Feasible Wrench Space and its Estimation for Isometric Haptic Interaction

Timothy Edmunds* Reinhard Gentner† Andrea d'Avella† Dinesh K. Pai*

*Department of Computer Science, University of British Columbia, Vancouver, Canada
†Santa Lucia Foundation, Neuromotor Physiology Laboratory, Rome, Italy

Abstract

Standard approaches to mapping force to displacement or force to velocity with multi-DoF isometric haptic devices typically ignore the directional variability in a user's feasible wrench. We demonstrate that such a directionally-uniform mapping tends to either over-sensitize the interaction in some directions or under-utilize the user's operational range. To increase the effective use of the user's operational range it is necessary to model that range across all directions; for high-dimensional devices that measure wrenches (i.e, forces and torques) the space of directions is non-trivial to model by sampling. We present an approach that uses in-depth measurement of the feasible wrench space of a small number of users to extract a generic model for a given device interaction context; the generic model can then be automatically fitted to other users through a small number of measurements. In a user study comparing our method against the standard directionally-uniform assumption we show that our method generates a significantly better estimation of a user's output range, while requiring only a few measurements.

1 Introduction

Isometric force exertion on a manipulandum is an effective mechanism for haptic interaction. It can be used in purely passive devices (such as a 2-DoF pointing Stick, or 6-DoF SpaceMouse™ [1]), or to extend the effective workspace of active devices with small physical workspaces [6]. Combining the isometric interaction with other output modalities allows for the creation of a perception of active force-feedback (or *pseudo-haptics* [7]). This kind of multi-modal isometric interaction can be combined with different control strategies (e.g., rate-control, or force-to-displacement mapping) to allow for effective haptic interaction with virtual environments.

In order for a haptic device to be most effective, the system should be calibrated to maximize the operational range; for most force-feedback haptic devices the operational range is determined by the capabilities of the device, as the limiting factor is the device's maximum force output. In general, grounded isometric devices do not have this constraint — a strain-gauge or other force sensor can be calibrated to measure forces well in excess of the maximum a human user can exert. In spite of this capability, in many isometric haptic devices (e.g., IBM Trackpoint™, Synaptics TouchStyk™, 3DConnexion SpaceMouse™ [1]) the operational range is usually fixed within a small subset of the user's feasible range. This represents a loss of interactivity-bandwidth that could be used to improve the efficiency of haptic task performance.

Different methods of mapping force to translation and torque to rotation have been studied [3, 9, 13], but in general all directions are treated uniformly: a single force-to-output transfer function for all translation [10] or all rotation [5], or a pair of transfer functions for translation and rotation in 6-DOF devices [1]. Due to the configuration of the musculoskeletal apparatus, a user's feasible output force is highly anisotropic across different directions (as we illustrate in our experimental results). When operating in a volume that is a small subset of the user's feasible range isotropic transfer functions cause few problems — if the force necessary for controlled interaction is sufficiently easy to produce then the relative easiness of force production in different directions is not a performance factor. However, when making fuller use of the user's feasible range a uniform transfer function will either significantly over- or under-emphasize certain directions.

A typical application is the use of isometric force-reaching tasks in a virtual environment to investigate human motor-control systems. How the central nervous system generates the appropriate contact forces for precise control of haptic interactions is still an open question in motor neuroscience [4]. A deeper understanding of the neural control of tool manipulation may be essential for the development of more usable and transparent interfaces. To investigate the neural control of isometric haptic interactions we have designed an apparatus in which subjects control a virtual 6-D cursor by applying wrenches to a mechanically grounded, instrumented handle attached to the wrist; the subjects attempt to reach visible

