2011 IEEE International Conference on Rehabilitation Robotics
Rehab Week Zurich, ETH Zurich Science City, Switzerland, June 29 - July 1, 2011

Limit-push training reduces motor variability

Ian Sharp
Department of Bioengineering
University of Illinois at Chicago
Chicago, IL 60607
isharp2@uic.edu

James L. Patton, PhD
Rehabilitation Institute of Chicago and
University of Illinois at Chicago
Chicago, IL 60607
pattonj@uic.edu

***Abstract*—Variability in human motor control has been a long observed phenomenon, which has come to be known by some as repetition without repetition. There are several explanations for this. One such explanation asserts that many equally optimal solutions exist for accomplishing the same task that naturally allows choices in how it can be successfully executed. The aim of this study was to determine whether variability could be conditioned within an invisible subspace, using visual and force feedback. We utilized a novel haptic-graphic boundary-oriented environment to condition motor variability. Subjects reduced the variability of their movements, such that action predominated within a subspace determined apriori; while the untreated group did not. These results show encouraging preliminary evidence that neural rehabilitative haptic-graphic interfaces can condition human motor variability. This type of training may benefit neurologically impaired individuals, who exhibit the commonly seen motor deficits of large trial to trial variability, such as victims of stroke and traumatic brain injury.**

I. Introduction

A number of recent studies have shown remarkable evidence that interactive environments can influence learning and the way individuals control movement. Robotic technology is among the most powerful of such tools, where practicing in the presence of provocative interactions with sensory feedback can: improve performance, facilitate adaptation, and even restore function to brain injured individuals. However, little attention has been given using these approaches to reducing motor variability, even though reducing the errors associated with variability would have a profound impact on areas such as surgical training, sports performance, piloting, and neurorehabilitation. This paper draws on some recent findings related to enhancing feedback in order to influence the consistency of movement patterns, and provides preliminary proof-of-concept for conditioning motor variability.

Most researchers agree that through repetitive practice, conditioning is achieved by forming an internal model of the environment. For example, someone may learn the characteristics of damping while reaching in a windy environment and this model may change if the same individual were submerged in water. Yet, what is still not known are the exact mechanisms by which this model becomes influenced and the meaning behind varability of learned movements.

Human movement scientists speculate that there may be benefits to variability. One thought maintains that when several paths exist to attaining a single outcome, some variability in well learned skills may not be just noise, but rather reflects meaningful exploration [11]. Still another benefit to optimal levels of variability may be that it aids in the creation of a system that is more adaptable to perturbations [17]. From a mechanical standpoint, reduced variability has been known to increase the probability of repetitive stress injury [19]. While variable repetitive strategies minimize sensory degradation and preserve motor control; in the example of research that trained owl-monkeys on a gripping task [1].

Naturally, movement variability carries inherent costs as well. Persistent variability increases energy expenditure and may reduce performance. Physiologically, variability is endogenous to motor units, the basal ganglia circuits, premotor cortex, and in reaching movements [2] [6] [12] [22]. For these reasons, variability and its costs are never fully eliminated. Some pathological states and special interactions result in variability that exceeds desired levels, disrupting the natural balance between benefits and costs.

It is well documented that stroke and victims of stroke and TBI have increased motor variability [16] [18] [7]. Fitts law equates the information capacity of the human motor system to a trade-off between speed and accuracy, known as the speed-accuracy trade off [4]. For some individuals, their variability could be due to incoordination, decreased strength, absent sensation, or a combination thereof [11]. Consequently, they cannot control their movements as accurately.

MicroSurgery is another area where variability is higher than desired, due to tremor and low signal-to-noise of the system [20]. Areas where microsurgery is of notable influence are reinnervation prosthetics and plastic surgery. Physiological sensors often do not have the fidelity to allow feedback to appropriately reject errors. Resulting surgical errors may then lead to severed vasculature. One compelling question is whether variability is something frozen by physiological constraints as some models dictate or whether it can be changed, perhaps through robotic training.

