Figure 1: Examples of "light touch" patient guidance

Safe Human-Robot Physical Interaction To Assist Balance

PI: Christopher G. Atkeson, Lead: Carnegie Mellon University

Our goal is to learn from current clinical practice how assistive robots (physical co-robots) **should** physically interact with patients with balance disorders, and develop robots with human level competence in assisting and providing therapy to patients. Caregivers provide physical support using large forces, provide guidance using light touch, and provide additional sensory feedback and references (Figure 1). These aids are sometimes provided intermittently on a schedule or based on patient progress to optimize patient recovery and learning (assist as needed (AAN) control). Patients may be standing, standing up from a chair, sitting down, walking without aids, or using crutches, a cane, a walker, an IV stand, or grasping hand rails or grab bars. Contact strategies include holding hands, hand to arm or body, full arm contact, and no contact. Grasping (push and pull) or unilateral contacts (pushing only) can be used.

Current robot assistance is typically provided by the robot doing all the work, and providing a rigid physical support for patients. This is quite different from what human caregivers provide. We will explore using safe soft robots to provide compliant guidance using light touch, and sensory feedback and references. Figure 2 shows some possible applications of our work. Figure 3 shows some of our work on safe soft robot arms. Both to instrument patient-caregiver interations and to provide sensors for our safe soft robots, we will utilize soft sensors we have developed (Figure 4).

We will model caregivers as providing multiple contact points, each providing a variable three dimensional force. The forces will be generated using time varying desired position, stiffness, and damping. We will learn appropriate parameters for these variables by instrumenting patient-caregiver interactions, learning appropriate strategies for support, guidance, and and other situations. We will implement, optimize, and evaluate these strategies using both external robot and wearable device experimental testbeds. These testbeds will be based both on our existing work on soft safe robots and wearable devices, as well as new versions of these devices customized for individual patients. Our basic research will lay the foundation for several demonstration systems built on open platforms for transfer to industry and for use in clinical studies.

PITT: *Paragraph on why focus on balance disorders needed here.*

Intellectual Merit: 1) We will collect data in order to characterize how human caregivers interact with patients. We will organize, analyze, and publish this data. 2) We expect modeling active humans will be a major challenge in this work. We will develop *new modeling techniques* to quickly tailor kinematic and dynamic models to individuals and their capabilities. We will also model tasks, including how patients and caregivers might react to or actively participate in a particular joint behavior, to choose effective individual behaviors and to plan for errors and possible failure. 3) We will improve our *inherently safe robot designs* using soft (non-rigid) struc-

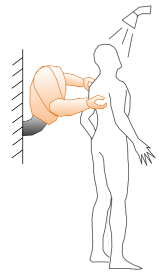

Figure 2: Our system aiding balance in a shower or bath. MORE EXAMPLE APPLICATIONS GO HERE

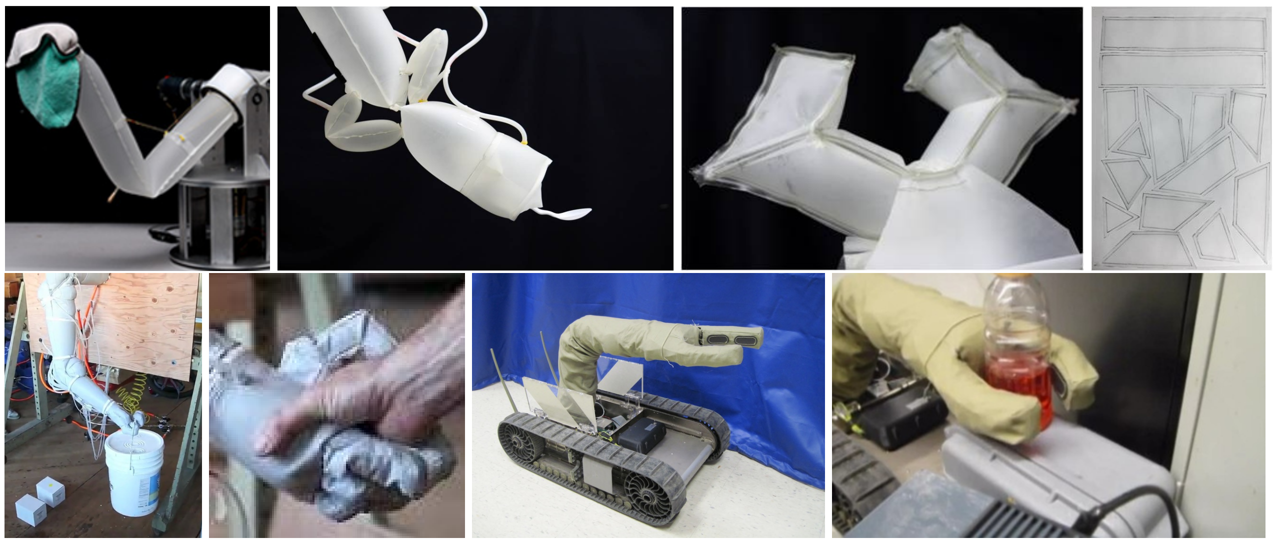

Figure 3: Previous work on soft safe inflatable arms [73]: **Top:** CMU Arm version 1, CMU wrist from version2, CMU gripper, pattern for making CMU gripper parts. **Bottom:** OtherLab Arm (2 views), IRobot Arm (2 views).

tures and mechanisms with smooth, pliable and reactive surfaces. 4) We will develop *new control techniques* that refine existing human strategies and invent new ones, based on cognitive optimization: general-purpose systems that learn to maximize reward or utility, using new, more powerful model-based forms of adaptive/approximate dynamic programming (ADP) and reinforcement learning. We will develop efficient multiple model methods for designing robust nonlinear and time varying task controllers for physical human-robot interaction where the robot, the patient, and possibly other caregivers all participate in the interaction. 5) We will develop *new machine perception techniques* using computer vision, ultrasound, inertial measurements (acceleration and angular velocity), and muscle activity measurements (electromyographic activity: EMG). Our computer vision techniques will allow accurate tracking of patients and caregivers and their body parts at a distance using real time surface tracking based on markers, structured light, and time of flight techniques. Our ultrasound-based sensing will allow tracking of a patient's tissue when the robot contacts the patient and provide measures of tissue movement as the interaction progresses. EMG measurements and accelerometers and gyros will be used in studying caregiver/patient interactions. During actual use these measures will be part of an early warning system for patient falls and other failures, and to allow the robot to more effectively cooperate with instrumented caregivers, by reacting more quickly to anticipated human movement.

Broader impacts I:

Physically interacting with patients is an important area for health care, but leads to injuries and occupational disability, and caregiving is statistically one of the most harmful of human occupations [**?**, **?**, **?**, **?**, **?**]. Many care facilities are instituting "no lift" policies, in which caregivers are no longer allowed to do large force physical manipulation of patients, but are required to use mechanical aids. There is currently a shortage of nursing staff for hospitals and nursing homes, compounded by injury and retirements due to disability. We also want to help older couples stay in their own homes longer, when one member has limited movement and the spouse is typically not very strong but would be cognitively able to operate a transfer aid. One goal is to reduce injuries

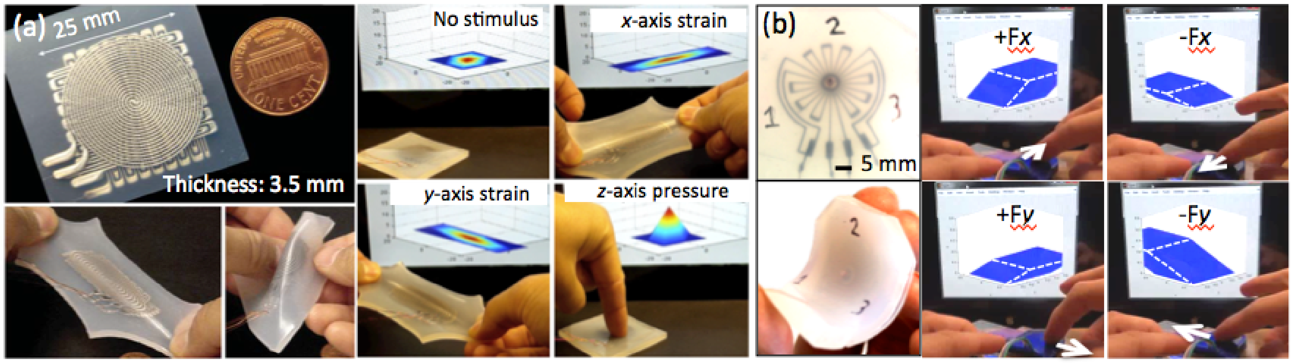

Figure 4: Soft artificial skin sensors. (a) Multi-axis strain and pressure sensing skin. (b) Soft multi-axis force sensor.

to caregivers. A second goal is to help patients recover more fully and recover faster. A third goal is to support the patient's ability to exercise and learn on their own both in the clinic and at home. A fourth goal is to provide better "customer service" and autonomy to patients, particularly in rehabilitation settings. Patients are often not allowed to go to the bathroom on their own for fear they might fall. Instead they must press a call button and sometimes wait a long time for a human aide.

