


Improving short-term retention after robotic training by leveraging fixed-gain controllers

Dylan P Losey¹ , Laura H Blumenschein², Janelle P Clark¹ and Marcia K O'Malley¹

Abstract

Introduction: When developing control strategies for robotic rehabilitation, it is important that end-users who train with those strategies retain what they learn. Within the current state-of-the-art, however, it remains unclear what types of robotic controllers are best suited for promoting retention. In this work, we experimentally compare short-term retention in able-bodied end-users after training with two common types of robotic control strategies: fixed- and variable-gain controllers.

Methods: Our approach is based on recent motor learning research, where reward signals are employed to reinforce the learning process. We extend this approach to now include robotic controllers, so that participants are trained with a robotic control strategy and auditory reward-based reinforcement on tasks of different difficulty. We then explore retention after the robotic feedback is removed.

Results: Overall, our results indicate that fixed-gain control strategies better stabilize able-bodied users' motor adaptation than either a no controller baseline or variable-gain strategy. When breaking these results down by task difficulty, we find that assistive and resistive fixed-gain controllers lead to better short-term retention on less challenging tasks but have opposite effects on the learning and forgetting rates.

Conclusions: This suggests that we can improve short-term retention after robotic training with consistent controllers that match the task difficulty.

Keywords

Robot-assisted rehabilitation, control systems, neurorehabilitation, motor learning, haptic device

Date received: 3 January 2019; accepted: 5 July 2019

Introduction

Robotic rehabilitation can be an effective therapy for end-users that are suffering from long-term motor impairments following stroke or spinal cord injuries.¹ During robotic rehabilitation, a robot physically interacts with the human to encourage the correct performance of repetitive movements. After the robot is removed, users should not revert to their impaired behavior: instead, these users should remember how they moved with the robot, and then replicate these correct motions during their everyday life. Accordingly, when developing devices and methodologies for robotic rehabilitation, we want to ensure that the robot not only helps the user to learn the correct motion but also to *retain what they have learned after the robot is removed*.

Promoting long-term retention requires that rehabilitation robots carefully select their behavior when interacting with humans. In other words, we must program rehabilitation robots with the right *control strategy*. Within the current state-of-the-art, control strategies for robotic rehabilitation can be

¹Department of Mechanical Engineering, Rice University, Houston, TX, USA

²Department of Mechanical Engineering, Stanford University, Stanford, CA, USA

Corresponding author:

Dylan P Losey, Mechatronics and Haptic Interfaces Lab, Department of Mechanical Engineering, Rice University, 6100 Main St., Houston, TX 77251, USA.

Email: dylan.p.losey@gmail.com



divided into three main categories²⁻⁴: control strategies that constantly assist the user, control strategies that constantly resist the user, and control strategies that adapt their amount of assistance or resistance based on the user's performance. Strategies that constantly assist or resist are *fixed-gain* controllers, because here the robot's feedback strategy is static, and does not change between tasks. By contrast, *variable-gain* controllers—such as assist-as-needed control—adapt the robot's feedback strategy over time in response to the end-user's behavior.

To illustrate the difference between fixed-gain and variable-gain control strategies for robotic rehabilitation, consider a simple robotic controller that behaves like a spring. When the human makes mistakes and deviates from the desired trajectory, this controller applies a force in proportion to the user's error: if the controller's spring constant is positive, the robot helps by pulling the human back towards the desired trajectory, and if the spring constant is negative, the robot resists by pushing the human farther away from their goal. Here, fixed-gain assistance corresponds to a positive spring constant while fixed-gain resistance corresponds to a negative spring constant. Importantly, for these fixed-gain controllers, the robot maintains the same spring constant across tasks, regardless of the user's performance. Alternatively, the robot may start the training with a positive spring constant to help users move correctly, and then—as the human improves—the robot can gradually decrease the spring constant to incrementally increase task difficulty. Changing the spring constant between tasks is an example of variable-gain control, where the robot adapts its strategy online to match the user's capabilities.

Although prior research has introduced several instances of fixed and variable gain controllers for robotic rehabilitation, it is not yet clear which controller type(s) are suitable for encouraging long-term retention. In this paper, we take a first step towards addressing this issue by experimentally comparing how these different control strategies stabilize motor adaptation in the *short-term*, immediately after the robot is removed. We find that effective controller strategies should be challenging but consistent:

Robotic controllers that match the task difficulty but provide fixed responses lead to humans with better short-term retention of their motor adaptation

Overall, our work suggests that rehabilitation robots should select a control strategy appropriate for the user's skill and the task's difficulty, and then maintain that same fixed-gain controller throughout the training process. More specifically, we make the following contributions (Note: Parts of this work have been

published at the Conference on Biomedical Robotics and Biomechatronics⁵):

Extending reward-based reinforcement with robotic

controllers: Recent experiments on motor learning⁶⁻⁸ indicate that reinforcing the user's adapted behavior through rewards can improve their short-term retention. We extend these prior works to now include robotic controllers during training, so that the user receives kinesthetic feedback (from the robot's control strategy) in addition to auditory feedback (rewarding the user when they complete the task successfully). We apply this multimodal approach for robotic rehabilitation.

Conducting user studies with fixed and variable-gain

controllers: We introduce a visuomotor offset between the robot's actual position and the robot's displayed position and teach this offset to participants using no robotic controller, a fixed-gain controller, or a variable-gain controller. We then compare how well participants retain the visuomotor offset immediately after all robotic feedback is removed. In our first study, we find that variable-gain control does not result in significantly different behavior than using no controller. In the second study, we investigate whether variable or fixed-gain control better stabilizes the participants' motor adaptation.

Comparing control strategies across tasks of different

difficulty: To assess how the task difficulty affects our results, we consider two different adaptation settings: one task where the target position remains constant and a second task where the target position changes at each iteration. We expect that the control strategy should be tuned to match the task difficulty. For the easier task, participants should better retain their motor adaptation after training with a challenging fixed controller (i.e., fixed-gain resistance), and, for the harder tasks, participants should benefit from an assistive fixed control strategy (i.e., fixed-gain assistance).

In summary, we experimentally compare control strategies for robotic rehabilitation and focus on short-term retention of the human's learned behavior. This paper is a first step towards developing effective control strategies for long-term retention following robotic therapy.