If physiological sensors are a key element of learning and adaptation, the next question to ask would be whether or not artificially increasing the magnitude of such feedback would promote faster or more complete learning/adaptation. Some researchers have found that applying greater sensory feedback forces or visual cues does indeed provide adequate neurological stimulus to promote higher levels of adaptation/learning. This is known as ”Error Augmentation” (EA). EA may be due to the fact that once outcomes of a motor control action deviate from an ideal; our internal model self-adjusts according to the magnitude of error. Consequently, as the brain becomes better at modelling the external environment, error in task execution decreases. Yet no matter how complete our learning is motor performance always carries with it some inalienable degree of error and variability due to sensory, motor and environmental noise.

A compelling question is therefore whether the end-point motor variability can be reduced, because it is one impairement that prevents stroke and TBI patients from taking care of themselves. If stroke and TBI patients have decreased end-point accuracy, this implies that the information capacity of the human motor system in stroke patients diminishes post stroke. However, if the information capacity of stroke patients can be increased via robotic interactions, the possibility exists that robotic-assisted rehabilitation would reduce variability in stroke patients. Evidence is accumulating that supports this specific possibility [14] [8] [3] [23]. In addition, studies demonstrate robot assisted therapy to be more effective, when compared to non-robot aided therapy over specific time periods [23]. However, to our knowledge no studies aimed to condition variability.

Although equally optimal solutions exist for many tasks, this area of research remains largely unexplored. One mechanism most suited to condition variability, that was used in this study, was a simple haptic-graphic boundary-oriented approach. We a region to determine the efficacy of EA in conditioning end-point variability toward one specific subregion, where many end-point locations within that subspace would be equally optimal in terms of completing the task.

II. Methods

Subjects consisted of 18 healthy individuals (9 control and 9 treatment), ages 18 to 64, with no history of neurological or motor impairment. All subjects signed an IRB and were randomly divided into 2 groups, described below. This experiment utilized a three-dimensional, large-workspace haptics/graphics system called the Virtual Reality and Robotic Optical Operations Machine (VRROOM). VRROOM is an integrated system combining display environment, robotic forces, and tracking of limb movement. VRROOMs visual display system, the Personal Augmented Reality Immersive System (PARIS), was developed in the Electronic Visualization Lab at the University of Illinois at Chicago. PARIS, described in more detail here [13] is currently the highest quality system available (Fig. 1). This virtual reality environment which provided stereovision utilized the Barrett WAM robot from Barrett Technology Inc, MA for haptic feedback. The WAM was wrapped into the HAPI and H3DAPI libraries, where the H3DAPI was used to build our virtual reality environment.

The virtual reality scene consisted of 3 objects in black space: a blue projectile (1kg virtual mass and 5x5cm), a red parallelepiped to represent the position of the subject's hand (5x5 cm), and a semi-transparent green workspace sphere (12.5cm radius). The projectile's home position remained constant throughout the experiment. The direction of projectile launch was randomized with constraints which guaranteed that the projectile intersected the semi-transparent workspace sphere at some point, and traveled at a constant velocity of 0.8 m/s. This is important, because all subjects will be asked to keep their hand within the semi-transparent sphere workspace.

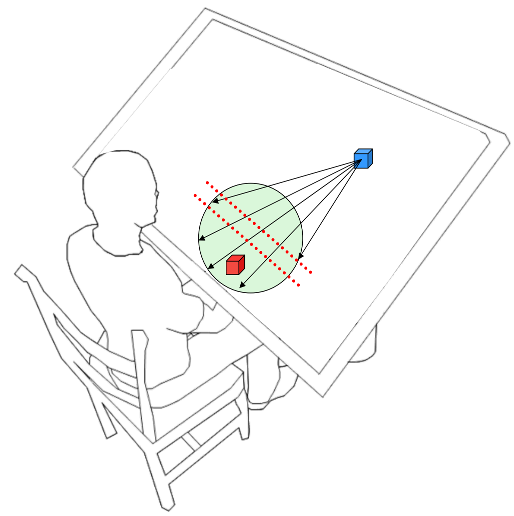

Fig. 1. Subject seated at the haptic/graphic apparatus. The red cube represents the subjects hand. The blue cube represents the projectile. Vectors from the blue cube represent possible launch directions of the projectile (always intersection the green workspace). The two dotted red parallel lines were not visible to the subject and represent the boundary region. This image is not drawn to scale.