Unfortunately, so far robotics has largely addressed this problem in a narrow and probably unacceptable way, using expensive stiff cumbersome approaches. We believe an approach that is more synergistic with current practice will be more useful, allow the use of weaker, cheaper, and safer robots, and will be more acceptable to patients and caregivers.

Our goal is a set of open demonstration systems that can lead to commercialization of this technology. We have already participated through formal collaboration and student internships with technology transfer of inflatable system with iRobot and OtherLab [32, 58] (Figure 3). Although we are proposing to do work on a narrow range of applications, the technology we are developing will have broad applicability, including prosthetics, orthotics, injury prevention, physical therapy and rehabilitation, sports, recreation, and entertainment, for example. Furthermore, future concepts and technologies on these lines will likely advance other systems that intimately interact with a human body like smart chairs and beds that may be used for therapeutic massage or to treat decubitus ulcers or bedsores (when bony prominences first penetrate onternal tissue and then skin; a very common condition for those suffering from limited mobility and affecting more than 2.5 million people in the United States each year).

The description of Broader Impacts is continued in Section 9.

1 Related Work In Robotic Balance Assistance

Most work in robotic balance assistance has focused on balance training devices or devices to provide balance, and include platforms a patient stands in, exoskeletons a patient wears, or adding sensing and actuation to walkers, canes, or other balance aids. These devices would complement the type of assistance we are focusing on.

PITT: *Need to add a nice review reference to the following paragraph:*

Platforms: Measurement and perturbation platforms are widely used in clinical balance assessment and therapy. Force platforms which measure foot forces and torques are common. Perturbations include the support platform translation and rotation, as well as push or pull devices located at various points, such as pushing or pulling at the hips, shoulders, or hands. Sensory perturbations are provided by virtual reality displays. Perhaps the most advanced platform-based balance trainer is Toyota's Balance Training Assist. The patient stands on a Segway-like wheeled robotic platform while playing a videogame. Action in the game affects the behavior of the platform in order to train the user to balance more effectively.

Gait machines: Large devices such as the Lokomat embed an exoskeleton into a larger support structure. Early designs were stiff and strong, and simply moved the patient's limbs along predetermined patterns. More recent designs are compliant and try to elicit more user effort.

Portable exoskeletons: Work on portable exoskeletons for both therapy and assistance are reviewed in [**?**, **?**]. A dominant design that has emerged from the laboratory is the exoskeleton with rigid frame structures [**?**]. Exoskeletons are often combined with high-torque electromechanical actuators for military applications [**?**, **?**]. They are also considered as rehabilitation and/or assistive devices by employing lightweight, compliant (or elastic) actuators [**?**, **?**, **?**].

While there are certain advantages that are derived from traditional exoskeleton design, such as relatively transparent force transmission and rigid mechanical body supports, they also carry several practical limitations, including bulkiness, mechanical constraints to host bodies, and safety issues when the wearers physically interact with other people. A new design, soft exoskeletons, overcome such limitations by removing rigid frame structures and mechanical joints. Such systems include active soft orthotic devices for assistance of ankle [60, 98] and knee [**?**, **?**] motions, a compliant exo-suit for improved metabolic energy efficiency through active gait assistance [**?**], and upper extremity rehabilitation with cable-driven [**?**] and pneumatic muscle [**?**] actuations. The design of soft exoskeletons is guided by the biomechanical properties of the human body so that synthetic components harness the elastic properties of soft tissue and the mechanical advantage of the skeletal levers. However, currently available actuation mechanisms limit further miniaturization and weight reduction of the devices, resulting in reduced wearability.

We have developed a programmable orthotic device that is intrinsically soft, wearable, pneumatically powered, and modular (Figure 6). When worn around a joint such as a wrist or knee, the orthotic will electronically monitor body motion and assist with motor function. In contrast to existing active orthotics [**?**, **?**, **?**, **?**] and exoskeletons [**?**, **?**], this device is soft, modular, multifunctional, and composed almost entirely of elastomer and embedded micro-fluidic channels. This builds on work in soft robotics [**?**, **?**, **?**, **?**, **?**] and embedded sensing [**?**, **?**, **?**, **?**]. These unique properties allow the soft orthotic to support a broad range of motor tasks and function reliably under extreme loads and body movements.

Instrumented and actuated walkers and canes

I need something here.

2 Proposed Work: What do caregivers actually do?

PITT: *This section (2) needs to come from Pitt folks.*

2.1 Summary of current practice

What are the balance disorders we will deal with? How are they treated? Support, Light touch, Sensory aids. Aid vs. teach?

2.2 How will we instrument interactions?

Motion capture, force plates, EMG, force gloves (Figure 5), ultrasound?

2.3 What experiments will we do?

Standing, sit to stand, sitting, down, and then walking on treadmill. Perturbations? Sensory interactions? Capture actual behavior in clinics?

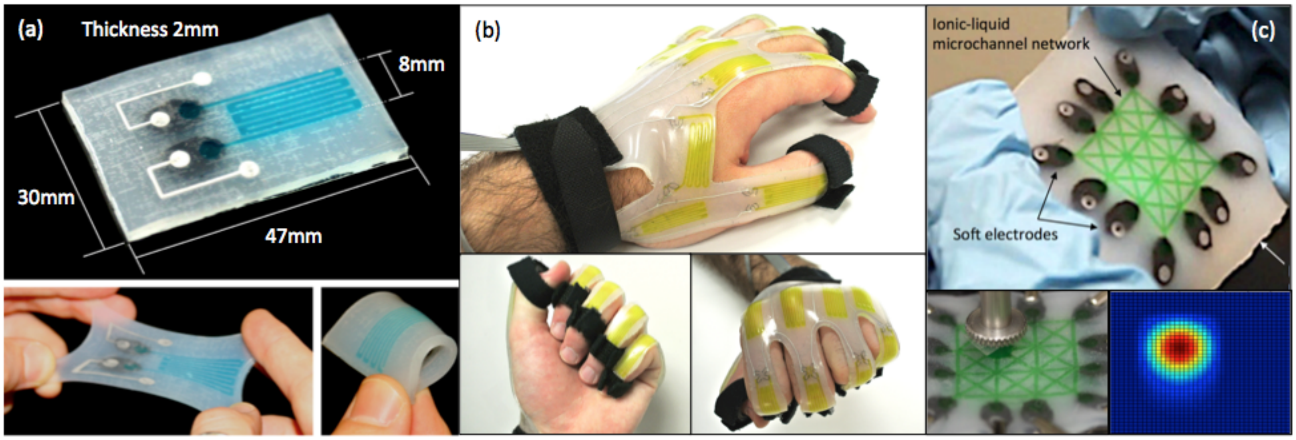

Figure 5: *This figure will illustrate how we will use force sensing "skin" for both patients and caregivers to instrument possible contacts.*

3 Proposed Work: Modeling and Controlling the Interaction

The goal our modeling work is to develop models that predict **when, where,** and **what** contact is used during balance assistance. Due to space limitations, this section will only be able to describe our modeling approach for assisting free standing balance with the patient in an "arms down" position. We will initially search for appropriate state variables to describe the interaction. We will start with the location and velocity of the estimated center of mass of the patient and the same variables for a single caregiver, as well as information describing current contacts. Data from human caregiver/patient interactions will be indexed by the state variables in a database.