Related work

Control strategies for robotic rehabilitation

Several control strategies have been proposed for upper-limb robotic rehabilitation.^{2,9} Here, we discuss three common types: fixed-gain assistance, fixed-gain

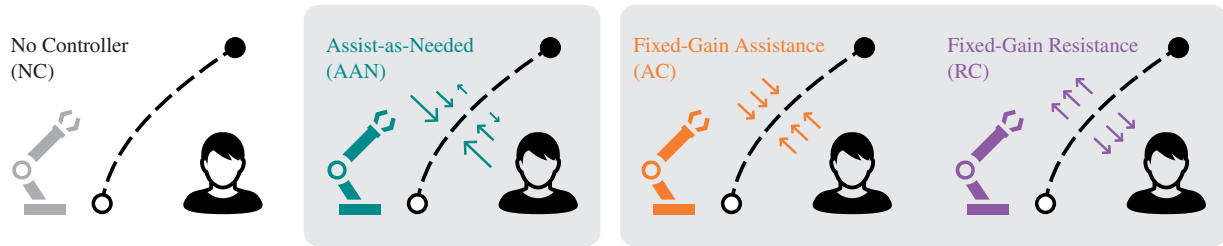


Figure 1. Fixed and variable-gain control strategies compared in our user studies. In the first user study, we trained participants with either NC or AAN (variable-gain control), and in the second user study, we trained participants with either NC, AC, or RC (fixed-gain control). The dashed line shows the desired path. Participants adapt to a visuomotor offset while receiving kinesthetic feedback from one of the listed controllers; we then remove all robot feedback and test the participant’s short-term retention.

resistance, and assist-as-needed control. Assist-as-needed control is an instance of a variable-gain controller (see Figure 1).

Fixed-gain assistance. These controllers guide and support the human throughout the task. Fixed-gain assistance can be implemented via impedance control with positive gains, so that—whenever the human makes a mistake, and deviates from the correct trajectory—the robot applies forces and torques to guide them back towards the right motion.^{10–12} Other works introduce a virtual tunnel around the correct path, enabling the human to complete the task with their own preferred timing¹³; if the human leaves the virtual tunnel, however, the robot begins to assist the user.¹⁴

Fixed-gain resistance. These controllers are instances of error augmentation, where the robot challenges the end-user by making their motion harder to complete. One method is to actively resist the human’s affected limb as a function of the movement velocity,^{15–17} so that the human must complete the task within a force field that exaggerates their mistakes. We note that both fixed-gain assistance and resistance lie along the same continuum, so that many of the approaches used for assistance can also be tuned to provide resistance, and vice-versa.^{18,19}

Assist-as-needed. In contrast to the prior strategies, assist-as-needed automatically modulates the robot’s interactions to correspond with the human’s capabilities. The robot trades off between maximizing human accuracy and minimizing robot effort: ideally, the human completes the task correctly, with as little robot assistance as possible.^{20–23} Finding the right amount of assistance has been addressed through optimization,²⁰ bounding tracking error,²² and learning human motion patterns.²³ By intervening only as necessary, the robot increases the user’s involvement, which is an important factor in facilitating recovery.³

Leveraging these strategies, we compare fixed-gain assistance, fixed-gain resistance, and assist-as-needed

control within a new paradigm, where we explore how the human maintains their motor adaptation after the robot is removed.

Motor learning and robotic rehabilitation

Fundamental research on motor learning provides several insights for robotic rehabilitation and controller selection. Below we introduce motor learning, and then overview how reward-based feedback and altering task difficulty can affect the human’s short-term retention.

Motor learning. Research on motor learning studies how humans make accurate movements: by using sensory feedback and prior experience, humans develop adaptive models of their body and the environment.²⁴ Understanding motor learning is important for neuro-rehabilitation in general and robotic rehabilitation specifically,^{25,26} since improved models of the recovery process can be applied to develop better training methodologies. We are particularly interested in methods to enhance the retention of motor learning.

Reward-based reinforcement. Some recent motor learning research suggests that a supervised approach—where the user receives a reward signal after they successfully perform the task—improves the retention of motor learning.^{6–8} Shmuelof et al.⁸ show that adapted behavior can be stabilized by training with binary reward feedback which notifies the user if they have succeeded. Galea et al.⁶ separately consider punishment and reward: negative feedback causes the user to learn faster, but positive feedback results in better retention after the feedback is withdrawn. We apply these experimental designs to robotic rehabilitation²⁷ and explore how adding kinesthetic feedback from the robot’s controller alters the human’s short-term retention. We recognize that—in practice—combining the kinesthetic signal from the controller with the auditory signal from the reward is an instance of multimodal feedback.⁴

Task difficulty. When applying robotic controllers to teach human users, the task difficulty can affect the resultant motor learning. Intuitively, the challenge should match the current user's capabilities¹⁸; prior work has shown that less skilled users better adapt to the task with fixed-gain assistance, and more skilled participants benefit from fixed-gain resistance.²⁸ Beyond challenge level, the task type can also be an important factor when determining the right controller,²⁹ so that the same control strategy may have different effects on motor learning when applied to new users and tasks.

Building on this prior work, we augment robotic control strategies with binary reward-based reinforcement to help humans retain what they have learned, and we test two task difficulties within our user studies to see how the user's short-term retention is influenced by the task difficulty.

Methods

Experimental overview

In this section, we describe two user studies that assess how fixed and variable-gain robot control strategies stabilize the participant's motor adaptation. We follow the experimental protocol introduced by Shmuelof et al.,⁸ where reward-based reinforcement is tested with able-bodied end-users. Within our user studies, participants physically interact with a kinesthetic haptic device and try to complete planar reaching motions in which they move the haptic device to the desired goal position. These motions are challenging, however, because we add a visuomotor offset between the robot's displayed position (that the user can see) and the robot's actual position (which is occluded). Participants must therefore adapt their behavior to compensate for the unknown visuomotor offset.

To adapt to this perturbation, participants train with a robotic rehabilitation control strategy: either no controller, fixed-gain assistance, fixed-gain resistance, or assist-as-needed. After their training is completed, we remove all robotic feedback—so that users are always told they are performing the task correctly—and explore whether the adapted behavior decays to the baseline behavior, or if the participants continue to follow the visuomotor offset. Overall, we manipulate both the control strategy and the task difficulty and assess how these factors relate to short-term retention of an artificial visuomotor offset with able-bodied end-users.

