Sensorimotor Cognition — the Earliest Knowledge Structures

Rod Grupen — grupen@cs.umass.edu
Paul Cohen — cohen@cs.umass.edu

Contents

1 Project Description

A common paradigm in artificial intelligence takes the position that intelligent behavior is derived from lots of domain-specific knowledge. However, our efforts at engineering knowledge work only in narrow contexts and produce, therefore, very special purpose artifacts. The tradition in the behavioral sciences, in contrast, appreciates the interactive aspects of knowledge and intelligence, but provide relatively little insight into mechanisms. Recent paradigms emerging in Philosophy, Psychology, Linguistics and Robotics offer a plausible framework for modeling the process of sensorimotor development. Such a framework can function both as a descriptive account of human sensorimotor development and as a design principle for intelligent machines. This proposal advances a computational model for sensorimotor development and explores the implications of such a model. This framework leads to mechanisms for organizing experience as behavioral policies, and a means of segmenting sensory feedback gathered over a sequence of motor activities into "objects." These data structures form the basis, we argue, of a conceptual structure that influences all subsequent cognitive development. This proposal gathers insights from developmental psychology, control theory, reinforcement learning, robotics and artificial intelligence to look for the mechanisms responsible for our very earliest knowledge structures.

1.1 Introduction

By Piaget's account, the sensorimotor stage in human infants lasts roughly twenty four months [83, 84]. In the first four months, reflexive responses begin to organize into coherent motor strategies, sensory modalities are coordinated and attentional mechanisms begin to emerge. Native reflexive responses like the primary walking reflex and the palmar grasp reflex [100] provide primitive sensorimotor behavior that manage appropriate musculo-skeletal structures to do relevant, sensory-driven work in the world. Bruner [102] refers to these types of behaviors as "preadaptation" primitives for learning skillful motor policies.

Policies for coordinating multiple sensory and motor modalities appear next in Piaget's account. Primary circular reactions form, which can be viewed as a primitive form of an identity judgment. From four to six months, these primary circular reactions are repeated (assimilation) and perturbed (accommodation) until the infant finds it possible to prolong certain interactions with the world. Between six and eighteen months Piaget believed that these primary circular reactions direct subsequent exploration and influence evolving behavioral models of the world. During this same period, *figurative schemas* appear to encode information regarding "types" or "classes" of behavioral relationships in a manner that depends less on specific resources.

Complex robot systems also require some preadaptive structure to organize the acquisition of sensorimotor behavior. Our model (introduced in Section 1.2) describes a mechanism for constructing behavior as sequences of control situations. Control situations, in turn, are combinations of controllers and system resources. Although the details differ significantly across human and robot subjects, we claim that both systems must initially acquire a policy for engaging resources in a coordinated fashion, and that these policies form the most basic model of the agent/world interaction.

The central objective of the research is to test the hypothesis that many situations can be expressed in terms of relatively few, schematic structures that are grounded in physical activity, which we call *physical schemas* [66, 67, 68, 73, 74, 52] and the related hypothesis that physical schemas are a sensorimotor foundation for conceptual knowledge. One implication of our theory is that once stable relationships in the world are discovered and captured as physical schemas, an agent may notice other instances of these relationships involving different entities. This is one form of *metaphorical extension* [66], whereby physical schemas are extended to other physical, and sometimes nonphysical, phenomena. For example, underlying the word "contain" is a physical schema of the activities and relationships we call CONTAINMENT. Explicit in this schema is a stable pattern of dynamic activity in which one entity is constrained by another. This, we argue, is learned by direct physical interaction with the world when, for instance, we grasp an object. The entailments of CONTAINMENT are preserved if we consider that water is contained in a cup — it is only the participants in the relationship that have changed. Physical schemas, from this perspective, are a foundation for an enormous range of plausible physical and mental phenomena. The idea of physical schemas is not new [66, 67, 68], nor is its role in theories of conceptual development [73, 74, 52], but we are proposing a computational account of physical schemas implemented on robotic platforms.

Philosophical Underpinnings The approach advocated in this proposal challenges the philosophical tradition of Artificial Intelligence and many other fields. Mark Johnson argues that the foundation of epistemology — the study of knowledge — has been *objectivist* [66]. On this account, all knowledge is grounded in objective features of the world such as color, mass, shape, and so on. Were this not so, say the objectivists, we could not have common experience. Concepts are then represented as combinations of these features and the meanings of concepts are given by truth conditions; that is, an object is a chair iff it satisfies all the properties of chairs. This edifice is falling apart. Since Wittgenstein, philosophers have been concerned about the objectivist account of meaning; since Rosch [91], psychologists have recognized that the features underlying human concepts are not all objective; since Piaget [83, 84, 23], we have known that human knowledge arises from interaction with the world and that many of the features on which knowledge is based are subjective, not objective. For example, an infant ontology may distinguish graspable and non-graspable objects, but this distinction is clearly subjective in the sense that graspability depends on the size and strength of the infant's hand, and on the infant's ability to control its hand.

Interactionist Accounts of Knowledge The *interactionist* account grounds human knowledge in *activity*. Concepts such as "chair" are defined not in terms of necessary and sufficient conditions, but in terms of activities: "chair," in the most general sense, denotes anything I can sit upon. The concept of a chair appears to represent a set of objective features only because we humans tend to sit in very much the same way. Our most fundamental concepts derive from our earliest activities. Activity during the sensorimotor period of infancy produces physical schemas that serve for a lifetime as a basis for all conceptual knowledge, finding application to nonphysical concepts through metaphorical extension.

Many in AI continue to view knowledge — its acquisition, organization and use — in terms of Boolean combinations of objective features. The reliance on features *per se* isn't worrisome. Rather, it is the fact that the features we chose are often only indirectly meaningful to the agent. In the worst case, we must define the correct system response to every string of features — we must engineer both knowledge and behavior. Thus, our approach is interactionist not merely because objectivism is untenable, but because ultimately, it is the only approach that will produce agents with intrinsically meaningful mental states, and as a consequence, agents that don't require constant human supervision in novel situations.

Evidence for Physical Schemas Much of the evidence for physical schemas in humans comes from linguistics (and much of it is summarized in [67]). Lakoff and Johnson [68] were the first to speculate that metaphor is evidence that we organize experience in terms of a relatively small number of schematic models. In later writing, Lakoff and Johnson called these models *image schemas* and they described their structure informally, how they are used, and what they entail [66, 67]. Langacker has developed a "cognitive grammar" around similar ideas [69]. Spelke seems to have evidence of an amodal notion of RIGIDITY in infants as young as four months [95]. Other infant researchers have found evidence of physical schemas in very young children; for example, Leslie [70] has demonstrated a notion of causality in six month old children, and Baillargeon, in a controversial experiment, claims to have demonstrated that eight month old children understand the entailments of being a physical object — namely, spatio-temporal permanence [7, 6].

So there is evidence for physical schemas and evidence that infants develop physical concepts and their entailments very early. What is the evidence that such schemas are a basis for conceptual knowledge in artificial agents? Perhaps the logically prior question is whether *human* knowledge can develop from such a basis. Mandler has developed a theory of conceptual development based on image-schemas [73, 74]. The basic idea is that we categorize objects by how we interact with them — so graspable objects are different from ungraspable ones, self-animate objects are different from those that must be pushed, and so on. Because infants' earliest activities are physical, physical concepts are the basis for what develops later. As to artificial agents, Schank's *conceptual dependency* representation [92] has its share of physical primitives (recall, e.g., the three forms of TRANS, very like the PATH schema discussed by Johnson [66]). Schank and Riesbeck even worked out a suite of inference methods — which Lakoff and Johnson would call entailments — to go with each of the primitives in conceptual dependency. Drescher developed a program that acquired the Piagetian object concept through interaction with its environment, but he didn't account for general conceptual structure, nor ground it in physical schemas [48].

Nativism and Constructivism in Sensorimotor Development The nativist/constructivist debate concerns the nature of the pre-adaptive structure required to produce cognitive facilities such as categorization and abstraction. Nativists believe we are born with knowledge specific to tasks such as language and logic, constructivists believe we are born with domain-general principles that are adapted to specific tasks. If humans do learn physical schemas during the sensorimotor period then the dynamics of this process may explain how and when infants attain some important intellectual milestones. Some psychologists think these milestone concepts are innate, others think they are learned [30, 50, 75]. For example, much controversial research has been devoted to the question of whether human infants might develop the concept of *object* (an entity that persists even when the infant is not in direct sensorimotor contact with it) or whether this concept must be innate [6, 33, 98]. We have a unique opportunity in this project to explore whether artificial agents need to be "born with" domain-specific knowledge, or can develop it by the application of general principles in specific environments.

