A Novel Haptic fMRI Interface for Five-Axis Force and Motion Neuroimaging Experiments

Samir Menon[1], Amaury Soviche[1], Jananan Mithrakumar[1], Alok Subbarao[1], and Oussama Khatib[1].

***Abstract*— In this paper, we demonstrate that it is feasible to conduct high-fidelity multi-axis force control motor neuroimaging experiments in humans using a novel five degree-of-freedom Haptic fMRI interface (HFI-5). We demonstrate that HFI-5 supports high fidelity force and torque control for all its axes. In addition, it minimizes the force-to-torque coupling induced by its closed chain kinematic structure. HFI-5 has low backdrive friction, which improves its haptic transparency. While the device is large and consequently flexible, a linear model can correct kinematic accuracies due to device deflection. Resulting position error measurements are similar in scale to human hand jitter, and are thus satisfactory. To test HFI-5's efficacy, we performed force perception and force discrimination tests with it on four human subjects. The test revealed that humans could readily perceive forces above 1N and could discern forces that were about 33% apart. As such, HFI-5's force range (x:±16.36N; y:±4.62N; z:±6.9N; passively gravity compensated) can support a wide variety of forces that are perceptible and distinguishable. Finally, we performed ten fMRI scans with a human subject who performed a visually guided force-control experiment. Temporal noise patterns in fMRI measurements matched the fMRI baseline (0.8%), which demonstrates that HFI-5's electromagnetic motors do not introduce substantial noise. HFI-5 is thus ready to be used in motor neuroimaging experiments.**

I. Introduction

Efforts to combine haptics and functional magnetic resonance imaging (fMRI) for studying the brain's motor system have led to the development of a variety of fMRI-compatible haptic interfaces [1], [2], [3], [4], [5], [6], [7], [8], [9], [10]. These *"Haptic fMRI"* interfaces allow researchers to overcome the workspace limitations of the confined MRI scanner by using virtual simulation [11], [12], [13], [14]. In addition, they allow precise measurement of human motor control, which is necessary to construct high fidelity models of the neural response and take advantage of high-resolution fMRI measurements (mm, sec) [15], [16].

Past research efforts in Haptic fMRI have been multi-faceted, focusing on low-level MRI-compatible actuation and sensing, on adapting existing haptic interfaces for fMRI, or on engineering novel devices. Actuator research has studied

*This work was supported by National Science Foundation National Robotics Initiative grant (IIS-1427396, O. Khatib and R. Bajcsy), and by a Stanford Center for Neurobiological Imaging Seed Grant (S. Menon).

[1]S. Menon, J. Mithrakumar, A. Subbarao, and O. Khatib are with the Stanford Robotics Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305, USA. [2]A. Soviche is with the Robotics and Autonomous Systems Department, Ecole Polytechnique Fédérale de Lausanne, Route Cantonale, 1015 Lausanne, Switzerland. Email: smenon@cs.stanford.edu, janamith@stanford.edu, asubbarao@stanfordhealthcare.org, ok@cs.stanford.edu, amaury.soviche@epfl.ch

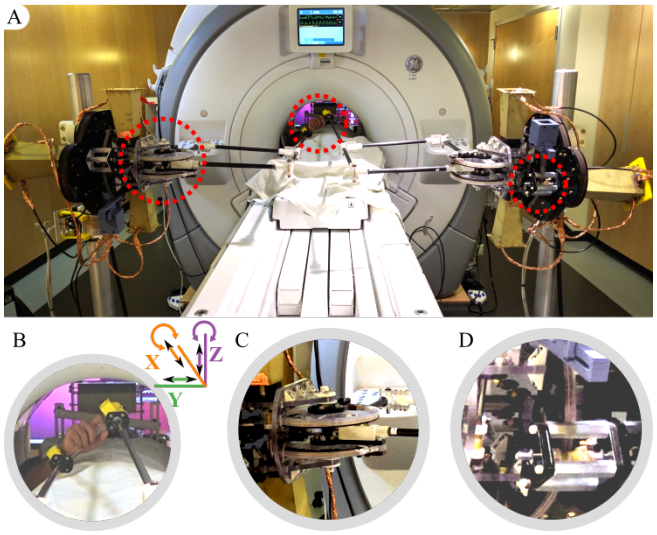

Fig. 1. Haptic fMRI Interface 5 (HFI-5): *(A)*. A novel fMRI-compatible haptic interface that realizes five degrees-of-freedom (DOF) by using two three-DOF arms connected with a gimbal at the end-effector (handle). Red circles outline zoomed insets. *(B)*. A subject grasping the end-effector gimbal while lying in the MRI scanner's bore. All axes except the handle-cylinder's axis are actuated. Actuated DOF are indicated. *(C)*, The actuation mechanism for the planar parallelogram structure of one arm. The brass Faraday cages that serve as electromagnetic shields and house the motors are visible. *(D)*, The actuation mechanism for the up-down axis. All drivetrains use kevlar cables to transmit forces, with capstans for gear reduction.