Independent variables

We varied the control strategy with four levels: *no controller* (NC), *assist-as-needed* (AAN), *fixed-gain*

assistance (AC), and *fixed-gain resistance* (RC). Both AC and RC are fixed-gain approaches, while AAN is a variable-gain controller.

We also varied the task difficulty: in *Easy* tasks, the human has the same goal position during every reaching trial, while in *Hard* tasks, we change the goal position between trials. The results of our experiments demonstrate that it is more challenging to adapt to the visuomotor offset when the target changes (i.e., the *Hard* task) as compared to a training with a stationary target (i.e., the *Easy* task).

Dependent measures

For every control strategy and task difficulty, we measured objective outcomes that captured each participant's training and retention.

Error. We measured the participant's actual hand direction at every trial and compared that to the 30° visuomotor offset (that we wanted the human to learn). Ideally, the difference is zero, indicating that the human has adapted to this offset. We computed the root-mean-squared error (RMSE) across trials to determine the participant's motion error. Thus, *error* here refers to the RMSE between the actual and desired hand directions.

Training. During the training portion of the user studies, we recorded the participant's *success rate* (i.e., how frequently they reached the desired goal) as well as their *learning rate* (i.e., how quickly they adapted to the visuomotor offset).

Retention. In addition to the error, we also used a *decay rate* to assess how rapidly participants reverted to some steady-state behavior when feedback was removed.

It is important to separate training and retention measures, since control strategies that lead to improved training may not result in better retention. We assess the *error* metric separately for training and retention blocks. To avoid including transient behavior in these measurements (where the human is still learning or decaying), we only compute the *error* over the final 80 training trials and final 50 error clamp trials. These thresholds were selected based on data from pilot users.

Control strategy

We controlled the 3 degree-of-freedom haptic device to track a desired path in task space.³⁰ Let $x \in \mathbb{R}^3$ be the robot's current position in task space, and let $x_d \in \mathbb{R}^3$ be the closest point along the straight line reaching motion between start and goal. The robot applies

task space forces $u \in \mathbb{R}^3$ according to the impedance control law¹¹

$$u = K \cdot (x_d - x), \quad K = \text{diag}(K_x, K_y, K_z) \quad (1)$$

where $K \in \mathbb{R}^3$ is a gain matrix. For the fixed-gain control strategies, K was constant throughout training, but—for the variable-gain approach—we updated K between iterations. Intuitively, increasing the magnitude of K causes the robot to provide more assistance, while negative values of K push the human away from the desired path. Note that we did not include a damping term in equation (1): our control law was a proportional controller, which was naturally damped by friction within the experimental setup. A damping term could be added as needed.

The fixed- and variable-gain control strategies leveraged in our user studies are depicted in Figure 1.

Variable-gain control

In our *first user study*, we compared NC to AAN during Easy and Hard tasks. We implemented the AAN case by updating the gain K in equation (1) based on the human’s capability during the previous trial. Let $t \in \mathbb{Z}^+$ be the trial number; then, under AAN, the variable-gain at the next iteration is

$$K(t+1) = K(t) + \alpha(t) \cdot \text{diag}(1,1,0) \quad (2)$$

Here, $\alpha(t) = 3$ N/m if the participant failed to reach the goal (and therefore needs more assistance) and $\alpha(t) = -1$ N/m if the participant successfully reached the goal (and therefore can accept more challenge). We bounded the magnitude of K_x and K_y such that $K_x, K_y \in [-20, +20]$ N/m, and K_z was fixed at 500 N/m to ensure that participants moved in a plane. We selected these values based on offline preliminary experiments, where they were tuned heuristically. Note that NC did not have any control feedback: $K_x = K_y = 0$ N/m.

Fixed-gain control

In our *second user study*, we compared NC to AC and RC during Easy and Hard tasks. We again controlled the robot using equation (1), but here we kept K constant, so that users receive consistent feedback. During AC, the robot continually assisted the user towards the goal: $K_x = K_y = +20$ N/m. By contrast, during RC, the robot exaggerated the human’s mistakes, pushing participants away from the straight line path: $K_x = K_y = -20$ N/m. We selected the gain values for AC and RC to be consistent with the bounding values of AAN from our first user study.

Experimental setup

Participants physically interacted with a haptic device (Touch X, 3D Systems) while observing a computer monitor. This monitor rendered a 1:1 representation of the participant’s reaching motion, and displayed the robot’s offset position, as well as the start and goal positions. Because a curtain occluded the participant’s view of the haptic device, they had to rely on the displayed positions when making decisions. This setup allowed us to introduce visuomotor offsets, where the displayed robot position was a rotation of its actual position in task space. Participants also wore headphones, which provided auditory reward-based reinforcement by playing a pleasant sound when the human reached the goal. We implemented the visualization and robot controllers with MATLAB/Simulink (MathWorks) and QUARC (Quanser). The experimental setup is shown in Figure 2.

Task and procedure

The experimental task consisted of repetitions of a reaching motion, where the participant moved the haptic device from a start to goal position. We refer to each reaching motion as a trial and the set of all trials as the task.

Reaching motion. At the beginning of each trial, the current participant held the robot at its start position. After a variable time interval (1.5 ± 0.5 s), a goal position appeared 80 mm from the start; participants then physically guided the robot towards this goal. If the displayed cursor intersected the goal, the




Figure 2. Experimental setup for our user studies. The participants perform reaching motions while grasping the haptic device, which employs a robotic control strategy (1). The screen displays the robot’s current position, with an added visuomotor offset (2). Participants receive auditory reinforcement that rewards them when they complete the task correctly (3).

trial was a success. The goal position was then erased, and, after another variable time interval (1.5 ± 0.5 s), the robot autonomously moved back to the start. Participants performed 40 unrecorded reaching motions before the task to become familiar with our setup.

Overall, the experimental task consisted of four blocks:

Baseline. A familiarization block with 20 recorded trials. We used this block to assess the participant's initial performance without any visuomotor offset.

Training. We next introduced a 30° visuomotor rotation, and participants adapted to this offset over a total of 140 trials. During these trials, participants received kinesthetic feedback from the robot controller and auditory reinforcement if they completed the trial successfully. Within the first 60 trials, we also provided visual feedback (so that participants could see the current robot position), and then we removed all visual feedback during the last 80 trials, so that users depended on the kinesthetic and auditory feedback to determine whether their reaching motion was successful.

Disturbance. To perturb the adaptation to the original offset, participants completed 30 trials with a 45° visuomotor offset and visual feedback.

