25

The Role of Haptics in Physical Rehabilitation

G. C. Burdea

While the majority of today's haptic interfaces and applications are targeted at the able-bodied user, a rapidly growing field of science studies the use of this technology in physical rehabilitation. There are many reasons the reader may wish to take a closer look at this application domain. One reason concerns societal impact, as there are about 70 million people with disabilities in the European Union [Bühler 97]. Such therapy is needed by various patient populations ranging from post-stroke survivors, to those with traumatic brain injury, cerebral palsy, spinal cord injuries, musculo-skeletal deficits, and others. The United States alone spends about $30 billion every year on physical rehabilitation [Patton et al. 06]. Of the above-mentioned costs, the majority represent labor costs (therapist time), and economic pressures tend to make rehabilitation interventions shorter than in prior years.

Rehabilitation science, in contrast to current rehabilitation practice, has recently shown that intense and longer physical therapy will benefit even chronic patients through the phenomenon of "brain plasticity." By repeating meaningful limb movements, similar to those done in *activities of daily living* (ADL), dormant neurons are recruited into new neural paths, and patients regain some of their lost function. Here robots are ideal, since they can train patients for the required long duration without tiring (unlike human therapists), and may eventually lead to a reduction in labor costs.

Robotic systems coupled with virtual reality simulations bring additional improvements to today's conventional physical therapy methods, since they introduce objective measures of performance. Data on total exercise time, speed and smoothness of movement, peak and average velocities, mechanical work, and endurance are among the variables that can be stored transparently and used to objectively gauge a patient's progress. This is a clear departure from the subjective therapist's evaluation of a patient, which is prevalent today.

When robotics is coupled with virtual reality, the resultant rehabilitation becomes fun, since patients can practice in the form of a video game play. They can also be challenged according to their specific abilities and can be given auditory or graphics rewards for their performance. The flexibility of virtual reality also means that a number of different simulations and haptic effects can be produced by the same hardware, thus creating variety and progression of therapeutic games difficulty to challenge each patient. It is intuitive that any therapy that motivates the patient will produce better outcomes, compared to approaches where the patient is disinterested, bored, and otherwise mentally detached from the task she/he is asked to perform.

A more subtle reason to look at haptic applications in physical therapy is the dual use of the same technology for able-bodied individuals. Such users will benefit from techniques presented in this chapter by augmenting their capabilities and thus improving their task performance in virtual reality or telerobotics applications. After all, disability is a question of degree, and we are all disabled to some extent.

This chapter starts with a review of robotic systems used in physical rehabilitation (Section 25.1), followed by a discussion of the specifics of haptics targeted at the disabled (Section 25.2). Safety issues are clearly important in systems, such as those described in this chapter, where users are in close proximity to the haptic interface or robot. Safety issues for the disabled, which are reviewed in Section 25.3, are even more important, since patients often have degraded hand-eye coordination or cognitive or reflex capabilities, and thus are at higher risk compared to able-bodied users. A look at the future use of haptics in physical rehabilitation concludes this chapter (Section 25.4).

25.1 Robotic Systems for Physical Rehabilitation

The terms *upper extremity* and *lower extremity* are commonly used by physical therapists to refer to either the upper or the lower limbs. Thus, upper extremity rehabilitation aims at improving the patient's shoulder, elbow, wrist, and fingers (and the patient's ADLs). Lower extremity training refers to exercising the patient's knee, ankle, foot, or the whole leg in walking. Robots have been used in physical rehabilitation for more than a decade, and they target all of the above areas of therapy.