electro-active polymers [17], pneumatics [1], hydraulics [18], [19], or long cables with dynamic models to improve performance [20]. Efforts to repurpose existing devices include using a shielded PHANTOM [5], [21]. Finally novel devices include a pneumatic planar manipulandum [1], a wrist device [10], and a family of electromagnetically actuated devices [7], [9]. A common problem faced by all approaches is that it is challenging to develop Haptic fMRI interfaces that allow high fidelity force control [9] while supporting natural motions [22], [7], [23] spanning the MRI workspace. Another common problem is precisely identifying the sources of noise introduced by Haptic fMRI devices [24], [9]. Finally, no existing device supports more than three degrees-of-freedom (DOF).

Here, we present a novel device, Haptic fMRI Interface 5 (HFI-5; Fig. 1), a five DOF electromagnetically actuated fMRI-compatible haptic interface that supports translation and rotation across the entire MRI workspace. HFI-5 supports high fidelity force (error, $F_{x-err} < 0.15$N, $F_{y-err}, F_{z-err} < 0.8$N) and torque (error, $\Gamma_{x-err} < 0.008$Nm, $\Gamma_{z-err} < 0.02$N) and effectively decouples forces and torques at the end-effector (force-to-torque coupling

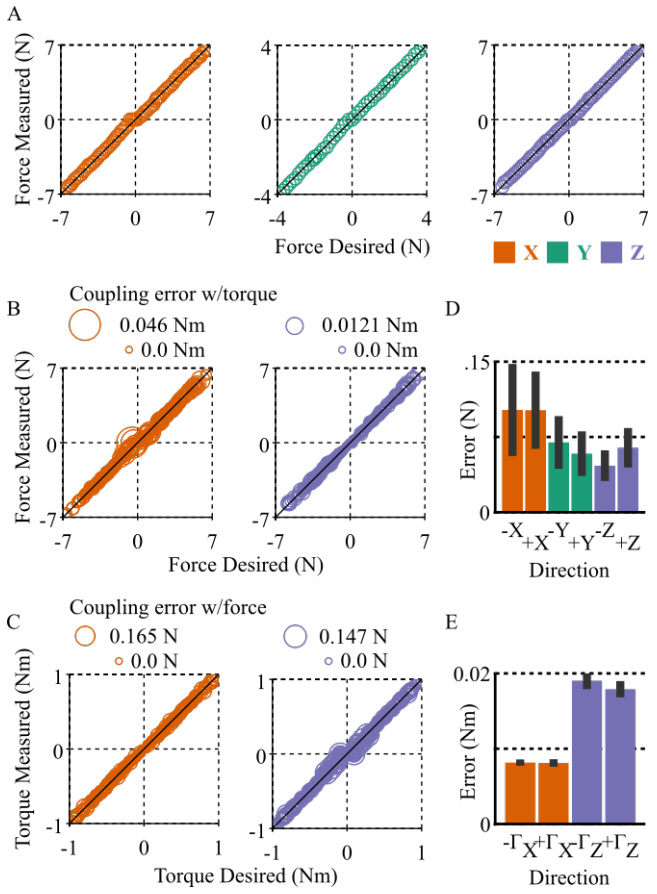

Fig. 2. Force Control for HFI-5: *(A)*. Forces measured at the end-effector match desired forces for the three translation axes (x, y, z), which indicates that force control is reliable and accurate. *(B)*. HFI-5 generates torques along two axes (x, z) by computing the associated end-effector forces. Force errors, as such, result in torque errors. The force-to-torque coupling error is indicated by circles of different sizes across a range of forces. The coupling error is small: $<$ 0.046 Nm for x-axis forces to z-axis torques, and $<$ 0.0121 Nm for z-axis forces t x-axis torques. *(C)*. Torque control is similarly reliable and accurate across a range of torques for the actuated axes (x, z). The torque-to-force coupling error is indicated with circles and is small: $<$ 0.165 N for x-axis torques to z-axis forces, and $<$ 0.147 N for z-axis torques to x-axis forces. *(D)*. The force control error for controlled axes. *(E)*. The estimated torque control error for controlled axes. The larger z-axis torque errors are a consequence of larger errors in x-axis force control. Error bars indicate 99%*ile* bootstrap confidence intervals (for D and E).

error $<$ 0.046Nm; torque-to-force coupling error $<$ 0.165N). In addition, HFI-5 has low backdrive friction at the end-effector (translation: x$<$ 0.4N, y$<$ 0.2N, z$<$ 0.17N; rotation: x$<$ 0.02Nm, z$<$ 0.04Nm). Since HFI-5's electromagnetic motors operate best beyond the MRI machine's 400 Gauss line [7], [9], the kinematic structure is large and consequently somewhat flexible. To compensate for HFI-5's flexibility, we used an empirical linear model to estimate and invert end-effector position offsets while applying forces. This limited deflection errors (x$<$ 0.2cm at 11.2N; y$<$ 0.3cm at 4N; z$<$ 0.67cm at 4N) to a level near steady-state human hand jitter (~0.5cm [23]).