The entire experiment consisted of 3 phases in sequence, where each phase was 200 trials long. In the first phase, both groups attempted to catch the virtual projectiles. In the second phase, referred to as the training phase, graphical cues and a haptic boundary were enforced when the treatment group moved their hand moved outside of the parallelepiped boundary. The visual cue consisted of the workspace sphere turning from green to red, while the haptic cue consisted of 4 Newtons of force pointing away from the center of the workspace. The control group did not receive haptic or graphic feedback during their training phase. The depth of the parallelepiped boundary was approximately 3.75cm, and the height extending ad infinitum. During the final phase, known as washout, feedback was removed from the treatment group, making conditions equivalent to the baseline phase.

Trials within this experiment were defined by the following finite state machine:

1) Wait - This state waited for 0.1 seconds before enabling the launch state.

2) Launch - this state launched the projectile by applying a series of pulse forces along a vector until the magnitude of the velocity reached 0.8m/s at which point all virtual forces acting on the projectile were turned off.

3) Caught - If the subject slowed the velocity of the projectile equal to or less than zero before the projectile exited the semi-transparent sphere workspace, the state projectile was

defined as caught otherwise it was defined as missed.

4) Stop-After-Launch - This state slowed the projectile to near zero velocity and then enabled the Chamber state.

5) Chamber - This state applied force on the projectile until it returns to the home position.

6) Stop-After-Chamber - This state slowed the projectile to near zero velocity and then enabled the Chamber state.

1) Error metrics: The error metrics used were: distance-to-edge, distance-to-center, fraction of time spent outside of the boundary, hand position at its closest point of contact along the anterior axis (ZPC), and time-to-edge (TTE). All error metric per phase were calculated by using a window of 20 trials. In order to determine whether subjects learned the boundary, we tested if they learned to keep safely away from the edge with the metric distance-to-edge. To determine whether subjects learned the center location of the region, we used the metric distance-to-center. In order to determine what percentage of the movement remained inside of the region we used the metric fraction of time spent outside of the region. Lastly, we explored whether subjects learned first order derivatives of the boundary region, which was tested by changes in the average time it would take the subject to reach the edge of the region. Variability was observed by analyzing the probability distribution of movement along the anterior axis. The fraction of time spent outside of the region was calculated by summing the total number of hand position observations that lied within the boundary for a given trial and dividing by the total number of hand position observations. Closest point of contact along the anterior axis was calculated by selecting the point at which the users hand was closest to the incoming projectile. The distance-to-edge was then calculated from the point of closest contact to the closest boundary edge. If the subject's hand was no longer within the boundary, the distance was marked as zero. Distance-to-center was calculated as the distance from the same point of closest contact to the center of the boundary. The mean time-to-edge was calculated for each trial by using the parametric equation of a line where, $P_t = P_0 + P_0 * t$ for every observation made of the subjects hand within every given trial. Lastly, differential entropy was calculated, described in more detail here in [9]. Entropy calculations were normalized with respect to the entropy of a uniform distribution and the bin number was determined using Izenman's nonparametric density estimation [5].

III. Results

As expected, during the baseline phase, both groups spent most of their movement outside of the boundaries and there was no difference between groups for any error metric. The control group spent 91% of their movement outside of the boundary whereas the treatment group spent 77% of their movement outside of the boundary. Both the distance-to-edge and distance-to-center during this phase remained low for both with no significant difference between groups (p = 0.2, p = 0.1, respectively). Although the time-to-edge was greater for the treatment group by .18 seconds (p = 0.03). Comparing mean-squared-errors on both the best-fit Gaussian and best-fit uniform distributions to the movement distribution along the anterior axis yielded no differences between groups (p = 0.08, p = 0.8). Within group comparisons showed that neither the control group nor the treatment group showed preference toward a best-fit of the uniform nor the normal distribution (p = 0.07, p = 0.3). Lastly, the entropy of the distribution of both groups also yielded no difference (p = 0.08), where the mean normalized entropy for the control group was 0.72 and the mean normalized entropy for the treatment group was 0.82.