Section 1.2 describes a computational framework for sensorimotor development based in control theory, discrete-event dynamic systems (DEDS) and Markov Decision Problems (MDPs); an example of learning to compose behaviors into a schema is presented in Section 1.2.4. Section 1.3 illustrates how concepts can be extracted from physical schemas and Section 1.4 discusses how the physical schema CONTAINMENT might be learned. Finally, the proposed work plan is discussed in Section 1.5.

1.2 A Computational Framework: Physical Primitives and Physical Schemas

Every agent starts life with significant structure, our objective is to understand what kind of structure is required, and how it develops into conceptual knowledge. The components of this prior structure are called *physical primitives.* The choice and representation of physical primitives is important for many reasons, not the least of which is that they are the field on which the nativist/constructivist debate plays out (see Section 1.1).

The design of physical primitives is well-motivated on psychological and engineering grounds. The agent's most primitive actions will be implemented in the form of closed-loop *controllers* that are asymptotically stable and that produce equilibria that are correct solutions to generic subtasks. This decision is consistent with current research in infant motor development and adult motor control [119, 120, 25, 29] and robotics [112, 110, 47].

Modeling primitive actions as controllers has three related consequences. First, controllers suppress local perturbations by virtue of their closed-loop structure. Some variations in the context of a control task are simply rejected by the action of the controller. Second, controllers provide a basis for abstraction. Instead of dealing with a continuous state space, a behavioral scheme need only worry about control activation and convergence events. The result is a discrete model of controller/world interactions. Third, time is abstracted; in particular, durations of actions are not prescribed. This is appropriate, because the duration of a physical schema depends on the task and the context.

Using higher level primitives dramatically reduces the complexity of behavioral representations but preserves flexibility. In our model, behavior is constructed as a sequence of controllers bound to system resources. When a control objective is met for a particular resource commitment, a predicate is asserted in an abstract model of the system behavior. The "state" of the system is a vector of such predicates each element of which asserts a possible property of convergence for some controller and resource combination. This predicate vector can be viewed as the universe of possible subgoals available to an agent with these native controllers.

This representation provides a complete description of all physically plausible behavioral states and the control transitions between them. Having assessed its state in the predicated space, an agent can use such a representation to enumerate the range of sensory and motor alternatives available from this state. Even one ply of this local "behavior graph" might be huge in practice, so additional guidance is required and a variety of formal techniques can be used to provide that guidance. If we view the structure as a Markov Decision Process (MDP), then we may apply a variety of reinforcement learning techniques in order to derive behavioral policies (Section 1.2.4). If we view the structure as a Discrete Event Dynamic System (DEDS) then there exist tools for analyzing and interacting with the processes for acquiring behavior. DEDS provide a means of proving that certain predicate states can't occur and as a consequence, a means of investigating the role of constraints as "bootstraps" for a learning system. For example, *shaping* and *maturational mechanisms* can be approximated by constraining exploration and/or biasing the behavior acquisition process.

1.2.1 An Example of Physical Primitives: The Control Basis

To structure control and to reduce the complexity of behavior composition, the control basis approach [111, 114] uses a small set of carefully designed feedback control laws to construct behavior on-line. Individual controllers in this approach are largely task and device independent and represent solutions to generic robot control problems. This allows a small set of these controllers to span a large range of task domains on a wide variety of robot platforms. In addition, these elemental controllers can easily be combined to form composite controllers without the loss of the predictability imposed in their design. Using these properties, the control basis can be used to reduce the size of the state space that has to be considered since they represent stable attractors which partition the continuous physical space into individual basins of attraction.

An example of a control basis that we have used in previous studies consists of solutions to three generic robot control problems, namely configuration space motion control, contact configuration control, and kinematic conditioning [114, 111].

Φ_0: Configuration space motion control.
Harmonic function path controllers [109] are used to generate collision-free motion through the configuration space of the robot. Formally, this approach minimizes collision probabilities for the robot system [108] by following the gradient of a harmonic potential.

Φ_1: Contact configuration control.
A contact configuration controller [107] is employed to move contacts based on the local geometry of the environment in order to minimize residual forces and moments about some reference coordinate frame.

Φ_2: Kinematic conditioning.
This controller optimizes the posture of an articulated structure while it is engaged in an interaction with the world.

This means that, at the lowest level, the agent is either *moving* (Φ_0), *pushing* (Φ_1), or *posturing* itself (Φ_2) for moving and pushing tasks.

Each of these control laws can be bound to subsets of the input and output resources of the system. Instances of controllers, $\Phi_i\frac{\sigma}{\tau}$, are generated by associating sets of possible input resources, σ, and output resources, τ, with a control law, Φ_i. Each candidate control law can be used to manipulate the predicate state of a sensorimotor machine. In complex systems, this can be quite a large range of control alternatives. Additional constraints such as safety, or the focusing provided by maturational effects, or supervision can be used to further prune the space of possible behavior — these constraints can be modified or relaxed at a future point in the development of the machine.

Throughout system operation, the discrete part of the hybrid architecture is used to activate the controllers while maintaining stability and safety within the limitations of the basis controllers. The continuous control modules, on the other hand, interpret all the low level sensory information, satisfy the reactivity requirements, and form the interface to the DEDS component by providing the symbolic activation and convergence events.

1.2.2 System Identification

Each element of the control basis can generate many instances of controllers distinguished by their respective sensory and motor resource allocations, as described in Section 1.2. A predicate s_i is defined that describes the state of the associated controller ϕ_i. The goal for controller ϕ_i is to achieve an equilibrium state for which the control objective is satisfied. Under these circumstances, predicate s_i is true, asserting that, for instance, a particular stance of a walking machine is stable, or that the robot arm has achieved a reference configuration.

In general, an agent's predicate state $\vec{s}$ reflects the current state of all the possible controllers. While control primitives are generally *robust* it will sometimes be the case that controller ϕ_i is unable to assert its predicate s_i in the current context. We denote by $p(s_i = T|\phi_i(\vec{s}))$ the probability that controller ϕ_i reaches a correct equilibrium from an initial predicate state $\vec{s}$.

The *system identification* task is to learn $p(s_i = T|\phi_i(\vec{s}))$ for all predicate states $\vec{s}$. These probabilities constitute a *model* of the agent interacting with its environment. Clearly, the number of unique predicate states is enormous even if we limit ourselves to a few controllers and a handful of resources. Assuming n control predicates then each controller may be invoked from 2^n predicate states. System identification, formulated this way, involves learning $(2^n)^2$ transition probabilities.

1.2.3 Reinforcement Learning

Autonomous systems in largely unknown environments and faced with changing objectives must be capable of reacting to novel situations without requiring external supervision. In principle, reinforcement learning is well suited for this type of learning problem: it is generally on-line, exploration-based and does not require an external supervisor [101]. The only feedback required is the relative reward resulting from a sequence of behavior. The learning algorithm solves the temporal credit assignment problem by associating credit with elements of a behavioral sequence that leads to a reward. One of the major drawbacks of reinforcement learning methods, however is the large number of trials they require in order to first find and then improve a given policy. A second problem in exploration-based learning techniques is that they have to take random actions in order to discover better control policies. This, however, can lead to catastrophic failures which are not acceptable for an autonomous system. The same problem arises in most of the behavior-based approaches since such behaviors do not allow general predictions about their outcome and can not help to guide the exploration.

The hybrid DEDS architecture can be an effective means of addressing these issues. First, reinforcement learning is applied at the level of control activations in a symbolic state space, thus dramatically reducing the size of the search space. In addition, the proposed architecture derives an explicit system model (in the transition probabilities). While not strictly required in reinforcement learning algorithms, this model can be used by the learning algorithm to improve its performance and it appears that transition probabilities transfer well across tasks in some domains. Moreover, the DEDS supervisor can automatically limit exploration to safe control policies and can thus avoid catastrophic failures.

1.2.4 Example: Learning to Walk

We have already demonstrated the use of a control basis with a DEDS supervisor in representative manipulation and walking tasks. This section presents the result of a reinforcement learning task on top of the DEDS supervisor that was designed to rotate our 4-legged walker counterclockwise in its own footprint.