25.1.1 Robots for Upper Extremity Physical Rehabilitation

One of the earliest applications of haptics in rehabilitation is the MIT MANUS system shown in Figure 25.1(a) [Krebs et al. 04]. It consists of

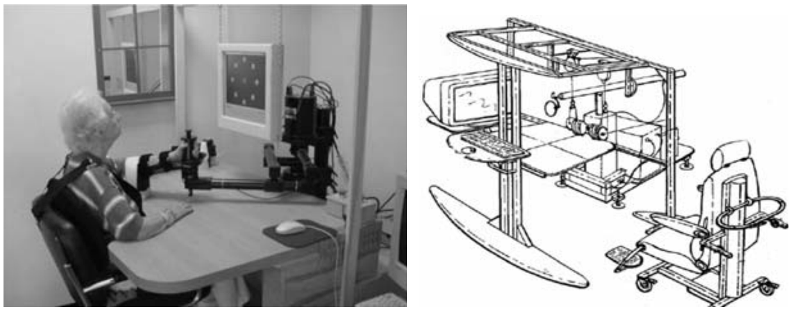

(a) (b)

Figure 25.1. Haptic systems for shoulder rehabilitation: (a) Commercial version of the MIT MANUS [Krebs et al. 04] (Open source material). (b) The Haptic Master [Loureiro et al. 04]. Reprinted by permission.

a direct-drive SCARA two-degree-of-freedom robot that trains the patient arm in a plane while monitoring forces at the end effector. The patient rests the forearm on a special support with safety coupling that detaches in case of excessive forces. The patient is strapped in a chair in order to prevent compensatory torso leaning and faces a monoscopic display controlled by a PC. The robot has its own controller, which implements a back-drivable impedance control aimed at increasing the patient's safety. More recent versions of the MIT MANUS allow the integration of modules for additional degrees of freedom.

Figure 25.1(b) [Loureiro et al. 04] illustrates the adaptation of the Haptic Master, a general-purpose haptic interface, for use in physical rehabilitation. The robot differs from the MIT MANUS, as it has three degrees of freedom and a cylindrical work envelope. Its control is also different, since the Haptic Master uses an *admittance controller* which moves the robot in response to forces applied by the patient on its end effector. Similar to the MIT MANUS setting, the patient is strapped in a chair and faces a monoscopic display showing graphics generated by a PC. These scenes are updated based on the data received by the PC from the Haptic Master. Since the work envelope and output forces of this robot are larger than those of the MIT MANUS, a much more complex apparatus is used to offload gravity-induced forces from the patient's extended arm.

Neither of the above robots is able to train the patient's fingers, which are essential in ADLs. The only commercially available haptic glove is the CyberGrasp (shown in Figure 25.2(a)) [McLaughlin et al. 05]. It consists of an exoskeleton worn on the back of the hand, and five actuators, which apply one degree of force feedback for each finger through a combination of cables and pulleys. Finger sensing is done by the CyberGlove on which

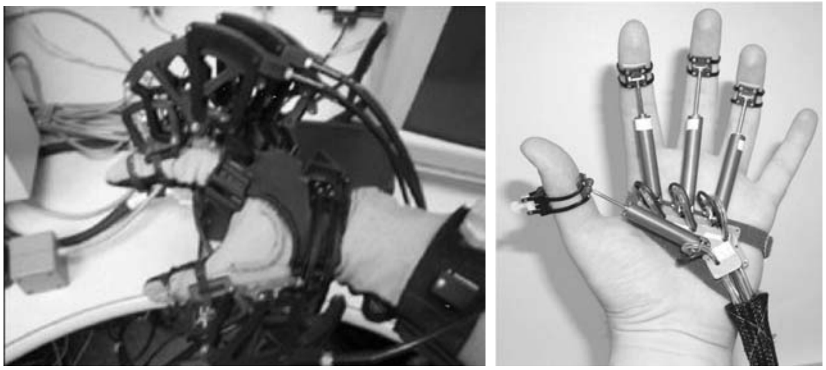

(a) (b)

Figure 25.2. Robots for finger rehabilitation: (a) The CyberGrasp [McLaughlin et al. 05]. Reprinted by permission; (b) the Rutgers Master II [Bouzit et al. 02](©Rutgers University). Reprinted by permission.

the CyberGrasp exoskeleton is retrofitted, and adjustments need to be made for various hand sizes, using mechanical stops on the exoskeleton cable guides. When applied in a physical rehabilitation setting, the weight of the CyberGrasp (about 400 grams) becomes a problem, since patients who need rehabilitation have a diminished arm weight-bearing capability. Furthermore, this weight is placed (by necessity) away from the body, which creates a mechanical amplifier effect.

