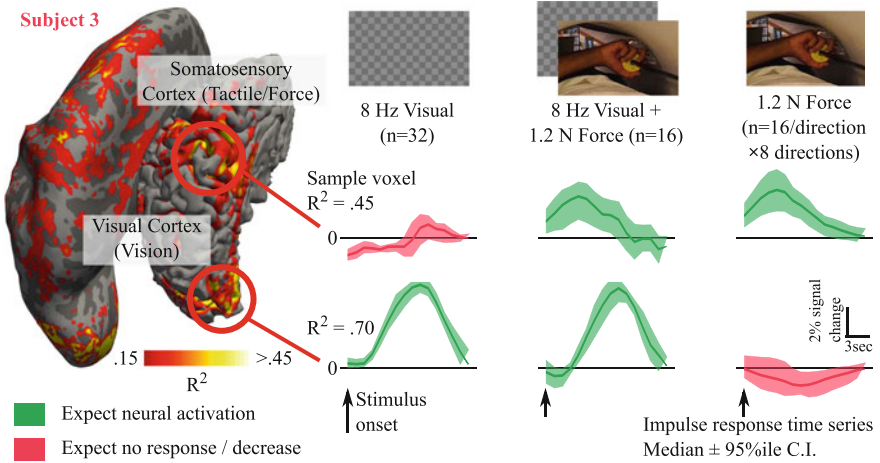
This experiment demonstrates that Haptic fMRI's spatial resolution can reliably segregate time series responses across brain regions that are a few centimeters apart. To obtain our results, we avoided spatial smoothing as it reduces noise but correlates nearby neural activation patterns. As an added advantage, our results also implicitly indicate the lack of motion or haptic artifacts, which create low spatial-frequency task-correlated noise and can make time series responses and variability look similar across nearby brain regions. We noted that the number of reliable voxels in high-level cortical regions was lower than in the low-level regions for subject two when compared to subject one. As such, we plotted time series estimates for the less reliable subject.

### ***3.2 Neural Activation for Vision and Haptic Perception***

Having tested neural activation for planning and movement, we proceeded to determine whether we could localize activation for visual and force perception in visual and somatosensory cortex, respectively. We did so using an experiment design that exposed subjects to a high contrast checkerboard visual stimulus, or a haptic force perception stimulus (see Appendix for details). Both stimuli were designed to be simple, and were each expected to reliably activate their own corresponding (low-level) sensory region without activating the other's. As such, monitoring changes in visual cortex activation while our haptic interface applied forces would reveal any radiofrequency noise due to motor operation.

Our experiment elicited reliable neural activation in visual, somatosensory, and motor cortex (Fig. 3). While we found little pre-motor activation (there was no planning task), we did find activation in parietal (between vision and motor) and supplementary (inner brain) cortex, which have been associated with object perception and motor control, respectively.

As expected, somatosensory cortex did not respond to the visual stimulus, and early visual cortex did not respond to the force stimulus. In addition, the presence