Figure 1 illustrates the control policy acquired in the form of a nondeterministic finite state automaton. It depicts a 5 dimensional predicate subspace in which a walking gait was constructed[113]. This graph shows all the possible transitions that can occur under the learned policy. The core of the policy is the cycle indicated by bold transition arrows. Transition probabilities within this cycle are greater than 95%, making it a stable attractor for this policy. The controllers are given on the right side of the figure for elements of the core cycle. Notice that these controllers are actually combinations of the primitive controllers. The interaction between controllers is managed using the "subject-to" constraint, $\triangleleft$, where $\phi_i \triangleleft \phi_j$ implies that a subordinate controller, ϕ_i, is executed

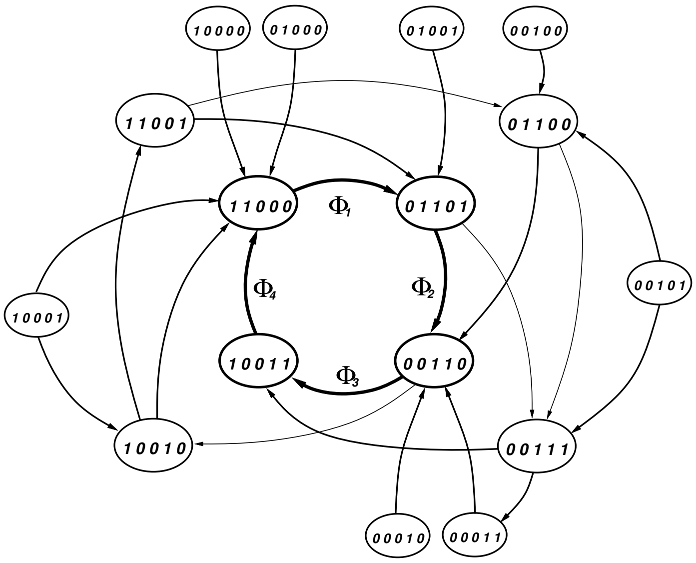

$$\Phi_1 = \phi_{1\,\underline{1}}^{\underline{0,1,2}} \lhd \phi_{2\,\underline{\varphi}}^{\underline{0,1,2,3}} \lhd \phi_{1\,\underline{1}}^{\underline{0,1,3}}$$

$$\Phi_2 = \phi_{2\,\underline{\varphi}}^{\underline{0,1,2,3}} \lhd \phi_{1\,\underline{2}}^{\underline{0,1,2}}$$

$$\Phi_3 = \phi_{2\,\underline{\varphi}}^{\underline{0,1,2,3}} \lhd \phi_{1\,\underline{3}}^{\underline{1,2,3}}$$

$$\Phi_4 = \phi_{1\,\underline{0}}^{\underline{0,2,3}} \lhd \phi_{1\,\underline{0}}^{\underline{0,1,3}}$$

Figure 1: Learned Control Policy for the Rotation Task

subject to constraints imposed by ϕ_j. In practice, this requires that actions produced by ϕ_i do not interfere with ϕ_j. This is similar to requiring that ϕ_i only produce actions within the null space of ϕ_j, although somewhat less restrictive.

The control-based architecture proposed exhibited several useful properties in this application:

1. The policy at predicate states not on the bold control cycle tends to direct the system back into the central control gait. This provides some **robustness** with respect to run-time conditions.
2. The system learned approximate **transition probabilities**, and thus a more precise model which can be transferred to different tasks.
3. Approximately 70% of the transitions included in the DEDS model actually never occurred, thus offering an **efficient** mechanism for focusing attention on the relevant portions of the predicate space.
4. The policy illustrated was acquired **on-line**, in a **single trial**, in about 30 minutes of experience.
5. The range of behavior was constrained to **safe policies**. This constraint permitted the robot to learn the gait without falling down.

Figure 1 really represents a family of control policies that are robust with respect to perturbations in the predicate state. If kinematic limits and irregular terrain, for example, were to cause the walker to fail to assert any of the four predicate states that comprise the core gait, the resulting predicate state is still guaranteed to be safe, and leads to a modified policy that is optimal in these contexts. This family of control policies represents a stable, long-term dynamic relationship between a pool of resources available and the environmental geometry. We can therefore, draw several similarities to the primary circular reactions proposed by Piaget. In Piagetian terms, these types of coordinated behavior are "burned in" by repetition (assimilation) and generalized by a kind of randomized search (accommodation) until an infant finds it possible to prolong an interaction with the world. The policy illustrated in Figure 1 shares these properties — it is acquired by repetitive, stochastic search and the robust behavior so acquired maintains the stability of the walking machine while accumulating a net rotation of the vehicle's heading. What is missing from our model is the Piagetian notion of the a *figurative schema* where declarative knowledge emerges from robust, primary circular reactions. The next section will show how sensorimotor policies can give rise to declarative knowledge, in particular, *categories*.

1.3 Categorization and Abstraction in Sensorimotor Schema

The supervisory policy illustrated in Figure 1 is a procedural knowledge structure describing the characteristic behavior of agent resources engaged in certain tasks in a variety of contexts. It is also a policy for interacting with a variety of possibly unknown plants. The term "plant" refers to the dynamical system formed by a particular combination of controllable resources and environmental constraints.

Figure 2 presents a closed-loop control scheme used to form stable grasps on the basis of instantaneous contact information. The plant consists of a k-contact grasp configuration on an unknown object geometry. Grasping is formulated as a *robust* control problem in which variations in object geometry and contact friction properties are modeled as *disturbances*. The grasp controller $\mathcal{C}^*$ is designed to be robust to these forms of variations in the plant. $\mathcal{C}^*$, is really a combination of two

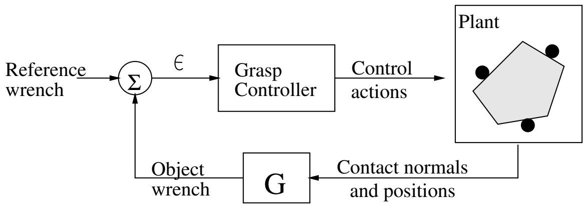

Figure 2: Grasp synthesis as a control problem.

independent controllers $\mathcal{C}_f$ and $\mathcal{C}_m$ (see [104, 103, 106, 105] for the definitions) that yield statically stable contact configurations for 2 and 3 contact grasps on any convex, regular polygon. Moreover, the history of errors in the $\mathcal{C}_{\{}$ and $\mathcal{C}_{\Updownarrow}$ controllers reveals some of the parameters of the plant and can be used to generalize the grasping behavior to other objects. A model of the interaction between $\mathcal{C}_f$ and $\mathcal{C}_m$ controllers called the *preimage* can be constructed by plotting the system trajectory on the (ϵ_f, ϵ_m) error plane as a grasp unfolds. Furthermore, grasp *quality* can be backed-up from equilibrium grasp configurations into precursor states and used to predict the eventual quality of grasp on an unknown object.

For example, consider the coefficient of friction, μ_0, required to stabilize the equilibrium grasp configuration as the grasp quality measure. The preimage can be consulted to determine whether the "worst-case" μ_0 expected exceeds the contact friction specified for the task, μ_{task}. A simple but effective policy for compensation involves superimposing a small, random control action onto the $\mathcal{C}^*$ controller. When the state of the system enters the preimage of a suitable attractor ($\mu_{task} < \mu_0(\epsilon_f, \epsilon_m)$), the random control action is deactivated and $\mathcal{C}^*$ achieves this equilibrium grasp. This procedure is represented schematically in Figure 3.

When multiple candidate objects are possible, a composite preimage is defined:

$$\mu_0^*(\epsilon_f, \epsilon_m) = \max(\mu_0^{obj}(\epsilon_f, \epsilon_m))$$

where $obj \in \{obj_1, obj_2, \ldots, obj_n\}$, which predicts the worst-case coefficient of friction required over all candidate objects. Figure 4 illustrates a typical grasp on a trapezoid for which $\mu_{task} = 0.1$. The left column depicts the composite preimage for the objects in the right column, given the current number of contacts. The initial grasp configuration has two contacts ($\theta = (0, \pi)$). At $t = 0$, the

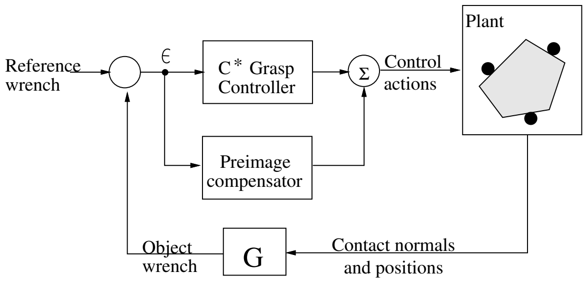

Figure 3: Compensating $\mathcal{C}^*$ using the preimage.

observed (ϵ_f, ϵ_m) state is inaccessible in the preimages of the triangles and so they are eliminated from the list of candidate objects.