Error clamp. We then removed the robot's feedback and measured if participants returned to the trained 30° offset or if they reverted to their baseline behavior (i.e., no offset). The error clamp block consisted of 100 trials, where, during each trial, the displayed robot position reached the goal regardless of the human's actual motion. Leveraging error clamp trials is common in studies examining the short-term retention or decay of learned motor behaviors.^{6-8,24}

The task procedure is outlined in Figure 3. Participants moved from baseline, to training, to disturbance, and, finally, to error clamp blocks. Our different control strategies (NC, AAN, AC, or RC) and task difficulties (Easy or Hard) were implemented during the *training* block. We tested the short-term retention during the *error clamp* block.

Easy and Hard tasks. In order to determine how the task difficulty affected retention, we tested two task types: Easy tasks and Hard tasks. The only difference between the Easy and Hard tasks was the location of the target during the baseline and training blocks. Within the Easy task, the target position was constant and did not change between trials (the target was always located at 135°). During the Hard task, the target position was randomly chosen at each trial (the target was randomly assigned between 90° and 180°). For both the Easy and Hard tasks, the target was at a radius of 80 mm from the start position. These Easy and Hard tasks are based on related works in motor learning.^{8,31}

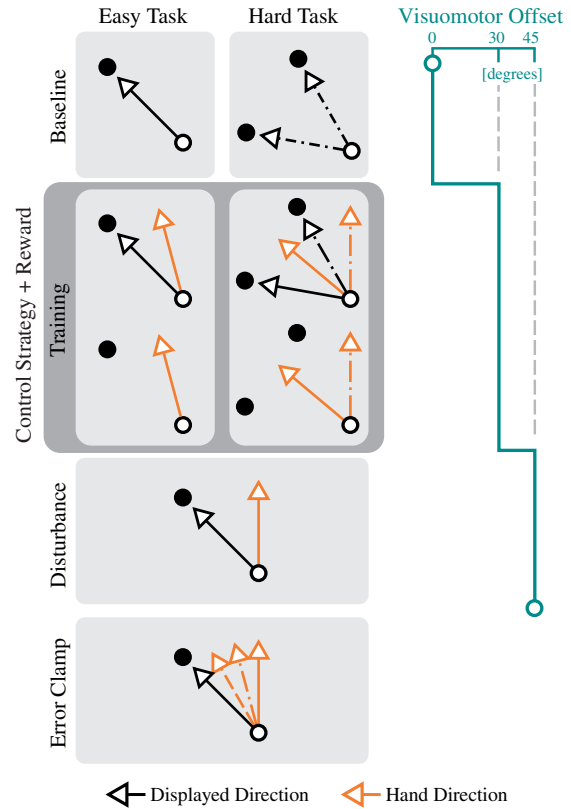


Figure 3. Outline of the experimental procedure, where participants complete four blocks of trials: baseline, training, disturbance, and error clamp. During the Easy task, the goal position is constant, but for the Hard task, the goal randomly changes at each iteration (during baseline and training). Training is split into two parts: first, participants get visual feedback in addition to the kinesthetic controller and the auditory reward, and then the visual feedback is removed.

Participants

Our participant pool consisted of 66 Rice University affiliates (aged 21.0 ± 3.4 years, 20 females) who provided informed written consent. Our user study was approved by the Rice University Institutional Review Board (IRB-FY2018-29). None of the participants had known neurological impairments, and all identified as right handed.

Of these 66 participants, 37 were involved in the first user study (aged 19.8 ± 0.7 years, 11 females), and the remaining 29 took part in the second user study (aged 22.6 ± 4.5 years, 9 females). The participants were divided into groups based on the robot's control strategy, i.e., NC, AAN, AC, or RC. Within our first user study, half of the participants trained on the Easy task and half trained with the Hard task. During the second user study, participants completed the experimental task twice: once with the Easy task and once with the Hard task. We counterbalanced the order of task

presentation (i.e., half started with the Easy task and half started with the Hard task), and we separated the two task sessions by a minimum of three days to mitigate between-task learning.

Learning and decay models

To measure the participant’s learning rate during training and decay rate during the error clamp block, we applied models to the human’s hand direction. These models are consistent with prior works on motor learning.^{7,6,32}

Learning model. Let $z \in [0,360^\circ)$ be the human’s estimate of the visuomotor rotation, and let $r = 30^\circ$ be the actual visuomotor rotation. We model the human’s adaptation as

$$z(t + 1) = Az(t) + B[r(t) - z(t)] \quad (3)$$

$$\hat{y}(t) = -z(t) \quad (4)$$

where $A \in [0,1]$ is a forgetting factor, $B \in [-1,1]$ is the learning rate, and \hat{y} is the predicted hand direction. We solved for the *learning rate* (B) offline by finding the model parameters that minimized the total squared error between the predicted (\hat{y}) and measured (y) hand directions.

Decay model. During the error clamp block, we similarly fit an exponential function to the human’s hand directions

$$\hat{y} = C_1 \cdot \exp(-\lambda t) + C_0 \quad (5)$$

where $C_0, C_1 \in \mathbb{R}$ are constants and $\lambda > 0$ is the decay rate. Like before, we solved for the *decay rate* (λ) by finding the model parameters that minimized the total squared error between the predicted (\hat{y}) and measured (y) hand directions. We point out that this decay model is leveraged in prior motor learning works,^{7,32} and equation (5) is equivalent to the learning model described by equations (3) and (4) for the specific case where $B=0$ and $C_0 = 0$.⁶ Hence, we can think of the decay model as a setting where the human stops learning (i.e., the learning rate is zero), and we are explicitly focused on determining the forgetting factor A , which we here refer to as the decay rate λ to prevent confusion.

Data analysis

We excluded a participant’s data if they either (a) noticed the error clamp or (b) decayed in the wrong direction (i.e., their hand direction became increasingly negative during the error clamp block). Only a single

participant reported noticing the presence of the error clamp (in the second user study), and a total of six participants had a negative decay rate (two in the first user study and four in the second). Hence, there were 35 subjects remaining in the variable-gain user study analysis and 24 subjects in the fixed-gain user study analysis. We performed our data analysis with SPSS (IBM).

