Toward Physical Interaction and Manipulation: Screwing in a Light Bulb

R. Grupen, J. Coelho, C. Connolly, V. Gullapalli,
M. Huber, K. Souccar
Department of Computer Science
University of Massachusetts, Amherst, 01003

1 Introduction

Manipulation tasks using dextrous robot hand/arm systems present formidable control challenges that are not amenable to traditional control methods. This is primarily because such systems are highly redundant, are applied to tasks composed of multiple goals, operate in uncertain and unstructured environments, and must capture a broad range of control contexts.

Interaction with the environment necessarily connotes a feedback process. The efficient exploitation of available feedback information is key to success in purposeful interaction with the environment. This document proposes a control structure for such systems, illustrates mechanisms for exploiting tactile and kinesthetic feedback, and presents a sequence of experiments designed to explore sensorimotor performance specifications in such systems.

2 Design Principles

The ultimate goal in designing manipulation systems is the autonomous execution of manipulatory tasks, given meaningful sensory information and limited computational power. In our work, we have adopted certain basic principles for the design of such systems:

1. Domain expertise can be encoded tacitly in control processes with formal stability and convergence results. The fundamental robustness of these controllers with respect to unmodeled disturbances can be exploited when operating in uncertain or unstructured environments.

2. Sensor evidence regarding relevant aspects of the process should be integrated into control at the highest possible rate.

3. The broad range of sensory contexts asks for adaptive control architectures, preferably expandable control structures with explicit mechanisms for skill acquisition.

3 Control Paradigms for Manipulation Systems

Classical **control theoretic** approaches require detailed process models. As a result, they only work as well as the process model predicts system performance. For most interesting problems, the process models are highly non-linear and non-stationary. **Motion planning** is fundamentally exponential in the number of degrees of freedom in the system and can be very sensitive to the complexity of the environment. Consequently, monolithic forms of model-referenced control or knowledge-based planning technologies do not scale to complex, unstructured systems.

Behavior-based approaches were also used in the implementation of manipulation systems. This paradigm partitions the control process into sets of subprocesses to address the complexity issue. And since models are expensive and often misleading as predictive devices, traditional models are avoided.

Behaviors embody mappings from specific sensorimotor apparatus into control policies for action. The result is a set of responsive, sensorimotor transformations. The coupling is hardwired and tuned carefully by the designer to yield acceptable results in the task domain.

Most often, these behavioral modules are designed to express elements of the target *psychology* for the robot — *wander, explore, follow, avoid.* Arrays of these behaviors capture qualitative performance as viewed by an external observer and speak only indirectly to the underlying sensorimotor controller. As a result, this approach leads to complex organizations of *ad hoc* control elements.

However, tuning behavioral interactions is itself an exponential process. In a behavioral network that captures context dependencies, the complexity of adapting behavioral interactions is exponential in the number of behaviors and in the range of relevant contexts. In order to respond to a broad range of operating contexts, the system designer must enumerate all sensorimotor alternatives, capture them as behavior, and tune the manner in which they interact.

This paper proposes that behavioral proliferation and its consequences are not necessary and that a much more efficient basis for sensorimotor control is possible without sacrificing system responsiveness. We propose that elemental control modules should be formally designed from a *declarative* specification of system performance, and implemented as asymptotically stable controllers (as opposed to procedural behaviors), amenable to *composition* into a predictable composite controller. Individual control modules are crafted to be responsive, and to be robust with respect to incomplete information. Furthermore, such control modules are designed to be a *basis* for a family of composite controllers.

Declarative control specifications express more abstract objectives, independent of sensors, mechanisms, the environment, or the task. Grasp stability, or more generally, imparting a force to the world through a mechanical contact, is such an objective. This goal can be expressed independent of the number of fingers, hands and arms employed, independent (largely) of the object geometry and task.

Correct nominal composition policies are difficult to derive in closed form, primarily because of the difficulty of modeling control module interactions and of predicting the range of control contexts. We propose a control architecture that combines control decomposition and learning. Our approach is based on using carefully designed sets of control modules together with learned control module activation policies to obtain near-optimal performance. Control decomposition can be used to *structure* the learning problem, ensuring that it remains tractable despite the high dimensionality and complexity of the control task [1, 2, 9, 6].