The preimage can be used to trigger the addition and removal of additional resources. Grasp configurations that are expected to lead to unsatisfactory grasps for all candidate plants cause resources (additional contacts) to be introduced. Consequently, at $t = 1$, the controller allocates a third contact (shown in white in Figure 4), at $\theta_2 \approx \pi$. If all $k-1$ contact configurations satisfy grasp friction constraints, then one of the contacts can be removed. At $t = 21$, the controller removes contact 0 at $\theta_0 \approx 0$ (marked with a circle in Figure 4), and subsequently eliminates the pentagon, hexagon, leaf, and the bulb object interpretations. The set of candidate objects now consists of the square and the trapezoid; the controller converges at $t = 87$ to a suitable grasp configuration for either of these, requiring a friction coefficient $\mu \leq \mu_{task}$.

This demonstration illustrates that evolving state of a control task can be a powerful perceptual cue regarding the ultimate quality of the motor policy. The preimage delineates regions where the system behavior is dependent on the particular plant thus affecting both recognition and control objectives. The preimage is an abstraction whose cost is amortized over the lifetime of more efficient grasping tasks.

1.4 Learning Physical Schemas

With the foundations of the previous sections we may now describe how a robotic agent may learn physical schemas. Whereas the previous walking example (Section 1.2.4) has been implemented, the following example is hypothetical. Briefly the problem is this: Once an agent has learned extended behaviors, such as learning to walk, or to reach and grasp, it is possible to learn abstract representations of those behaviors called physical schemas. The problem, then, is to go from robust *ground instances* of behavior to abstract schematic representations of behavior.

We can view behaviors as generating multivariate time series of descriptive variables. The task is to identify features in this time series that establish *roles* in the on-going pattern of interaction. For example, Figure 5 is a *semantic* characterization of CONTAINMENT. Any time series of descriptive variables that matches Figure 5 is potentially an instance of CONTAINMENT. This abstract schema may carry additional information, not represented in Figure 5. For example, the probabilities of agents' actions in Figure 5, p and q, are also important: In a containment relationship we expect $p >> (1 - p)$, that is, when Agent1 "escapes," Agent2 re-establishes the containment with high probability. Moreover, we may choose to require that Agent 1 and Agent 2 constitute disjoint sets of resources. Our claim is that the core semantic characterization of a physical schema is a trajectory through the predicate space, as shown in Figure 5.

Several methods could be adapted to learning schemas like CONTAINMENT from time series data. The representation of CONTAINMENT as a state-transition diagram suggests methods for learning Markov models and hidden Markov models. If we view physical schemas as essentially *grammatical*,

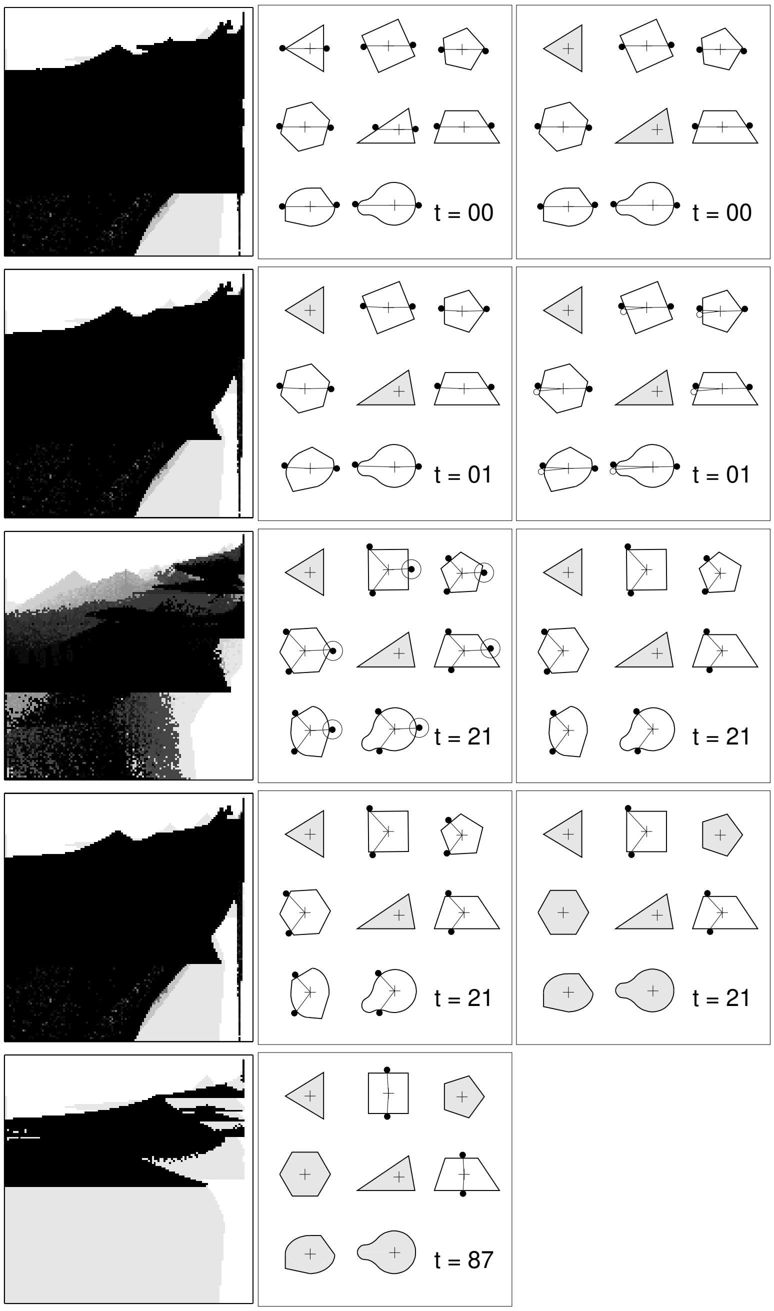

Figure 4: Grasping a trapezoid with two contacts, with $\mu_{task} = 0.1$.

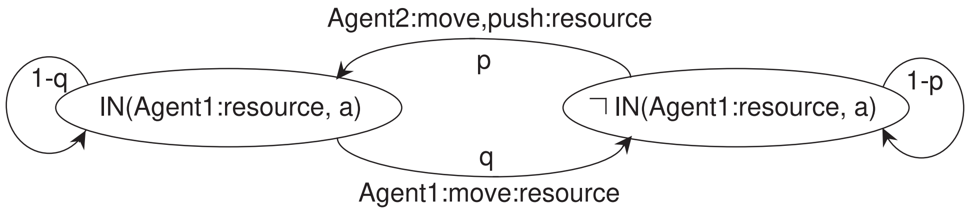

Figure 5: The dynamic relationship between agents' actions and the IN predicate in the CONTAINMENT schema.

that is, structured arrangements of tokens over time, then we might apply methods for learning probabilistic grammars e.g., [35]). In any case, a learning method must be able to lift patterns out of time series of potentially huge state descriptions, where most of the state information is irrelevant.

The Multi-stream Dependency Detection (MSDD) algorithm efficiently finds patterns in time series of large state vectors [81]. For example, we currently use MSDD to find patterns in the behavior of computer networks [80, 65]. Although MSDD will require some modifications to learn physical schemas, it can lift regular patterns, like Figure 5, in time series. We must develop methods to postprocess the dependencies, to gather up and fit together rules that describe pieces of a schema like the CONTAINMENT model in Figure 5. MSDD in its current form will not suffice, but even in its current form, it can find predictive patterns in time series and abstract them away from background noise variables. Thus, if a physical schema is characterized by a predictive pattern that relates agent actions to predicates schema — then MSDD or something like it will find this characterization.