To determine when to initiate an interaction, where to make contacts, and what type of contacts to make, we will continuously use the current state as an index to look up in the database what the human caregiver did in this situation. If there is not a close match, we will interpolate among a set of nearby matches. In machine learning, this approach is known as memory-based learning. Key technical issues are choosing a state, choosing a similarity measure, actually finding relevant experience using nearest neighbor lookup, and interpolating among sets of experiences.

We will model caregivers as providing multiple contact points, each providing a variable three dimensional force $\mathbf{f} = (f_x, f_y, f_z)$. The forces will be generated using time varying desired position $(\mathbf{x}_0)$, stiffnesses $\mathbf{k} = (k_x, k_y, k_z)$, and damping $\mathbf{b} = (b_x, b_y, b_z)$. For example:

$$f_x = k_x(t)(x - x_0(t)) + b_x(t)\dot{x} + f_{0_x}(t) \tag{1}$$

We will initially manually estimate $\mathbf{x}_0(t), \mathbf{k}(t), \mathbf{b}(t)$, and $\mathbf{f}_0(t)$. Later in the study we will develop methods to automatically estimate $\mathbf{x}_0(t)$. We will initially use EMG to guide manual selection of $\mathbf{f}_0(t)$, $\mathbf{k}(t)$ and $\mathbf{b}(t)$, and then develop more automatic methods to estimate these variables.

We will initially use straightforward control approaches, involving having the robot play back the desired positions $\mathbf{x}_0(t)$, stifnesses $\mathbf{k}(t)$, damping $\mathbf{b}(t)$, and nominal force $\mathbf{f}_0(t)$.

Refinements to the control approach in later years will involve manually identifying "behavioral primitives" or "strategies" used by the caregivers for support, guidance, and and other situations, such as "move until contact" or "maintain roughly constant force", and perhaps doing all control in terms of these higher level behaviors. We would initally identify behavioral parameters manually, and then, explore can they be automatically estimated from instrumented patient-caregive data.

A later section will describe how to optimize these behaviors using simulation.

4 Proposed Work: Robotic Testbeds

To evaluate our behavior measurements, models, and control approaches we will implement two types of robot balance assist systems. The first involves an external robot that reaches out and touches the patient, much like a human caregiver would. The second approach uses a safe soft wearable device to simulate (as much as possible) the mechanical effect of a caregiver's touch. It is possible that a lightweight robot arm mounted on the patient reaches out and touches the enviroment, for example [**?**]. A wearable device could use reaction wheels to generate forces.

External robots: Our current plan for external devices will be refinements of our inflatable arms shown in Figure 3. We are focusing on standing balance, sit-to-stand transitions, sitting down, and walking on a treadmill so that we do not need to intially develop extensive mobility. Developing legs or a wheeled base is part of our future work. We expect systems to need fast mobility to handle fast patient foot support changes or collapse such as knees buckling, so we expect final systems to need omnidirectional wheels or legs.

We will attach our inflatable arms to nearby walls or furniture (an advantage of these lightweight robots is that they can easily be deployed and moved around in a home environment.

We have experimented with both pneumatic actuators such as McKibben actuators [28, 29], and other forms of tendon driven actuation, such as tendons attached to rotary or linear electric actuators. It is not yet clear which type of actuator is superior for this application, so we will compare current state of the art systems. A concern with inflatable robots is payload. One of our pneumatic arm prototypes is capable of lifting 7kg (70N). We expect this is more than enough for light interpersonal touch. However, If higher force levels are needed, we may explore hydraulically actuated systems, which will also allow higher stiffness, but will weigh more. Hydraulically actuated systems have several advantages over pneumatically actuated systems. The system response times are typically much faster, energy losses are much smaller, and for incompressible fluid forces per unit area may be much larger allowing for more compact design of system with the same peak force dynamical output. We will explore systems with mixed actuation, where pneumatics, hydraulics or electric motors are used as needed. We will also explore other hybrid actuation schemes, where different types of actuation are used on the same degree of freedom. Macro-Mini control [107] demonstrated that electric motors could provide accurate and fast control, and pneumatic muscles could generate high force output, for example. This hybrid actuation concept has been tested in our previous work on a human safety robot [80, 81].

Wearable Systems: A natural extension of our work on lightweight inflatable robots is extending it to wearable systems. Inflatable struts can provide bracing and support. Pneumatic actuation can be used, as well as other types of lightweight actuation.

We are currently exploring how to embed pneumatic actuators in two dimensional sheets or three dimensional volumes [54, 93] for increased actuation force and controllability. We have developed pneumatic actuators that can be arrayed in the fabric of a worn co-robot (Figure 6). Our miniaturized pneumatic artificial muscle arrays are composed of elastomer air chambers with embedded Kevlar fibers. When the air chambers are filled, the constrained length of the Kevlar fibers reduces the axial length of the chamber resulting in axial contraction. The modular design of the muscle network makes the system easily resizable and reconfigurable. This two-dimensional design of the muscle will not only simplify the fabrication process but also make the actuation material thin and light to be wearable or as actuation for a lightweight robot. Figure 6 shows a preliminary prototype of a flat pneumatic artificial muscle array. Our prior work has successfully demonstrated that this type of actuators can be implemented for actuation of wearable assistive devices [59, 65, 31] generating relatively high contraction force for active assistance of knee flexion and extension motions. The fiber path can be designed for the expected co-robot behaviors.

5 Proposed Work: Sensing

There are three key innovations we will pursue in our work on sensing for action: 1) We will develop a whole-body vision system that uses optical sensing from all robot skin, not just localized eyes. 2) We will develop soft

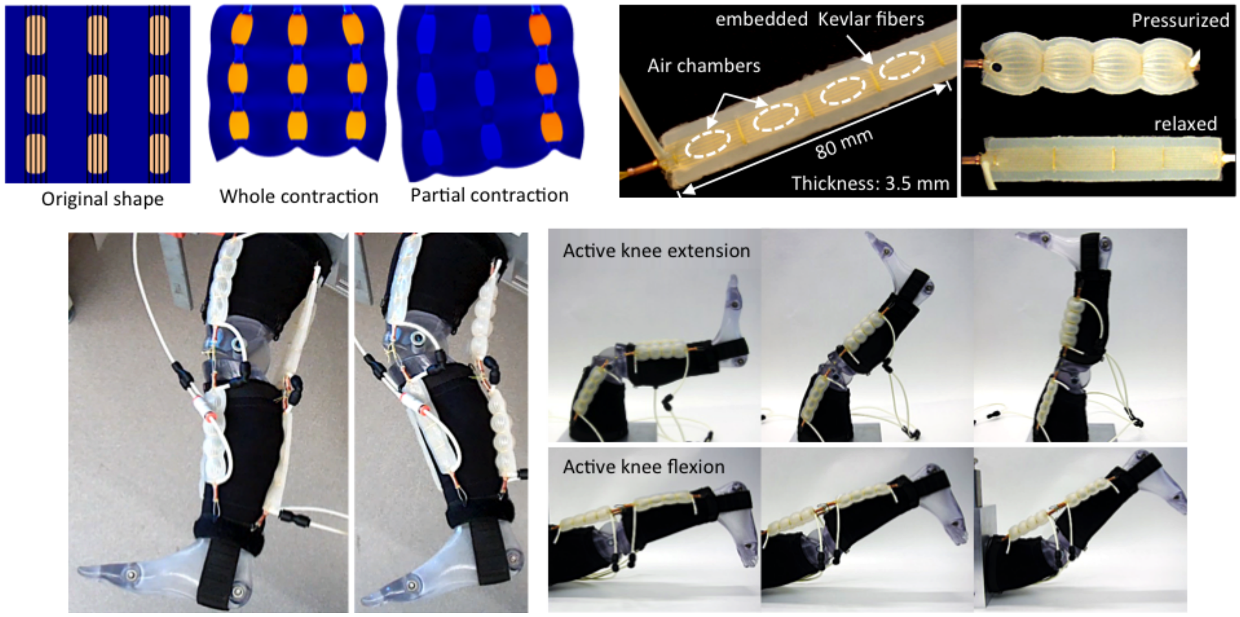

Figure 6: Flat pneumatic artificial muscle array and its application: **Top Left:** Design concept. **Top Right:** Actual prototype with four muscle cells. **Bottom:** Implementation to a wearable soft robotic device

skin sensors that work well on inflatable arms. 3) We will develop ultrasound-based tracking of nearby patient tissue to avoid injury and pain. A recent overview of robot skin and robot skin-based sensing is provided by [**?**].