The requirement for reduced weight is addressed in the prototype Rutgers Master II glove shown in Figure 25.2(b) [Bouzit et al. 02], which weighs about 100 grams. Similar to the CyberGrasp, the Rutgers Master II has an exoskeleton that provides one degree of force feedback per finger (less the pinkie). However, it does not require a separate sensing glove, as its exoskeleton incorporates non-contact position sensors. The glove uses a direct-drive configuration and compressed air, such that each fingertip is resisted in flexion with up to 16 N force. The lack of a separate glove makes its donning faster and easier than the CyberGrasp.

25.1.2 Robots for Lower Extremity Physical Rehabilitation

While robots for upper extremity rehabilitation have existed for over a decade, those used to train the patient's walking and ankle control are more recent. Among them, the best known (and commercially available) is the Lokomat [Frey et al. 06, Riener et al. 06] shown in Figure 25.3(a), used for gait training. Patients with spinal cord injury or post-stroke patients have diminished weight-bearing capacity, which hampers walking. There-

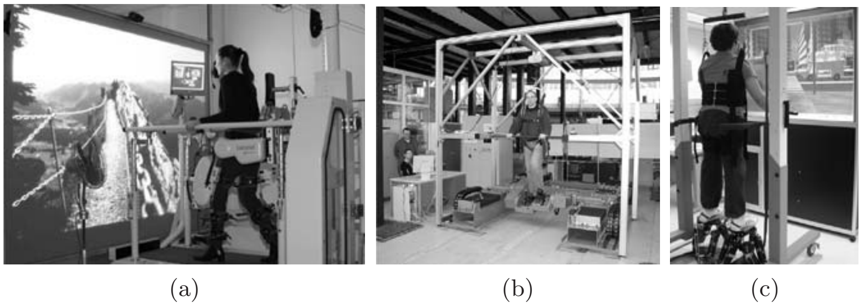

(a) (b) (c)

Figure 25.3. Robotic systems for walking rehabilitation: (a) the Lokomat [Riener et al. 06] (© IEEE). Reprinted by permission. (b) The HapticWalker [Schmidt et al. 05] (© ACM). Reprinted by permission. (c) The Mobility Simulator [Boian 05] (© Rutgers University). Reprinted by permission.

fore, therapists use treadmills and passive *body weight supports* (BWSs) in the form of a harness and elastic element to reduce the weight the patient's legs have to support by 60 to 80%. The Lokomat uses the same treadmill + BWS approach, but adds two important elements. The first is a pair of leg exoskeleton robots, which assist the gait cycle with speeds up to about 3 km/h. The robots greatly reduce the therapist's physical effort and thus allow longer therapy than otherwise possible. The second improvement over non-robotic approaches to gait training is the addition of an active (actuator) based BWS in addition to the passive one. The combination of passive + active BWS results in much more uniform weight unloading during walking, and optimal gait training. Recently, the Lokomat has added advanced biofeedback, which immerses the patient in a virtual environment. The patient views the scene of a hiking trail and obstacles that need to be negotiated. If the foot is not lifted high enough, haptic and sound feedback of the collision with the obstacle are produced. A fan provides tactile feedback (in the form of wind) proportional with the patient's walking speed. Thus the patient trains in a meaningful environment, which is adjustable to his/her performance and helps highlight proper walking patterns.

Treadmill training cannot realistically reproduce walking on uneven terrain, such as up and down the stairs. A system that addresses this limitation is the HapticWalker seen in Figure 25.3(b) [Schmidt et al. 05]. Similar to the Lokomat, the HapticWalker consists of two exoskeleton robots that move the patient's legs, coupled with a BWS. The robots incorporate direct-drive electric motors capable of assisting walking up to a speed of 5 km/h. The HapticWalker design uses hybrid serial (large workspace) and parallel (large payload) kinematics. Two actuators connected in parallel move the foot either up/down or front/back. A third actuator is used to tilt the foot.