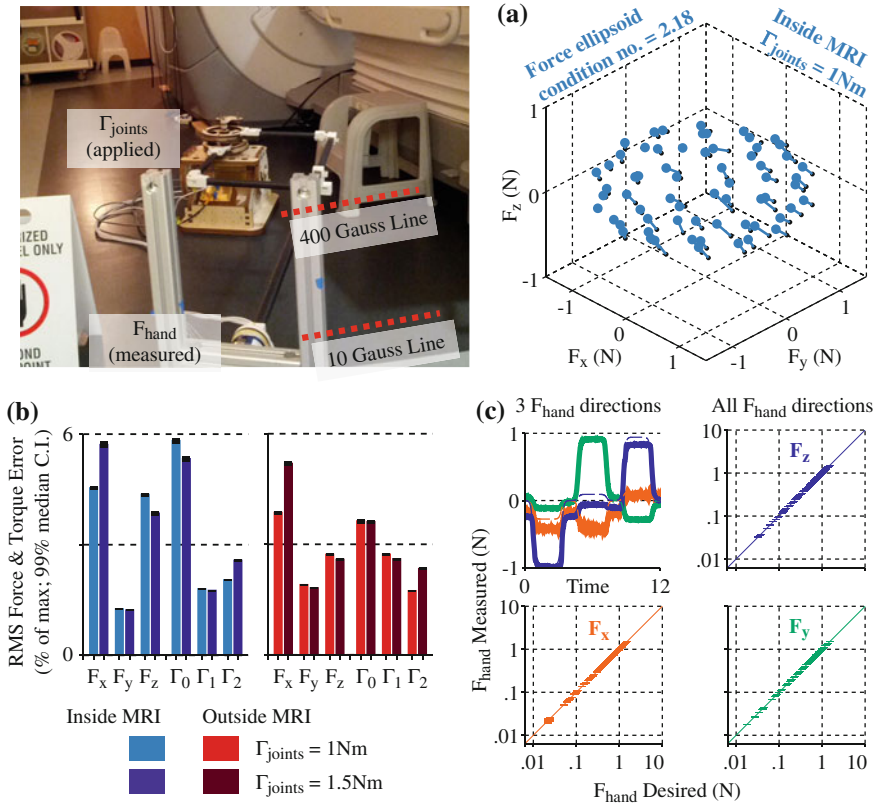
**Fig. 3 Neural activation for visual and tactile perception:** Vision and haptic-force perception elicit reliable responses in visual and somatosensory cortex, respectively. Impulse response time series for exemplar voxels are similar for simultaneous vision and force stimuli, as well as for each individual stimulus. The lack of spurious activation in visual regions during the force-only stimulus indicate that neural signals dominate any potential low temporal-frequency task-correlated noise due to HFI’s electromagnetic actuation

of reliable activation in visual cortex (high  $R^2$ ; see exemplar voxel’s time series) indicates low radiofrequency interference when motors actively apply haptic forces.

It is noteworthy that visual cortex responses to the checkerboard are more reliable than somatosensory cortex responses to force perception. There are a few reasons for this. First, the high-contrast checkerboard is a well tested visual stimulus that is known to strongly activate early visual cortex [41]. Since no past fMRI studies have estimated time series activation for multi-axis force perception while the arm is unconstrained, our experiment was likely suboptimal and may be greatly improved in the future. Second, our force perception stimulus was oriented along one of eight random directions (see Appendix for details). We did so to probe visual cortex while actuating multiple motors in HFI with a variety of currents. While this increased the chance that HFI would generate interfering radiofrequency noise, it also led to less-reliable neural activation patterns for force perception.

### 3.3 Force Generation in the MRI Scanner

Having elicited expected neural activation for eight force directions in our perception experiment, we proceeded to test HFI’s ability to apply forces in any direction (Fig. 4). Our force sensing rig, however, was not MRI-compatible and had to be placed at the 10 Gauss line, far away from the MRI machine. As such, we placed HFI at the 400 Gauss magnetic field line—its typical position—but inverted its direction to point



**Fig. 4 Force generation in an MRI scanner:** HFI reliably produced desired forces at the end-effector while operating at an MRI scanner’s 400 Gauss line—its typical position during Haptic fMRI experiments. **a** Measured end-effector forces (light blue) closely track desired forces (black). Directions were uniformly sampled on a sphere in joint-torque space. **b** Root mean square force tracking errors for end-effector forces and joint torques for two applied joint torque levels (1, 1.5 nm; light and dark bars) shown for measurements made inside (blue) and outside (red) the MRI scan room. Noise levels increase marginally in the MRI scan room. **c** Raw end-effector force time series shown for three exemplar directions (top-left). Measured forces (solid lines) closely track desired forces (dashed). Breaking forces along all directions into three individual axes (x:bl, y:br, z:tr) demonstrates reliable force generation across three log decades. Error-bars are vertical (and small)

away from the scanner bore instead of into it. We then attached the end-effector to the force sensor and directly measured generated forces.