Kinematic sensing: We will use joint angle measurments, augmented by link angular velocities, linear accelerations measured by MEMS gyros and accelerometers, to estimate where the robot is. In addition, we intend to have a largely transparent skin, and use simple IR range finders and embedded cameras to sense objects near to the arm and more accurately localize arm parts withing the space.

Contact force sensing: Average contact forces across a hand can be estimated using six axis wrist force torque sensing, which we will build into our hands. Local force sensing to detect local concentrations of force will be done by measuring skin strain in various directions. We will explore the use of existing tactile imaging devices, such as Tekscan material, which we currently use, but it only senses compression forces in the normal direction (perpendicular to the skin/sensor). We think obtaining estimates of local shear forces (parallel to the skin/sensor) will be important in reducing the risk of skin damage to the patient. In addition to estimating shear forces from robot skin strain (stretch) measured by traditional strain gages, strain gages printed directly on the robot skin, strain measured by liquid metal strain gages (Figure 4), or strain measured optically by cameras inside the skin or arm that we are developing, along the lines of [**?**, **?**, **?**, **?**, **?**, **?**].

Ultrasound sensing: We will explore using ultrasound transducers mounted in the skin of the robot to see into the person being touched, to track any visible bones, tissue movement, and estimate tissue strain. We can use real time ultrasound measurements to locate bones to guide manipulation. We can also use ultrasound to monitor tissue deformation (displacement and strain) to guide forces and detect when a manipulation is failing and should be stopped or adjusted. Using phase-sensitive 2D speckle tracking [**?**, **?**], tissue deformation can be accurately measured in real time at high spatial (< 1 mm) and temporal (< 10 ms) resolution with a large imaging depth of a few cm at typical clinical ultrasound imaging frequencies (2-10 MHz). In year 1, we will do feasibility tests using a commercially available linear array transducer, connected to a commercial ultrasound research platform (Verasonics US Engine, Verasonics Inc, WA, USA), which will be mounted in the robot skin. In year 2, a set of off-the-shelf single element transducers will be mounted on the palm and fingers of the robot hand for optimal sensing, imaging and feedback. Based on the design parameters obtained during Y2, a custom built array transducer system incorporated in the skin will be developed and evaluated in Y3. We expect substantial revision and improvements in our volumetric sensing in response to how well it does in experimental use, including development of better image processing algorithms, as we go. In principle, the

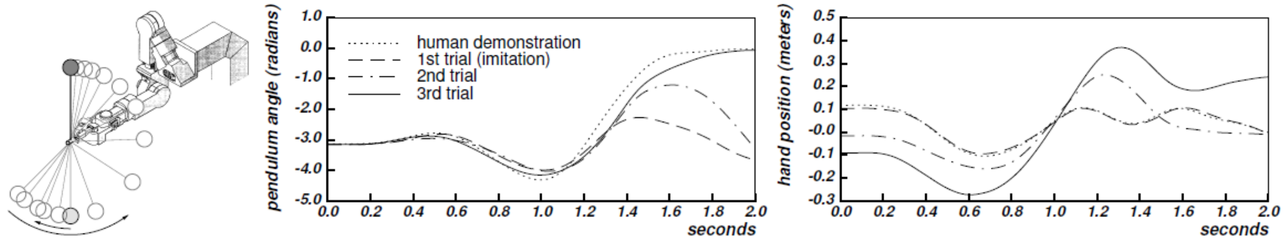

Figure 7: **Left:** The robot swinging up an inverted pendulum by moving the hand side to side. **Right:** The pendulum and hand motion during robot learning from demonstration and practice using a non-parametric model.

concept of mounting ultrasound transducers in robot skin is very similar to the detection of prostate cancers from an ultrasound probe mounted on a doctor's finger in the rectum to measure the elastic properties of nearby tissue [**?**, **?**]. The integration of an ultrasound linear array transducer onto the robot hand in Y1 will also be similar to the approaches of a previous study of muscle fatigue in the forearms [**?**]. Building a custom ultrasound transducer opens the possibility of distributing the transducer across the skin, or putting elements of the transducer in support surfaces such as a bed or chair, and doing transmission ultrasound imaging as well as reflective imaging. Transmission imaging may give us better images at greater depth.

6 Proposed Work: Programming Safe Soft Robots

For our approach to be successful, we need to make the proposed robots easy to program. The mechanical design of our co-robot presents new control and programming challenges: 1) Our robots do not necessarily have well-defined rotary or linear joints, but instead may have deformations spread over a structure. This has implications for defining and estimating robot state and planning behaviors. 2) Our robots are better described as maintaining an impedance around an equilibrium configuration, rather than moving to a position. Other aspects of control, such as a focus on force control when in contact with a human or objects, is common to many co-robot approaches. 3) We expect most manipulations of deformable objects such as humans with soft skin to be difficult to model.

We propose using policy optimization and learning to control the robot and improve its performance over time instead of standard model-based optimal control. We will explore learning complex policies for physically interacting with humans by watching humans execute them or by programming them directly based on our understanding of these procedures, and then refining those policies using optimization and learning. We will explore behavior and control based on explicit object trajectories and force control, policy optimization, discriminant or predicate based policies, and matching learned sensory templates. We will develop new control techniques that refine existing human strategies and invent new ones, based on cognitive optimization: general-purpose systems that learn to maximize reward or utility, using new, more powerful model-based forms of adaptive/approximate dynamic programming (ADP) and reinforcement learning. We will develop efficient multiple model methods for designing robust nonlinear and time varying task controllers for physical human-robot interaction where the robot, the patient, and possibly other caregivers all participate in the interaction.

In addition to new programming methods we will develop, we will support several existing methods to specify shapes and contact forces: 1) direct specification of shapes and forces, 2) programming by demonstration, and 3) programming by direct manipulation. All of these specification phases will be followed by an optimization phase to optimize a policy that achieves the specification. In direct specification, the programmer uses current programming languages and tools to develop robot policies to achieve particular tasks. We use this approach for almost all of our current robot programming.

Programming by demonstration and programming by direct manipulation are closely related. In programming by demonstration the programmer manipulates their own body or tools to demonstrate how to do the task. Body or tool motion is captured by a motion capture system or a computer vision system. Programming by direct manipulation involves grasping the robot and physically moving it to desired positions. This is possible since our robots are soft and will allow a lot of deformation. We will build on our previous work on learning parametric and non-parametric robot and task models [8, 51, 53, 52, 78, 76, 79, 77]. We will also build on learning from demonstration [**?**, **?**]. In our past work we used inverse models of the task to map task errors to command corrections [2, 1]. We have found that optimization is more effective than trying to track a learned reference movement, especially with non-minimum phase plants. Figure 7 shows an example of a robot learning to do a nonlinear, unstable, and non-minimum phase task (pendulum swing up) from watching a human do it (once) [9, 10, 6]. The figure shows the teacher's demonstration of how to manipulate a jointed deformable object by moving the hand side to side and the robot's practice trials. On the first attempt the robot imitates its perception of the teacher's movement, which fails to swing the jointed object upright. The robot then uses optimization and an updated model of the task dynamics to adapt its hand motion. The 2nd trial is better, the model is updated again, and the 3rd trial succeeds. We have found that optimization greatly speeds up this type of learning. However, this type of learning sometimes gets stuck because updating the model with new data causes only a very slow change in the policy because the planned movement is in a different area of state space from the new data. We use policy optimization to solve this problem. We have also implemented direct policy learning to allow a robot to learn air hockey and a marble maze task from watching a human [25, 23, 24]. Other prior work on policy learning and optimization and learning includes [3, 85, 86, 87]. We will develop tools to make it easy for human programmers to "morph" demonstrations in response to feedback.