We tested HFI-5 with two experiments involving human subjects, a force perception experiment and an fMRI visually guided force-control experiment. First we studied how

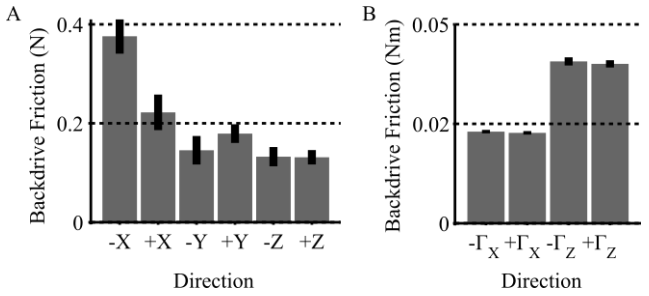

Fig. 3. Backdrive Friction for HFI-5: *(A)*. The minimum force required to produce a translation along each axis direction. *(B)*. The estimated minimum torque required to produce a translation along each axis direction. Error bars indicate 99%*ile* bootstrap confidence intervals for both plots.

humans perceived forces generated by HFI-5 by applying a sequence of force tests where one force was larger than the other. Humans reliably perceived forces larger than ~1N, and could distinguish between forces that were about 33% different in magnitude. HFI-5's steady-state force-generation abilities (x:$\pm$16.36N; y:$\pm$4.62N; z:$\pm$6.9N; passively gravity compensated) indicate that the device supports a wide variety of forces that are perceptible and distinguishable. Finally, we tested HFI-5's performance in a series of ten fMRI scans where the subject performed a visually guided force-control task. Recorded temporal noise patterns were close to the fMRI baseline (~0.8% noise-to-signal), indicating that HFI-5 is suitable for studying force control in the human brain.

II. Five Axis Force Control

HFI-5 uses electromagnetic motors (z-axis, Maxon 370356; x- and y-axis, Maxon 148877) and supports open-loop force control through analog-amplifier current control (Copley Controls, model 4122). Since placing any electronic subsystems inside the MRI-room causes strong electromagnetic interference, the control system requires large analog current lines ($> 4m$ long). This can cause instabilities in the current controller and thus reduce the fidelity of the force control. As such, we conducted extensive tests to calibrate force control and made suitable adjustments that ensured low force-generation error.

We tested HFI-5's force-generation ability with a sequence of axis-by-axis tests (Fig. 2). Forces were measured with a precise single-axis force sensor (resolution, ~10mN; see Appendix for details). Measurements indicated that force output was accurate and highly linear for all translation axes. We estimated torque errors by sampling errors from the bootstrap linear force error distribution [25], and transformed those errors into torque errors with a closed-chain kinematic model at the end-effector gimbal. This allowed us to identify how force errors produced torque errors (see Fig. 2. B), and estimate torque generation (see Fig. 2. C). As with force-generation, torque-generation was accurate and linear.

We note that the x-axis force-generation error is larger than the other axes because of the device's kinematic structure. The long arms deflect along the z-axis under their own weight, which makes the actual x-axis direction differ

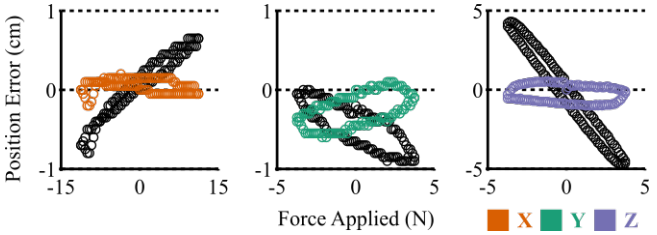

Fig. 4. Device Flexibility and Mechanical Stiffness: A linear device-flexibility model effectively compensated for device deflection while applying forces across all three axes. Black traces indicate position-error before the model is applied, and colored traces indicate position-error after correction for deflection for each axis. The device is naturally stiffest along the x-axis ($\pm$0.7cm deflection for $\pm$11.2N), less stiff along the y-axis ($\pm$0.45cm deflection for $\pm$4N), and least stiff along the z-axis ($\pm$4.35cm deflection for $\pm$4N). The mechanical stiffness is similarly high along the x-axis (1600N/m), lower along the y-axis (890N/m), and lowest along the z-axis (92N/m). Applying the flexibility model greatly reduced position errors along the x-axis ($\pm$0.2cm deflection for $\pm$11.2N), y-axis ($\pm$0.3cm deflection for $\pm$4N), and z-axis ($\pm$0.67cm deflection for $\pm$4N).

from the estimated direction. As a consequence, the x-axis Jacobian is somewhat error prone. Errors in x-axis forces are reflected in z-axis torques due to the kinematics. Finally, force-generation may vary by axis direction due to anisotropic friction and motor winding effects. As such, we computed error estimates and confidence intervals for each axis and each direction separately (see Fig. 2. D, E).