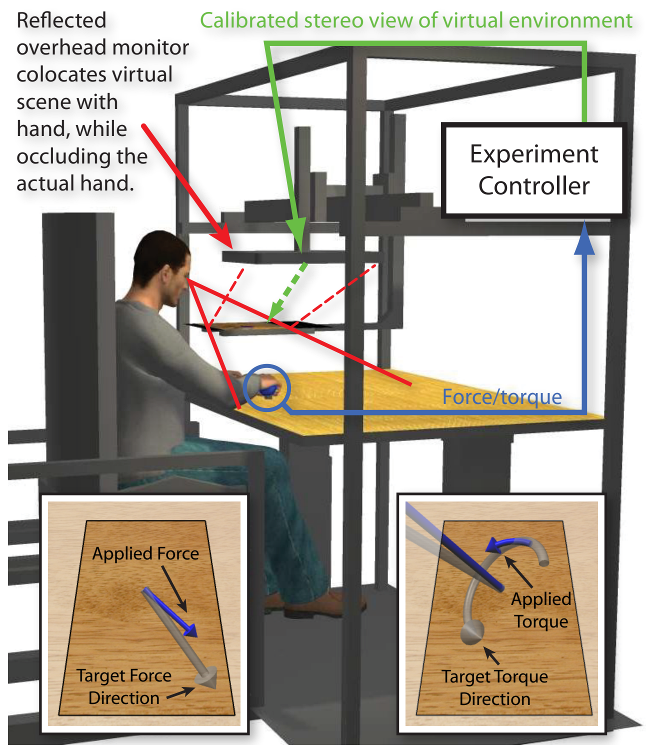

Figure 1: Artist's rendition of the apparatus used for sampling a subject's Maximum Voluntary Wrench. Inset left/right: the subject's view during a force/torque sampling trial.

*email:{tedmunds,pai}@cs.ubc.ca
†e-mail:{r.gentner,a.davella}@hsantalucia.it

 DOI: 10.1109/HAPTIC.2012.6183809

targets while various perturbation schemes are applied to the mapping from applied force/torque to cursor translation/rotation. In order to allow for reasonable comparison of performance — both between subjects and between different directions for a single subject — it is necessary to assign targets that are appropriately scaled for a subject's feasible output. Likewise, to present meaningful visual feedback to the subject, we need to scale a subject's output wrench to displayed translation/rotation in such a way that the wrench-region-of-interest for the experiment maps well to the visual workspace (i.e, not so large as to move offscreen, but not so small as to make differences in wrench difficult to perceive).

Human factors studies have gauged the force output capability of individual joints [11], and literature in the biomechanics field has addressed the problem of modelling human neuromuscular output (see [12] for a review), but constructing and calibrating such a model to accurately estimate the range of endpoint wrenches a user can exert would be prohibitively expensive. Our proposed solution is to leverage the similarity of the musculoskeletal plant across subjects by constructing a data-driven low-dimensional model of a generic user's operational range from a small set of users; we can then calibrate this generic model for new users using measurements of the maximum force capability in just a few directions.

The remainder of this paper is organized as follows. In Section 2 we present some theoretical background to our approach to characterizing a subject's feasible wrench space in a 6-DOF setting. Section 3 describes the experiment we performed to obtain a low-dimensional space in which to characterize a subject's feasible wrench space. In Section 4 we describe how we use this low-dimensional characterization to calibrate an estimate a subject's feasible wrench space using only a small number of samples; we compare the effectiveness of our calibration approach with the typical uniform force/torque normalization. In Section 5 we draw conclusions about the utility of our approach for calibrating the operational range of isometric haptic devices.

2 Theoretical Background

In the idealized view of our experimental setup, the subject's musculoskeletal apparatus is completely fixed in a particular posture, and isometric end-point wrenches are generated by the linear combination of wrenches from n muscles. In this view, the subject's Maximum Voluntary (endpoint) Wrench (MVW) space is the interior of an m-vertexed polytope in $\mathbb{R}^6$, with $m \leq 2^n$ (i.e., the convex hull of the wrenches generated by every possible combination of the maximum exertion of each muscle). Assuming independent activation of each muscle, the MVW is a reasonable first-approximation of the volume of $\mathbb{R}^6$ in which the subject can perform controlled manipulation of the isometric haptic device. (Whether the human motor control system makes use of the full independence-of-activation of the muscles is an open question that is a main focus of the motivating investigation, but the MVW space is a useful classification of the subject's range.) A "perfect" normalization to MVW would consist of characterizing any endpoint wrench $w \in \mathbb{R}^6$ as $w' = \frac{w}{\hat{w}}$, where $\hat{w}$ is the intersection of the ray $\vec{w}$ with the MVW polytope. Thus any target could be expressed as a fixed percentage of MVW *in a particular direction*; likewise, any output wrench could be scaled for display based on the MVW in that output direction. However, run-time calculation/lookup of $\hat{w}$ is relatively complex, and per-subject sampling of every vertex in a subject's MVW polytope is virtually impossible. Sampling *an estimate* of the MVW polytope is feasible in *subspaces* of $\mathbb{R}^6$ (e.g., sampling a number of directions in the x-y force plane), but the number of (physically tiring) samples needed to even sparsely sample a 6-D hypersphere is prohibitive for our application (as it would be for most others).