We expect our recently developed policy optimization approach to be more efficient and effective than our previous work, because we have developed very efficient first and second order gradient methods to do policy optimization [7]. Unlike inverse dynamics for torque-based robots, which can be unstable, our deformable robots are dominated by stable passive dynamics, so calculating actuator commands should be well conditioned. To facilitate this type of programming, we will develop a fast approximate simulator that uses an approximate rigid body skeletal model of the robot with thick deformable skin, and a slow accurate simulator that uses finite element models of the deformable parts of the robot and task, Our task models must necessarily include models of humans being interacted with. We take a multiple model optimization approach to try to capture user variability, in which a set of models is used to capture possible human responses. Our goal is to support rapid programming development, and to develop automatic tools that enable a robot program that works well on the fast but approximate model to be converted to work well on the slow but accurate model and on the real hardware.

We will also explore control approaches based on model-based approximate dynamic programming (ADP) and reinforcement learning.

7 Results from Prior NSF Support Related to this Work

PITT: *I need a section in this format from Pitt*

Co-PIs Coros and Park have had no prior NSF support.

Atkeson: (a) NSF award number: EEC-0540865; **amount:** $29,560,917; **period of support:** 6/1/06 - 5/31/15. **(b) Title:** NSF Engineering Research Center on Quality of Life Technology (PI: Siewiorek). **(c) Summary of Results:** This NSF Engineering Research Center supported during the period 2007-2013 a research project (lead C. Atkeson) on soft robotics. We only report on this part of the award.

Intellectual Merit. This project explored a design extreme: Can we build a useful robot arm with no rigid parts? We converged on an inflatable design, and built several prototype inflatable arms. We developed new torsional actuators, and new ways to make inflatable arms strong. We showed that inflatable arms could be fast and strong, dispelling common misconceptions.

Broader Impacts. This work inspired the inflatable robot character Baymax in the Disney movie Big Hero 6 [**?**]. As a result, the plot of the movie strongly supported STEM education and is having a huge positive impact.

Development of Human Resources. The project trained one graduate student, Siddharth Sanan. Siddharth is now a postdoctoral fellow at The Wyss Institute & School of Engineering and Applied Sciences (SEAS), Harvard University.

(d) Publications resulting from this NSF award: [74, 75, **?**].

(e) Other research products: Web pages www.cs.cmu.edu/~cga/soft and www.cs.cmu.edu/~cga/bighero6

(f) Renewed support. This proposal is not for renewed support.

8 Research Plan

The team has extensive experience in the areas covered by this proposal. Atkeson has been developing inflatable robots for several years, and has extensive experience in learning and optimal control, particularly of humanoids. Redfern is an expert in human balance control. (**PITT**: *Need Pitt person here*) is an expert in clinical assessment of balance and patient-caregiver interaction. Park is a leader in soft sensing and actuation. Coros is an expert in co-design, co-control, co-adaptation, and co-learning, as well as designing robot behaviors.

The postdoc will serve as the major coordinator between CMU and the University of Pittsburgh. One CMU student will focus on modeling human-robot physical interaction, based on the data from the University of Pittsburgh. Another CMU student will focus on robot fabrication. A third CMU student will focus on robot control.

Year 1: Year 1 will focus on aiding standing balance control and initiating the iterative robot design process. Pilot data will be taken from subjects, and initial subject models will be created. The existing robot arms and wearable devices will be used for initial conrol design, while new arm designs are fabricated. Quantitative performance, cost, and safety goals will be created, as well as simulations including the patient and the robots.

Year 2: The goal of the second year is completed systems for the aid of standing balance. This will involve more data and probably more accurate modeling of subjects. Task specific external robot arms and wearable devices will be fabricated and controlled. Preliminary data, models, and system for aiding balance during sit to stand, sitting down, and treadmill walking will be initiated.

Year 3: The goal of the third year is completed systems for the aid of balance during sit to stand, sitting down, and treadmill walking. We will collect more extensive data and refine our models. Walking specific external robot arms and wearable devices will be fabricated and controlled. Implementation and evaluation of all systems will be completed, with actual patients. Technology transfer activities based on open systems such as ROS will be increased.

In all years we will ask human subjects to interact with different versions of our systems, to get their input on design decisions, robot performance, and user acceptability issues throughout the work. We have performed this type of study previously. Most recently, we performed an IRB-approved study where an inflatable arm fed subjects. The benefits and risks of such studies is discussed in the supplementary document on human subject protection.

We will coordinate through a shared computer file system, weekly online meetings, and periodic physical visits.

9 Broader Impact II

PITT: *We need Pitt additions to this.*

This section extends Section **??**.

We believe this work will lead to benefits to society in terms of enabling people with disabilities and older adults to live independently, as well as producing economically useful robots for more general applications. We have already had productive interactions with two companies, iRobot and Otherlab.

We will take advantage of the low cost and ease of manufacture of the proposed co-robots for possible STEM education applications. We will explore how such robots can be widely and cheaply distributed, either by distributing instructions, producing kits, or selling actual robots. We will seek partners to develop appropriate K-12 curriculum.

Another form of outreach will be to transform our results and products to exciting educational materials that can be presented to K-12 students in various forms with an easy access. Examples of these materials include class notes, videos, and simplified prototypes that students can easily try and have increased interests and understanding of wearable robots, assistive robots, soft robots, humanoid robots etc. We will participate in CMU/Pitt NSF Engineering Research Center on Quality of Life Technologies outreach activities including the QoLT Ambassador program, TechLink, and BodyScout activities. The Carnegie Mellon Robotics Institute already has an aggressive outreach program at the K-12 level, and we will participate in that program, as well as other CMU programs such as Andrew's Leap (a CMU summer enrichment program for high school students), the SAMS program (Summer Academy for Mathematics + Science: a summer program for diversity aimed at high schoolers), Creative Tech Night for Girls, and other minority outreach programs.

Unplanned Broader Impact. Often the broader impacts of our work are serendipitous, and not planned in advance. Examples of such ad hoc broader impacts from our recent work include: 1) Our technologies being demonstrated on entries in the DARPA Robotics Challenge. 2) A graduate student was a participant on a Discovery Channel TV series, "The Big Brain Theory: Pure Genius". One purpose of the TV series is getting people excited about engineering. 3) A graduate student helped run a group of all female high school students in the robot FIRST competition. 4) Our work on soft robotics inspired the soft inflatable robot Baymax in the Disney movie Big Hero 6 [**?**]. We have participated in extensive publicity as a result. An explicit goal of this movie was to support STEAM. We expect to be involved with sequels to Big Hero 6, and to support Disney's STEAM effort (which is quite large and well funded). We expect similar unplanned broader impacts to result from this work, especially based on dramatic videos of agile and responsive robots.

Development of Human Resources. Students working on this project will gain experience and expertise in teaching and mentoring by assisting the course students in class projects. Students working on this project will also have the opportunity to train their communication and inter-personal skills, as they will actively participate in the dissemination of the research results at conferences and in related K-12 outreach programs of the Robotics Institute.

Participation of Underrepresented Groups. We will attract both undergraduate and graduate students, especially those from underrepresented groups. We will also make use of existing efforts that are part of on-going efforts in the Robotics Institute, and CMU-wide efforts, as well as similar efforts at WPI and NCSU. These efforts include supporting minority visits, recruiting at various conferences and educational institutions, and providing minority fellowships. CMU is fortunate in being successful in attracting an usually high percentage of female undergraduates in Computer Science. Our collaboration with the Rehabilitation Science and Technology Department of the University of Pittsburgh in the area of assistive robotics is a magnet for students with disabilities and students who are attracted by the possibility of working on technology that directly helps people.

Outreach. Our track record in this area is also discussed in the NSF-supported Prior Work section.

A major outreach program will be a collaboration with the Pittsburgh Carnegie Science Center to build soft robot kits for use in the museum and at home. Beginner kits will involve cutting and gluing materials, suitable for young kids who can be trusted with a pair of scissors and a supply of glue. Basically they would be making models that can be moved by hand. Intermediate kits will involve sensors, actuators, and computers. Advanced kits will involve using Maker technology to fabricate parts etc.

