A Factorization Approach to Manipulation in Unstructured Environments

Dov Katz and Oliver Brock

Abstract We propose factorization as a concept to analyze and solve manipulation problems in unstructured environments. A factorization is a decomposition of the original problem into factors (sub-problems), each of which can be solved much more easily than the original problem. The appropriate composition of these factors results in a robust and efficient solution to the original problem. Our assumption is that manipulation problems live in lower-dimensional subspaces of the high-dimensional state space associated with unstructured environments. A factorization identifies these subspaces and therefore permits finding simple and robust solutions to the factors. In this paper, we examine the effects of factorization in the context of our recent work on manipulating articulated objects in unstructured environments.

1 Introduction

Mobile manipulation in unstructured environments[1] remains an important challenge in robotics. Even after several decades of research, our ability to endow robotic systems with general manipulation skills remains limited. What is the key to making tangible progress in this domain?

In this paper, we hypothesize that fundamental progress in autonomous manipulation can only be achieved through an understanding of how to adequately compose simple perception, control, planning, and learning skills so that they incrementally realize increasingly complex manipulation behavior.

Dov Katz and Oliver Brock
Robotics and Biology Laboratory, School of Electrical Engineering and Computer Science, Technische Universität Berlin, Germany

[1] When using the term "unstructured" we refer to environments that have not been modified to accommodate limitations of the robot. We consider providing *a priori* models to the robot as a way to accommodate these limitations.

This hypothesis may seem obvious. We believe, however, that it contrasts with some common believes and practices applied in much of the current research in robotics. Should the hypothesis prove to be correct, there would be important implications for how research in autonomous mobile manipulation should be conducted.

Within the manipulation community, some researchers argue that a break-through in manipulation will be triggered by better sensing technologies. Such advances, so the assumption, will lead to the availability of highly accurate models, even in unstructured environments. The problem of acquiring those models is therefore deferred to the sensor community and research proceeds under the assumption that accurate models are available. This view is commonly taken, for example, in motion planning and grasp planning. In contrast, we believe that the perceptual problems associated with obtaining adequate models for task execution will remain very challenging, irrespective of advances in sensor technology.

Another popular view is that the challenges of manipulation in unstructured environments may be addressed using ever-increasing computational power. Considering the recent successes of sampling-based motion planning and POMDP-based approaches to planning, it seems credible that soon we will be able to plan in complex environments while taking sensing and actuation uncertainties into account. In contrast, we believe that the combinatorial explosion associated with problems in unstructured environments will leave general planning for real-world environments out of our computational reach for some time to come.

Why then do we claim that advances in manipulation can be achieved through a suitable composition of perception, control, planning, mechanisms, etc.?

At a high level, our argument is about appropriate decompositions of high-dimensional state spaces. The goal of decomposition is to find sub-problems that can be solved easily and whose composition solves the original, more difficult problem. We refer to such a decomposition as a *factorization*, emphasizing that the decomposition leads to simpler components (factors) that, when combined (multiplied), solve the original problem (equal the product). As an example consider the expression $a^2 - 2ab + b^2$, which can be decomposed as $a^2(1 - \frac{2b}{a} + \frac{b^2}{a^2})$ or as $(a-b)(a-b)$. Clearly, both expressions are equivalent—they compute the same number (or achieve the same functionality)—but the latter one is much simpler and requires less computation. We only refer to the latter decomposition as a factorization.

Solving complex problems by decomposition is hardly a new idea. In fact, the fragmentation of robotics into sub-fields such as vision, control, planning, grasping, and manipulation represents a particular decomposition of the "robotics" problem. However, we believe that factorizations, i.e. "good" decompositions, do not naturally coincide with the boundaries imposed by the traditional academic sub-fields. Instead, we hypothesize that a factorization typically exploits synergies that arise when these very boundaries are crossed.