Our approach is to construct a low-dimensional model with which to describe the MVW space. The MVW that a user can generate in one direction may be significantly different from that of another direction, and may also differ from the the MVW of another user, but due to the similar arrangement of joints and muscles with which the different users apply the wrenches there will be similarity in the *shape* of their MVW spaces (i.e., the *relative* magnitudes of MVW in different directions). We can take advantage of this similarity by performing an in-depth characterization of the shape of the MVW space for a small number of users and identifiying a sparse set of sampling directions that can be used to calibrate a generic shape for subsequent users. Without knowing the direction of the MVW space's boundary vertices (i.e., knowing the effective direction of each muscle's activity in a particular posture) a full characterization of the MVW shape can be impractical to sample even for a single subject; consider that just to sample all combinations of positive/negative axial directions in 6-D involves measuring a user's maximum effort in 728 directions ($3^6 - 1$), and that even that much provides somewhat sparse coverage of the 6-D hypersphere. We address this problem of dimensionality by settling for sampling an approximation of the full MVW space. By considering force and torque separately, we can sample two 3-D spaces in sufficient directions to provide a reasonably confident description of an approximate MVW space, with the observation that the more thorough one can afford to be in the sampling of the preliminary set of users the more robust an estimate of the true MVW space one will be able to calibrate for subsequent users.

3 Methods

In order to apply and test our feasible wrench space calibration method, we conducted a user study in which we measured the MVW space of 5 subjects; 4 of these subjects' results were used to estimate a generic MVW shape, while the remaining subject's results were used to validate the estimate. Below we describe the apparatus we used to measure the subjects' exerted wrenches, and the protocol for sampling their MVW.

3.1 Apparatus

The central component of our apparatus is a fixed handle instrumented with a 6-D force/torque sensor [2]; subjects grasp and manipulate the handle, and the exerted wrench is recorded. In order to maintain a consistent posture throughout the experiment, the subject's torso is restrained by shoulder straps, and the subject's forearm is held by an inflated cuff inside a tube; the tube is attached to the handle (and the f/t sensor) so the subject is effectively exerting wrench with the arm — wrist forces are internalized (though they do play a minor role, in that the user's grasp of the handle prevents rotation of the forearm within the cuff). The subject's view of the actual hand and handle is obscured by a mirror that reflects the display of an overhead monitor; commercial shutter-glasses [8] are used to provide a stereoscopic visual display. In this framework the experiment is structured as an interaction with a co-located haptic/graphic virtual environment (see Figure 1).

3.2 Experimental Protocol

For each subject the experiment consists of a series of trials in which a visual target is presented and the subject attempts to exert a maximal force/torque in that target direction. For clarity, we chose to only present targets that are purely force or purely torque; we present pure force targets in each of the axial and cross-axial directions (26 directions), and likewise present rotation targets about each of these directions. Force targets are presented as a straight arrow originating at the actual position of the subject's hand, while torque targets are presented as an axis of rotation and a curved arrow (see Figure 1 insets).