To test HFI’s force generation, we uniformly sampled joint torque vectors on a unit sphere, applied each torque vector sequentially to the motors, and measured the resulting end-effector forces. HFI’s kinematics transform the (unit) joint torque sphere to an end-effector force ellipsoid characterized by the Jacobian’s eigenvectors. Our force measurements (see Fig. 4a) matched the theoretical Jacobian [12], and show that HFI’s force generation is close to isotropic (condition number, 2.18).

The applied forces and torques have a negligible root mean square error (0.06–0.3 N; see Fig. 4b), which is below the human detection threshold [11]. Moreover, the force and torque errors are similar inside and outside the MRI machine. This indicates that HFI’s motors are sufficiently far from the MRI machine and their performance is not unaffected by magnetic interference.

Finally, the measured end-effector forces match theoretical predictions over three decades on a log scale, indicating a high dynamic range (see Fig. 4c). The force generation becomes less reliable at very low forces ( $<0.08$  N) due to a combination of sensor noise and device friction.

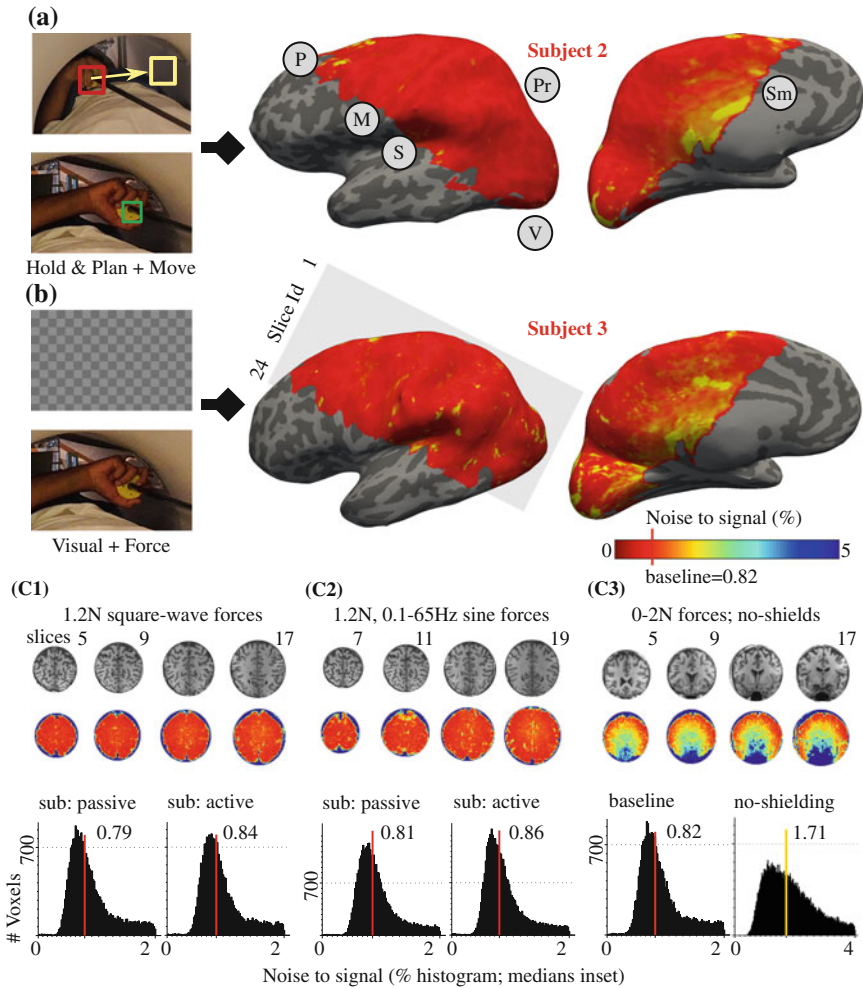
### ***3.4 Noise During Haptic Force Transmission***

Our final experiment quantitatively characterized the temporal noise in our fMRI measurements for the planning and moving tasks, as well as visual and force perception (Fig. 5). Since our force perception task involved square wave forces, which might not represent all possible haptic force interactions, we also added a noise test where HFI applied sine wave forces across a large frequency range (0.1–65 Hz). As a reference, we also compared temporal noise measurements across these conditions against noise levels with HFI’s radiofrequency shields removed, and against MRI scanner baseline noise levels with no device and passive subjects.