III. Friction Estimation

While the open-loop force error upper bounds friction from the motors to the end-effector, completely characterizing friction for a haptic device also requires testing the *"backdrive"* friction levels associated with producing motion at the end-effector (Fig. 3). We did so by applying a force at the HFI-5 end-effector gimbal with a linear-spring deflection scale (see Appendix for details) and noting the point at which either (or both) HFI-5 arms moved. To compensate for changes in grasp position, we sampled the backdrive friction at many different points across the end-effector (see Appendix for details).

The measured end-effector backdrive friction was, surprisingly, substantially larger than the open-loop force-generation errors (see Fig. 3). We attribute this to the device's closed-chain kinematics and flexibility. While backdriving the device, we noted that the elastic deformation due to flexibility caused the force application plane to be misaligned with the device drivetrain (capstan) plane. This, combined with the closed chain gimbal, meant that internal forces in the device kinematics compensated part of the applied end-effector forces. This effect, however, can not be avoided without substantially improving z-axis flexibility (a future goal). In contrast, the actual drivetrain produces no internal forces since the base is not flexible, which leads to lower force-generation errors.

Finally, our manual backdrive-and-observe method introduces errors in both force application direction as well as measurement. The linear spring had a lower resolution than our force sensor. However, we could not use the force sensor for measuring backdrive friction since motion transients interfered with its measured force. Given experimental constraints, we suggest that the estimated backdrive friction values should primarily be interpreted as upper bounds. It is highly likely that the actual backdrive friction is much lower.

IV. Device Flexibility Modeling

HFI-5 is large, lightweight, and made of non-ferromagnetic metals, plastics, and composites. As such, it is flexible along the x-, y-, and z-axes. Device flexibility introduces systematic errors in estimating end-effector kinematics while applying forces. We empirically characterized this deflection as a function of applied force (Fig. 4).

We found that the flexibility along the x- and y-axes was substantially lower than the z-axis. This upper-bounds the stiffness the device can project in a haptic simulation (see Fig. 4). In addition, when applying forces across HFI-5's range, the device's flexibility causes position errors that are larger than the steady-state hand-position-holding jitter for humans (~0.5cm [23]). The errors in measurement accuracy thus could prove to be an experimental confound for motor neuroimaging experiments.

To overcome device flexibility, we developed a linear deformation model that corrects the kinematic estimates while applying forces (see Appendix for details). Our model successfully lowered the position error for all three axes (see Fig. 4, colored circles). While x- and y-axes exhibited little deflection to begin with, the z-axis position estimates improved dramatically. We note that improving this estimate also helps improve the estimate of the force projection Jacobian, which consequently reduces force-generation error.

V. Force Perception

Having calibrated HFI-5's force-generation, identified friction properties, and corrected for device flexibility, we proceeded to an experiment that tested how subjects perceived haptic forces (Fig. 5). We selected a paired-force experiment design where subjects held HFI-5 and were subjected to a sequence of force-pairs (see Appendix for details). The experiment had two goals: first, to identify the threshold at which subjects start perceiving applied forces; and second, to determine how well subjects could distinguish between applied forces. A randomized sequence of force-pairs helped ensure that subjects were not overtrained, which can render measured results uninterpretable. Force magnitudes were selected along a range that allowed subjects to perform the entire sequence (300 force-pairs) in one sitting, and without getting excessively tired.

We conducted the paired-force experiment on four human subjects (see Appendix for details) and found that all subjects had similar patterns of force perception (see Fig. 5. A). It is clear that forces above 1N can be readily perceived by all subjects. As such, experiments conducted with HFI-5 should involve forces $>$ 1N.

Finally, we tested the ability of human subjects to distinguish between the two forces in a pair. This is important