As the subject applies wrenches to the handle in response to the target stimulus, the f/t sensor attached to the handle samples the wrench (at 1 kHz), and the sampled wrench is graphically rendered

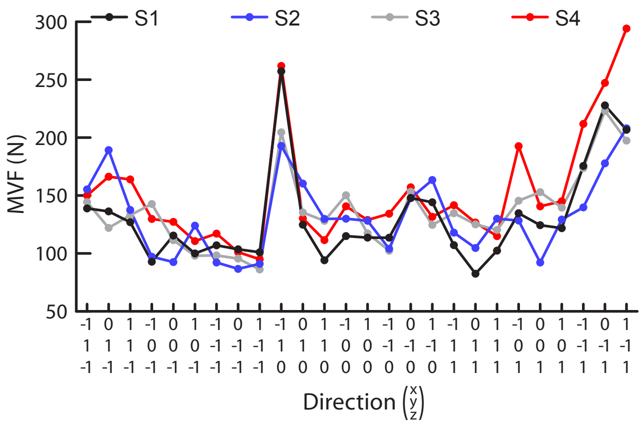

(a) Maximum Voluntary Force

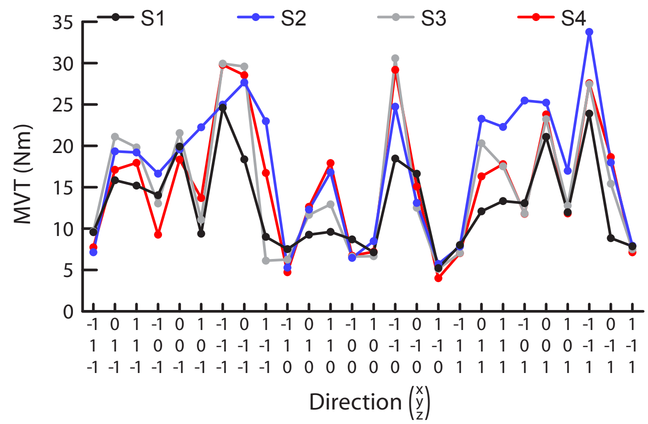

(b) Maximum Voluntary Torque

Figure 2: Maximum Voluntary Force/Torque for the 4 subjects used to estimate the generic MVW shape.

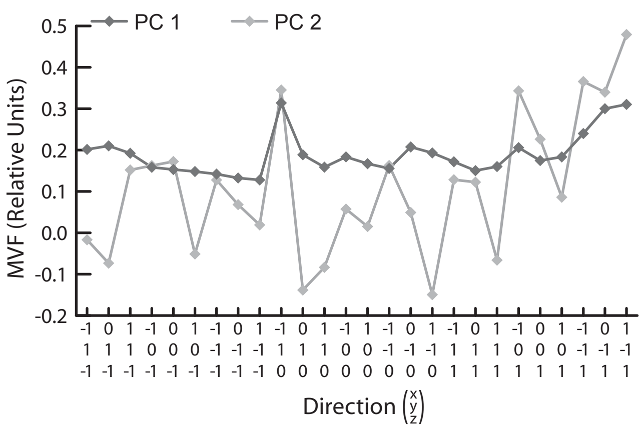

(a) Force PCA

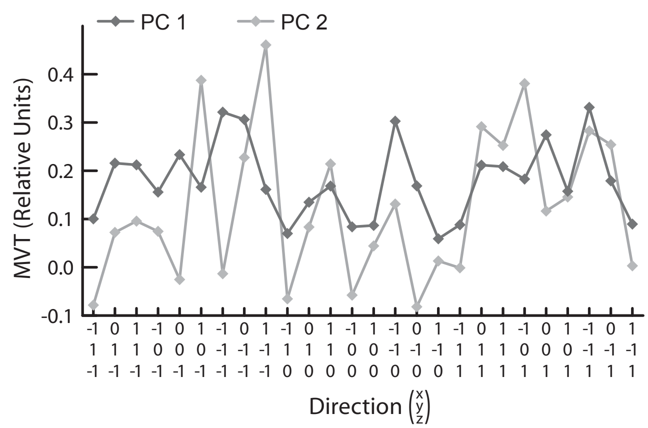

(b) Torque PCA

Figure 3: Dominant principal components for (a) force and (b) torque.

in the virtual environment. For trials involving pure-force targets the subject's wrench is displayed as an arrow originating at the actual hand position and radiating in the direction of applied force (with a length scaling of $\frac{1}{1600}\ \frac{\text{m}}{\text{N}}$). For trials with pure-torque targets, the subject's applied torque is visualized by rotating a small curved arrow about the axis of the torque through the centre of the handle ($\frac{\pi}{25}\ \frac{\text{rad}}{\text{Nm}}$); the axis of the applied torque is also displayed as a rod.