The resulting structure expresses the coupling between control modules. Adaptive composition schemes can exploit this structure by isolating the exploration of control strategies to relatively small subsets of controllers. Adaptation focuses on relevant control actions and is distributed over independent sets of sensorimotor controllers. Exploration is limited to safe, predictable variations of existing control compositions.

The system retains the baseline robustness and error suppression built in the underlying control modules as learning proceeds. Overall, systems exploiting the structure of a complex control problem should in principle acquire motor skill at significantly higher rates than systems exploring atomic action sequencing.

4 Task Design: Screwing in a Lightbulb

This demonstration task was chosen because it is representative of a large class of everyday manipulation tasks: turning a knob, picking up objects, performing insertions (tape in VCR, cube in square), flipping a light switch, (un)screwing a lid on a jar. In this section, our goal will be to derive conditions for executing this class of tasks autonomously in response to sensory information. We are not interested in special purpose solutions for a single task, but instead we wish to capture the characteristics for a broad class of simple manipulation tasks.

4.1 Assumptions

1. **Quasistatic Environments:** We focus on the role of information in generating solutions rather than focusing on the characteristic rate of that information.
2. **Task Specification:** We will not require that the robot knows what a lightbulb is, but simply knows which end goes in the socket and that the lightbulb engages the socket as result of a clockwise rotation about the principal axis of the lightbulb following a successful insertion. We consider the case when no other specific geometric information is available a priori.
3. **Control Structure:** In Figure 1, leaves in the structure establish the control basis. Each bifurcation is an independent control composition problem and defines the relevent working set of control. Working sets of control are sets of formally stable and convergent controllers that are composed based on relevant state information. In this paradigm, tasks are expressed as "gaits" over working sets of control, and each element of a gait is a dynamic composition of feedback control elements.

The task of screwing in a lightbulb is decomposed into the following sequence of control experiments:

1. "pick-and-place" — we will discuss Cartesian space resolution requirements for collision free reaching and demonstrate its use.
2. "sensory-guided insertion" — we will identify specifications for the effective use of an adaptive admittance control methodology.
3. "kinesthetic contact localization" — we will try to get a handle on specifications for practical contact observers.
4. "manipulation" — we will suggest that the proposed control design methodology yields re-usable manipulation controllers, mixed and matched to respond to specific tasks.

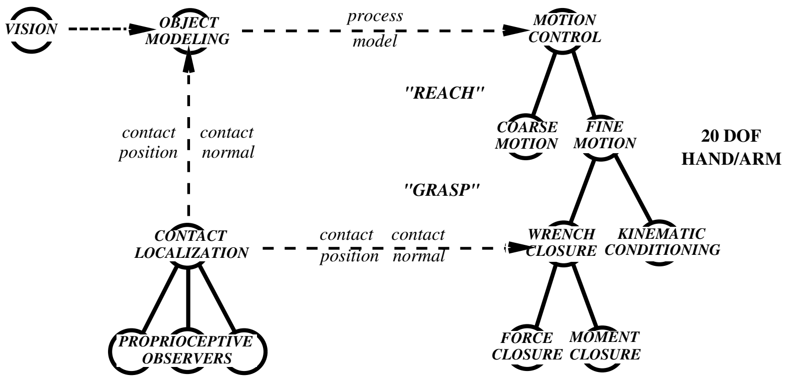

Figure 1: Structure of sensorimotor control for learning and skill acquisition in reaching and grasping tasks.

5 Experiments: Toward Physical Interaction And Manipulation

5.1 Experiment: Pick-and-Place

The first experiment consists of picking up a wooden block from a cluttered table and placing it at a different location. The target system for this design is our 20 DOF hand/arm system. This platform consists of a Utah/MIT dextrous hand (16 DOF) mounted on a GE P50 industrial robot (4 DOF). A single spatially calibrated overhead vision sensor was also used. The object and the point on the table to which it should be delivered are user specified by clicking a mouse on an overhead image of the cluttered table. The objects used are prisms with polygonal bases, approximately 20 cm in height, and with a maximum width between edges of 5 cm.