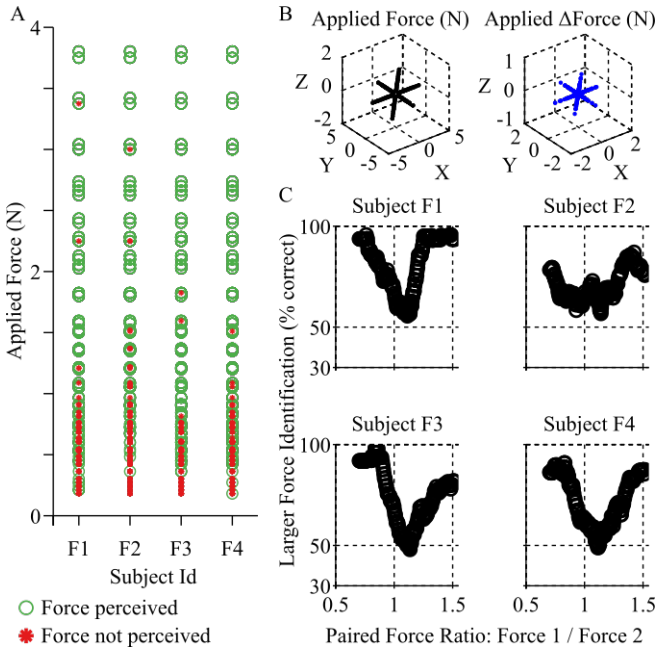

Fig. 5. Human Force Perception with HFI-5: *(A)*. Data from four human subjects indicating haptic forces that were perceived (green circles) or not (red stars) while holding HFI-5. The experiment involved feeling two different forces in a sequence. Subjects could indicate both, one, or none perceived. Subjects can perceive forces above 1N, though they occasionally detect forces even at the force-error threshold (> 0.2N). *(B)*. Applied forces (left; black) were sampled uniformly along the shown directions with different magnitudes. The difference between successive forces (right; blue) were similarly sampled along the directions shown. *(D)*. Subject performance at determining which force in the pair was larger as a function of the ratio of the two forces. Subject performance converges to chance (50%) when the forces are similar. Subjects perform fairly well ($> 75\%$) when forces differ more than 33%. It is noteworthy that subjects have a bias towards feeling that the second force is larger (the graph minima are above 1).

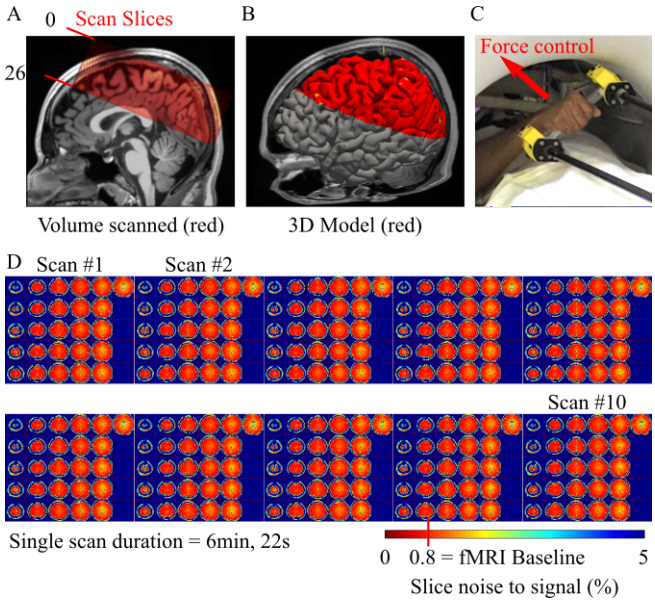

Fig. 6. fMRI Noise Tests for a Force Control Experiment: *(A)*. A volume of a subject's brain was scanned. The volume was divided into 26 fMRI slices (2.5mm thick). *(B)*. A volumetric reconstruction of the subject's brain shows the regions scanned in three dimensions. *(C)*. The experiment involved controlling x-axis forces (3.25, 4.5, 5.75N) at the end-effector. *(D)*. Temporal noise levels shown in a heatmap for ten fMRI scans. All slices are shown in an image grid (pixel resolution, $2.5 \times 2.5\,mm^2$). The baseline temporal noise level (0.8%, when a subject is sleeping) is marked on the colorbar.

since while designing a force-based motor neuroimaging experiment, the experimenter should select a set of force levels that are discriminable. To analyze discriminability, we sorted force pairs by the ratio of the first to second force, then binned forces using a sliding window that contained 60 samples from the sequence, and finally computed the success rate for each bin. We found that subject force discriminability approached chance (50%) when the forces were close to each other (see Fig.5. C). Forces that were well separated were perceptible with a higher chance rating. We do note that while three subjects (F1, F3, F4) had very similar response profiles, one subject's (F2) responses did not reach a very high level of force discrimination.

VI. fMRI Compatibility: Noise Testing

Having satisfactorily tested all aspects of HFI-5's force generating abilities, we proceeded to conduct fMRI brain-scanning experiments with a human subject (Fig. 6). Our goal was to determine whether HFI-5 was properly shielded and truly fMRI compatible: that its electromagnetic motors did not introduce any temporal noise into fMRI measurements. Based on our experience with past Haptic fMRI devices [7], [24], [9], we decided that an experiment that involved a human applying large forces with random force transients was most likely to interfere with fMRI measurements. As such, we chose a visually guided force control experiment.

Our experiment design required subjects to hold HFI-5's handle and stare at a computer screen (through a mirror) while in the MRI. By default the screen was grey, which indicated to subjects that they should relax and stay still. At random, the sreen would turn light blue indicating that subjects should apply a +x force (towards their chin from their navel). As subjects applied force, the screen turned deeper blue, and finally green (stable point). If the force applied was too large, the screen turned red. Forces were applied with an impedance field (262.5N/m). Subjects randomly applied a variety of such forces over the course of ten scans, each of which was six and a half minutes long.