We found limited noise in the haptic planning and movement task (Fig. 5a) across our regions of interest in the cortex. Temporal noise levels increased marginally in visual cortex for the vision and force perception task, potentially due to a random combination of higher head motion (up to 1.25 mm), scanner calibration drift, or interfering motor radiofrequency noise (Fig. 5b). This, however, was one of the worst runs of the experiment; there are usually scanner-calibration dependent changes across runs. The noise levels were much lower in other runs (see a second run in Fig. 5c1; head motion  $<0.4$  mm). We did note a large systematic increase in noise near the center of the brain, which could be attributed to partial volume artifacts caused by a proximity to ventricles and other non-brain regions. Moreover, we used a 32 channel head coil, which surrounds the head and provides higher signal-to-noise near the surface of the skull.

Our final noise testing condition involved continuously applying a mixture of sine wave forces with HFI. While potentially unrealistic—applying forces non-stop will limit any statistical analysis—it does quantify worst-case noise patterns (Fig. 5c2). Temporal noise levels even in this extreme case were not significantly different from baseline, and are likely to be dominated by common fMRI artifacts related to head- and hand-motion, partial voluming, poor scanner calibration, or very high spatial resolution scanning. With shields removed, however, HFI’s radiofrequency interference reaches unacceptable levels (Fig. 5c3).





**Fig. 5 Low noise during HFI operation:** Temporal noise (high-frequency) for brain regions measured as a percent of the median time series signal. **a** fMRI measurements for the planning and movement task had little noise (<1 %) in pre-motor (*P*), motor (*M*), somatosensory (*S*), visual (*V*), parietal (*Pr*), and supplementary motor (*Sm*) cortex. **b** The force perception task induced greater temporal noise in the inner brain regions (slices match C.1). **c** Temporal noise histograms for three task types while the subject actively operated HFI or was passive (an operator held HFI). **c.1** Brain slices from a second run of the vision and force perception task. **c.2** HFI applied sine wave forces of 0.1–35 Hz at the end-effector instead of square waves. **c.3** Subjects randomly interacted with a haptic simulation with HFI’s radiofrequency shields removed. Note the dramatic increase in temporal noise (*bottom right*)

## 4 Experimental Insights

We demonstrate that it is feasible to conduct three degree-of-freedom Haptic fMRI experiments involving force perception as well as unconstrained motion. Our experiments used an electromagnetically actuated interface, HFI, whose design helps it achieve high fidelity force generation in the MRI scan room, and whose radiofrequency shields prevent imaging artifacts. Our experiments used simple stimuli and elicited neural activation patterns consistent with past research. This demonstrates that we avoid the numerous artifacts possible in fMRI experiments, and sets the stage for detailed haptic studies to map complex motor coordination patterns, tactile perception, and visuo-motor integration.

The primary challenge facing Haptic fMRI is to now engineer a transparent six degree-of-freedom haptic interface that is cost-effective yet achieves robust motion tracking and force generation over the long-term. For instance, we have used HFI for more than eighty-five Haptic fMRI scans spread across sixteen sessions over more than 1 year. Moreover, the device is compact, costs less than ten thousand dollars to fabricate, and has a short setup time of 10–15 min. Thus, HFI minimizes both fixed fabrication costs as well as recurring MRI scanning costs. A six degree-of-freedom successor must meet or exceed HFI's metrics.

A second challenge is to develop experiment designs that leverage fMRI's ability to simultaneously image multiple brain regions at a high spatial resolution, while accommodating its slow temporal responses. Ideal experiments would ensure that overlapping sensory or motor task conditions elicit neural activation in anatomically distinct brain regions. Achieving this, however, requires improving upon past experiments that predict a fractured somatotopic organization where functionally related limbs are anatomically co-localized [7, 42]. Whole body control frameworks [24, 25] applied to subject-specific musculoskeletal models [43] can provide a theoretical basis to explain this complex organization; testing their predictions with Haptic fMRI experiments is an immediate future goal.