By the end of training, the control group had a higher fraction of their movement outside of the boundary than the treatment group (p = 0.002). The fraction of time spent outside of the boundary was 62% for the control group, whereas the treatment group spent 17% of their movement space outside of the boundary. The treatment groups distance-to-edge was also greater than the control group (p = 0.0007). The control group had a distance-to-edge of 5.8 mm., while the treatment groups distance-to-edge was 7.6 mm. The treatment groups distance-to-center was less than the control group (p = 8e-5). The control group had a distance-to-center of 3.2 cm while the treatment had a distance-to-center of 1.5 cm. Lastly, at the end of training, the time-to-edge for the treatment group was less than the control group (p = 8e-5). The control group time-to-edge was 0.45 seconds, while the treatment group time-to-edge was 1.1 second. Within group comparisons showed that the control group showed no preference toward any distribution type, while the treatment group showed preference toward a uniform distribution (p = 0.07, p = 0.01); where the MSE for the uniform distribution of the treatment group was $2.6 * 10^4$, and the MSE for the normal distribution of the treatment group was $7 * 10^4$. The mean normalized entropy for the treatment group was 0.62 while mean normalized entropy of the control group was 0.78. By the end of training the treatment group has 0.16 less entropy (p = 0.02), more closely representing a normal distribution.

Comparing the end of baseline phase to the end of training phase determined whether groups improved after training. The control group did not show improvement in any of the following: fraction of time spent outside boundary, distance-to-edge, distance-to-center, time-to-edge (p = 0.8, p = 0.6, p = 0.7, p = 0.8). In contrast, the treatment group displayed improvement for all metrics. The fraction of time spent outside of the boundary decreased by 56% (p = 0.002), their distance-to-edge increased by 3.7 mm (p = 0.005), their distance-to-center decreased by 1.2 cm (p = 0.005), and their time-to-edge increased by 0.8 seconds (p = 0.004). Comparing mean-squared-errors on both the best-fit Gaussian and best-fit uniform distributions to the movement distribution along the anterior axis data yielded no differences in changes for the control group (control p = 0.5, p = 0.9). However, the treatment group became both less gaussian and uniform (p = 0.01, p = 0.03). The mean MSE of the treatment group's best fit gaussian distribution started at $1.3 * 10^4$ and increased to $7 * 10^4$, while the mean MSE of the treatment group's best fit uniform started at $1.1 * 10^4$ and increased to $2.6 * 10^4$. Entropy did not change for the control group (p = 0.29), yet decreased for the treatment group from 0.8 to 0.6 (p = 0.01).

Both visual and haptic feedback was then removed from the treatment group, where the subjects entered the onset of the washout phase. During the onset of washout, differences between groups persisted for the fraction of time spent outside of the boundary (p = 0.05), distance-to-edge (p = 0.05), and distance-to-center (p = 0.04). Compared to the control group, the treatment group spent 9% more of their movement inside of the boundary, was 2.5 mm farther from the edge, and 1.1 cm closer to the center. Differences did not persist for the time-to-edge (p = 0.1). The control group's distribution of movement along the anterior axis did not favor the best-fit of a Gaussian nor a uniform distribution (p = 0.4). However, the treatment group favored the uniform distribution over the Gaussian (p = 0.02), where the mean MSE of the gaussian distribution was $4.2*10^4$ and the uniform distribution was $1.8*10^4$. Lastly, the entropy of the treatment group was less than the control, where the control group had a mean entropy of 0.81 and the treatment group had a mean entropy of 0.70. Differences between groups achieved significance (p = 0.02).

Comparing the end of training phase to the end of washout phase determined whether groups show any evidence of learning. Neither the control group nor the treatment group showed any significant change for distance to edge, distance to center, fraction of time spent inside of the region, nor time to edge (control p =0.6, p =0.9, p =0.8, p =0.8) (treatment p =0.5, p =0.4, p =0.3, p =0.2). With respect to mean squared error of best-fits, the control group did not show a change in preference toward a uniform, nor a normal distribution (p = 0.2, p = 0.5). The treatment group, showed no change in the mean squared error for the best fit of a gaussian distribution either (p = 0.2), however the mean squared error decreased from $1.8 * 10^4$ to $1 * 10^4$, in favor of a more uniform distribution. Entropy did not change for the control group (p = 0.2), yet increased for the treatment group from a mean of 0.7 to 0.8 (p = 0.04).