This object is modeled as a polygon using the Hough edge detection algorithm[1]. All other blocks and the table are mapped as obstacles into a 3 dimensional cylindrical space, which is in-turn mapped into a 4 dimensional configuration space used by the path planning. The selected object is mapped to two goal configurations; one on top, and one from the side.

The basis for controlling this system consists of four main components: harmonic function path planning [4], manipulability-based posture controller [5], and a wrench closure grasp controller [3].

This experiment introduced some evidence regarding the resolution required in the configuration space to accomplish collision free motions of the hand/arm system. The underlying representation for this control task consists of a Cartesian and configuration space bitmap. Each node in this mesh is marked as an obstacle or as an element of freespace. Obviously, the performance of this approach will degrade exponentially for both the mapping between bitmaps and the control computation with respect to dimension and resolution. We attempted to find an acceptable working solution for this domain.

A coarse discretization of the 4D configuration space (16x8x8x8), was used to span a 3D Cartesian workspace. The volume of the workspace mapped in this fashion was roughly $1 \times 1 \times 1$ m^3. In addition to controlling the endpoint position of the hand/arm, we also control its orientation to allow for both top and side grasps.

Figure 2 shows a sequence of pictures of the robot executing a harmonic controller to reach for a top grasp.

The coarse reach converges to a goal bin in a discrete, regular configuration space. Therefore, the Cartesian space resolution of the controller varies over the workspace. In practice, the cspace resolution was chosen to be equivalent to about one half the span between index finger and thumb on average over the workspace. After the coarse reach, it is assumed that there are no remaining obstacles en route to the final grasp. Finger controllers position each contact onto and subsequently over the object's surface. The evolving hand configuration is reflexively adjusted toward isotropic configurations by the action of a concurrent arm controller.

At this point, the user is prompted to select a delivery point. Again, this selection is a mouse click on the overhead image. After mapping the delivery goal to configuration space, a collision free path is computed for the arm/hand/objec system (illustrated in Figure 3). The same volume of the workspace and the same resolutions were found suitable during this operation.

[1] The approach was designed to use only contact information, but due to inadequate contact detection (see section 5.3), we simulated the local contact information with the help of a visual model.

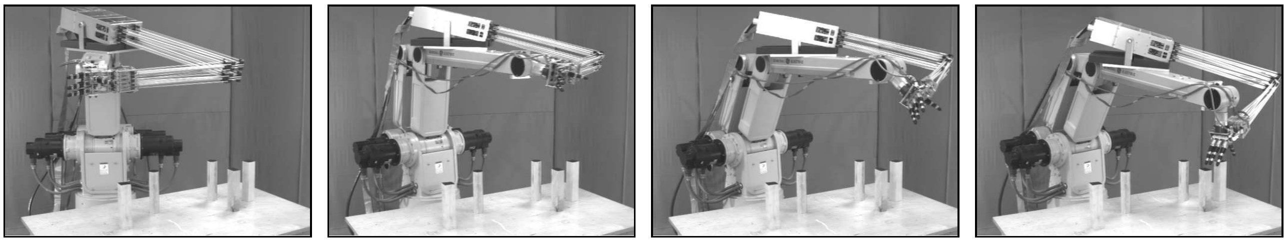

Figure 2: Harmonic Function Reach for a Top Grasp

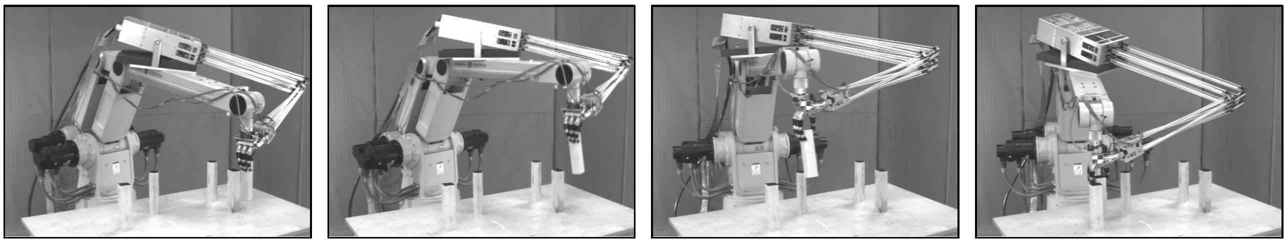

Figure 3: Robot System Delivering Object on a Cluttered Table

5.2 Experiment: Force-Guided Insertion

In this experiment, we attempted to get a handle on the specifications for a sensorimotor system that interprets a pattern of sensory feedback during an assembly task to compensate for uncertainty. We employed the approach described in [7] for learning a dynamic admittance relationship which maps sensed contact forces to compensatory motions during an insertion task.