Factorization will enable progress in manipulation for two reasons. First, they lead to simple, efficient, and robust solution to manipulation problems, because factorizations identify the low-dimensional subspace of the high-dimensional state space within which the solution to the problem lies. Second, factorizations enable

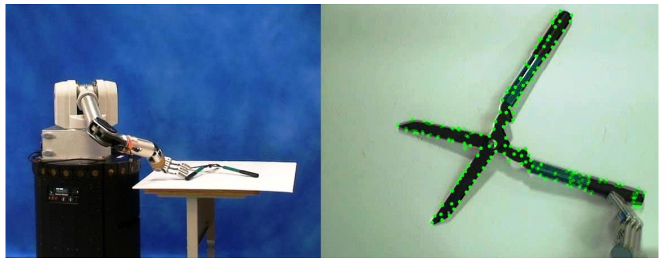

Fig. 1 UMan (UMass Mobile Manipulator) performs a manipulation task without prior knowledge about the manipulated object. The right image shows the scene as seen by the robot through an overhead camera; dots mark tracked visual features.

the *incremental* development of increasingly complex skills. Once the "right" factor has been split off, the remaining product itself becomes easier to factorize. Choosing a poor decomposition, however, may leave us with parts that are as hard to solve as the original problem.

Let us consider for instance the problem of grasping. Grasping is often decomposed into perception, planning, execution, and mechanism design. Planning methods assume accurate models and determine force closure based on this information. They have rarely, if at all, scaled to real-world environments. In contrast, consider the impressive real-world performance of the shape-deposition manufacturing (SDM) hand design by Dollar and Howe [9]. The hand has a single actuated degree of freedom and is able to robustly and repeatably pick up objects of greatly differing geometries, assuming the hand is positioned appropriately relative to the object.

A closer look at the decompositions used by these two approaches reveals why, in our view, one is much more successful than the other. The classical approach decomposes the problem along the boundaries of existing sub-fields into sensing to build an accurate model and planning a grasp. Both of these sub-problems have proved to be very difficult. Dollar and Howe pursue a different approach: they decompose grasp planning into determining a hand placement and closing the hand around the object. We believe that these two factors are much easier to solve than those of the classical decomposition, while achieving the same objective. Dollar and Howe show that the second factor (closing the hand to form a stable grasp) can be achieved easily for a variety of objects. They do so by leveraging compliance in the hand design. Compliance provides resilience to uncertainty and eliminates the necessity for complex perception and accurate models. The solution chosen by Dollar and Howe for the second factor thus has the potential of greatly facilitating the solution to the first factor—a sign of a good factorization. (A similar example of a good factorization is RHex [2], the biologically inspired hexapod robot. Both of these examples can be viewed as instances of the more general concept of morphological computation [27].)

In this paper, we evaluate our hypothesis in light of our recent work on manipulation in unstructured environments [16, 17, 18] (see Figure 1). We show that by translating our metaphor of factorization into a practical method for manipulation, a robot can autonomously obtain general domain knowledge for manipulation. Our work is preliminary and not intended as a conclusive validation of our hypothesis. However, we hope to be able to initiate a discussion about the most suitable way for making progress towards autonomous manipulation capabilities in unstructured environments.

2 Related Work

In our discussion of related work, we do not intend to survey research in manipulation (please refer to [6, 19, 26]). Instead, we examine the relationship of factorization to other areas of research and to prior work in robotics.

In the eighties, the psychologist Gibson [13, 14] questioned the separation of action and perception. He argued that perception is an active process and highly coupled with motor activities. Motor activities are necessary to perform perception—and perception is geared towards detecting opportunities for motor activities. He called these opportunities "affordances." This view of perception stands in contrast with the classical take on computer vision as "inverse optics", proposed by David Marr in 1982.

Gibson's theories continue to be relevant in psychology, cognitive science, and philosophy [24]. In a recent book, the philosopher-turned-cognitive-scientist Alva Noë describes an "enactive" approach to perception. He argues that perception is an embodied activity that cannot be separated from motor activities and that can only succeed if the perceiver possesses an understanding of motor activities and their consequences [23].