One form of outreach we have pursued is an aggressive program of visiting students and postdocs. This has been most successful internationally, with visitors from the Delft University of Technology (4 students) [4, 46, 103], the HUBO lab at KAIST (1 student and 1 postdoc) [15, 27, 35], and the Chinese Scholarship Council supported 5 students [42, 104]. We welcome new visitors, who are typically paid by their home institutions during the visit. We are currently experimenting with the use of Youtube and lab notebooks on the web to make public preliminary results as well as final papers and videos. We have found this is a useful way to support internal communication as well as potentially create outside interest in our work. We will continue to give lectures to visiting K-12 classes. The Carnegie Mellon Robotics Institute already has an aggressive outreach program at the K-12 level, and we will participate in that program, as well as other CMU programs such as Andrew's Leap (a CMU a summer enrichment program for high school students), the SAMS program (Summer Academy for Mathematics + Science: a summer program for diversity aimed at high schoolers), Creative Tech Night for Girls, and other minority outreach programs.

Dissemination Plan. For a more complete description of our dissemination plan, see our Data Management Plan. We will maintain a public website to freely share our simulations and control code, and to document research progress with video material. We will present our work at conferences and publish it in journals, and use these vehicles to advertise our work to potential collaborators in science and industry.

Technology Transfer. Our research results and algorithms are being used by iRobot, Otherlab/Pneubotics, and Disney Research and through the Disney Research technology transfer path will eventually be used in entertainment and education applications, and will be available to and inspire the public. Three former students work at Boston Dynamics/Google transferring our work to industrial applications, several students have done internships at Disney Research, and two former postdocs work there full time.

Benefits to Society. This is discussed in Section **??**.

Enhancement of Infrastructure for Research and Education. We will make our testbeds available to other researchers, and will encourage their replication at other sites.

Relationship to DRC funding. We are currently part of one of the funded teams in the DARPA Robotics Challenge. This support will end in the summer of 2015. The DRC focuses on reliability, implementation issues, and logistics. We are not able to develop the proposed ideas in the time frame of the DRC, so longer term NSF funding complements our soon to end DRC funding.

10 Curriculum Development Activities

We will develop curriculum based on this work. We will develop a new course on soft robots, as well as augmenting our existing courses on robot design, optimization, and humanoid robotics. Our robotics classes currently largely focus on rigid body dynamics, and there is almost no discussion of how to deal with soft objects, soft robots, or physical interaction with humans.

We will develop course materials on skin, tactile sensing, robot control, and biologically inspired approaches. These materials will directly be influenced by the planned activities of this proposal. The CMU PIs currently teach several courses that will benefit from this material. We make these course materials freely available on the web. We teach a course designed to attract undergraduates into the field, *16-264: Humanoids.* Park will create *Robotics Capstone,* a senior design course for the robotics second major undergraduate students, and *Mechanism Design,* a course that discusses various sensing and actuation mechanisms for robots for both undergraduate and graduate students. Both these courses will reflect the proposed work. Majidi is developing *MEG24-673: Special Topics on Soft Robotics* This interdisciplinary course examines the central role of material constitutive properties in robotics and includes hands-on projects that are performed in a dedicated laboratory workspace. Students are challenged to compare different paradigms for artificial muscle, soft sensors, and stretchable electronics. Majidi also participates in the ETH Summer School on Soft Robotics. He gave a series of lectures on the mechanics of soft-matter technologies. This included background on nonlinear elasticity and inflation mechanics and their application to the design and operation of pneumatic "artificial muscle" actuators.

References

[1] E. Aboaf, S. Drucker, and C. Atkeson. Task-level robot learning: Juggling a tennis ball more accurately. In *IEEE International Conference on Robotics and Automation*, pages 1290–1295, Scottsdale, AZ, 1989.

[2] Chae H. An, C. G. Atkeson, and John M. Hollerbach. *Model-Based Control of a Robot Manipulator*. MIT Press, Cambridge, MA, 1988.

[3] S. O. Anderson, J. K. Hodgins, and C. G. Atkeson. Approximate policy transfer applied to simulated bongo board balance. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2007.

[4] S. O. Anderson, M. Wisse, C. G. Atkeson, J. K. Hodgins, G. J. Zeglin, and B. Moyer. Powered bipeds based on passive dynamic principles. In *Proceedings of the 5th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2005)*, pages 110–116, 2005.

[5] C. Atkeson and J. Santamaria. A comparison of direct and model-based reinforcement learning. In *International Conference on Robotics and Automation*, 1997.

[6] C. G. Atkeson. Nonparametric model-based reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 10, pages 1008–1014. MIT Press, Cambridge, MA, 1998.

[7] C. G. Atkeson. Efficient robust policy optimization. In *American Control Conference*, 2012.

[8] C. G. Atkeson, A. W. Moore, and S. Schaal. Locally weighted learning. *Artificial Intelligence Review*, 11:11–73, 1997.

[9] C. G. Atkeson and S. Schaal. Learning tasks from a single demonstration. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation (ICRA97)*, pages 1706–1712, 1997.

[10] C. G. Atkeson and Stefan Schaal. Robot learning from demonstration. In *Proc. 14th International Conference on Machine Learning*, pages 12–20. Morgan Kaufmann, 1997.

[11] C. G. Atkeson and B. Stephens. Random sampling of states in dynamic programming. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 38(4):924–929, 2008.

[12] V. Azhmyakov, V. Boltyanski, and A. Poznyak. The dynamic programming approach to multi-model robust optimization. *Nonlinear Analysis, Theory, Methods & Applications*, 72(2):1110–1119, 2010.

[13] D. Bagnell and J. Schneider. Autonomous helicopter control using reinforcement learning policy search methods. In *IEEE International Conference on Robotics and Automation*, 2001.

[14] F. J. Bejarano, A. Poznyak, and L. Fridman. Min-max output integral sliding mode control for multiplant linear uncertain systems. In *American Control Conference*, pages 5875–5880, 2007.

[15] D. Bentivegna, C. G. Atkeson, and J.-Y. Kim. Compliant control of a hydraulic humanoid joint. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2007.

[16] D. C. Bentivegna and C. G. Atkeson. Using primitives in learning from observation. In *Humanoids 2000: The First IEEE-RAS International Conference on Humanoid Robotics*, 2000.

[17] D. C. Bentivegna and C. G. Atkeson. Learning from observation using primitives. In *Proceedings, IEEE International Conference on Robotics and Automation*, 2001.

[18] D. C. Bentivegna and C. G. Atkeson. Learning how to behave from observing others. In *SAB'02-Workshop on Motor Control in Humans and Robots: on the interplay of real brains and artificial devices*, Edinburgh, UK, 2002.

[19] D. C. Bentivegna, A. Ude, C. G. Atkeson, and G. Cheng. Humanoid robot learning and game playing using PC-based vision. In *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2449–2454, 2002.

[20] D.C. Bentivegna, C.G. Atkeson, and G. Cheng. Learning from observation and practice at the action generation level. In *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2003.

[21] D.C. Bentivegna, C.G. Atkeson, and G. Cheng. Learning from observation and practice using primitives. In Nakamura, editor, *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Los Vegas, NV, 2003.

[22] D.C. Bentivegna, C.G. Atkeson, and G. Cheng. Learning to select primitives and generate sub-goals from practice. In *Humanoids 2003: The Third IEEE-RAS International Conference on Humanoid Robotics*, 2003.

[23] D.C. Bentivegna, C.G. Atkeson, A. Ude, and G. Cheng. Learning tasks from observation and practice. *Robotics and Autonomous Systems*, 47:163–169, 2004.

[24] D.C. Bentivegna, C.G. Atkeson, A. Ude, and G. Cheng. Learning to act from observation and practice. *International Journal of Humanoid Robotics*, 1(4):585–611, 2004.

[25] D.C. Bentivegna, G. Cheng, and C.G. Atkeson. Learning from observation and from practice using behavioral primitives. In *11th International Symposium on Robotics Research*, Siena, Italy, 2003.

[26] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.

[27] D. Choi, C. G. Atkeson, S. J. Cho, and J. Y. Kim. Phase plane control of a humanoid. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 145–150, 2008.

[28] C.-P. Chou and B. Hannaford. Measurement and modeling of McKibben pneumatic artificial muscles. *IEEE Trans. Rob. Autom.*, 12(1):90–102, 1996.

[29] F. Daerden and D. Lefeber. Pneumatic artificial muscles: actuators for robotics and automation. *Eur. J. Mech. Environ. Eng.*, 47:10–21, 2002.