We conducted the experiment for a healthy right handed male human (Subject F3), and performed ten scans of a region of the brain that included pre-motor (planning), motor [26], [16], [27], somatosensory (tactile perception) [28], parietal (grasping) [29], and visual [30] cortex (see Fig. 6. A, B). On analyzing the temporal noise patterns (see Appendix for details) of all the recorded voxels, we found that noise levels were close to baseline noise (0.8% of signal) for all ten scans. This indicated that HFI-5 does not indicate with fMRI measurements. We also plotted the relevant voxels on an inflated 3D model of the brain for clarity (Fig. 7; see Appendix for plotting details).

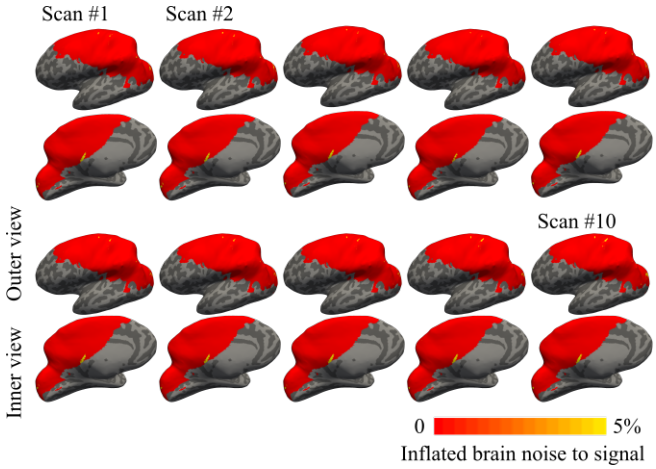

Fig. 7. fMRI Temporal Noise Maps: Inflated brain models help observe noise levels across the entire brain's surface (including hidden grooves). Temporal noise maps for the left half of the brain are shown for ten scans (seen earlier in Fig. 6). View from the outside and inside are shown in alternate rows. The noise is low ($< 1\%$) across all the brain.

VII. Conclusions

To summarize, we present a novel electromagnetically actuated five-DOF Haptic fMRI interface, HFI-5. We present results indicating that HFI-5 supports high fidelity forces at its five axes, that it has low friction, and that a model based approach can correct kinematic estimation errors due to device flexibility. We used HFI-5 to perform force perception and force discrimination tests on four human subjects. The results we obtained can help guide fMRI experiment design. Finally, we performed ten human visually guided force-control fMRI scans to test HFI-5's fMRI compatibility and found that HFI-5 did not interfere with fMRI measurements. Our tests follow an established testing methodology developed for a family of past Haptic fMRI interfaces [7], [23], [8], [24], [9]. HFI-5 is thus ready to be used in motor neuroimaging experiments.

APPENDIX

Friction Estimation: The mount for friction estimation consisted in having the robot at the zero potion (robot's end effector perfectly static in the air, without the need of any support). The goal here is to determine, for both robot's arm and every axis (x, y, z) and direction (positive, negative), the minimum amount of force to apply in order to have the robot move. To achieve this goal, we would push the robot's arms with a spring deflection-based scale. When the robot started to move, we recorded the force given by the scale. We applied the latter procedure for every axis and every direction, leading to a total of 180 measures. The spring deflection-based scale is Jomard GPP-8, with a range of 200g * 5g and an accuracy of +-1 graduation (= 7.08 g)

Force Calibration: The mount for the force calibration consisted in having a scale (Spirit digital scale with precision of 0.1g, weight capacity of 3000g, platform of 2.54 x 2.54 cm) attached to a tripod. The tripod's head orientation could be re-adjusted both in linear translation along the z axis (floor - ceiling axis) and in angle. The latter mount allowed us to measure forces applied from the robot at the end effector along any of the main axes (x, y, z) and directions (positive, negative). The procedure was a two step iterative approach. First, without any force calibration, we had the robot applied forces and collected the measured forces from the scale. We then, using linear interpolation, found the coefficient k between the desired forces and the applied forces ($F_{desired} = k * F_{measured}$). Once this iterative process was done (found coefficient k), we could implement the force calibration in our robot's server: $F_{sent-to-robot} = (1/k) * F_{task}$. We would not use any constant coefficient to avoid jerk within low forces. Forces along the x and z axis were contained inside +/-6.93N. Forces along the y axis were contained inside +/-4.21N. Along any direction, we decided to take force steps of 0.17N, leading to 80 measurements along x and z axis and 51 along y axis. The y axis has less measures because this axis saturates faster than others due to motor contraints associated to our robot's Jacobian.