Similar "enactive" theories have been proposed for the development of cognitive capabilities [31, 12]. These "enactive" theories, be it in psychology, cognitive science, or philosophy, reject a functional separation of perception, thinking, and acting (as in sense, plan, act). Such a separation is at odds with experimental evidence in psychology and neuroscience and cannot hold up to the theoretical scrutiny of philosophers. This evidence might suggest that the development of advanced manipulation capabilities in the context of robotics will greatly benefit from the reorganization, and possibly the convergence, of existing sub-disciplines.

The trend towards eliminating the separation between perception, action, and cognition has long been present in the robotics community. Brooks' behavior-based robotics [8] exhibits conceptual parallels with the theory of behavioral psychology [25] (behaviorism). Both are, viewed simplistically, reactive paradigms. Based on this paradigm, behavior-based robotics already departs from the sense-plan-act paradigm and replaces it with hierarchies of reactive behavior [7], thereby overcoming the separation between action and perception.

Psychologists have criticized behaviorism as it does not account appropriately for the deliberate actions of an individual. The "enactive" perspective [31, 23] responds to this limitation by emphasizing the role these actions play in perception and cognition. It arrives at the conclusion that action, perception, and the associated cognitive processes cannot be separated from embodiment.

We believe Noë's theories lend support to our view that the simplest solution to manipulation in unstructured environments does not necessarily have to—or maybe even: must not—follow a strict separation between sensing, thinking and acting.

There is, of course, much work in robotics that is consistent with our hypothesis. Active vision [1] and visual servoing [15], for example, tightly couple perception and action (for a special set of skills). And we already discussed the use of compliance in embodiment [2, 9] to replace aspects of traditional computation with morphological computation [27], effectively crossing the boundary between embodiment and action.

The work of Edsinger and Kemp in mobile manipulation [11] also follows similar ideas. They "let the body do the thinking," an idea similar to morphological computation. They also emphasize the importance of task-relevant perception, a consequence of close coupling between action and perception.

There are also a number of ongoing mobile manipulation projects that in our view proceed along a different direction. These projects demonstrate impressively what robotic systems can accomplish today, based on the integration of technologies that has been developed within the boundaries of existing sub-fields. Among them are the STAIR project at Stanford University [3], El-E at Georgia Tech [21, 22], and HERB at Intel/CMU [4, 5]. All of these projects share one goal: they want to develop robots that can perform manipulation tasks in everyday environments. Whether the right path towards that goal will prove to be the integration of existing technologies or the factorization of specific manipulation problems remains an open question.

3 Factorizing a Manipulation Skill

We now present a case study of factorization for manipulation skills in unstructured environments. The specific skill we are interested in concerns the manipulation of articulated objects. To reflect the fact that the robot operates in an unstructured environment, it initially has no specific knowledge about the objects it interacts with. The robot plays with an articulated object until it has understood the object's kinematic structure. Based on the acquired information, the robot then manipulates the object into a given configuration. In the current scenario, illustrated in Figure 1, we restrict the class of objects to planar kinematic chains, as the ones shown in Figure 2.

The ability to manipulate kinematic objects is elementary for a wide range of manipulation tasks (all prehensile manipulation tasks with rigid objects). We therefore believe that the skill discussed here can serve as a sensorimotor foundation for more complex manipulation tasks. We believe this manipulation skill represents a "good" factor and will therefore facilitate the factorization of the other, more complicated factors, i.e. it will be useful for the development of more complex capabilities. This, of course, remains to be demonstrated in future research.

Fig. 2 Two examples of kinematic structures: scissors with a single revolute joint and a wooden toy with a prismatic joint and two revolute joints.

3.1 Action and Perception

Determining the kinematic structure of a planar articulated object is difficult based on visual clues alone. It is equally difficult based on haptic interactions alone. When using visual clues and interactive abilities together, however, the task becomes very simple.

The key idea is simple: we use the embodiment of the robot to create a visual signal that facilitates the identification of the kinematic structure of articulated objects. This means that the robot pushes the object and observes the resulting changes in the scene (Figure 1). The observation consists of tracking the motion of visual features in the scene. The motion of these features can be analyzed to determine the kinematic structure of the articulated object and the approximate extent of the links.