During the final trials of the experiment, there was no difference when comparing groups for any metric. The after-effects washed out for the fraction of time spent outside of the boundary (p = 0.1), distance-to-edge (p = 0.07), distance-to-center (p = 0.09), and time-to-edge (p = 0.1). The control group spent 81% of their movement outside of the boundary whereas the treatment group spent 65% of their movement outside of the boundary. The control group distance-to-edge was 2.5 mm, whereas the treatment group's distance-to-edge was 5 mm of their movement outside of the boundary. The control group distance-to-center was 0.04 m, whereas the treatment group's distance-to-center was 0.025 m. The control group time-to-edge was 0.2 s, whereas the treatment group's time-to-edge was 0.3 s. The control group's distribution of movement did not favor a uniform or Gaussian distribution (p = 0.1). However the treatment group's distribution of movement appeared to favor a uniform distribution (p = 0.04), where the mean MSE of the uniform distribution was 1.0^4 and the mean MSE of the Gaussian distribution was 2.3^4. The control group did not have a movement distribution that was significantly more Gaussian when compared to the treatment group (p = 0.8), nor did they have a movement distribution that was significantly more uniform (p = 0.4). No significant difference in normalized entropy of the movement distribution was noted between the control and the treatment group (p = 0.4), where the control group's normalized entropy was 0.80 and the treatment group's normalized entropy was 0.83.

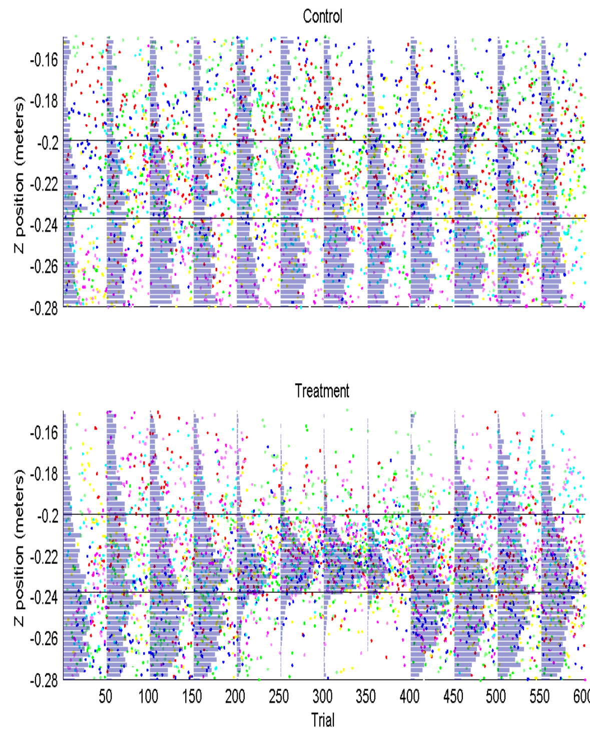

Fig. 2. The control group is on the top subplot while the treatment group is on the bottom subplot. Each colored dot represents a different subject. The location of the dot shows where the subject's came closest to contacting the projectile. The two horizontal lines display the location of the boundary. Blue semi-transparent histograms are overlayed on top of the raw data to display the movement distribution of the group along the anterior axis. The baseline phase is located within trials 1 to 200, the training phase is located within trials 201 to 400, and the washout phase is located within trials 401 to 600.

IV. Discussion

This study suggests that boundaries and safety margins can play an important role in learning. At the onset of training, subjects learn the boundary of the region. Movement outside of the boundary elicits an error augmenting force. Therefore subjects reduce their variability and tend to stay within the limits of safety. However, once haptic-graphic feedback are eliminated and subjects learn that being outside of the sub-space no longer imposes a punishing force on their appendage, their variability starts to increase outside the boundaries again.

Researchers such as Todorov and Jordan have proposed that optimal control theory of motor control allows variability to build in task-irrelevant dimensions of space [21]. Furthermore, Scholz and Schoener explain their uncontrolled manifold