Results

Variable-gain control

The objective results from our first user study are displayed in Figures 4 and 5. We interpret these results below:

Error. We leveraged a mixed analysis of variance (ANOVA), where controller strategy and task difficulty

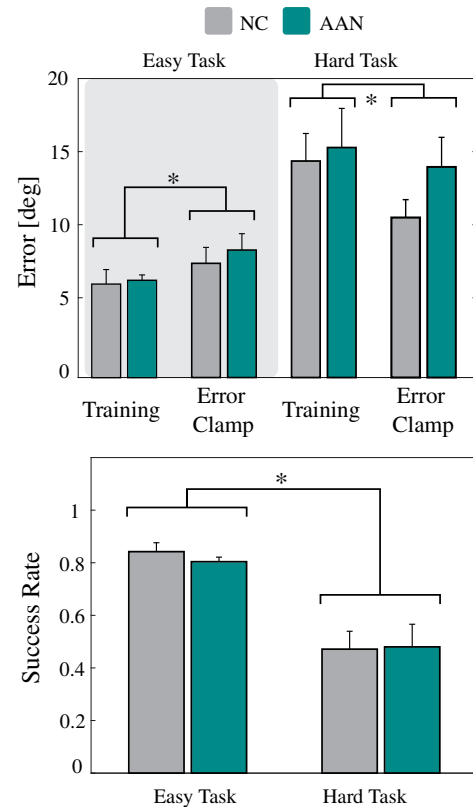


Figure 4. Results from our user study with no controller (NC) and an assist-as-needed variable controller (AAN). Top: root-mean-squared error from the desired visuomotor offset during the final trials of the training and error clamp blocks. Bottom: success rate during training. Viewed together, participants found the Hard task more challenging (increased error in training and lower success rate). Error bars show standard error about the mean (SEM), and an asterisk denotes statistical significance ($p < .05$).




Figure 5. Variable gain term with the assist-as-needed control strategy during the training block. This gain is updated with equation (2), and we here plot $|K|$. Positive gain values indicate assistance, while at negative gain values the robot resists the human. Based on the human's demonstrated capability, the robot updated its gain to increase the challenge during the Easy task, but decrease the challenge during the Hard task. Shaded areas show SEM.

were between-subject factors, and the block (training or error clamp) was the within-subject factor. Control strategy did not have a significant main effect ($F(1,31) = 1.8$, $p = .19$). However, there was a statistically significant interaction between task difficulty and block ($F(1,31) = 15.9$, $p < .001$). We thus reran our ANOVA with simple main effects: participants in the Easy task had significantly increased error during the error clamp as compared to training ($p = .05$). The opposite trend occurred in the Hard task, with less error during the error clamp as compared to training ($p < .01$).

Success rate. Control strategy did not significantly affect success, but within the Hard task, participants found it more challenging to reach the target, as evidenced by their lower success rate ($F(1,31) = 38.4$, $p < .001$).

Controller gain. The gain values for AAN also indicate that the Hard task was more challenging (see Figure 5). We observe that participants on the Easy task converged towards resistive feedback—i.e., a negative gain—while participants on the Hard task required positive assistance ($p < .001$). This result is in line with challenge point theory in motor learning literature, where learning is best at a specific challenge level.^{18,28} The assistive control gains allow better learning and retention for the Hard task because they match the harder difficulty, while resistive control gains add some challenge for the Easy task that may improve learning and retention.

Learning and decay. There were no statistically significant results for learning rate or decay rate.

Summary (NC vs. AAN). During both tasks, AAN and NC led to similar short-term retention; although AAN resulted in higher error across the board

(also see Table 1), this increase was not statistically significant. Based on the results for error, success rate, and AAN controller gain, we confirm that the Hard task was more challenging than the Easy task.

Fixed-gain control

The results from our second user study are summarized in Figure 6, and a side-by-side comparison of the error between our first and second user studies is listed in Table 1. We separately discuss these results below:

Error. We performed a mixed ANOVA, with control strategy and presentation order as between-subject factors, and task difficulty and block as within-subject factors. We found that the control strategy had a significant main effect ($F(2,18) = 5.1$, $p < .05$). For the Easy task, there was also a significant interaction between the control strategy and the task block ($F(2,18) = 8.2$, $p < .01$). Contrasts are used for the primary hypotheses of this experiment, i.e., that AC leads to improved retention for the Hard task and RC results in improved retention for the Easy task. Participants who trained with AC or RC had significantly less error than those trained with NC during error clamp trials for the Easy task ($p < .05$ for both). During the Hard task, neither controller led to significantly different error than NC.

Similar to the first user study, there was also an interaction between the task difficulty and block ($F(1,18) = 10.9$, $p < .01$). For the Hard task, participants had less error during the error clamp trials than during the training trials ($p < .01$).

Success rate. Control strategy did not have any effect on success during training; however, the success rate during the Hard task was significantly lower than during the Easy task ($F(1,18) = 78.7$, $p < .001$).

Learning rate. The control strategy had a significant effect on the participants' learning rate ($F(2,18) = 8.9$, $p < .01$). Contrasts showed that RC resulted in slower learning than NC ($p < .001$) and AC ($p < .01$).

We additionally found that the overall learning rate was higher for the Easy task when compared to the Hard task ($F(1,18) = 8.8$, $p < .01$).

Decay rate. The control strategy also affected the decay rate during error clamp trials ($F(2,18) = 3.7$, $p < .05$). Here, contrasts determined that AC decayed faster to the steady-state than NC ($p < .05$) across the board.

Summary (NC vs. AC and RC). Control strategy had a significant effect on how participants adapted and maintained the visuomotor offset. During the Easy task, subjects that trained with the RC or AC controllers