The first step of our algorithm analyzes the motion of features to identify all rigid bodies observed in the scene. The algorithm builds a graph $G(V,E)$ from the feature trajectories obtained throughout the interaction. Every vertex $v \in V$ in the graph represents a tracked image feature. An edge $e \in E$ connects vertices (v_i, v_j) if and only if the distance between the corresponding features remains smaller than some threshold throughout the observed interaction. Features on the same rigid body are expected to maintain approximately constant distance between them throughout the entire observation. In the resulting graph, all features that lie on a single rigid body form a highly connected component (see Figure 3). To separate the graph into these components we use the min-cut algorithm. Identifying the highly connected sub-graphs is analogous to identifying the object's different rigid bodies.

The min-cut algorithm we use has worst case complexity of $O(nm)$, where n represents the number of nodes in the graph and m represents the number of clusters [20]. Most objects possess only few joints, making $m \ll n$. Thus, for practical purposes, we consider clustering to be linear in the number of tracked features.

This procedure of identifying rigid bodies is robust to the noise present in the feature trajectories. Unreliable features randomly change their relative distance to other features. This behavior places such features in small clusters, most often of size one. In our algorithm, we discard connected components with three of fewer features. This is a very simple and effective way to filter out sensor and tracking noise. The remaining connected components consist of features that were tracked reliably throughout the entire interaction. Each of these components corresponds to a rigid body in the scene.

The second step of our algorithm identifies the kinematic relationship among rigid bodies. We will discuss this for revolute joints, prismatic joints and other algorithmic aspects of the method are presented in detail in [16]. To find revolute joints, our algorithm examines all pairs of rigid bodies identified in the previous step. Based on their relative motion, it classifies their kinematic relationship as either revolute, prismatic, or disconnected.

To find revolute joints, we exploit the information captured in the graph G. Vertices that belong to one connected component must have maintained constant distance from all vertices in their cluster. This property holds for features on or near revolute joints, connecting two or more rigid bodies. To find all revolute joints, we simply search the entire graph for vertices that belong to two clusters (see Figure 3).

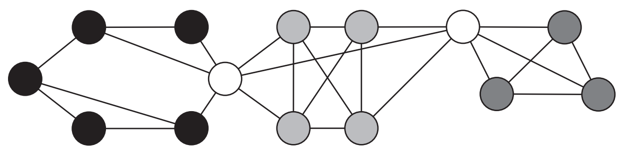

Fig. 3 Graph for an object with two revolute degrees of freedom. Highly-connected components (shades of gray) represent the links. Vertices of the graph that are part of two components represent revolute joints (white vertices).

After all pairs of rigid bodies represented in the graph have been considered, our algorithm has determined appropriate explanations for their relative motions. Using this information, we build a kinematic model of the object using Denavit-Hartenberg parameters. This is illustrated for a real-world object in Figure 4. Note that both the tool and the table have wood texture and color, making the vision problem difficult for any color- or texture-based algorithm.

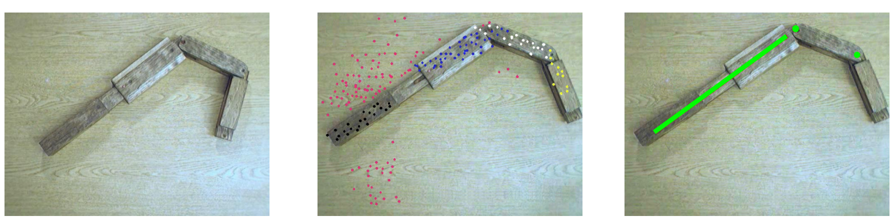

Fig. 4 Experimental results from [16] showing the extraction of the kinematic properties of a wooden toy (length: 90cm) using interactive perception: The left image shows the object in its initial pose. The middle image shows the object after the interaction. The color-coded clusters of visually tracked features correspond to the rigid bodies of the toy. The right image shows the detected kinematic structure (line marks the prismatic joint, dots mark the revolute joints).