Deflection Calibration: The mount for the deflection calibration consisted in having the robot's end effector at the zero position, attached to a tripod, so that it could not move along any axis nor direction. The procedure was again a two step iterative approach. We first ran run an automated script to have the robot apply forces along every axis (x, y, z) and every direction (positive, negative). Along every axis, forces were contained inside +-3N, with increments of 0.1N. For every force sent to the robot, we automatically saved the force sent and the estimated end-effector position by the robot into a text file. We could then, for every axis, plot the relationship $Position - error = k * F_{applied}$. By using a linear regression technique, we extracted the coefficient k and corrected the position error.

Force Perception: For this experiment, we reproduced the same environment as if patients were inside the fMRI. The patients laid down on a table with pillows to make the experiment more comfortable. We then applied 300 pairs of forces, with a 3 second break in between and each force applied for 4 seconds, and another 2 second force ramp up and cancellation (sigmoid type). For every pair, the subjects were given five options: 1. "could not feel force 1", 2. "could not feel force 2", 3. "could not feel any force" 4. "force 1 was stronger than force 2", 5. "force 2 was stronger than force 1". The subjects had a keyboard next to them in order to give their answer after every pair of forces. Subjects' answers were then collected into a text file to conduct further data analysis.

fMRI Scanning: All fMRI scans were conducted at Stanford University's Center for Cognitive and Neurobiological Imaging on a GE Discovery MR750 3 Tesla MRI scanner, with a 32 channel Nova Medical head coil. The scan protocol was gradient echo EPI with a 16cm field of view sampled at a 64x64 resolution (2.5x2.5x2.5 mm^3 voxels), a 1.57s repetition time, a 28ms echo time, and a 72^o flip angle. All scan runs were preceeded by 2^{nd}-order polynomial shimming and were sandwiched by fieldmap scans. After scanning, the fMRI images were slice time corrected, motion corrected

(SPM), spatially undistorted using fieldmaps, and analyzed to compute temporal signal-to-noise. A subject-customized bite-bar minimized head motion. All runs (265 sec each) had frame-to-frame head motion >0.1mm or overall head motion >1mm.

fMRI Analysis: Temporal noise-to-signal computations used the median neural response distribution obtained by regressing out a line from each voxels 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. Surface registration was done using Freesurfer, and all surface images were plotted using Freeview.

Haptic Data: Motions were right handed. The haptic control loop was ~7450Hz. As part of the haptic control interface, haptic trajectories were filtered with a second order Savitzky-Golay filter (programmed in C++) to remove high frequency noise and estimate derivatives. Trajectories were down-sampled to 100Hz using a cubic spline to simplify figure plotting.

Human Subjects: Subjects were healthy right-handed males with no history of motor disorders: F1, 26yr, 148lb, 5'7"; F2, 23yr, 166lb, 6'0"; F3, 19yr, 137lb, 5'8"; F4, 22yr, 115lb, 5'7". Informed consent was obtained in advance on a protocol approved by the Institutional Review Board (IRB) at Stanford University.

ACKNOWLEDGMENTS

We gratefully acknowledge Jihee Hwang, Fahad Al-Shaibani, Deeksha Goyal, Valerie Hau, and Sahaana Suri, for their assistance in constructing and calibrating the electrical control system for HFI-5. We acknowledge Jack Zhu for his assistance with designing the stand for HFI-5, and for helping construct the prototype D-1 device. Finally, we thank Kendrick Kay, Laima Baltusis, and Robert Dougherty for helping develop fMRI scanning and data processing protocols.

REFERENCES

[1] J. Diedrichsen, Y. Hashambhoy, T. Rane, and R. Shadmehr, "Neural correlates of reach errors," *The Journal of Neuroscience*, vol. 25, no. 43, pp. 9919–9931, 2005.

[2] J. Hidler, T. Hodics, B. Xu, B. Dobkin, and L. Cohen, "Mr compatible force sensing system for real-time monitoring of wrist moments during fMRI testing," *Journal of neuroscience methods*, vol. 155, no. 2, pp. 300–307, 2006.

[3] R. Gassert, L. Dovat, O. Lambercy, Y. Ruffieux, D. Chapuis, G. Ganesh, E. Burdet, and H. Bleuler, "A 2-dof fMRI compatible haptic interface to investigate the neural control of arm movements," in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*. IEEE, 2006, pp. 3825–3831.

[4] R. Gassert, R. Moser, E. Burdet, and H. Bleuler, "Mri/fMRI-compatible robotic system with force feedback for interaction with human motion," *Mechatronics, IEEE/ASME Transactions on*, vol. 11, no. 2, pp. 216–224, 2006.

[5] A. Hribar, B. Koritnik, and M. Munih, "Phantom haptic device upgrade for use in fMRI," *Medical and Biological Engineering and Computing*, vol. 47, pp. 677–684, 2009.