## Appendix

### *MRI Protocol*

All fMRI scans were conducted at Stanford University's Center for Cognitive and Neurobiological Imaging on a GE Discovery MR750 3T MRI scanner, with a thirty-two channel Nova Medical head coil. The scan protocol was gradient echo EPI with a 16 cm field of view sampled at a  $64 \times 64$  resolution ( $2.5 \times 2.5 \times 2.5 \text{ mm}^3$  voxels), a 1.57 s repetition time, a 28 ms echo time, and a  $72^\circ$  flip angle. Each scan run was preceded by 2nd-order polynomial shimming and was sandwiched by spiral fieldmap scans ( $2.5 \times 2.5 \times 5 \text{ mm}^3$  voxels). Fieldmap scans were conducted within 10 s of each

scan run's start and end. After scanning, the fMRI images were slice time corrected, motion corrected (SPM [44]), spatially undistorted using fieldmaps, and analyzed to compute temporal noise-to-signal.

## *fMRI Analysis*

Temporal noise-to-signal computations used the median neural response distribution obtained by regressing out a line from each voxel's time series, computing the absolute value of the difference between successive time points, computing the median of these absolute differences, dividing the result by the mean of the original time series, and then multiplying by 100. Cortex segmentation used Freesurfer's Desikan-Killiany atlas [45]. Surface registration was done using Freesurfer, and all surface images were plotted using Freeview. Freeview smoothed the surface plots while rendering (2 steps).

## *Estimating fMRI Impulse Response Time Series and $R^2$*

fMRI measures changes in blood oxygenation induced by neural metabolic activity [1, 2], which have a slower time course than neural computation and persist long after sensory stimuli and motor tasks terminate. Such persistent responses cause raw fMRI measurements to overlap in experiments where consecutive task conditions are not be separated by large time-intervals. Separating task conditions by large time-intervals, however, makes fMRI runs very long, which can induce a variety of unwanted artifacts related to MRI scanner calibration drift, neural adaptation, or subject attention lapses, microsleep and exhaustion. Instead, we optimized our experiments to ensure reliable motor task execution [13], which caused fMRI measurements for different task conditions to overlap.

We segregated neural activation for individual tasks using a finite impulse response (FIR) model (implemented using GLMdenoise [46]). The FIR model works by associating each task type with a unique time course and segregates time courses while assuming that overlapping responses sum linearly. fMRI signal linearity, however, is an active area of research [1, 41, 47]. As such, we randomized inter-task delays and randomly ordered tasks, which made the model's time series match anatomical expectations based on past research (see Figs. 2 and 3; read [6] for an overview). When tasks were closely spaced in time, as with planning and motion, this method was noisy. The parts of planning that overlap with motion are thus less reliable and the confidence interval for the planning time series estimates is larger after motion starts (but still above zero; see Fig. 1).

We computed 95 % confidence intervals by bootstrapping [48] runs (400 bootstraps), fitting FIR models to each, and taking the median percentile estimates across

the estimated bootstrap time series. Finally, we computed  $R^2$  values for each voxel by comparing the time series variance with the variance after regressing out median FIR model estimates.

### ***Data Collection Protocols***

See [13] for precise specifications of the motion protocol. Subjects executed one practice run inside the MRI scanner, and then executed at least eight scan runs (S1, 10; S2, 8). Each run was 630 s long.

The force and visual perception experiment protocol involved fixed duration stimuli instances with visual, motor, or visual and motor sensory input. The experiment was divided into runs, and each run was divided into blocks. During each block, the subject started with their hand at rest. Next, they were instructed to move their hand into free space. After a randomized delay period of 3–5 s, the subject experienced two randomly selected stimuli instances. Each stimulus instance was 3–5 s long and was separated from the other by a randomized delay 3–5 s. Finally, the subject was required to rest their hands for a random time interval (4–20 s), and then restarted the process. The subject executed four scan runs with multiple blocks. Each run was 459 s long.