Noisy force sensations compound the problem of control when the robot is interacting with external objects. Moreover, hybrid position/force controllers often exhibit instabilities near the contact boundary. The second source of uncertainty is due to the dynamics of the robot itself. Because of friction, gear backlash, etc., giving the same command in the same starting configuration can result in different motions of the robot. Our target platform for this experiment is our Zebra Zero robot. This robot is equipped with a wrist force sensor, and can also sense positions of the joints of the arm through position encoders. Endpoint precision of our beat-up Zero can be as bad as $\pm 5mm$ due primarily to gear backlash.

Figure 4 shows a representative performance on one test case in terms of the evaluation received over the course of insertions.

In addition to the obvious speedup in insertion time, another characteristic of the learning behavior can be observed from this figure by contrasting the curves showing performance after 100 and after 500 training runs. The presence of jitter in the evaluation after 100 training runs indicates that the force threshold is being violated often, causing drops in the evaluation. This jitter is not present in the curve showing performance after 500 training runs, indicating that the controller has learned to perform insertions without causing excessive forces to act on the peg. This approach has since been successfully applied to square pegs inserted into chamferless holes down to $0.001in$ clearances.

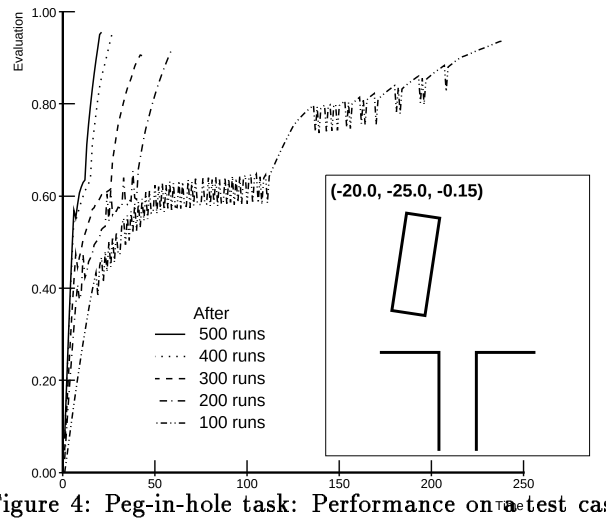

Figure 4: Peg-in-hole task: Performance on a test case after various intervals of training. The inset shows the initial position and orientation of the peg for this test case.

The force feedback in this system came from a wrist mounted, 6-axis force sensor. The object was bolted into the gripper. Although coupling inertial and gravitational loads to the signal, this sensor configuration promoted good signal to noise in the force information. The learning procedure was successful at rejecting high frequency noise in the sensor with amplitudes of approximately 0.1 N and 2.5 Nm, respectively. The sensor provides 10 bit A/D yielding a force resolution of approximately $0.05N$.

5.3 Experiment: Kinesthetic Contact Localization

Packaging tactile sensors has and continues to be a formidable engineering challenge. Many research groups have contributed to this technology and to the use of such information in modeling and control. In this experiment, we explored the use of proprioceptive information for contact localization. Patterns of manipulator position, velocity and torque feedback from the hand is used in the context of a Kalman filter to estimate the position and orientation of contacts[8]. An example of the system is included in Figure 5.

One limitation of this approach is the quality of the torque feedback provided by the hand, namely the tendon tensions which require contact forces of approximately $2 - 3$ N in order to unambiguously detect a contact. For effective manipulation and touch guided grasp formation, however, tolerable contact forces are more likely to be in the range of $0.1 - 0.5$ N to avoid movements of the object.