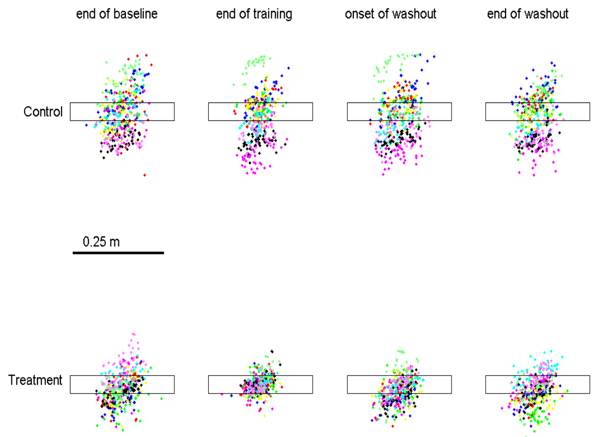

Fig. 3. Displays the point of closest contact in the XZ plane (view from above). The rectangle in each phase displays the safety region that treatment subjects were attempting to stay inside of while performing the task. Note that feedback was not delivered to the treatment subjects at the end of baseline, onset of washout, or end of washout.

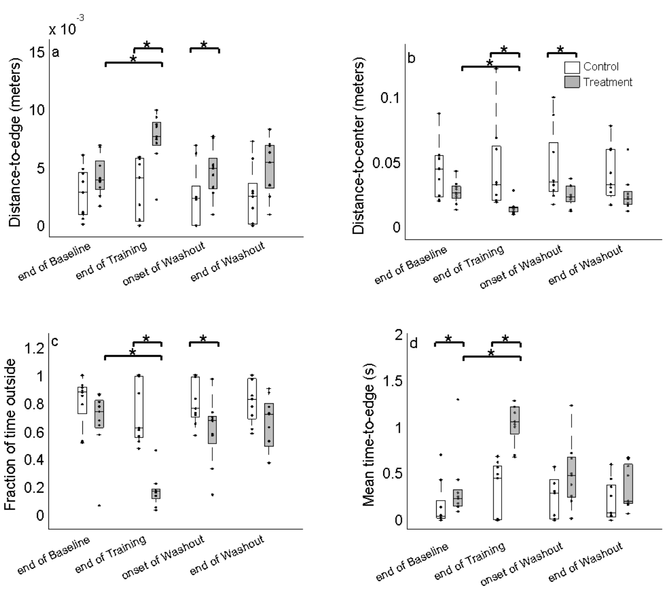

Fig. 4. Four different views of how Limit-push changes movement. a. Distance-to-edge. Displays the average distance-to-edge of the boundary at the closest point of interception. The treatment group moves farther from the edge by the end of training. b. Distance-to-center. Displays the average distance-to-center of the boundary at the closest point of interception. The treatment group moves closer to the center by the end of training. c. Fraction of time spent outside boundary. Displays what percentage of the movement was spent outside of the boundaries of the boundary. The treatment group spent more time inside the boundary by the end of training. d. Mean time-to-edge. Displays the average time-to-edge for each phase. The treatment group has an increased time-to-edge by the end of training. Asterisks show significance at p <= 0.05.

theory where a range joint postures result comparable task completion [15]. Similarly, this experiment shows that we may

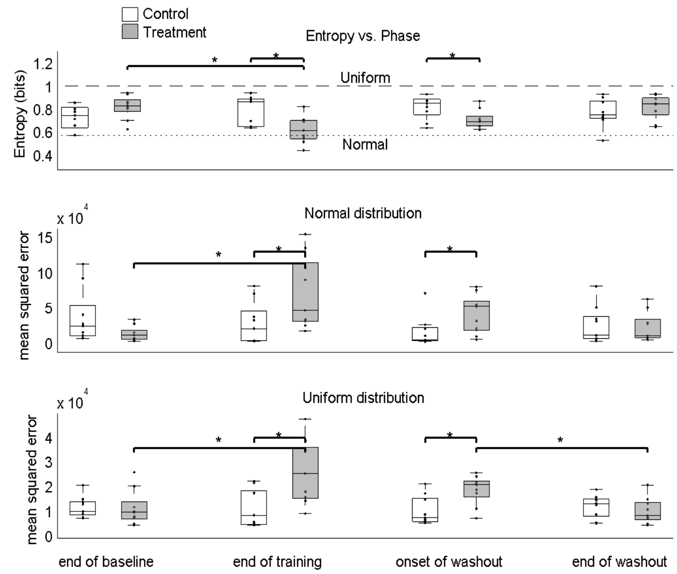

Fig. 5. Movement distributions of uniformity and normality change over the course of training. a. Each box and whisker plot displays the distribution of each subjects normalized entropy. The dotted horizontal line at the top represents the normalized differential entropy of a uniform distribution, and the dotted horizontal line at the bottom represents the normalized differential entropy of a normal distribution with a variance equal to the width of the region. b. Each box and whisker plot displays the mean squared error of the best-fit of a normal distribution for each subject's movement distribution. c. Each box and whisker plot displays the mean squared error of the best-fit of a uniform distribution for each subject's movement distribution.

manipulate variability in one dimension. However, position with respect to boundaries may not be the only parameter involved in this type of learning.