[6] N. Yu, N. Estévez, M. Hepp-Reymond, S. Kollias, and R. Riener, "fMRI assessment of upper extremity related brain activation with an mri-compatible manipulandum," *International journal of computer assisted radiology and surgery*, vol. 6, no. 3, pp. 447–455, 2011.

[7] S. Menon, G. Brantner, C. Aholt, K. Kay, and O. Khatib, "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*, July 2013, pp. 4137–42.

[8] S. Menon, A. Stanley, J. Zhu, A. Okamura, and O. Khatib, "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*, August 2014.

[9] S. Menon, H. Ganti, and O. Khatib, "Using haptic fMRI to enable interactive motor neuroimaging experiments," *Springer Tracts in Experimental Robotics*, vol. 81, 2014.

[10] A. Erwin, M. O'Malley, D. Ress, and F. Sergi, "Kinesthetic feedback during 2dof wrist movements via a novel mr-compatible robot," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2016.

[11] K. S. Hale and K. M. Stanney, "Deriving haptic design guidelines from human physiological, psychophysical, and neurological foundations," *Computer Graphics and Applications, IEEE*, vol. 24, no. 2, pp. 33 –39, March-April 2004.

[12] J. K. Salisbury, F. Conti, and F. Barbagli, "Haptic rendering: Introductory concepts," *IEEE Computer Graphics and Applications*, vol. 24, no. 2, pp. 24–32, March 2004.

[13] F. Conti and O. Khatib, "Spanning large workspaces using small haptic devices," *Proc. of the Eurohaptics on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 183–188, Mar. 2005.

[14] ——, "A new actuation approach for haptic interface design," *The International Journal of Robotics Research*, vol. 28, no. 6, pp. 834–848, 2009.

[15] N. K. Logothetis, "What we can do and what we cannot do with fMRI," *Nature*, vol. 453, no. 7197, pp. 869–878, Jun 2008.

[16] J. D. Meier, T. N. Aflalo, S. Kastner, and M. S. A. Graziano, "Complex organization of human primary motor cortex: A high-resolution fMRI study," *Journal of Neurophysiology*, vol. 100, pp. 1800–1812, 2008.

[17] J. Vogan, A. Wingert, J. Plante, S. Dubowsky, M. Hafez, and D. Kacher, "Manipulation in mri devices using electrostrictive polymer actuators: with an application to reconfigurable imaging coils," in *IEEE International Conference on Robotics and Automation*, 2004, pp. 2498–2504.

[18] E. Burdet, R. Gassert, G. Gowrishankar, and H. Bleuler, "fMRI compatible haptic interfaces to investigate human motor control," *Experimental Robotics IX*, vol. 21, pp. 25–34, 2006.

[19] S. Klare, A. Peer, and M. Buss, "Development of a 3 dof mr-compatible haptic interface for pointing and reaching movements," *Lecture Notes in Computer Science*, vol. 6192, pp. 211–218, 2010.

[20] D. Chapuis, R. Gassert, G. Gowrishankar, E. Burdet, and H. Bleuler, "Investigation of a cable transmission for the actuation of mr compatible haptic interfaces," in *Proc. of Biomedical Robotics and Biomechatronics*, 2006, pp. 426–431.

[21] T. Massie and J. Salisbury, "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. IOS Press, 1994, pp. 295–300.

[22] M. Kostic, D. Popovic, and M. Popovic, "Influence of planar manipulandum to the hand trajectory during point to point movement," in *IEEE International Conference on Rehabilitation Robotics*, July 2011, pp. 1–4.

[23] S. Menon, M. Yu, K. Kay, and O. Khatib, "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*, August 2014.

[24] S. Menon, P. Quigley, M. Yu, and O. Khatib, "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*, August 2014.

[25] B. Efron and R. Tibshirani, *An introduction to the bootstrap*. Chapman & Hall/CRC, 1994, vol. 57.

[26] M. L. e.a., "Activation of cortical and cerebellar motor areas during executed and imagined hand movements: An fMRI study," *Journal of Cognitive Neuroscience*, vol. 11, no. 5, pp. 491–501, 1999.

[27] R. Shadmehr, M. A. Smith, and J. W. Krakauer, "Error correction, sensory prediction, and adaptation in motor control," *Annual Review of Neuroscience*, vol. 33, no. 1, pp. 89–108, 2010.

[28] H. Hämäläinen, J. Hiltunen, and I. Titievskaja, "fMRI activations of si and sii cortices during tactile stimulation depend on attention," *Neuroreport*, vol. 11, no. 8, pp. 1673–1676, 2000.

[29] J. Connolly, R. Andersen, and M. Goodale, "fMRI evidence for a'parietal reach region'in the human brain," *Experimental Brain Research*, vol. 153, no. 2, pp. 140–145, 2003.

[30] S. Courtney, L. Ungerleider, *et al.*, "What fMRI has taught us about human vision." *Current opinion in neurobiology*, vol. 7, no. 4, p. 554, 1997.