All trials begin with a "fixation" cue — a spherical target is displayed at the origin (the actual location of the subject's hand) until the returns to a relaxed state (force magnitude less than 20 N, after biasing for gravity). The target is then displayed and the subject attempts to exert a maximal sustained force/torque in/about the target direction/axis. In addition to the quantitative visual representation of the subject's applied force/torque, the target changes colour (from transparent grey to transparent yellow) to signal when the subject is within 10° of the target direction. The subject has up to 20 s to attempt to generate and briefly hold a maximum voluntary wrench as close to the target direction as possible — once the subject believes the MVW has been achieved he/she returns to relaxation.

The entire experiment for each subject consists of two blocks of 52 trials each; each target direction is presented (in randomized order) once per block. To prevent fatigue, there is a fixed 5 s break between trials, and a subject-controlled break after every fourth trial.

3.3 Generic Model Extraction

Based on the maximum voluntary wrenches that a subject applied over the course of the experiment, we characterized that subject's MVW space by a 52-dimensional vector, with each dimension corresponding to the MVW in one of the target directions. To compute our estimate of the subject's MVW in a particular direction we considered every force/torque sample throughout the experiment; we rejected all samples that had a magnitude below 25 N or 2 Nm, or were more than 22.5° away from the target direction. Having culled the many samples from the relaxed portions of the trials or from exertion in the direction of other targets, we projected the remaining samples onto the target direction and selected the sample that was in the 99^{th} percentile by magnitude (to reject noise or unsustainable voluntary wrenches).

We used the MVW samples of 4 subjects to build a generic model of the MVW shape for interaction with an isometric haptic device in this posture. To deal with the inherent difference in units we treated force and torque directions separately, decomposing our 52-dimensional shape description into two 26-dimensional components. We applied principal component analysis (Matlab's `princomp` function) to the two 4x26 matrices containing the sub-

(a) Force Reconstruction

(b) Torque Reconstruction

Figure 4: Validation subjects' actual MVW and reconstruction calibrated from the minimal probing measurements.

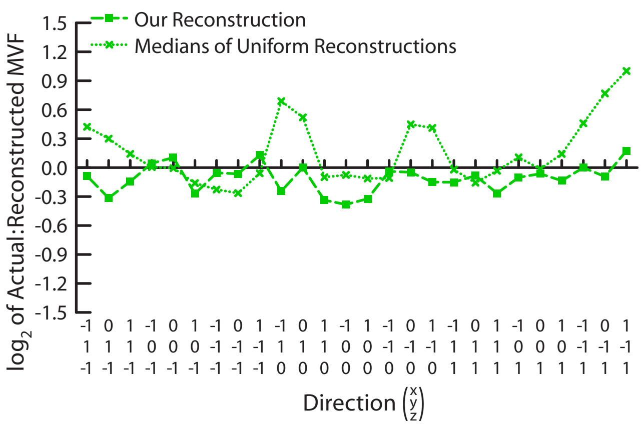

(a) Force Reconstruction Error

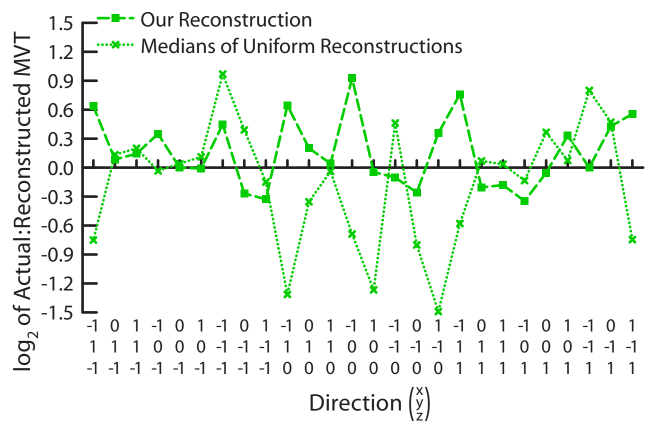

(b) Torque Reconstruction Error

Figure 5: Validation subject's actual-to-reconstructed MVW ratio; medians of results if one measurement is used as a uniform MVF/MVT estimate.

jects' Maximum Voluntary Force/Torque (MVF/MVT) samples; selecting a small set of principal components (those before the elbow of the variance scree plot) provided a basis in which to express MVW shapes.