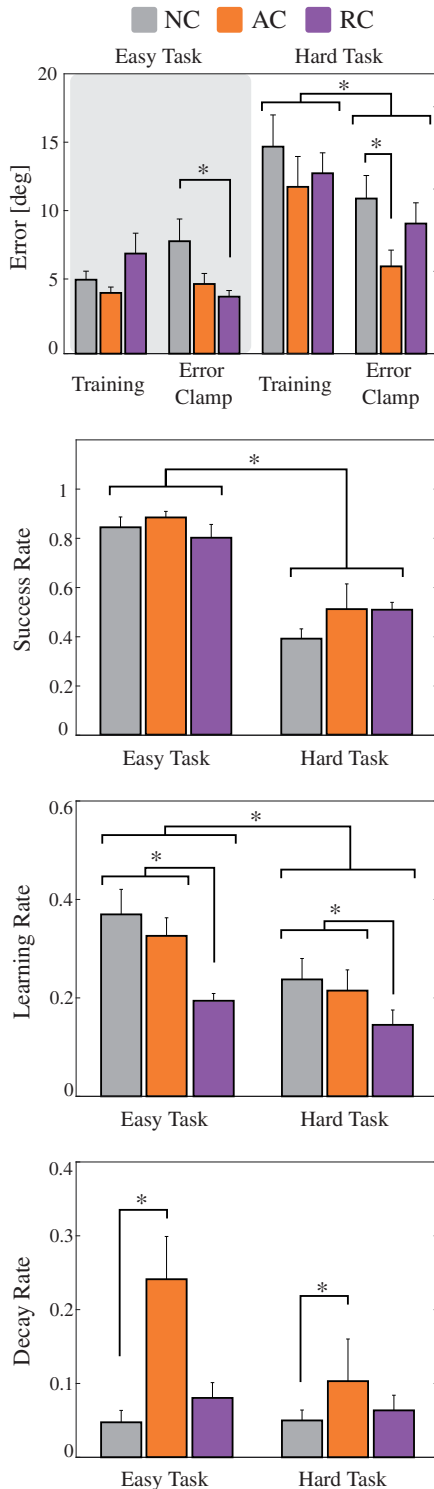


Figure 6. Results from our user study with fixed-gain controllers: no controller (NC), fixed-gain assistance (AC), and fixed-gain resistance (RC). Training with either RC or AC led to better short-term retention than NC during the Easy task. RC had a lower learning rate across the board, and AC decayed faster to steady-state than NC. The Hard task was more challenging, as evidenced by the decreased success rate and learning rate, and the increased error during training as compared to error clamp.

Table 1. Error across control strategies and task difficulties for both of our user studies: variable-gain control and fixed-gain control.

	Variable gain		Fixed gain		
	NC	AAN	NC	AC	RC
Easy task					
Training	5.69	5.93	5.18	4.34	7.15
Error Clamp	7.29	8.14	7.72	4.98	4.07
Hard task					
Training	14.44	15.15	14.12	11.92	12.89
Error Clamp	10.91	14.09	11.01	6.24	9.30

Note: We list the average root-mean-squared error (in degrees) between the desired and actual visuomotor offset at the end of the training and error clamp blocks. Error clamp scores display the users' short-term retention after training with the listed controller. NC: no controller; AAN: assist-as-needed; AC: fixed-gain assistance; RC: fixed-gain resistance.

had better short-term retention than the NC baseline. We also found that participants learned the visuomotor offset more slowly with RC, but—after all robotic feedback was removed—they more quickly reverted to their adapted behavior after AC. As before, the Hard task was more challenging than the Easy task: participants had lower success and learning rates and higher training error in the Hard task.

Discussion

The results of our user studies suggest that training able-bodied subjects with the right robot control strategy improves short-term stabilization of motor adaptation after the robot is removed. These control strategies should be consistent but challenging: fixed-gain controllers that match the task difficulty outperform variable-gain controllers that adjust based on the user's capability.

For the less difficult task, subjects who trained with fixed-gain controller (assistance or resistance) had better performance during the retention trials than those trained without a robotic controller. During the same task, training with a variable-gain controller (assist-as-needed) led to similar retention as training without a controller. For the challenging task, the training control strategy did not have a statistically significant effect on short-term retention: however, when compared to training without a controller, we found that the fixed-gain approaches led to lower error, while the variable-gain controller increased error (see Table 1). We note that—for each of these control strategies and task difficulties—the robot also provided auditory reinforcement throughout training, which rewarded users when they completed the motion correctly.

Overall, our findings demonstrate an interplay between multimodal feedback, controller type, and task difficulty during robotic training. We separately discuss each of these issues below.

Robotic control and reward-based reinforcement

During our user studies, participants trained while receiving kinesthetic feedback from the robotic controller in addition to auditory feedback that reinforced successful reaching movements. These experiments extend prior works on motor learning⁶⁻⁸: in the previous works, the authors compare learning with visual feedback plus reinforcement to learning with reinforcement alone. Importantly, their results suggest that including visual feedback can harm short-term retention. The human may become overly dependent on the presence of visual feedback, since this feedback continually informs them about their error from the correct motion.³³ We point out that—within our user studies—kinesthetic feedback from the robotic controller is similar to this visual feedback, as it continually provides the user with directional information about their distance from the straight line path.

Based on our results, however, we conclude that training with kinesthetic feedback has a different effect on short-term retention than training with visual feedback. Even though the control strategies provided continual information during each reaching motion, users trained with the best-suited controller outperformed users trained with reinforcement alone (i.e., the no controller case). Hence, not only can we combine robotic control strategies with reward-based reinforcement, but we can also improve short-term adaptation by leveraging robotic controllers instead of auditory reinforcement alone. This result is in line with works on multimodal feedback. While auditory and visual feedback can cause detriments to learning when combined—as they share many cognitive pathways—kinesthetic and auditory feedback have been effectively combined to reduce cognitive workload and convey more complex feedback.⁴ We might expect similar results if visual feedback was used for the reward instead of auditory, as long as the addition does not confuse the visual information already being received.

Fixed and variable-gain controllers

Our user studies suggest that fixed-gain robotic controllers are better suited for short-term retention than variable-gain control. Indeed, in the Hard task, using variable-gain control resulted in worse retention than a no controller baseline.

In order to understand why the variable-gain approaches were less effective, we recall that fixed-

gain controllers follow the same feedback direction at each iteration while maintaining a constant gain, and variable-gain controllers not only have variable control gains, but are also capable of switching the direction of the force feedback based on the human's performance. Users training with fixed-gain assistance or fixed-gain resistance received the same feedback when they made the same mistakes. By contrast, users training with the variable-gain control strategy interacted with an adaptive robot, so that the same mistakes could result in different kinesthetic feedback during different trials. Overall, we suggest that this variability is one explanation for why assist-as-needed control is not as well suited for promoting short-term retention.