Even more degrees of freedom may be needed for realistic haptics and purposeful training. For example, quick horizontal translations overimposed to gait are needed to simulate walking on ice. A robot that can reproduce such haptic effects is the *mobility simulator* prototype seen in Figure 25.3(c) [Boian 05]. Similar to the Lokomat and the HapticWalker, this robot incorporates a BWS system. However, each foot sits on top of a Rutgers Mega Ankle Stewart Platform with direct-drive pneumatic actuators [Boian et al. 04]. Thus, each foot is moved in six degrees of freedom, which allows training for walking on even or uneven terrain (mud, gravel, ice). To provide gait training associated with ADLs, the patient faces a large (monoscopic) display showing a street crossing. The patient has to cross at the pedestrian stop light under various surface, time to cross, and visibility conditions. Distractions, in the form of street noises (honking) or impatient drivers pushing onto the street, are provided for additional training difficulty. Due to the fact that the bases of the two Rutgers Mega Ankle platforms are fixed and their dimensions are more compact than those of the HapticWalker, the step length is smaller that normal values, which is a drawback of the current design.

25.2 Specifics of Haptic Feedback for the Disabled

Haptic feedback used in physical therapy is different from that provided to able-bodied users due to the force and motor coordination deficits of the disabled. In domains not related to rehabilitation, haptic feedback is usually in the form of resistive forces which complement graphics and other simulation modalities. Such resistive forces are required to more realistically simulate object compliance, weight, inertia, and surface properties (roughness, stickiness, and friction).

Haptic feedback in physical therapy is more demanding, since it needs to adapt to each patient's functioning level and each therapy session. Furthermore, certain types of haptic feedback (such as vibrations) that adversely affect normal training can prove beneficial in physical therapy. The discussion here is focused on two aspects that play a central role in haptic feedback for physical therapy, namely *assistive haptics* and disturbances.

25.2.1 Assistive Haptics

Dues to the weakened upper or lower extremities of various patient populations, such as those with neurological disorders (stroke, spinal cord injury,

(a) (b)

Figure 25.4. Assistive haptics used to train ankle strength in children with cerebral palsy: (a) System view showing the Rutgers Ankle robot. (b) Screen image highlighting the ideal trajectory the robot is using to pilot the plane while patient is passive (© Rutgers University). Reprinted by permission.

cerebral palsy), the haptic interface needs to assist the patient in performing the simulated task. An example is the use of the Rutgers Ankle robot [Girone et al. 01] in the training of patients with cerebral palsy. Patients sit facing a PC display while their foot is strapped on the mobile platform of the Rutgers Ankle Stewart Platform-like robot. The simulation depicts an airplane that has to fly through a series of hoops while under patient control. In prior studies done with stroke patients, the robot provided purely resistive spring-like forces [Mirelman et al. 06]. This is not possible with children with CP, since at the start of each rehabilitation session their ankle needs to be stretched and moved over its range of motion, with the patient being passive. While in conventional therapy, this is done manually by the physical therapist: here the robot pilots the airplane over an ideal sinusoidal path (see Figure 25.4(a)–(b). During this time, the patient is completely passive. Subsequently, the patient is asked to progressively exert more torques to tilt the foot up/down while the robot creates a “haptic tunnel.” Small corrective forces are applied to keep the airplane within an acceptable (threshold-determined) neighborhood of the ideal path. In subsequent rehabilitation sessions, while the patient’s ankle exertion capability increases, the robot will switch off assistance and eventually apply resistive forces, which will challenge the patient more.