1.5 Experimental Agenda

As a consequence of our commitment to interactionist models of development, we plan to build integrated robot systems to test our hypotheses. Our objective is not to invest too heavily in the details of the techniques presented in previous sections, but instead to prove that interaction, through sensorimotor primitives and policies, leads to conceptual structure. In this section, we will present a tentative experimental agenda for demonstrating our objectives. We will embed a set of robot resources (see Section **??**) in a common work area. For example, a stereo pair of cameras and our integrated hand/arm system will interact with objects on a table top. In all such experimental designs, our goal is to document a developmental sequence from physical primitives, through physical and figurative schema, and toward a conceptual structure with predictive semantics.

Consider the following sequence for acquiring concepts in the *blocks world* domain. Our robot configuration consists of a dextrous hand, a robot arm, an articulated stereo vision system, and a binaural auditory system. The robot will interact with the blocks world environment in order to build domain rules for this class of tasks. This expertise will necessarily proceed from coordinated sensorimotor policies (spatial sense), to models describing the relationship between the robot and environmental features (categorization), and finally to declarative models of the interactions between extrinsic features (*stability*).

Spatial Sense Two of a child's earliest accomplishments are auditory localization and coordinated binocular vision. Somewhat later (approximately 26 weeks), an infant begins to reach to visual or auditory cues. All three of these sensory modalities contribute to the infant's emerging sense of space. Consider the acquisition of a policy for coordinated control of an articulated stereo vision system. Treated independently, the correspondence problem between two visual images is under constrained and leads, typically, to a number of plausible three dimensional reconstructions. But, if corroborating evidence is found in other sensory data, then the inherent ambiguity of stereo matching can be reduced. There seems strong evidence that these modalities interact to a significant degree very early in development. If this is true, then our very earliest sensorimotor concepts may be found in the cross correlations between sensing modalities. **Our framework for engaging sensory and motor resources will be extended in scope to represent multimodal sensory alternatives. We will use the resulting behavioral model to learn how to coordinate multiple spatial observers. The patterns of co-occuring spatial features provide the robot with its first clues about "objects" in the world.**

Categorization Our approach to categorization will be guided by the intuition afforded by examples like that presented in Section 1.3, that is, important contexts can be encoded in the details of the trajectory observed during interactions between the robot and the world. However, that example was hand-crafted — we defined the space in which the preimage was constructed. The robot must be capable of doing this for itself. **We will generalize the notion of the haptic preimage (introduced in Section 1.3) to incorporate visuomotor information so that correlated haptic and visual signatures of "objects" can be constructed.** These multimodal abstractions are the precursor of sensory and motor predictions. A visual feature space can be used to index into a haptic model and vice versa. This marks the beginning of predictive semantics as the agent learns to predict what the object "feels" like on the basis of visual observation. With auditory feedback as well, we may consider rich preimages of behavior.

STABILITY In an experiment first proposed by Whitehead and Ballard[8], humans were asked to copy a planar arrangement of blocks. The experiment exposed some of the character of the motor schemas underlying the task, specifically, it showed how the visual attention of the subject varied during the copying task. This comes close to identifying a sequence of visual operators (visual resources) preferred by the human subjects. Another task proposed by Brand [32] produced a sequence of visual operations necessary to establish the stability of a three dimensional stack of blocks. While they expressed their stability relations in the form of production rules, we offer a plausible mechanism for acquiring stability schema during interaction with the world, and to subsequently use this behavioral knowledge to both establish the conditions of stability and to explain why a particular stack of blocks is stable.

Consider a hybrid experimental platform in which the robot is to copy three dimensional arrangements of blocks. The goal is for the robot to learn to assemble a stack of blocks identical or similar to a reference stack. However, rather than a set of production rules regarding stability (as in [32], **we propose that the robot interact with blocks on a table in order to learn the relationships underlying stable grasps and to generalize them to stable interactions *between* blocks. This form of metaphorical extension signals the transformation from procedural to declarative representation**. And it is not a such a great leap — all that is required is that the robot "notice" the block in question is stable regardless of whether it is due to a contact configuration generated by fingers or by other blocks.

For example, if asked to copy the reference stack, a sequence of sensorimotor behavior would ensue to construct a physical copy of the stack. This much is a direct extrapolation of the examples in Sections 1.2.4 and 1.3. However, when asked to explain why a particular block in the stack is stable, the same policy that computed grasps in one context may now be used to direct a visual search in order to identify interactions between blocks that serve to stabilize it. If queried about the role of a block in a stable stack, the robot could reply by engaging a sequence of visual operators to determine what interactions this block has with its neighbors and how these interactions serve to stabilize the stack. With figurative schema such as STABILITY in place, our developing robot is capable of predicting and explaining. Experimental data regarding the formation of figurative schema such as this would contribute significantly to the on-going debates regarding sensorimotor development and the origins of cognitive structure.

2 References Cited

References

[1] J. S. Albus. Mechanisms of planning and problem solving in the brain. *Mathematical Biosciences*, 45:247–293, 1979.

[2] G. Anderson and C. R. Beal. Children's recognition of inconsistencies in science texts: Multiple measures of comprehension monitoring. *Applied Cognitive Psychology*, 9:261–272, 1995.

[3] Scott D. Anderson. *A Simulation Substrate for Real-Time Planning*. PhD thesis, University of Massachusetts at Amherst, February 1995. Also available as Computer Science Department Technical Report 95–80.

[4] Scott D. Anderson, Adam Carlson, David L. Westbrook, David M. Hart, and Paul R. Cohen. Tools for experiments in planning. In *Proceedings of the Sixth International Conference on Tools with Artificial Intelligence*, pages 615–623, Los Alamitos, CA, 1994. IEEE, IEEE Computer Society Press.

[5] Araujo, E., Grupen, R., "Learning Control Composition in a Complex Environment," submitted to *from Animals to Animats*, 1996.

[6] R. Baillargeon. How do infants learn about the physical world? *Current Directions in Psychological Science*, 3:133–140, 1994.

[7] R. Baillargeon, E. Spelke, and S. Wasserman. Object permanence in five-month-old-infants. *Cognition*, 20:191–208, 1985.

[8] D.H. Ballard. Reference frames for animate vision. In *IJCAI89*. IJCAI, 1989.

[9] Avron B. Barr, Paul R. Cohen and Edward A. Feigenbaum. 1989. *The Handbook of Artificial Intelligence, Volume IV*. Reading, Mass: Addison Wesley.

[10] A. G. Barto, S. J. Bradtke, and S. P. Singh. Learning to act using real-time dynamic programming. *Artificial Intelligence*, 72:81–138, 1995.

[11] A.G. Barto, R.S. Sutton and C. Watkins, "Learning and sequential decision making," in *Learning and Computational Neuroscience*, M. Gabriel and J.W. Moore (Eds.), Cambridge, MA: MIT Press, 1990, pp. 539–602.

[12] A.G. Barto, "Connectionist learning for control: An overview," in *Neural Networks for Control*, T. Miller, R.S. Sutton, and P.J. Werbos (Eds.), Cambridge, MA: MIT Press, 1990, pp. 5–58.

[13] A.G. Barto. Reinforcement learning and adaptive critic methods. In D. A. White and D. A. Sofge, editors, *Handbook of Intelligent Control: Neural, Fuzzy, and Adaptive Approaches*, pages 469–491. Van Nostrand Reinhold, New York, 1992.

[14] A. G. Barto. "Reinforcement learning." In *Handbook of Brain Theory and Neural Networks,* M.A. Arbib (Ed.), Cambridge: MIT Press, 1995.

[15] J. Bates. The role of emotion in believable agents. *Communications of the ACM*, 1994. In Press.

[16] C. R. Beal. Development of knowledge about the role of inference in text comprehension. *Child Development,*, 61:1011–1023, 1990.

[17] Beal, C. R. (1990). The development of text evaluation and revision skills. *Child Development*, 61:247-258.

[18] C. R. Beal. The role of comprehension monitoring in children's revision. *Educational Psychology Review*, 1996. In Press.

[19] Beal, C. R., Garrod, A. C., & Bonitatibus, G. J. (1990). Fostering children's revision skills through training in comprehension monitoring. *Journal of Educational Psychology*, 82:275-280.

[20] C. R. Beal and S. L. Belgrad. The development of message evaluation skills in young children. *Child Development*, 61:705–712, 1990.

[21] Beal, C. R., & Bonitatibus, G. (1994, July). Theory of mind in middle childhood: Children's appreciation of text interpretation. Paper presented at the Annual meeting of the American Psychological Society, Washington DC.