[30] L. A. Feldkamp and G. V. Puskorius. Fixed-weight controller for multiple systems. In *International Conference on Neural Networks*, pages 773–778, 1997.

[31] E. C. Goldfield, Y.-L. Park, B. Chen, W.-H. Hsu, A. Wessendorf, D. Young, M. Wehner, D. Stephen, L. Stirling, D. Newman, R. Nagpal, E. Saltzman, K. G. Holt, C. Walsh, and R. J. Wood. Bio-inspired design of soft robotic assistive devices: The interface of physics, biology, and behavior. *Ecological Psychology*, 24(4):300–327, 2012.

[32] iRobot. Advanced platforms, 2014. www.irobot.com/us/cool_stuff/Research/Advanced_Platforms.aspx.

[33] S. M. Joshi. Optimal output feedback control of linear systems in presence of forcing and measurement noise. Technical Memorandum TM X-72003, NASA, Langley Research Center, Hampton VA, 1974. DTIC ADA280162; handle.dtic.mil/100.2/ADA280162.

[34] K. Jung, K. J. Kim, and H. R. Choi. A self-sensing dielectric elastomer actuator. *Sens. Actuators, A*, 143(2):343–351, 2008.

[35] J.-Y. Kim, C. G. Atkeson, J. K. Hodgins, D. Bentivegna, and S. J. Cho. Online gain switching algorithm for joint position control of a hydraulic humanoid robot. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2007.

[36] J. Z. Kolter. *Learning and control with inaccurate models*. PhD thesis, Stanford, August 2010.

[37] R. Kramer, C. Majidi, R. Sahai, and R. J. Wood. Soft curvature sensors for joint angle proprioception. In *Proc. IEEE Int. Conf. Intell. Rob. Syst. (IROS'11)*, pages 1919–1926, 2011.

[38] K. Kure, T. Kanda, K. Suzumori, and S. Wakimoto. Flexible displacement sensor using injected conductive paste. *Sens. Actuators, A*, 143(2):272–278, 2008.

[39] S. Kuriyama, M. Ding, Y. Kurita, J. Ueda, and T. Ogasawara. Flexible sensor for McKibben pneumatic artificial muscle actuator. *Int. J. Autom. Technol.*, 3(6):731–740, 2009.

[40] F. L. Lewis, G. Lendaris, and D. Liu. Special issue on approximate dynamic programming and reinforcement learning for feedback control. *IEEE Trans. Syst. Man, Cybern – B: Cybernetics*, 41(1):896–897, 2008.

[41] F. L. Lewis and D. Vrabie. Reinforcement learning and adaptive dynamic programming for feedback control. *IEEE Circuits and Systems Magazine*, 9(3):32–50, 2009.

[42] C. Liu and C. G. Atkeson. Standing balance control using a trajectory library. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3031 — 3036, 2009.

[43] V. J. Lumelsky, M. S. Shur, and S. Wagner. Sensitive skin. *IEEE Sens. J.*, 1(1):41–51, 2001.

[44] C. Majidi, R. Kramer, and R. J. Wood. Non-differential elastomer curvature sensors for softer-than-skin electronics. *Smart Mater. Struc.*, 20(10):105017, 2011.

[45] P. M. Mäkilä and H. T. Toivonen. Computational methods for parametric LQ problems — a survey. *IEEE Transactions on Automatic Control*, AC-32(8):658–671, 1987.

[46] T. Mandersloot, M. Wisse, and C.G. Atkeson. Controlling velocity in bipedal walking: A dynamic programming approach. In *Proceedings of the 6th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2006.

[47] M. McNaughton. CASTRO: robust nonlinear trajectory optimization using multiple models. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 177 – 182, 2007.

[48] Y. Menguc, Y.-L. Park, E. Martinez-Villalpando, P. Aubin, M. Zisook, L. Stirling, R. J. Wood, and C. J. Walsh. Soft wearable motion sensing suit for lower limb biomechanics measurement. In *Proc. IEEE Int. Conf. Rob. Autom. (ICRA'13)*, Karlsruhe, Germany, May 2013.

[49] Y. Menguc, Y.-L. Park, H. Pei, D. Vogt, P. Aubin, E. Winchell, L. Fluke, L. Stirling, R. J. Wood, and C. J. Walsh. Wearable soft sensing suit for human gait measurement. *Submitted to Int. J. Rob. Res.*, 2013.

[50] W. T. Miller, R. S. Sutton, and P. J. Werbos. *Neural Networks for Control*. MIT Press, 1990.

[51] J. Morimoto and C. G. Atkeson. Improving humanoid locomotive performance with learnt approximated dynamics via Gaussian processes for regression. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007.

[52] J. Morimoto and C. G. Atkeson. Nonparametric representation of an approximated Poincaré map for learning biped locomotion. *Autonomous Robots*, 27(2):131–144, 2009.

[53] J. Morimoto, S. H. Hyon, C. G. Atkeson, and G. Cheng. Low-dimensional feature extraction for humanoid locomotion using kernel dimension reduction. In *IEEE-RAS Conference on Robotics and Automation*, pages 2711–2716, 2008.

[54] T. Nakamura and H. Shinohara. Position and force control based on mathematical models of pneumatic artificial muscles reinforced by straight glass fibers. In *Proc. IEEE Int. Conf. Rob. Autom. (ICRA'07)*, pages 4361–4366, 2007.

[55] A. Y. Ng and M. Jordan. Pegasus: A policy search method for large MDPs and POMDPs. In *Uncertainty in Artificial Intelligence, Proceedings of the Sixteenth Conference*, 2000.

[56] R. H. Nyström, J. M. Böling, J. M. Ramstedt, H. T. Toivonen, and K. E. Häggblom. Application of robust and multimodel control methods to an ill-conditioned distillation column. *Journal of Process Control*, 12:39–53, 2002.

[57] J. O'Reilly. Optimal low-order feedback controllers for linear discrete-time systems. In C. T. Leondes, editor, *Control and Dynamic Systems*, volume 16, pages 335–367. Academic Press, 1980.

[58] Otherlab. Otherlab, 2014. otherlab.com.

[59] Y.-L. Park, B. Chen, C. Majidi, R. J. Wood, R. Nagpal, and E. Goldfield. Active modular elastomer sleeve for soft wearable assistance robots. In *Proc. IEEE/RSJ Int. Conf. Intell. Rob. Syst. (IROS'12)*, pages 1595–1602, Vilamoura, Portugal, October 2012.

[60] Y.-L. Park, B. Chen, N. Perez-Arancibia, D. Young, L. Stirling, R. J. Wood, E. C. Goldfield, and R. Nagpal. Design and control of a bio-inspired soft wearable robotic device for ankle-foot pathologies. *Bioinspiration & Biomimetics*, 9(1), 2014.

[61] Y.-L. Park, B. Chen, and R. J. Wood. Soft artificial skin with multi-modal sensing capability using embedded liquid conductors. In *Proc. IEEE Sens. 2011 Conf.*, pages 81–84, Limerick, Ireland, October 2011.

[62] Y.-L. Park, B. Chen, and R. J. Wood. Design and fabrication of soft artificial skin using embedded microchannels and liquid conductors. *IEEE Sens. J.*, 12(8):2711–2718, 2012.

[63] Y.-L. Park, B. Chen, D. Young, L. Stirling, R. J. Wood, E. Goldfield, and R. Nagpal. Bio-inspired active soft orthotic device for ankle foot pathologies. In *Proc. IEEE/RSJ Int. Conf. Intell. Rob. Syst. (IROS'11)*, pages 4488–4495, San Francisco, CA, September 2011.

[64] Y.-L. Park, C. Majidi, R. Kramer, P. Bérard, and R. J. Wood. Hyperelastic pressure sensing with a liquid embedded elastomer. *J. Micromech. Microeng.*, 20(12), 2010.

[65] Y.-L. Park, J. Santos, E. C. Goldfield, and R. J. Wood. A soft wearable robotic device for active knee motions through development of compact flat pneumatic artificial muscles. In *Proc. IEEE Int. Conf. Rob. Autom. (ICRA'14) - to appear*, Hong Kong, China, May 2014.