More specifically, learning the visuomotor offset while simultaneously adapting to the robot's changing controller gains may have confused users within the variable-gain group and distracted these users from internalizing the visuomotor offset. For example, users may have associated the offset with the variable controller, so that, after we removed the kinesthetic feedback, users also forgot the offset.³⁴ Alternatively, participants with variable-gain control may have been unsure about which kinesthetic forces would result in reward (since the strength and direction of the robot's force field changed over time). Especially when switching between assistance (positive) and resistance (negative) gains, any learned response to the controller force field will likely lead to fewer rewards, not more. Explicitly connecting actions to reward is important for effective reward-based reinforcement,⁶⁻⁸ and the variable-gain controller may have inhibited this connection.

By contrast, the fixed-gain controllers may have improved retention because they provided consistent feedback, which was easily understood by the users. As expected, the fixed-gain resistance strategy resisted users during training and resulted in lower initial learning rates: users had to overcome this resistance before completing the task correctly. Fixed-gain assistance appeared to strongly reinforce the user's adapted behavior, where users more quickly reverted to what they had learned after the robot was removed. We suggest that fixed-gain assistance led to faster reversion since users with this robotic assistance had less variable movements during training.

Our findings for fixed and variable-gain controllers are supported by some recent studies.³⁵⁻³⁷ Within these works, variable-gain controllers do not outperform fixed-gain controllers and fail to increase the user's retention.^{36,37} Our results complement these more recent works and also highlight the connection between fixed-gain control and short-term retention. We recognize, however, that there are situations where variable-gain controllers are better suited than fixed-gain

controllers, particularly when the robot needs to maintain user participation over long-duration training.

Matching controllers to task difficulty

Within our user studies, the subjects trained on tasks of two different difficulties: an Easy task and a Hard task. We verified that training on the Hard task was more challenging, where subjects had lower success rates and higher error. Interestingly, we found a relationship between task difficulty and the best control strategy. Specifically, short-term retention improved on the Easy task after training with either fixed-gain controller, while fixed-gain resistance led to slower learning and fixed-gain assistance led to faster decay across the board.

Similar to previous research on robotic training,^{18,28} we suggest that resistive control strategies were better suited for less challenging tasks, while assistive control strategies appear to lead to increased retention on more challenging tasks (also see Table 1). One key novelty here, however, is that this pattern applies to short-term motor adaptation, where the users are trained through reward-based reinforcement. In practice, we recommend that the designer use the subject's baseline behavior to tune the control strategy, so that—after a few trials without any robotic intervention—the designer can select the controller that corresponds to the user's ability in the presented task.

Applying our results to robotic rehabilitation

The experiments in this paper were performed on able-bodied end-users; however, we are also interested in how our results may translate to robotic rehabilitation, where the users suffer from long-term motor impairments. Several related works on motor learning and retention have similarly conducted user studies on able-bodied subjects with implications for the rehabilitation of impaired users.^{6,8,28,33,34} Although we recognize that extending results from training to rehabilitation is not always straightforward,²⁵ we point out that some existing studies with impaired users appear to support some of our findings. For instance, Frullo et al.²¹ recently used a variable-gain assist-as-needed controller to perform robotic rehabilitation on subjects with incomplete spinal cord injuries. Like in our experiments, their results showed that the variable-gain controller did not lead to improved retention when compared to the baseline. In this work, we extend this concept to also compare the baseline and variable-gain controller to two fixed-gain controllers.

Replicating our experiments on subjects with motor impairments is a topic of future work.

Another topic for future work is how retention changes over time: here we have focused on short-term retention, but it is also important to explore how control strategies affect retention over longer time durations.

Conclusion

In this paper, we presented an experimental comparison of robotic rehabilitation control strategies on able-bodied subjects, with a focus on short-term stabilization of the human's motor adaptation. We first extended a recent reinforcement approach to include robotic control, where the end-user trains while receiving kinesthetic feedback from the controller as well as auditory reward signals when they complete the task successfully. We then performed user studies to compare the effects of fixed and variable-gain control strategies: no controller, assist-as-needed, fixed-gain assistance, and fixed-gain resistance. Overall, we found that the fixed-gain control strategies led to better short-term retention than either the variable-gain controllers or the no controller baseline. These improvements were broken down by task difficulty: during the less difficult task fixed-gain assistance and resistance resulted in better short-term retention, while fixed-gain resistance had a lower learning rate and fixed-gain assistance decayed more rapidly across the board. When applied to robotic rehabilitation, our results suggest that designers should promote short-term retention by selecting fixed-gain robotic controllers that match the user's perceived difficulty. This work is only a first step towards developing effective control strategies for long-term retention following robotic therapy: future works should explore how our results transfer to subjects with motor impairments, as well as retention over longer time durations.

Declarations of Conflicting Interests

The author(s) declared the following potential conflicts of interest with respect to the research, authorship, and/or publication of this article: MKO is an employee of Rice University and is a consultant at TIRR-Memorial Hermann. MKO is a co-founder of Houston Medical Robotics. MKO has received grants from NSF, ONR, NASA, NIH, Facebook, and TIRR Foundation.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by NSF GRFP-1450681 and NSF CNS-1135916.

Guarantor

MKO.

Contributorship

DL and LB researched literature, conceived, and conducted the study. JC was involved in the data analysis. DL wrote the first draft of the manuscript. MKO supervised the researchers throughout the study. All authors reviewed and edited the manuscript and approved the final version of the manuscript.