To cope with these requirements we are investigating the use of a hybrid system employing a low resolution tactile sensor to provide additional, improved force information for the kinesthetic localization.

5.4 Experiment: Object Reorientation

In the last sequence of experiments we attempt to get a handle on how re-useable our structured control architecture might be by applying it to several simple manipulation tasks. First, sequences of the pick-and-place control gait described in Section 5.1 were used to rotate a rectangle through a sequence of stable intermediate states. Figure 6 illustrate the idea.

The new objective is an extension of the pick-and-place task. User input in the reach and deliver modules is replaced by setting only one goal (forward or top, according to the direction of rotation) in *Coarse Reach*, and the other goal in *Deliver Turned*. This leads to a steady clockwise or counterclockwise reorientation. One counterclockwise rotation is shown in Figure 6.

The same set of control modules are composed in an alternative gait to manipulate objects while maintaining a stable grasp. The task described here is the rotation of a cylindrical object. In particular, the rotation of a soda can in the air, and (of course) screwing in a light bulb are presented. The control gait in these examples cycles through a succession of stable 3 fingered grasps to "walk" over the object's surface.

This experiment suggests that there is room for hope that one can exploit declarative force control specifications to address broad classes of tasks and with varying kinematic sub-chains in the mechanism. We feel that this conjecture could be important for building flexible, autonomous robots by mitigating the scalability problem in such systems. We are currently implementing a legged platform to test the reusability of these declarative control modules in another domain.

References

[1] R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, March 1986.

[2] R.A. Brooks. Intelligence without representation. In *Proc. of the Workshop on the Foundations of AI.* MIT Press, 1987.

[3] J.A. Coelho Jr. and R.A. Grupen. Constructing effective multifingered grasp controllers. In *Proc. of the 1994 Conf. on Robotics and Automation*, San Diego, USA, May 1994. IEEE Press.

[4] C.I. Connolly and R.A. Grupen. The applications of harmonic functions to robotics. *Journal of Robotic Systems*, 10(7):931–946, October 1993.

[5] R. Grupen and K. Souccar. Manipulability-based spatial isotropy: A kinematic reflex. In *Proc. of Int. Workshop on Mechatronical Comp. Systems for Perception and Action*, pages 157–163, Halmstad, Sweden, June 1993.

[6] R.A. Grupen. *General Purpose Grasping and Manipulation with Multifingered Robot Hands.* PhD thesis, University of Utah, Merrill Engineering Building, Salt Lake City, UT 84112, August 1988.

[7] V. Gullapalli, R. Grupen, and A. Barto. Learning reactive admittance control. In *Proc. of the 1992 Conference on Robotics and Automation*, pages 1475–1480, Nice, FRANCE, May 1992. IEEE.

[8] M. Huber and R.A. Grupen. 2d contact detection and localization using proprioceptive information. *IEEE Trans. on Rob. and Automation*, to appear.

[9] M.H. Raibert. *Legged Robots that Balance.* MIT Press, Cambridge, MA, 1986.