[22] Beal, C. R., Schmitt, K., & Dekle, D. J. (1995). Eyewitness identification of children: Effects of absolute judgments, nonverbal response options, and event encoding. *Law and Human Behavior*, 19:197-216.

[23] H. Beilin. Piaget's enduring contribution to developmental psychology. *Developmental Psychology*, 28:191–204, 1992.

[24] B.I. Bertenthal, J. J. Campos, and R. Kermoian. An epigentic perspective on the development of self produced locomotion and its consequences. *Current Directions in Psychological Science*, 3:140–145, 1994.

[25] Neil Berthier. Learning to reach: A mathematical model. *Developmental Psychology*, 1995. In Press.

[26] Berthier, N.E. (in press). Analysis of reaching for stationary and moving objects in the human infant. In J. Donohoe (Ed.) *Neural Network Models of Complex Behavior—Biobehavioral Foundations.*

[27] Berthier, N.E., Singh, S.P., Barto, A.G. & Houk, J.C. (1992). A cortico-cerebellar model that learns to generate distributed motor commands to control a kinematic arm. In: *Advances in Neural Information Processing Systems 4*, J.E. Moody, S.J. Hanson, R.P. Lippmann (Eds.) San Mateo, CA: Morgan Kaufmann, 611-618.

[28] Berthier, N.E., Clifton, R.K., Gullapalli, V., McCall, D. & Robin, D. (in press). Visual information and the control of reaching. *Journal of Motor Behavior.*

[29] N.E. Berthier, S.P. Singh, A.G. Barto, and J.C. Houk. Distributed representations of limb motor programs in arrays of adjustable pattern generators. *Journal of Cognitive Neuroscience*, 5:56–78, 1993.

[30] J. N. Bohannon and L. Stanowicz. The issue of negative evidence: Adult responses to children's language errors. *Developmental Psychology*, 24:684–689, 1988.

[31] G. J. Bonitatibus and C. R. Beal. Finding new meanings: Children's recognition of interpretive ambiguity in text. *Journal ofExperimental Child Psychology*, 1996. In Press.

[32] M. Brand. *Making Sense of the World with Causal Analysis.* PhD thesis, Institute for the Learning Sciences, Northwestern University, 1994.

[33] S. Carey and R Gelman. *The epigenesis of mind: Essays on biology and cognition.* Hillsdale NJ: Erlbaum, 1991.

[34] David Chapman. *Vision, Instruction and Action.* MIT Press, 1991.

[35] Eugene Charniak. *Statistical Language Learning.* MIT Press, 1993.

[36] Y. Chauvin and D. E. Rumelhart. *Backpropagation: Theory, Architectures, and Applications.* Erlbaum, Hillsdale, NJ, 1995.

[37] Clifton, R.K., Rochat, P., Robin, D.J. & Berthier, N.E. (1994). Multimodal perception in the control of infant reaching. *J Exp Psychology: HPP*, *20*, 876-886.

[38] J. Coelho and R. Grupen. Online grasp synthesis. In *Proceedings of the Conference on Robotics and Automation*, Minneapolis, MN, April 1996. IEEE.

[39] Coelho, J., Grupen, R., "A Control Basis for Multifingered Grasps," submitted to the IEEE Journal of Robotic Systems, August, 1995.

[40] Paul Cohen. *Heuristic Reasoning About Uncertainty*. London:Pitman Publishing, 1983.

[41] Paul Cohen, editor. *Acquiring Conceptual Structure from Sensorimotor Interaction*. Seminar readings, unpublished collection, 1996.

[42] Paul R. Cohen. *Empirical Methods for Artificial Intelligence*. MIT Press, 1995.

[43] Paul R. Cohen, Scott D. Anderson, and David L. Westbrook. Simulation for ARPI and the air campaign simulator. In A. Tate, editor, *Advanced Planning Technology*, Menlo Park, CA, 1996. AAAI Press.

[44] Paul R. Cohen, Michael L. Greenberg, David M. Hart, and Adele E. Howe. Trial by fire: Understanding the design requirements for agents in complex environments. *AI Magazine*, 10(3):32–48, Fall 1989.

[45] Paul R. Cohen, Tim Oates, and Marc S. Atkin. Preliminary evidence that conceptual structure can be learned by interacting with an environment. 1996.

[46] Paul R. Cohen, Tim Oates, Marc S. Atkin, and Carole R. Beal. Building a baby. To appear in Proceedings of the Cognitive Science Conference, 1996.

[47] Thomas Dean and Michael Wellman. *Planning and Control*. Morgan Kaufmann, 1991.

[48] Gary L. Drescher. *Made-Up Minds*. MIT Press, 1991.

[49] B. Falkenhainer, K. D. Forbus, and D. Gentner. The structure mapping engine: Algorithm and examples. *Artificial Intelligence*, 41:1–63, 1989.

[50] K. W. Fischer and T. Bidell. Constraining nativist inferences about cognitive capacities. In S. Carey and R. Gelman, editors, *The epigensis of mind*. Hillsdale NJ: Erlbaum, 1991.

[51] D. Gentner and A. B. Markman. Structural alignment in comparison: No difference without similarity. *Psychological Science*, 5:152–158, 1994.

[52] G. J. Gibbs and H. L. Colston. The cognitive psychological reality of image schemas and their transforms. *Cognitive Linguistics*, 6–4:347–378, 1995.

[53] E. J. Gibson. Detection of the traversability of surfaces by crawling and walking infants. *JEXP: Human Perception and Performance*, 13:533–544, 1987.

[54] E. J. Gibson. Has psychology a future? *Psychological Science*, 5:69–76, 1994.

[55] Eleanor J. Gibson. *An Odyssey in Learning and Perception*. MIT Press, 1990.

[56] C.E. Granrud. Visual perception and cognition in infancy. *Carnegie Mellon Symposia on Cognition*, 23, 1993.

[57] R. Grupen and R. Weiss. Integrated control for interpreting and manipulating the environment. *Robotica*, 12(2):165–175, March 1994.

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

[59] Lawrence A. Hirschfeld and Susan A.Gelman, editors. Cambridge University Press, 1994.

[60] D. Holmes. The korean scenario for air camaign planning. Technical report, ISX Corporation, 1995.

[61] Adele E. Howe and Paul R. Cohen. Understanding Planner Behavior. *Artificial Intelligence Journal*, Special Issue on Planning Systems, Vol. 76, Nos. 1 & 2, pp. 125-166, 1995.

[62] Huber, M. and Grupen, R., "2D Contact Detection and Localization Using Proprioceptive Information." *IEEE Transactions on Robotics and Automation*, 10(1):23-33, 1994.

[63] M. Huber, W. MacDonald, and R. Grupen. A control basis for multilegged walking. In *Proceedings of the Conference on Robotics and Automation*, Minneapolis, MN, April 1996. IEEE.

[64] Ray Jackendoff. *Semantic Structures.* MIT Press, 1993.

[65] D. Jensen, T. Oates, and P. Cohen. Fault identification, requirements and design. Technical report, University of Massachusetts,Amherst, Department of Computer Science, 1996.

[66] Mark Johnson. *The Body in the Mind.* University of Chicago Press, 1987.

[67] George Lakoff. *Women, Fire, and Dangerous Things.* University of Chicago Press, 1984.

[68] George Lakoff and Mark Johnson. *Metaphors We Live By.* University of Chicago Press, 1980.

[69] Ronald Langacker. *Foundations of Cognitive Grammar, Vol. 1.* Stanford University Press, 1987.

[70] Alan M. Leslie. The necessity of illusion: perception and thought in infancy. In L. Weiskrantz, editor, *Thought Without Language.* Oxford University Press (Clarendon), 1988.

[71] M. Lewis, S. M. Alessandri, and M. W. Sullivan. Violation of expectancy, loss of control, and anger expressions in young infants. *Developmental Psychology*, 26:745–751, 1990.

[72] R. Maclin and J. W. Shavlik. Incorporating advice into agents that learn from reinforcements. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, 1994.

[73] Jean M. Mandler. How to build a baby: On the development of an accessible representational system. *Cognitive Development*, 3:113–136, 1988.

[74] Jean M. Mandler. How to build a baby: Ii. conceptual primitives. *Psychological Review*, 99(4):587–604, 1992.