3.4 Calibration of Generic Model

Having reduced the dimensionality of the representation of the MVW shape to a basis P whose columns are the k identified selected principal components, we can now fit the generic shape to a given subject with only k probing measurements by solving the linear equation:

$$P_p x_p = b_p \tag{1}$$

where P_p consists of the rows of P corresponding to a set of probing measurement direction indices p, and b_p are the subject's MVW results in those directions. Using the data from the same set of subjects that were used to extract the principal components, we identified the k sampling directions that minimized the overall error in the reconstructed MVW shape. For the purposes of comparing errors in MVW reconstruction, we consider an overestimate of 100% to be as erroneous as an underestimate of 50%. Therefore, to choose the reconstruction with a minimum overall error of, we minimize the norm of the log of the errors in each direction:

$$\underset{p}{\operatorname{argmax}} \| \log_2 (\frac{P x_p}{y}) \| \tag{2}$$

where y is the full set of MVW samples.

4 Results

The MVW shapes of the subjects that were used to construct the generic MVW model are shown in Figure 2. This figure illustrates the degree of directional asymmetry in the MVW that results from the structure of the human musculoskeletal system, but it is also suggestive of the cross-subject consistency in MVW shape that we seek to leverage. Using these MVW shapes resulted in a generic MVW model characterized by just 4 principal components (2 each for force and torque), illustrated in Figure 3. For these principal components the probing directions that minimized Equation 2 were:

- force in the (y) direction
- force in the (-x,-y,z) direction
- torque about the (-z) direction
- torque about the (-x,-y,z) direction

4.1 Validation

We used the MVW samples of the remaining subject to test the effectiveness of the calibration procedure determined above. Solv-

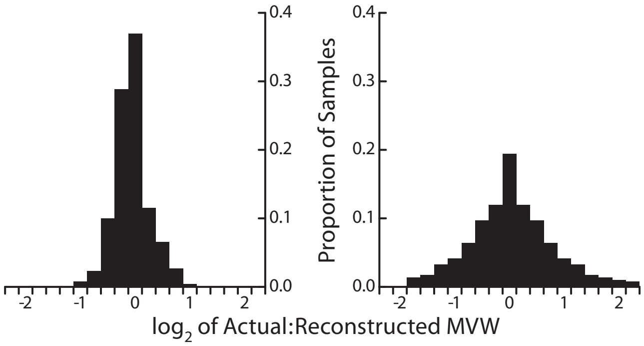

Figure 6: Histograms of errors over all choices of validation subject. Left: Our estimation/calibration approach. Right: Uniform force/torque scaling factors.

ing Equation 1 for just the 4 directions listed above provides a set of weights that express the subject's estimate MVW shape in the generic shape basis. We compare this estimated shape to the actual measured shape for each of the remaining subjects in Figure 4.

To quantify the error in our reconstruction and compare it with the standard normalization scheme (uniform scaling factors for force and torque), we evaluated how much the actual MVW was over-/underestimated in each direction. A log-plot of that ratio for our reconstruction is shown in Figure 5. As a comparison, we considered the effect of using the validation subject's maximum in a single direction as the uniform maximum estimate. For each direction there were thus 26 different estimates; we plotted the median in each direction (also in Figure 5). Note that in all directions our estimate was within the range $[2^{-0.38}, 2^{0.93}]$ times actual value whereas the uniform scaling approach had some cases where the *median* estimate-to-measured ratio was as low as $2^{-1.49}$ or as high as $2^{1.00}$.

To verify the tendency of our approach to result in more accurate estimation of a subject's MVW we reperformed the above analysis, each time withholding a different subject as the validation subject. The aggregated error distributions across all directions for all validation subjects are compared in Figure 6.