Force magnitudes were set to evenly spaced directions along the x-y plane, with a magnitude of 1.2 N. The force vectors used were (1.2, 0.0), (0.0, 1.2), (−1.2, 0.0), (0.0, −1.2), (0.85, 0.0), (0.0, 0.85), (−0.85, 0.0), and (0.0, −0.85).

### ***Haptic and Force Measurement Details***

Haptic experiments were conducted with Haptic fMRI Interface [12], a three degree-of-freedom fMRI-compatible device. All motions were right handed, and the haptic control rate was 350 Hz. The reaching task spanned the entire workspace (see [13] for more details), but avoided arm motion artifacts [14].

Visual stimuli were displayed on a 30 in. diagonal (76 cm, 16 : 10 aspect ratio) flat panel display custom built by Resonance Technology. Subjects viewed visual stimuli through a dual-mirror setup. The visual distance from screen to mirror-2 is 184.4 cm, from mirror-2 to mirror-1 is 6 cm, and from mirror-1 to the eye is about 15 cm, for a total viewing distance of about 205 cm. The visual field of view is about 30°, making each visual checkerboard square span about one and a half degrees of the visual field. The display has a native resolution of 2560 × 1600 but stimuli were displayed at 1280 × 800. The display has a 7 ms temporal response, and 10-bit color rendering. The maximum luminance of the display is 329 cd/m<sup>2</sup> (red is 88, green is 117, and blue is 124 cd/m<sup>2</sup>).

Forces were measured using a JR3 85M35A-U560 63N4S force sensor. The raw sensor data was sampled at 1 KHz, resampled to match HFI's control rate, and was finally filtered using a 75 Hz low pass filter to remove high frequency sensor noise.

## *Human Subjects*

Subjects were healthy right-handed males with no history of motor disorders: S1, 29y, 185lb, 5'9"; S2, 19y, 170lb, 6'2"; S3, 21y, 160lb, 5'8". Informed consent was obtained in advance on a protocol approved by the Institutional Review Board (IRB) at Stanford University.

## **References**

1. Logothetis, N.K., Wandell, B.A.: Interpreting the bold signal. *Annu. Rev. Physiol.* **66**, 735–769 (2004)
2. Logothetis, N.K.: What we can do and what we cannot do with fMRI. *Nature* **453**, 869–878 (2008)
3. Hale, K.S., Stanney, K.M.: Deriving haptic design guidelines from human physiological, psychophysical, and neurological foundations. *IEEE Comput. Graph. Appl.* **24**, 33–39 (2004)
4. Kostic, M., Popovic, D., Popovic, M.: Influence of planar manipulandum to the hand trajectory during point to point movement. In: *IEEE International Conference on Rehabilitation Robotics*, pp. 1–4 (2011)
5. Khatib, O.: Inertial properties in robotic manipulation: an object-level framework. *Int. J. Robot. Res.* **14**, 19–36 (1995)
6. Rosenbaum, D.: *Human Motor Control*. Academic Press, San Diego (2009)
7. Meier, J.D., Aflalo, T.N., Kastner, S., Graziano, M.S.A.: Complex organization of human primary motor cortex: a high-resolution fMRI study. *J. Neurophysiol.* **100**, 1800–1812 (2008)
8. Churchland, M.M., Yu, B.M., Ryu, S.I., Santhanam, G., Shenoy, K.V.: Neural variability in premotor cortex provides a signature of motor preparation. *J. Neurosci.* **26**, 3697–3712 (2006)
9. Aflalo, T.N., Graziano, M.S.A.: Relationship between unconstrained arm movements and single-neuron firing in the macaque motor cortex. *J. Neurosci.* **27**, 2760–2780 (2007)
10. Kay, K.N., Winawer, J., Mezer, A., Wandell, B.A.: Compressive spatial summation in human visual cortex. *J. Neurophysiol.* **110**, 481–494 (2013)
11. Mesa-Munera, E., Ramirez-Salazar, J., Boulanger, P., Bischof, W.F., Branch, J.W.: Estimation of vibration and force stimulus thresholds for haptic guidance in MIS training. *Rev. Ing. Biomed.* **5**, 17–22 (2012)
12. Menon, S., Brantner, G., Aholt, C., Kay, K., Khatib, O.: Haptic fMRI: combining functional neuroimaging with haptics for studying the brain's motor control representation. In: *Proceedings of the 13th Annual Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 4137–4142 (2013)
13. Menon, S., Yu, M., Kay, K., Khatib, O.: Haptic fMRI: accurately estimating neural responses in motor, pre-motor, and somatosensory cortex during complex motor tasks. In: *Proceedings of the 14th Annual Conference of the IEEE Engineering in Medicine and Biology Society* (2014)
14. Menon, S., Quigley, P., Yu, M., Khatib, O.: Haptic fMRI: using classification to quantify task-correlated noise during goal-directed reaching motions. In: *Proceedings of the 14th Annual Conference of the IEEE Engineering in Medicine and Biology Society* (2014)