[66] Y.-L. Park, D. Tepayotl-Ramirez, R. J. Wood, and C. Majidi. Influence of cross-sectional geometry on the sensitivity of liquid-phase electronic pressure sensors. *Appl. Phys. Lett.*, 101(19), 2012.

[67] Y.-L. Park and R. J. Wood. Smart pneumatic artificial muscle actuator with embedded microfluidic sensing. In *Proc. IEEE Sens. 2013 Conf.*, pages 689–692, Baltimore, MD, November 2013.

[68] Y.-L. Park, D. Young, B. Chen, R. J. Wood, R. Nagpal, and E. C. Goldfield. Networked bio-inspired modules for sensorimotor control of wearable cyberphysical devices. In *Proc. Int. Conf. Comput. Netwroking Commun. (ICNC'13)*, pages 92–96, San Diego, January 2013.

[69] Y. Piguet, U. Holmberg, and R. Longchamp. A minimax approach for multi-objective controller design using multiple models. *International Journal of Control*, 72(7-8):716–726, 1999.

[70] W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley-Interscience, 2nd edition, 2011.

[71] D. V. Prokhorov. Backpropagation through time and derivative adaptive critics — a common framework for comparison. In *Handbook of Learning and Approximate Dynamic Programming* [83], pages 381–404.

[72] D. V. Prokhorov. Training recurrent neurocontrollers for robustness with derivative-free Kalman filter. *IEEE Transactions on Neural Networks*, 17(6):1606–1616, 2006.

[73] S. Sanan. *Soft Inflatable Robots for Safe Physical Human Interaction*. PhD thesis, Carnegie Mellon University, 2013.

[74] S. Sanan, J. Moidel, and C. G. Atkeson. Robots with inflatable links. In *IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, pages 4331–4336, 2009.

[75] S. Sanan, M. H. Ornstein, and C. G. Atkeson. Physical human interaction for an inflatable manipulator. In *IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 7401–7404, 2011.

[76] S. Schaal and C. G. Atkeson. Constructive incremental learning from only local information. *Neural Computation*, 10(8):2047–2084, 1998.

[77] S. Schaal and C. G. Atkeson. Learning control for robotics. *IEEE Robotics & Automation Magazine*, 17(2):20–29, 2010.

[78] S. Schaal, C. G. Atkeson, and S. Vijayakumar. Real-time robot learning with locally weighted learning. In *Proceedings, IEEE International Conference on Robotics and Automation*, 2000.

[79] S. Schaal, C. G. Atkeson, and S. Vijayakumar. Scalable locally weighted statistical techniques for real time robot learning. *Applied Intelligence*, 16(1), 2002.

[80] D. Shin, I. Sardellitti, Y.-L. Park, O. Khatib, and M. R. Cutkosky. Design and control of a bio-inspired human- friendly robot. *Springer Tracts in Adv. Rob.*, 54(43-52), 2009.

[81] D. Shin, I. Sardellitti, Y.-L. Park, O. Khatib, and M. R. Cutkosky. Design and control of a bio-inspired human-friendly robot. *Int. J. Rob. Res.*, 29(5):571–584, 2010.

[82] J. Shinar, V. Glizer, and V. Turetsky. Solution of a linear pursuit-evasion game with variable structure and uncertain dynamics. In S. Jorgensen, M. Quincampoix, and T. Vincent, editors, *Advances in Dynamic Game Theory - Numerical Methods, Algorithms, and Applications to Ecology and Economics*, volume 9, pages 195–222. Birkhauser, 2007.

[83] J. Si, A. G. Barto, W. B. Powell, and D. Wunsch II. *Handbook of Learning and Approximate Dynamic Programming*. Wiley-IEEE Press, 2004.

[84] T. Someya, T. Sekitani, S. Iba, Y. Kato, H. Kawaguchi, and T. Sakurai. A large-area, flexible pressure sensor matrix with organic field-effect transistors for artificial skin applications. *PNAS*, 101(27):9966–9970, 2004.

[85] M. Stolle, H. Tappeiner, J. Chestnutt, and C. G. Atkeson. Transfer of policies based on trajectory libraries. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2007.

[86] Martin Stolle and Christopher G. Atkeson. Knowledge transfer using local features. In *Proceedings of the IEEE Symposium on Approximate Dynamic Programming and Reinforcement Learning (ADPRL)*, 2007.

[87] Martin Stolle and Christopher G. Atkeson. Finding and transferring policies using stored behaviors. *Autonomous Robots*, 29(2):169—200, 2010.

[88] H. T. Su and T. Samad. Neuro-control design: Optimization aspects. In O. Omidvar and D. L. Elliott, editors, *Neural Systems For Control*, pages 259–288. Academic Press, 1997.

[89] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.

[90] V. L. Syrmos, C. T. Abdallah, P. Dorato, and K. Grigoriadis. Static output feedback — a survey. *Automatica*, 33(2):125–137, 1997.

[91] R. Tajima, S. Kagami, M. Inaba, and H. Inoue. Development of soft and distributed tactile sensors and the application to a humanoid robot. *Adv. Rob.*, 16(4):381–397, 2002.

[92] K. Takei, T. Kakahashi, J. C. Ho, H. Ko, A. G. Gillies, P. W. Leu, R. S. Fearing, and A. Javey. Nanowire active-matrix circuitry for low-voltage macroscale artificial skin. *Nature Mater.*, 9(10):821–826, 2010.

[93] H. Tomori and T. Nakamura. Theoretical comparison of mckibben-type artificial muscle and novel straight-fiber-type artificial muscle. *Int. J. Autom. Technol.*, 5(4):544–550, 2011.

[94] A. Varga. Optimal output feedback control: a multi-model approach. In *IEEE International Symposium on Computer-Aided Control System Design*, pages 327–332, 1996.

[95] D. Vogt, Y.-L. Park, and R. J. Wood. A soft multi-axis force sensor. In *Proc. IEEE Sens. 2012 Conf.*, Taipei, Taiwan, October 2012.

[96] D. Vogt, Y.-L. Park, and R. J. Wood. Design and characterization of a soft multi-axis force sensor using embedded microfluidic channels. *IEEE Sens. J.*, 13(10):4056–4064, 2013.

[97] F.-Y. Wang, H. Zhang, and D. Liu. Adaptive dynamic programming: An introduction. *IEEE Computational Intelligence Magazine*, 4(2):39–47, 2009.

[98] M. Wehner, Y.-L. Park, C. J. Walsh, R. Nagpal, R. J. Wood, T. Moore, and E. Goldfield. Experimental characterization of components for active soft orthotics. In *Proc. IEEE Int. Conf. Biomed. Rob. Biomechatron. (BioRob'12)*, pages 1586–1592, Roma, Italy, June 2012.

[99] P. J. Werbos. Neurocontrol and supervised learning: An overview and evaluation. In *Handbook of Intelligent Control: Neural, Fuzzy, and Adaptive Approaches* [101], pages 65–89.

[100] P. J. Werbos. *The Roots of Backpropagation: From Ordered Derivatives to Neural Networks and Political Forecasting*. Wiley, 1994.

[101] D. A. White and D. A. Sofge. *Handbook of Intelligent Control: Neural, Fuzzy, and Adaptive Approaches*. Van Nostrand Reinhold, 1992.

[102] R. J. Williams. Adaptive state representation and estimation using recurrent connectionist networks. In *Neural Networks for Control* [50], pages 97–114.

[103] M. Wisse, C. G. Atkeson, and D. K. Kloimwieder. Swing leg retraction helps biped walking stability. In *Proceedings of the 5th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2005.

[104] D. Xing, C. G. Atkeson, J. Su, and B. Stephens. Gain scheduled control of perturbed standing balance. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4063–4068, 2010.

[105] I. V. Yannas and J. F. Burke. Design of an artificial skin. I. Basic design principles. *J. Biomed. Mater. Res.*, 14(1):65–81, 1980.

[106] S. Yao and Y. Zhu. Wearable multifunctional sensors using printed stretchable conductors made of silver nanowires. *Nanoscale*, DOI: 10.1039/C3NR05496A, 2014.

[107] M. Zinn, B. Roth, O. Khatib, and J. K. Salisbury. A new actuation approach for human –friendly robot design. *Int. J. Rob. Res.*, 23(4-5):379–398, 2004.