ORCID iD

Dylan P Losey  <https://orcid.org/0000-0002-8787-5293>

References

- Lo AC, Guarino PD, Richards LG, et al. Robot-assisted therapy for long-term upper-limb impairment after stroke. *New Engl J Med* 2010; 362: 1772–1783.
- Marchal-Crespo L and Reinkensmeyer DJ. Review of control strategies for robotic movement training after neurologic injury. *J Neuroeng Rehabil* 2009; 6: 20.
- Blank AA, French JA, Pehlivan AU, et al. Current trends in robot-assisted upper-limb stroke rehabilitation: Promoting patient engagement in therapy. *Curr Phys Med Rehabil Rep* 2014; 2: 184–195.
- Sigrist R, Rauter G, Riener R, et al. Augmented visual, auditory, haptic, and multimodal feedback in motor learning: a review. *Psychonom Bull Rev* 2013; 20: 21–53.
- Losey DP, Blumenschein LH and O'Malley MK. Improving the retention of motor skills after reward-based reinforcement by incorporating haptic guidance and error augmentation. In: *IEEE RAS/EMBS international conference on biomedical robotics and biomechanics (BioRob)*, 2016. pp. 865–871. Piscataway, NJ: IEEE.
- Galea JM, Mallia E, Rothwell J, et al. The dissociable effects of punishment and reward on motor learning. *Nat Neurosci* 2015; 18: 597–602.
- Huang VS, Haith A, Mazzoni P, et al. Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron* 2011; 70: 787–801.
- Shmuelof L, Huang VS, Haith AM, et al. Overcoming motor “forgetting” through reinforcement of learned actions. *J Neurosci* 2012; 32: 14617–14621.
- Proietti T, Crocher V, Roby-Brami A, et al. Upper-limb robotic exoskeletons for neurorehabilitation: a review on control strategies. *IEEE Rev Biomed Eng* 2016; 9: 4–14.
- Krebs HI, Hogan N, Aisen ML, et al. Robot-aided neurorehabilitation. *IEEE Trans Rehabil Eng* 1998; 6: 75–87.
- Hogan N. Impedance control: an approach to manipulation. *J Dyn Syst Meas Control* 1985; 107: 1–24.
- Celik O, O'Malley MK, Boake C, et al. Normalized movement quality measures for therapeutic robots strongly correlate with clinical motor impairment measures. *IEEE Trans Neural Syst Rehabil Eng* 2010; 18: 433–444.
- Duschau-Wicke A, von Zitzewitz J, Caprez A, et al. Path control: A method for patient-cooperative robot-aided gait rehabilitation. *IEEE Trans Neural Syst Rehabil Eng* 2010; 18: 38–48.
- Mao Y and Agrawal SK. Design of a cable-driven arm exoskeleton (CAREX) for neural rehabilitation. *IEEE Trans Robot* 2012; 28: 922–931.
- Lee J and Choi S. Effects of haptic guidance and disturbance on motor learning: Potential advantage of haptic disturbance. In: *IEEE haptics symposium*, 2010, pp. 335–342. Piscataway, NJ: IEEE.
- Lum PS, Burgar CG, Shor PC, et al. Robot-assisted movement training compared with conventional therapy techniques for the rehabilitation of upper-limb motor function after stroke. *Arch Phys Med Rehabil* 2002; 83: 952–959.
- Patton JL, Stoykov ME, Kovic M, et al. Evaluation of robotic training forces that either enhance or reduce error in chronic hemiparetic stroke survivors. *Exp Brain Res* 2006; 168: 368–383.
- Guadagnoli MA and Lee TD. Challenge point: a framework for conceptualizing the effects of various practice conditions in motor learning. *J Motor Behav* 2004; 36: 212–224.
- Rauter G, Sigrist R, Marchal-Crespo L, et al. Assistance or challenge? Filling a gap in user-cooperative control. In: *IEEE/RSJ international conference on intelligent robots and systems (IROS)*, 2011, pp. 3068–3073. Piscataway, NJ: IEEE.
- Emken JL, Bobrow JE and Reinkensmeyer DJ. Robotic movement training as an optimization problem: Designing a controller that assists only as needed. In: *IEEE international conference on rehabilitation robotics*, 2005, pp. 307–312. Piscataway, NJ: IEEE.
- Frullo JM, Elinger J, Pehlivan AU, et al. Effects of assist-as-needed upper extremity robotic therapy after incomplete spinal cord injury: a parallel-group controlled trial. *Front Neurobot* 2017; 11: 26.
- Pehlivan AU, Losey DP and O'Malley MK. Minimal assist-as-needed (mAAN) controller for upper limb robotic rehabilitation. *IEEE Trans Robot* 2016; 32: 113–124.
- Wolbrecht ET, Chan V, Reinkensmeyer DJ, et al. Optimizing compliant, model-based robotic assistance to promote neurorehabilitation. *IEEE Trans Neural Syst Rehabil Eng* 2008; 16: 286–297.
- Shadmehr R, Smith MA and Krakauer JW. Error correction, sensory prediction, and adaptation in motor control. *Ann Rev Neurosci* 2010; 33: 89–108.
- Krakauer JW. The applicability of motor learning to neurorehabilitation. In: Dietz V and Ward N (eds) *Oxford textbook of neurorehabilitation*. Oxford, UK: Oxford University Press, 2015, pp.55–63.
- Reinkensmeyer DJ, Burdet E, Casadio M, et al. Computational neurorehabilitation: Modeling plasticity and learning to predict recovery. *J Neuroeng Rehabil* 2016; 13: 42.
- Hasson CJ, Manczurowsky J and Yen SC. A reinforcement learning approach to gait training improves retention. *Front Hum Neurosci* 2015; 9: 459.
- Milot MH, Marchal-Crespo L, Green CS, et al. Comparison of error-amplification and haptic-guidance training techniques for learning of a timing-based

- motor task by healthy individuals. *Exp Brain Res* 2010; 201: 119–131.
29. Heuer H and Luetzgen J. Robot assistance of motor learning: a neuro-cognitive perspective. *Neurosci Biobehav Rev* 2015; 56: 222–240.
 30. Spong MW, Hutchinson S and Vidyasagar M. Robot modeling and control (vol. 3). New York, NY: John Wiley & Sons, 2006.
 31. Taylor JA, Krakauer JW and Ivry RB. Explicit and implicit contributions to learning in a sensorimotor adaptation task. *J Neurosci* 2014; 34: 3023–3032.
 32. Orban de Xivry JJ and Lefèvre P. Formation of model-free motor memories during motor adaptation depends on perturbation schedule. *J Neurophysiol* 2015; 113: 2733–2741.
 33. Ronsse R, Puttemans V, Coxon JP, et al. Motor learning with augmented feedback: modality-dependent behavioral and neural consequences. *Cereb Cortex* 2010; 21: 1283–1294.
 34. Kluzik J, Diedrichsen J, Shadmehr R, et al. Reach adaptation: what determines whether we learn an internal model of the tool or adapt the model of our arm? *J Neurophysiol* 2008; 100: 1455–1464.
 35. Li Y, Patoglu V and O'Malley MK. Negative efficacy of fixed gain error reducing shared control for training in virtual environments. *ACM Trans Appl Percept* 2009; 6: 3.
 36. Powell D and O'Malley MK. The task-dependent efficacy of shared-control haptic guidance paradigms. *IEEE Trans Haptics* 2012; 5: 208–219.
 37. Huegel JC and O'Malley MK. Progressive haptic and visual guidance for training in a virtual dynamic task. In: *IEEE haptics symposium*, 2010, pp. 343–350. Piscataway, NJ: IEEE.