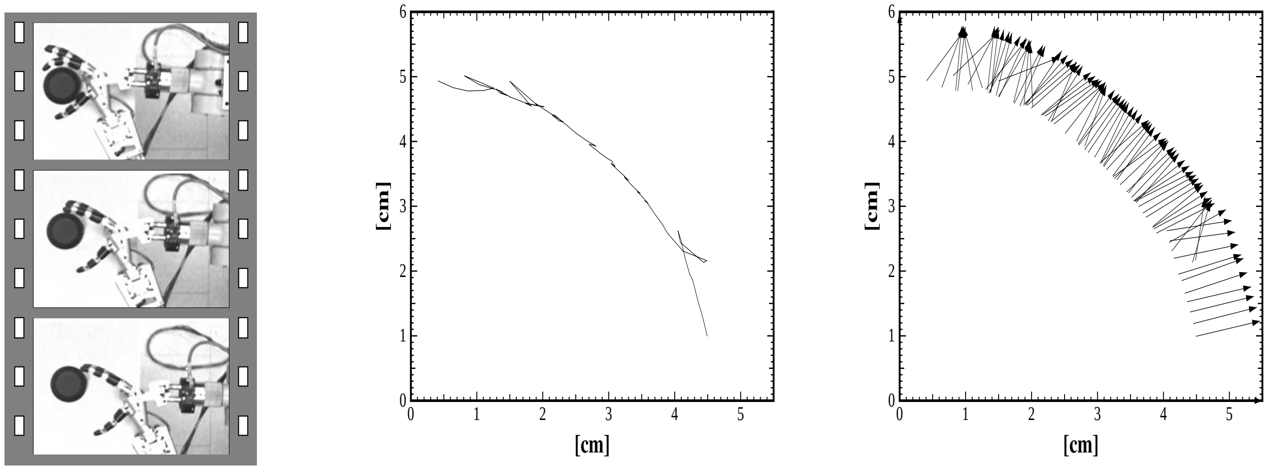

Figure 5: Contact Traces for Sweep over Circular Object

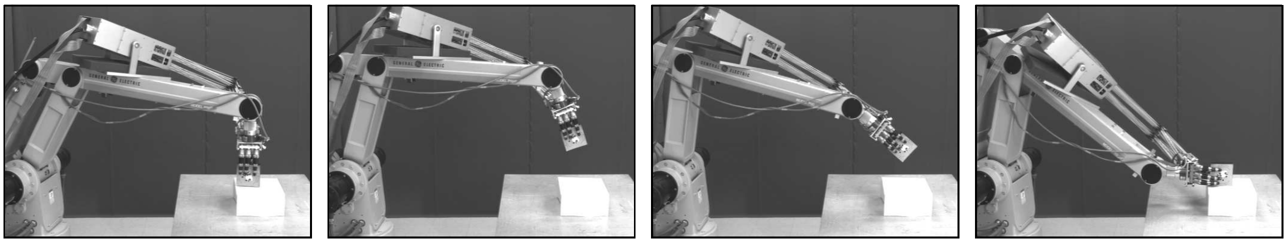

Figure 6: Robot System Performing Reorientation Task

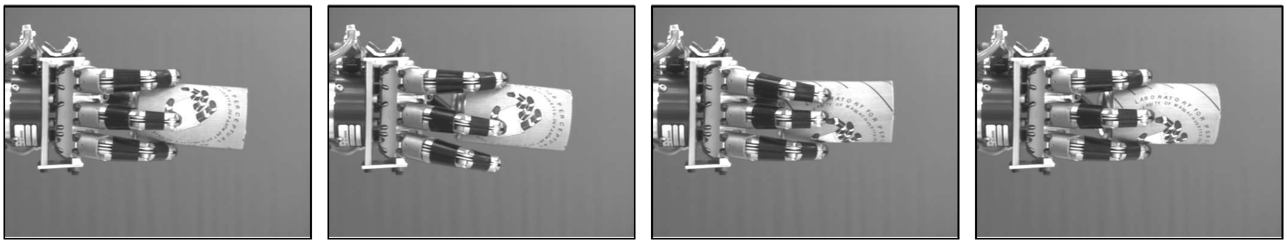

Figure 7: Utah/MIT hand Rotating Soda Can in Free-space

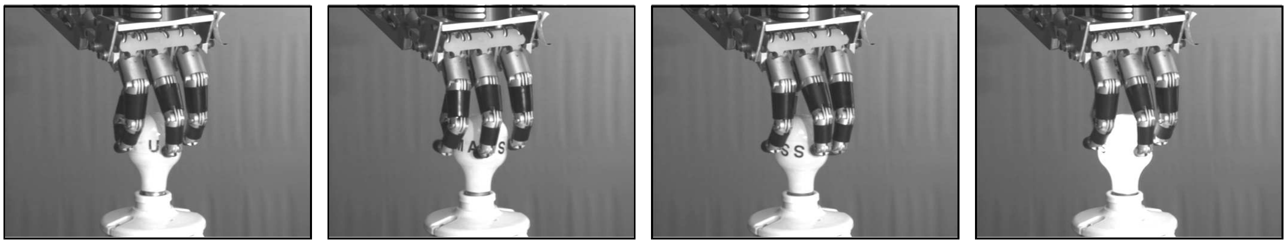

Figure 8: Utah/MIT hand Screwing in a Light Bulb