Another example of graded assistance by a robot is the upper extremity training provided by the MIT-MANUS system. As seen in Figure 25.1(a), the patient is asked to move the robot handle in a plane, such that a corresponding cursor on an associated display moves to a highlighted dot out of eight possible targets [Hogan and Krebs 04]. The robot implements an impedance control, which calculates a point that moves on an ideal path to the target while monitoring the position of the end effector. A spring-like force attempts to minimize the distance between the handle position and the moving ideal location on the ideal path. Tests showed this therapy to be useful; however, it did not adapt sufficiently to each patient's condition. This lack of adaptation was due to the fact that the speed of the ideal point on the nominal path was kept constant. A subsequent improvement in the haptic feedback provided by the MIT-MANUS was an adaptive impedance controller which implements a "virtual slot" running between the ideal position and the target position. The walls of the virtual slot are "springy" to provide assistance in case of inappropriate movements away from the ideal path. Furthermore, the back wall of the virtual slot moves to the target with a velocity that assures a fixed duration for a minimum-jerk trajectory. This back wall assists the patient if he or she lags behind the ideal position on the path. However, if the patient can move faster than the virtual slot back wall, he or she is free to do so (while getting no assistance from the robot). The duration of the ideal movement is set automatically based on the patient's past performance. If the patient was able to consistently move faster than the back wall of the virtual slot, then the simulation is made faster, requiring faster arm movements to stay ahead of the robot. Tests showed that this improved therapeutic haptic feedback which was between four to ten times more efficacious than the fixed impedance controller initially used.

25.2.2 Haptic Disturbances to Help Motor Control and Recovery

Haptic disturbances are effects overlaid in the simulation in order to increase therapy difficulty or induce desired after effects. Air turbulence was simulated when piloting the airplane during a storm by oscillating the Rutgers Ankle in the horizontal plane [Boian et al. 03]. Progressively more turbulence determined gradually faster swaying of the robot, while the amplitude of the vibrations was kept fixed. Tests showed that patients gradually learned to cope with these haptic disturbances, eventually being able to clear 100% of the target hoops. This is indicative of improved ankle control, which results in diminished reinjury due to accidents or falls.

Another type of haptic disturbance is illustrated by the graphs in Figure 25.5 [Patton et al. 04]. The curves represent planar arm-reaching movements towards one of six targets while holding a robot arm. Initial undis-

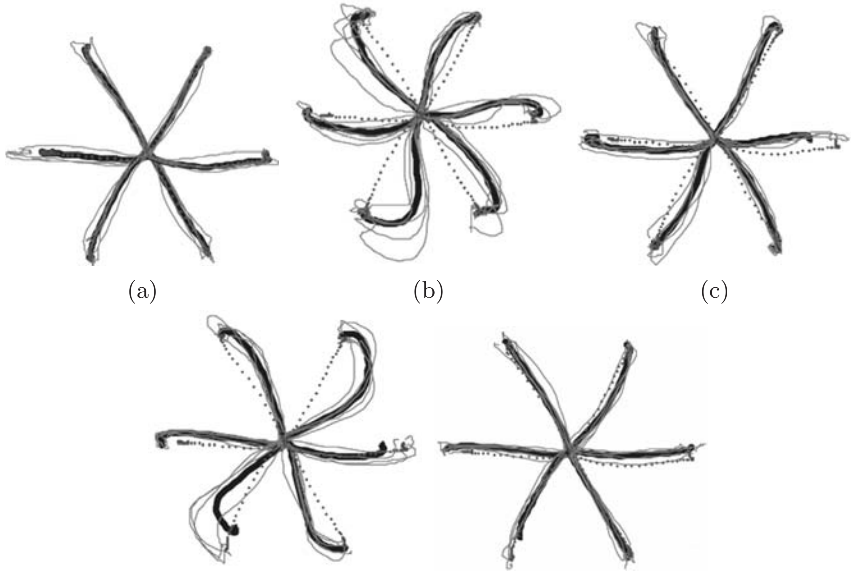

(a) (b) (c)

(d) (e)

Figure 25.5. Hand trajectories in horizontal plane illustrating aftereffects of systematic haptic disturbances: (a) unperturbed baseline; (b) early training with disturbance; (c) final training; (d) aftereffects when disturbance was removed; (e) final washout. Dotted lines are the initial baseline; bold lines represent average movements [Patton et al. 04] (© IEEE). Reprinted by permission.

turbed "baseline" reach movements for a healthy user are plotted in Figure 25.5(a), followed by subject's movements when first confronted with a steady lateral force. Gradually the subject learns to cope with these forces, such that by the end of training (Figure 25.5(c)), the arm moves in straight lines again despite the presence of disturbances. Figure 25.5(d) illustrates the aftereffects of haptic disturbances, as soon as the lateral forces are removed. It can be seen that the arm moves over trajectories, which curve in the opposite direction to the previously applied lateral forces. With continuing repetitions, the trajectory straightens out again, such that aftereffects disappear (or "wash out"). While washing out of learned movements is common with able-bodied users, this is not the case for the disabled [Matsuoka et al. 04]. For the disabled, the effects induced by haptic disturbances do not wash out, because the training leads the patient to activate different sets of muscles. Once the distorting haptic effects disappear at the end of training, the disabled continue to use the new coordinated movements that they learned, using the muscles that had previously been unused.