[75] G. F. Marcus. Negative evidence in language acquisition. *Cognition*, 46:53–85, 1993.

[76] A. N. Meltzoff. The development and neural bases of higher cognitive functions. In A. Diamond, editor, *Annals of the New York Academy of Sciences: Vol. 608.* New York: New York Academy of Sciences., 1990.

[77] Moore, J.W., Desmond, J.E. & Berthier, N.E. (1989). Adaptively timed conditioned responses and the cerebellum: A neural network approach. *Biological Cybernetics*, *62*, 17-28.

[78] Moore, J.W., Berthier, N.E. & Blazis, D.E.B. (1990). Classsical eye-blink conditining: Brain systems and implementation of a computational model. In *Learning and Computational Neuroscience: Foundations of Adaptive Networks.* M. Gabriel and J. Moore (Eds.), Cambridge, Mass.: MIT Press.

[79] K. Nelson. Structure and strategy in learning to talk. *Monographs of the Society for Research in Child Development*, 38((1-2,Serial No. 149)), 1973.

[80] T. Oates, D. Jensen, and P. Cohen. Technologies for fault management. Technical report, University of Massachusetts, Amherst, Department of Computer Science, 1995.

[81] Tim Oates and Paul R. Cohen. Searching for structure in multiple streams. In *Proceedings of the Thirteenth International Conference on Machine Learning*, 1996. To appear.

[82] Tim Oates, Matthew D. Schmill, and Paul R. Cohen. Parallel and distributed search for structure in multivariate time series. Submitted to *IEEE Transactions on Knowledge and Data Engineering, Special Issue on Database Mining*, Year=1996.

[83] Jean Piaget. *The Origins of Intelligence in Childhood.* International Universities Press, 1952.

[84] Jean Piaget. *The Construction of Reality in the Child.* Basic Books, 1954.

[85] George Polya. *Patterns of Plausible Inference.* New Jersey: Princeton University Press, 1968.

[86] P.K. Pook and D.H. Ballard. Recognizing teleoperated manipulations. In *Proceedings of IEEE Intl. Conf. on Robotics and Automation.* IEEE, 1993.

[87] E. Riloff and W.G. Lehnert. Automated dictionary construction for information extraction from text. In *Proceedings of the Ninth IEEE Conference on AI for Applications*, pages 93–99, 1994.

[88] Ricci, C. M., & Beal, C. R. (1996, April). The influence of prior knowledge on children's eyewitness memory, identifications, and suggestibility. Paper presented at the biennial Conference on Human Development, Birmingham AL.

[89] Ricci, C. M., Beal, C. R., & Dekle, D. J. (in press). The effect of parent versus unfamiliar interviewers on young witnesses' memory and identification accuracy. *Law and Human Behavior.*

[90] Robin, D., Berthier, N., & Clifton, R.K. (in press). Infants' predictive reaching for moving objects in the dark. *Developmental Psychology.*

[91] E. Rosch and C. B. Mervis. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7:573–605, 1975.

[92] Roger Schank. *Conceptual Information Processing.* North Holland: Amsterdam, 1975.

[93] K. Souccar and R. Grupen. Distributed motion control for multiple robotic manipulators. In *Proceedings of the Conference on Robotics and Automation*, Minneapolis, MN, April 1996. IEEE.

[94] E. S. Spelke, K. Breinlinger, J. Macomber, and K. Jacobson. Origins of knowledge.. *Psychological Review*, 99:605–632, 1992.

[95] Elizabeth S. Spelke. The origins of physical knowledge. In L. Weiskrantz, editor, *Thought Without Language.* Oxford University Press (Clarendon), 1988.

[96] E. Thelen, D. Corbetta, K. Kamm, J. Spencer, K. Schneider, and R. Zernicke. The transition to reaching: Mapping intention and intrinsic dynamics. *Unpublished*, 1992.

[97] E. Thelen and L. Smith. *A Dynamic Systems Approach to the Development of Cognition and Action.* MIT Press, Cambridge, MA, 1994.

[98] K. Wynn. Evidence against empiricist accounts of the origins of numerical knowledge. *Mind and Language*, 7:315–332, 1992.

[99] W. Zhang and T. G. Dietterich. A reinforcement learning approach to job-shop scheduling. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1114–1120, 1995.

[100] A.E. et. al. Aronson. *Clinical Examinations in Neurology.* W.B. Saunders Co., Philadelphia, PA, 1981.

[101] A. G. Barto, S. J. Bradtke, and S. P. Singh. Learning to act using real-time dynamic programming. Technical Report 93-02, University of Massachusetts, Amherst, MA, 1993.

[102] J.S. Bruner. Organization of early skilled action. *Child Development*, 44:1–11, 1973.

[103] J. Coelho and R. Grupen. Constructing effective multifingered grasp controllers. In *Proceedings of the 1994 Conference on Robotics and Automation*, San Diego, CA, May 1994. IEEE.

[104] J. Coelho and R. Grupen. Optimal multifingered grasp synthesis. In *Proceedings of the 1994 Conference on Robotics and Automation*, San Diego, CA, May 1994. IEEE.

[105] J. Coelho and R. Grupen. Online grasp synthesis. In *Proceedings of the Conference on Robotics and Automation*, Minneapolis, MN, April 1996. IEEE.

[106] J. Coelho and R.A. Grupen. Effective multifingered grasp synthesis. Technical Report CS Technical Report 94-12, CS Department, University of Massachusetts, February 1994.

[107] Jefferson A. Coelho Jr.. and Roderic A. Grupen. Effective multifingered grasp synthesis. Technical Report 94-12, CMPSCI Dept., Univ. of Mass., Amherst, February 1994.

[108] Christopher I. Connolly. Harmonic functions and collision probabilities. In *Proc. Int. Conf. Robotics Automat.*, San Diego, CA, May 1994. IEEE.

[109] Christopher I. Connolly and Roderic A. Grupen. The applications of harmonic functions to robotics. *J. Robotic Sys.*, 10(7):931–946, October 1993.

[110] R. Grupen and R. Weiss. Integrated control for interpreting and manipulating the environment. *Robotica*, 12(2):165–175, March 1994.

[111] Roderic A. Grupen, Manfred Huber, Jefferson A. Coelho Jr., and Kamal Souccar. Distributed control representation for manipulation tasks. *IEEE Expert*, 10(2):9–14, April 1995.

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

[113] M. Huber, W. MacDonald, and R. Grupen. A control basis for multilegged walking. In *Proceedings of the Conference on Robotics and Automation*, Minneapolis, MN, April 1996. IEEE.

[114] Manfred Huber, Willard S. MacDonald, and Roderic A. Grupen. A control basis for multilegged walking. In *Proc. Int. Conf. Robotics Automat.*, pages Vol.4 2988–2993, Minneapolis, MN, April 1996. IEEE.

[115] P. Maes and R. Brooks. Learning to coordinate behaviors. In *Proceedings of the 1990 AAAI Conference on Artificial Intelligence*. AAAI, 1990.

[116] S. Mahadevan and J. Connell. Automatic programming of behavior-based robots using reinforcement learning. *Artificial Intelligence*, 55:311–365, 1992.

[117] S. Singh, C. Connolly, R. Grupen, and A. Barto. Robust reinforcement learning in motion planning. In *Advances in Neural Information Processing Systems 6 (NIPS)*, 1994.

[118] Richard S. Sutton. First results with Dyna, an integrated architecture for learning, planning and reacting. In W. Thomas Miller III, Richard S. Sutton, and Paul J. Werbos, editors, *Neural Networks for Control*, pages 179–189. MIT Press, 1990.

[119] E. Thelen, D. Corbetta, K. Kamm, J. Spencer, K. Schneider, and R. Zernicke. The transition to reaching: Mapping intention and intrinsic dynamics. *Unpublished*, 1992.

[120] E. Thelen and L. Smith. *A Dynamic Systems Approach to the Development of Cognition and Action.* MIT Press, Cambridge, MA, 1994.

[121] C. J. C. H. Watkins. *Learning from Delayed Rewards.* PhD thesis, Cambridge University, Cambridge, England, 1989.

[122] C.J. Watkins and P. Dayan. Technical note: Q-learning. *Machine Learning*, 8:279–292, 1992.