Our results suggest that derivative space within the boundary may also play a role in learning tasks, namely time-to-edge; rather than the nervous system relying soley on boundaries in Cartesian space. Lee demonstrated how time-to-edge could be used by drivers who were on a collision course (rather than just distance, speed, or acceleration/deceleration), and proposed implications of safe speeds and following distances [10]. It may be that time-to-edge plays a complimentary role in learning and unlearning paradigms. However, time-to-edge does not show evidence of after-effects. Yet, if error augmentation could modulate time-to-edge, it may be beneficial to observe the effects of modulating time-to-edge during the washout phase and whether that would affect variability.

A common challenge in therapeutic rehabilitation is maintaining motor performance once training is over and the subjects return to everyday tasks. Our data suggest that the effects of training were retained for at least 20 trials but had washed out by trial 200. Yet, this is typical for healthy subjects to 'de-adapt' when there is little motivation to keep altered states [13]. It remains to be seen if stroke and TBI patients persist. One way to add meaning to our results would be to decipher how to increase washout time. Neurological parameters which dictate rate of washout are still being elucidated. Such issues may be probed by attempting to decipher the mechanism of error adjustment and the meaning of optimal learning.

Whether our nervous system is optimal or not remains an area of debate between motor control scientists, and the definition of optimal is not agreed upon. Therefore, any debate over this topic would not provide researchers with any consistent logic. For example, one side may argue that a tightly packed Gaussian distribution of points would suggest optimality, because this indicates the nervous systems attempt to attain a single point in space, surrounded by noise. Another may argue to the contrary that a uniform distribution of points within a boundarys subspace would show greater optimality, since this minimizes computational effort - controlling a hand into a region rather than to a single point.

In this study, although the treatment groups distribution may have qualitatively looked more normal by the end of training, these distributions failed Gaussian classification under the kolmogorov-smirnov test. According to the best-fit of mean-squared error, distributions for the treated group were more uniform. Yet, according to entropy calculations, movement distributions more closely resembled Gaussian curves. These data may suggest that our nervous system may be neither purely normal nor uniform; neither center-seeking nor region-seeking, or that two objectives are being met.

V. Conclusion

This study study demonstrated preliminary evidence that humans were capable of learning an invisible region's subspace via force and visual feedback within a haptic-graphic environment. Through haptic and graphic biofeedback methods condition variability. Lastly, we wanted to know whether the distribution of movement within a redundant task workspace would appear more normally or uniformly distributed. Subject movement distributions tended to more closely approximate uniform distributions according to mean-squared-error best-fit approximations. However, in terms of information transfer (via entropy), the treatment group appeared to more closely resemble a gaussian distribution in terms of information transfer. The results of these data suggest that for this type of task, feedback may assist in conditioning variability and learning boundary subspaces . Furthermore, the conditioned variability resemble gaussian distributions with respect to entropy, while in cartesian space appear to more closely approximate a uniform distribution. What is clear is that training treatmens such as that demonstrated in this study provide encouraging new evidence for prospects of variability conditioning using haptic and graphic environments.

References

[1] N.N. Byl, M.M. Merzenich, S. Cheung, P. Bedenbaugh, S.S. Nagarajan, and W.M. Jenkins. A primate model for studying focal dystonia and repetitive strain injury: effects on the primary somatosensory cortex. *Physical therapy*, 77(3):269, 1997.

[2] M.M. Churchland, A. Afshar, and K.V. Shenoy. A central source of movement variability. *Neuron*, 52(6):1085–1096, 2006.

[3] S.E. Fasoli, H.I. Krebs, J. Stein, W.R. Frontera, and N. Hogan. Effects of robotic therapy on motor impairment and recovery in chronic stroke. *Archives of physical medicine and rehabilitation*, 84(4):477–482, 2003.