25.3 Safety Issues in Haptics for Rehabilitation

While the haptic interface mediates interactions with virtual environments, the forces applied on the user are real. Robots designed for industrial applications, capable of high output forces and large accelerations, pose a real risk when used as haptic interfaces. Even robots designed from the start for physical rehabilitation applications may be dangerous to the patient, since they need to apply large enough forces and torques to make therapy meaningful.

The start of this chapter pointed out that the user's safety is even more important for the disabled. Their slower defensive reflexes, diminished awareness of surroundings, diminished sensory capability (blurred vision, degraded proprioception), and diminished cognitive capacity put the disabled at increased risk when involved in haptics-assisted rehabilitation. It is thus important to look at ways to design computerized physical rehabilitation systems that address the patient's safety concerns mentioned here.

The first line of defense, commonly used in industrial applications, is the provision of safety switches that disable the robot in case of danger. In rehabilitation settings, there should be several such manual switches, one for the patient and one for the attending therapist, who can stop the simulation in case of danger.

Manual switches, however, are not sufficient in a rehabilitation application, due to the slow human response. Additional measures are the integration of sensors and limit switches in the haptic interface itself. This is the approach taken in the design of the HapticWalker patient's foot attachment, as seen in Figure 25.6(a) [Schmidt et al. 04]. The patient wears a shank strap connected to an ankle goniometer through a lever. If the ankle dorsiflexion angle exceeds a prescribed limit, the controller monitoring the goniometer executes an emergency shutdown. Additional safety measures are the thrust pieces that snap in holes that incorporate emergency stop switches. These are built in the supporting plate under the foot, both front and back, and excessive forces detach the thrust pieces and thus trigger a shutdown of the robot.

The above example illustrates the redundancy principle used in good safety design. Several layers of safety measures are necessary in case one layer fails, and designers have to foresee such sensor failures. [Roderick and Carignan 05] describe how they improved the exoskeletons designed for shoulder therapy in order to incorporate redundant layers of safety. Their preliminary analysis identified hazards related to the movement of the patient's arms outside safe position ranges with excessive velocity, or hazards due to excessive torques applied to the patient. Their initial hardware design used an incremental encoder to measure joint values and provide feedback to the servo controller for that joint haptic feedback motor. This

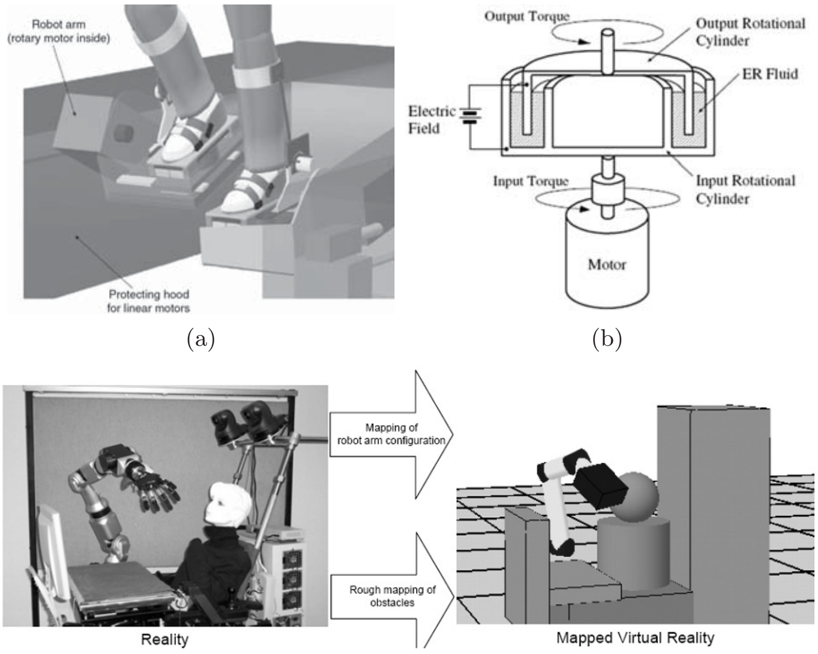

(a) (b)