15. Vogan, J., Wingert, A., Plante, J., Dubowsky, S., Hafez, M., Kacher, D.: Manipulation in MRI devices using electrostrictive polymer actuators: with an application to reconfigurable imaging coils. In: IEEE International Conference on Robotics and Automation, pp. 2498–2504 (2004)
16. Diedrichsen, J., Hashambhoy, Y., Rane, T., Shadmehr, R.: Neural correlates of reach errors. *J. Neurosci.* **25**, 9919–9931 (2005)
17. Menon, S., Stanley, A., Zhu, J., Okamura, A., Khatib, O.: Mapping stiffness perception in the brain with an fMRI-compatible particle-jamming haptic interface. In: Proceedings of the 14th Annual Conference of the IEEE Engineering in Medicine and Biology Society (2014)
18. Burdet, E., Gassert, R., Gowrishankar, G., Bleuler, H.: fMRI compatible haptic interfaces to investigate human motor control. *Exp. Robot.* IX **21**, 25–34 (2006)
19. Klare, S., Peer, A., Buss, M.: Development of a 3 DoF MR-compatible Haptic Interface for Pointing and Reaching Movements. *Lecture Notes in Computer Science*, vol. 6192, pp. 211–218. Springer, Berlin (2010)
20. Chapuis, D., Gassert, R., Gowrishankar, G., Burdet, E., Bleuler, H.: Investigation of a cable transmission for the actuation of MR compatible haptic interfaces. In: Proceedings of Biomedical Robotics and Biomechanics, pp. 426–431 (2006)
21. Hribar, A., Koritnik, B., Munič, M.: Phantom haptic device upgrade for use in fMRI. *Med. Biol. Eng. Comput.* **47**, 677–684 (2009)
22. Massie, T., Salisbury, J.: The phantom haptic interface: a device for probing virtual objects. In: Proceedings of the ASME Winter Annual Meeting, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, vol. 55(1), pp. 295–300. IOS Press (1994)
23. Khatib, O., Warren, J., Sapiro, V.D., Sentis, L.: Human-like motion from physiologically-based potential energies. In: *Advances in Robot Kinematics*, pp. 149–163. Springer, The Netherlands (2004)
24. Sapiro, V.D., Warren, J., Khatib, O., Delp, S.: Simulating the task-level control of human motion: a methodology and framework for implementation. *Vis. Comput.* **25**, 289–302 (2006)
25. Demircan, E., Besier, T., Menon, S., Khatib, O.: Human motion reconstruction and synthesis of human skills. In: Lenarčič, J., Stanisic, M. (eds.) *Advances in Robot Kinematics*, pp. 283–292. Springer, Berlin (2010)
26. Evarts, E.V.: Relation of pyramidal tract activity to force exerted during voluntary movement. *J. Neurophysiol.* **31**, 14–27 (1968)
27. Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., Massey, J.T.: On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* **2**, 1527–1537 (1982)
28. Georgopoulos, A.P., Schwartz, A.B., Kettner, R.E.: Neuronal population coding of movement direction. *Science* **233**, 1416–1419 (1986)
29. Georgopoulos, A.P., Ashe, J., Smyrnis, N., Taira, M.: The motor cortex and the coding of force. *Science* **256**, 1692–1695 (1992)
30. Schwartz, A.B., Kettner, R.E., Georgopoulos, A.P.: Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. relations between single cell discharge and direction of movement. *J. Neurosci.* **8**, 2913–2927 (1988)
31. Georgopoulos, A.P., Caminiti, R., Kalaska, J.F.: Static spatial effects in motor cortex and area 5: quantitative relations in a two-dimensional space. *Exp. Brain Res.* **54**, 446–454 (1984)
32. Kettner, R.E., Schwartz, A.B., Georgopoulos, A.P.: Primate motor cortex and free arm movements to visual targets in three-dimensional space. iii. positional gradients and population coding of movement direction from various movement origins. *J. Neurosci.* **8**, 2938–2947 (1988)
33. Cheney, P.D., Fetz, E.E., Palmer, S.S.: Patterns of facilitation and suppression of antagonist forelimb muscles from motor cortex sites in the awake monkey. *J. Neurophysiol.* **53**, 805–820 (1985)
34. Kakei, S., Hoffman, D.S., Strick, P.L.: Muscle and movement representations in the primary motor cortex. *Science* **285**, 2136–2139 (1999)
35. Holdefer, R., Miller, L.: Primary motor cortical neurons encode functional muscle synergies. *Exp. Brain Res.* **146**, 233–243 (2002)