5 Conclusions

We have demonstrated that a user's feasible force or torque is not well represented by a single value across all directions, and that while thorough measurement of a single subject's feasible range across directions is too expensive for everyday use, an accurate estimate of a user's MVW space can be established by combining a few measurements with a generic model pre-built around thorough measurement of a small group of users. Our investigation dealt exclusively with a one-handed haptic interaction in a fixed posture, but the overall approach can be extended to other situations. As long as there is sufficient similarity in musculoskeletal apparatus across users, a generic model of user capability can be extracted by careful measurement of a relatively small number of users; increased postural freedom will increase the number of principal components needed to describe the feasible space, but even that number will likely be small enough to allow rapid calibration for new users.

The main limitation of our approach is that for high dimensions (e.g., 6) a rigorous sampling of the full wrench space for even a small number of users is time consuming and fatiguing. We evaded some of this difficulty by treating force and torque in isolation (reducing a single 6-D space to two 3-D spaces), but for some applications this simplification could be unacceptable; in constrained conditions where a user has few degrees of freedom there could be certain narrow directions in which the feasible wrench is significantly lower than in nearby directions. An area of future work is to incorporate rough models of the user's musculoskeletal apparatus to guide the choice of sampling directions to best characterize the MVW space in a given posture.

Acknowledgements

This work is funded by the Human Frontier Science Program (RGP11/2008), with additional support from the Canada Resarch Chairs Program, NSERC, CFI, CIHR, NIH, and the Peter Wall Institute for Advanced Studies. Apparatus rendering courtesy of Daniele Borzelli.

References

[1] *SpaceMouse*. 3Dconnexion. URL `http://www.3dconnexion.com`. 323

[2] *F/T Sensor: Delta*. ATI Industrial Automation. URL `http://www.ati-ia.com`. 324

[3] G. Casiez and D. Vogel. The effect of spring stiffness and control gain with an elastic rate control pointing device. In *SIGCHI '08*, pages 1709–1718, 2008. doi: http://doi.acm.org/10.1145/1357054.1357321. 323

[4] A. d'Avella and D. K. Pai. Modularity for sensorimotor control: Evidence and a new prediction. *Journal of Motor Behavior*, 42(6):361–369, 2010. doi: 10.1080/00222895.2010.526453. 323

[5] B. Froehlich, J. Hochstrate, V. Skuk, and A. Huckauf. The GlobeFish and the GlobeMouse: Two new six degree of freedom input devices for graphics applications. In *SIGCHI '06*, pages 191–199, 2006. doi: http://doi.acm.org/10.1145/1124772.1124802. 323

[6] R. Hollis and S. E. Salcudean. Lorentz levitation technology: a new approach to fine motion robotics, teleoperation, haptic interfaces, and vibration isolation. In *International Symposium for Robotics Research*, Oct. 1993. 323

[7] A. Lécuyer, S. Coquillart, A. Kheddar, P. Richard, and P. Coiffet. Pseudo-haptic feedback : Can isometric input devices simulate force feedback? In *IEEE VR '00*, pages 83–90, 2000. doi: 10.1109/VR.2000.840369. 323

[8] *3D Vision*. NVIDIA. URL `http://www.nvidia.com`. 324

[9] I. Poupyrev, S. Weghorst, and S. Fels. Non-isomorphic 3D rotational techniques. In *SIGCHI '00*, pages 540–547, 2000. doi: http://doi.acm.org/10.1145/332040.332497. 323

[10] J. D. Rutledge and T. Selker. Force-to-motion functions for pointing. In *INTERACT '90: Human-Computer Interaction*, pages 701–706, 1990. 323

[11] H. Z. Tan, M. A. Srinivasan, B. Eberman, and B. Cheng. Human factors for the design of force-reflecting haptic interfaces. In *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 353–359, 1994. 324

[12] F. J. Valero-Cuevas, H. Hoffmann, M. U. Kurse, J. J. Kutch, and E. A. Theodorou. Computational models for neuromuscular function. *IEEE Reviews in Biomedical Engineering*, 2: 110–135, 2009. doi: 10.1109/RBME.2009.2034981. 324

[13] S. Zhai. *Human Performance in Six Degree of Freedom Input Control*. PhD thesis, University of Toronto, 1995. 323