(c)

Figure 25.6. Safety methods used when applying haptics in physical rehabilitation: (a) Sensors and mechanical limit switches incorporated in the foot support of the HapticWalker [Schmidt et al. 04] (© IEEE 2004). Reprinted by permission. (b) Electro-rheologic actuator couplings incorporated in a haptic interface for arm rehabilitation [Furusho et al. 05] (© IEEE 2005). Reprinted by permission. (c) Predictive real-time modeling used to prevent patient-robot collisions [Feuser et al. 05] (© IEEE 2005). Reprinted by permission.

design would not prevent motion outside safe ranges if the encoder failed. Thus, the improved design added a second position sensor (an absolute encoder) at each joint. The divergence between the values reported by the two position sensors is monitored to detect failure. The same hardware is used in joint velocity monitoring; thus redundancy is assured in order to prevent excessive joint velocities. In order to build redundancy in force control, the design adds a power amplifier thus senses the power draw of the feedback actuator motor. A motor power divergence check is done in software to detect when the requested output set by the servo controller does not correspond to the motor actual current draw.

Figure 25.6(b) illustrates another approach to increase the safety of a robot used in arm rehabilitation [Furusho et al. 05]. Instead of connecting the actuator directly to the robot joint, the designers use an *electrorheologic* (ER) coupling. The ER fluid changes its viscosity in proportion to the electrical field applied, which in turn is controlled by the robot controller. Hence it is possible to modulate slippage, thus limiting the potentially dangerous output torques. In case of power loss, the link is decoupled and the robot arm becomes completely back-drivable. In order to further improve safety, haptic interface arm inertia (which does not disappear even when power is lost) is minimized by placing the actuators at the base of the robot and passively counterbalancing the robot arm with weights.

A departure from the previous designs, which relied on robot actuators and internal sensors to improve the patient's safety, is the system illustrated in Figure 25.6(c) [Feuser et al. 05]. It uses a pair of cameras to create a simplified model of the environment consisting of 3D primitives (sphere, cylinder, prism). The robot is modeled as a series of linked 3D objects, and obstacles (including the patient) are also modeled with primitives. Such a simplified model facilitates real-time updates that are performed any time a new object is added or the patient moves. The robot control software performs collision detection using vertex-to-vertex distance calculation (it is thus necessary to convert the primitives to a sparse vertex lattice) [Gilbert et al. 88]. Once the real-time collision detection determines that distances in the updated virtual model fall below a threshold, the real robot is stopped before colliding with the patient.

25.4 Looking at the Future

It is expected that haptics will play an increasing role in physical rehabilitation in the years to come. Based on initial study data, it is expected that the technology will prove efficacious, especially when robotics is coupled with game-like virtual reality training. The penetration of the technology into widespread clinical use will benefit from lower cost hardware, such as game consoles and cheaper haptic interfaces.

Another direction of future growth is the nascent area of telerehabilitation, where therapy is provided at a distance (eventually in the patient's home). It is common in today's rehabilitation practice for the physical therapist to manually manipulate (move, stretch, warm up) the patient's affected limbs. Doing so at a distance will make at-home exercises more meaningful for the patient, without requiring the physical therapist to be co-located. Innovative approaches are clearly required to overcome the problems due to current network limited quality of service (jitter, time delays) in order to implement remote touch.

Acknowledgments

The author's research reported here was supported by grants from the National Science Foundation (BES-9708020 and BES-0201687), from the National Institutes of Health (5R21EB00653302), and the New Jersey Commission on Science and Technology (R&D Excellence Grant).