36. Caminiti, R., Johnson, P.B., Urbano, A.: Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. *J. Neurosci.* **10**, 2039–2058 (1990)
37. Reina, G.A., Moran, D.W., Schwartz, A.B.: On the relationship between joint angular velocity and motor cortical discharge during reaching. *J. Neurophysiol.* **85**, 2576–2589 (2001)
38. Fu, Q.G., Suarez, J.I., Ebner, T.J.: Neuronal specification of direction and distance during reaching movements in the superior precentral premotor area and primary motor cortex of monkeys. *J. Neurophysiol.* **70**, 2097–2116 (1993)
39. Churchland, M.M., Santhanam, G., Shenoy, K.V.: Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *J. Neurophysiol.* **96**, 3130–3146 (2006)
40. Ashe, J., Georgopoulos, A.P.: Movement parameters and neural activity in motor cortex and area 5. *Cereb. Cortex* **4**, 590–600 (1994)
41. Boynton, G.M., Engel, S.A., Glover, G.H., Heeger, D.J.: Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* **16**, 4207–4221 (1996)
42. Graziano, M.S.A., Aflalo, T.N.: Mapping behavioral repertoire onto the cortex. *Neuron* **56**, 239–51 (2007)
43. Delp, S.L., Anderson, F.C., Arnold, A.S.: OpenSim: open-source software to create and analyze dynamic simulations of movement. *IEEE Trans. Biomed. Eng.* **54**, 1940–1950 (2007)
44. Friston, K.J., Ashburner, J.T., Kiebel, S.J., Nichols, T.E., Penny, W.D.: *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Elsevier, London (2006)
45. Desikan, R.S., Killiany, R.J., et al.: An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* **31**, 968–980 (2006)
46. Kay, K., Rokem, A., Winawer, J., Dougherty, R., Wandell, B.: GLMdenoise: a fast, automated technique for denoising task-based fMRI data. *Front. Neurosci.* **7** (2013)
47. Lin, A.L., Fox, P.T., Hardies, J., Duong, T.Q., Gao, J.H.: Nonlinear coupling between cerebral blood flow, oxygen consumption, and ATP production in human visual cortex. *Proc. Natl. Acad. Sci.* **107**, 8446–8451 (2010)
48. Efron, B., Tibshirani, R.: *An Introduction to the Bootstrap*, vol. 57. Chapman & Hall, New York (1994)