3 Biographical Sketches

3.1 Roderic A. Grupen, Associate Professor

Department of Computer Science
University of Massachusetts
Amherst, MA

EDUCATION

Ph.D. Computer Science, University of Utah	1988
M.S. Mechanical Engineering, Pennsylvania State University	1984
B.A. Physics, Franklin and Marshall College	1980
B.S. Mechanical Engineering, Washington University	1980

PROFESSIONAL EXPERIENCE

1995 - present: Associate Professor, Department of Computer Science, University of Massachusetts, Amherst, Massachusetts.

1988 - 95 Assistant Professor, Department of Computer Science, University of Massachusetts, Amherst, Massachusetts.

1988 - present Co-Director of the Laboratory for Perceptual Robotics, Department of Computer Science, University of Massachusetts, Amherst, Massachusetts.

1984 - 88: Research and Teaching Assistant, Departments of Computer Science and Mechanical Engineering, University of Utah, Salt Lake City, Utah.

1981 - 82: Systems Design Engineer, General Electric Company Edison Engineering Program, Louisville, Kentucky.

1980: Washington University Technical Associates, St. Louis, Missouri.

RESEARCH INTERESTS

My research focus is the integration of Artificial Intelligence and Control Theory for autonomous robot control. I have been investigating the use of sensor-rich component automation as a means of enhancing the flexibility of robots applied to a diverse range of tasks. A central focus of this work is computationally tractible controllers for complex, highly redundant systems. The approach entails on-line composition of a control basis for manipulation tasks. Constraints on the expression and composition of control have been established which preserve stability and convergence characteristics in the integrated system. We have applied this paradigm to a 20 DOF hand/arm system, to a homebuilt legged platform and to a non-holonomic mobile platform as well. Also underway is a research program to characterize and exploit the underlying predictability of these control architectures to facilitate process scheduling. This work is focussed on evaluating process control alternatives with respect to quality, timeliness, and throughput in real-time, flexible automation. In the future, this work will be applied to autonomous robot control, to “agile” low volume manufacturing systems, and to control for intelligent vehicles and highway systems.

SELECTED PUBLICATIONS

1. Huber, M. and Grupen, R., "2D Contact Detection and Localization Using Proprioceptive Information." *IEEE Transactions on Robotics and Automation*, 10(1):23-33, 1994.
2. Connolly, C., and Grupen, R., "On the Applications of Harmonic Functions to Robotics," In *Journal of Robotic Systems*, 10(7):931-946, October, 1993.
3. Grupen, R., Huber, M., Coelho, J,. and Souccar, K., In "Distributed Control Representation for Manipulation Tasks," IEEE Expert, Special Issue on Intelligent Robotic Systems, Ed. R. Arkin, 10(2):9-14, 1994.
4. Coelho, J. and Grupen, R., "Optimal Multifingered Grasp Synthesis," the IEEE Conference on Robotics and Automation, San Diego, CA, May, 1994.
5. Grupen, R. and Souccar, K., "Manipulation-Based Spatial Isotropy: A Kinematic Reflex," In MCPA'93 - Workshop on Mechatronical Computer Systems for Perception and Action, Halmstad, Sweden, June 1-3, 1993.
6. Gullapalli, V., Barto, A., and Grupen, R., "Learning Admittance Mappings for Force-Guided Assembly," the IEEE Conference on Robotics and Automation, San Diego, CA, May, 1994.
7. Singh, S., Connolly, C., Grupen, R., and Barto, A., "Robust Reinforcement Learning in Motion Planning," Neural Information Processing Systems (NIPS), 1993.
8. Connolly CI, Grupen RA, "Nonholonomic Path Planning Using Harmonic Functions," the 1994 IEEE Conference on Intelligent Robots and Systems (IROS '94), Munich, Germany.
9. Coelho, J., Grupen, R., "A Control Basis for Multifingered Grasps," submitted to the Journal of Robotic Systems, August, 1995.
10. Stan, M., Burleson, W., Connolly, C., and Grupen, R., "Analog VLSI for Robot Path Planning," accepted by *Journal of VLSI Signal Processing*, Special issue on Analog Computation, 1993.

COLLEAGUES/RECENT COLLABORATORS

M.S. Students		**Ph.D. Students**		**Professional Associations**
David Oskard	1990	Christopher Connolly	1994	Ronald Arkin
King Shaw	1990	Gordon Dakin	1994	Thomas C. Henderson
Aditi Dubey	1991	Brian Pinette	1994	Edward M. Riseman
		Inigo Thomas	1993	Robin J. Popplestone
		Satinder Singh	1993	Richard Weiss
		Vijaykumar Gullapalli	1992	Wayne Burleson
		Rakesh Kumar	1992	Andrew Barto
		Jonathon Bachrach	1991	Richard Sutton
		Robert Jacobs	1990	
		Yanxi Liu	1990	

3.2 Paul R. Cohen, Professor of Computer Science

EDUCATION

Ph.D. Computer Science and Psychology, Stanford University	1983
M.A. Psychology, Univ. of California, Los Angeles	1978
B.A. Psychology, Univ. of California, San Diego	1977

PROFESSIONAL EXPERIENCE

1995 - present: Professor, Computer Science Department, University of Massachusetts, Amherst, MA.

1993 - 1996: Director, Computer Science Graduate Program, University of Massachusetts, Amherst, MA.

1989 - 1995: Associate Professor, Computer Science Department, University of Massachustts, Amherst, MA.

1989 - 1993: Visiting Associate Professor of Engineering, Thayer School of Engineering, Darmouth College, Hanover, NH.

1983 - 1989: Assistant Professor, Computer Science Department, University of Massachusetts, Amherst, MA.

1983 - present: Director, Experimental Knowledge Systems Laboratory, Computer Science Department, University of Massachusetts, Amherst, MA.

RESEARCH INTERESTS

Professor Cohen's primary research interests over the last five years have focused on empirical methods, the development of conceptual structure from sensorimotor interaction, campaign planning, simulation, and machine learning.

SELECTED PUBLICATIONS

1. Dawn E. Gregory and Paul R. Cohen. January 1997. Intelligent Assistance for Computational Scientists: Integrated Modelling, Experimentation, and Analysis. To appear in *Preliminary Papers of the Sixth International Workshop on Artificial Intelligence and Statistics.* Ft. Lauderdale.

2. Paul R. Cohen, Marc S. Atkin, Tim Oates, and Carole R. Beal. 1996. Neo: Learning Conceptual Knowledge by Sensorimotor Interaction with an Environment. To appear in *Proceedings of the First International Conference on Autonomous Agents.*

3. Marc S. Atkin, and Paul R. Cohen. 1996. Monitoring Strategies for Embedded Agents. In *Journal of Adaptive Behavior*, Volume 4, Issue 2.

4. Paul R. Cohen, Tim Oates and Marc S. Atkin. 1996. Preliminary Evidence that Conceptual Structure Can Be Learned by Interacting with an Environment. To appear in *Working Notes of the AAAI Fall Symposium on Embodied Cognition and Action.*

5. Paul R. Cohen, Tim Oates, Marc S. Atkin, and Carole R. Beal. 1996. Building a Baby. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*, L. Erlbaum, pp. 518-522.
6. Tim Oates and Paul R. Cohen. 1996. Searching for Planning Operators with Context-Dependent and Probabilistic Effects. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence.* AAAI Press. pp. 863-868.
7. Tim Oates and Paul R. Cohen. 1996. Searching for structure in multiple streams. In *Proceedings of the Thirteenth International Conference on Machine Learning*, Morgan Kaufman, Inc., pp. 346-354.
8. Robert St. Amant and Paul R. Cohen. 1996. A Planner for Exploratory Data Analysis. In *Proceedings of the Third International Conference on Artificial Intelligence Planning Systems (AIPS-96).* AAAI Press. pp. 205-212.
9. Paul R. Cohen. *Empirical Methods for Artificial Intelligence.* MIT Press, 1995.
10. Paul R. Cohen, Marc S. Atkin, Tim Oates, and Dawn E. Gregory. 1995. A Representation and Learning Mechanisms for Mental States. In *Working Notes of the Symposium on Representing Mental States and Mechanisms*, AAAI-95 Spring Symposium Series, pp. 15-21.

COLLEAGUES/RECENT COLLABORATORS

Ph.D. Students	**Professional Associations**
Scott Anderson	Andrew Barto
Adele Howe	Neil Berthier
Robert St. Amant	James Kurose
	Wendy Lehnert
	Carole Beal
	Gary Marcus
	Richard Sutton
	Beverly Woolf
	Nathaniel Martin