[4] P.M. Fitts. The information capacity of the human motor system in controlling the amplitude of movement. *J of*.

[5] A.J. Izenman. Recent developments in nonparametric density estimation. *Journal of the American Statistical Association*, 86(413):205–224, 1991.

[6] K.E. Jones, A.F.C. Hamilton, and D.M. Wolpert. Sources of signal-dependent noise during isometric force production. *Journal of neurophysiology*, 88(3):1533, 2002.

[7] D.G. Kamper, A.N. McKenna-Cole, L.E. Kahn, and D.J. Reinkensmeyer. Alterations in reaching after stroke and their relation to movement direction and impairment severity. *Archives of physical medicine and rehabilitation*, 83(5):702–707, 2002.

[8] HI Krebs, JJ Palazzolo, L. Dipietro, M. Ferraro, J. Krol, K. Rannekleiv, BT Volpe, and N. Hogan. Rehabilitation robotics: Performance-based progressive robot-assisted therapy. *Autonomous Robots*, 15(1):7–20, 2003.

[9] AV Lazo and P. Rathie. On the entropy of continuous probability distributions (corresp.). *Information Theory, IEEE Transactions on*, 24(1):120–122, 2002.

[10] D.N. Lee. A theory of visual control of braking based on information about time-to-collision. *Perception*, 5(4):437–459, 1976.

[11] D.B. Matchar, D.C. McCrory, H.J.M. Barnett, and J.R. Feussner. Medical treatment for stroke prevention. *Annals of internal medicine*, 121(1):41, 1994.

[12] B.P. Olveczky, A.S. Andalman, and M.S. Fee. Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLoS Biol*, 3(5):e153, 2005.

[13] J. Patton, Y. Wei, C. Scharver, R.V. Kenyon, and R. Scheidt. Motivating rehabilitation by distorting reality. In *Biomedical Robotics and Biomechatronics, 2006. BioRob 2006. The First IEEE/RAS-EMBS International Conference on*, pages 869–874. IEEE, 2006.

[14] R. Riener, T. Nef, and G. Colombo. Robot-aided neurorehabilitation of the upper extremities. *Medical and Biological Engineering and Computing*, 43(1):2–10, 2005.

[15] J.P. Scholz and G. Schoner. The uncontrolled manifold concept: identifying control variables for a functional task. *Experimental Brain Research*, 126(3):289–306, 1999.

[16] T. Sibylle, T. Phil, K. Laurence, H. David, and G. John. Movement variability in stroke patients and controls performing two upper limb functional tasks: a new assessment methodology. *Journal of Neuro-Engineering and Rehabilitation*, 6.

[17] N. Stergiou, R.T. Harbourne, and J.T. Cavanaugh. Optimal movement variability: a new theoretical perspective for neurologic physical therapy. *Journal of Neurologic Physical Therapy*, 30(3):120, 2006.

[18] C.D. Takahashi and D.J. Reinkensmeyer. Hemiparetic stroke impairs anticipatory control of arm movement. *Experimental brain research*, 149(2):131–140, 2003.

[19] J.B. Talmage and JM Mehorn. Why staying at work or returning to work is in the patients best interest. *Talmage JB, Melhorn JM. A Physicians Guide to Return to Work. Chicago, Ill: American Medical Association*, pages 1–6, 2005.

[20] R. Taylor, P. Jensen, L. Whitcomb, A. Barnes, R. Kumar, D. Stoianovici, P. Gupta, Z.X. Wang, E. Dejuan, and L. Kavoussi. A steady-hand robotic system for microsurgical augmentation. *The International Journal of Robotics Research*, 18(12):1201, 1999.

[21] E. Todorov and M.I. Jordan. Optimal feedback control as a theory of motor coordination. *Nature neuroscience*, 5(11):1226–1235, 2002.

[22] R.J. van Beers, P. Haggard, and D.M. Wolpert. The role of execution noise in movement variability. *Journal of Neurophysiology*, 91(2):1050, 2004.

[23] M. Van der Loos. Robot-assisted movement training compared with conventional therapy techniques for the rehabilitation of upper-limb motor function after stroke. *Arch Phys Med Rehabil*, 83, 2002.