The robustness of this skill was demonstrated in dozens of real-world experiments. The algorithm makes no assumptions about the kinematic structure of objects in the scene (except for the kinematic structure being planar); it can handle serial chains as well as planar branching mechanisms and kinematic loops. The algorithm does not require prior knowledge of the objects, is insensitive to lighting conditions and specularities, succeeds irrespective of the texture and color of the object's parts, works reliably even with low-quality video, and is computationally efficient. At the same time, the algorithmic components of this interactive perception skill are very basic (feature tracking, pushing, graph min cut). Nevertheless, the "right" composition of these simple ingredients results in an interactive manipulation skill for unstructured environments that is extremely robust and computationally efficient.

3.2 Learning Effective Interactive Perception

So far we assumed that the robot's interactions with objects in the environment were scripted. Now we show how a robot can learn how to interact with its environments in the most effective manner. In the process of performing a number of such self-observed interactions, the robot gathers domain knowledge and is able to use this knowledge to extract complete kinematic models with fewer and fewer interactions.

To enable learning in the domain of planar articulated objects, we capture the robot's experience in a relational representation. This representation is critical to the success of our learning-based approach to manipulation. Using a finite set of relations, we describe an infinite number of states and actions. It thus becomes feasible to represent and reason about situations that a propositional representation cannot handle. For example, a robot may encounter many types of scissors, varying in color, shape, and size. All scissors, however, have the same kinematic structure. A single relational formula can capture this structure for *all* scissors, irrespective of other physical characteristics. Therefore, a single relational action can be applied to *all* such objects (see Figure 5). Furthermore, experience gathered with one object can be applied to *all* objects that contain the same kinematic substructure. Propositional representations, in contrast, require a proposition for every link in the kinematic structure, and one for every action. The relational representation avoids this combinatorial explosion and makes learning possible.

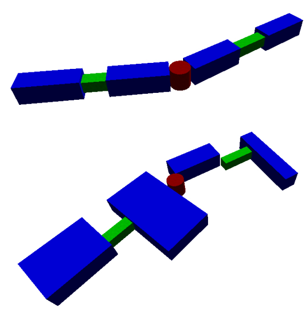

Fig. 5 Two objects with different physical properties but identical kinematic structure: because their relational representation is identical, experience acquired with one can be used directly to interact with the other

Our relational representation for kinematic models of articulated objects captures links, link properties, and kinematic relationships between links. Figure 2 shows two examples of planar kinematic structures. The scissors have a single revolute degree of freedom and the wooden toy is a serial kinematic chain with a prismatic joint (on the left of the figure) and two revolute joints. Our relational representation uses predicates $R(\cdot)$, $P(\cdot)$, and $D(\cdot)$ to describe that rigid bodies are connected by a revolute joint, a prismatic joint, or are disconnected, respectively.

The predicates are n-ary, with $n \geq 2$, to capture branching kinematic structures. The rigid body passed as the first argument to the relation is the one in relationship with all other arguments. For example, $R(x,y,z)$ is equivalent to $R(x,y) \wedge R(x,z)$. Using these relations, we can represent the kinematic structure of the scissors as $D(l_b, R(l_1, l_2))$, where l_1 and l_2 represent the two links of the scissors and l_b is a disconnected background link. The kinematic structure of the wooden toy if Figure 2 can be represented as $D(l_b, R(l_4, R(l_3, P(l_1, l_2))))$. Note that this representation is not unique. The wooden toy could also be represented as

$$D(P(l_4, R(R(l_1, l_2), l_3)), l_b).$$

Which of these representations is used by the robot depends on the order of discovery of the links. The most deeply nested relation is discovered first.

By extending our atomic representation of links to m-ary relations $L(\cdot)$, $m \geq 1$, we can include link properties in our description of kinematic chains. We will limit ourselves to a single property, the size of the link. The wooden toy can now be represented as

$$D(l_b, R(L(s, f_4), R(L(s, f_3), P(L(s, f_1), L(s, f_2))))),$$

where s stands for the property *small* and the sets of visual features f_i spatially identify links in the physical world. The extension to an arbitrary number of link properties is straightforward.

We also use a relational representation for the actions performed by the robot. Actions apply pushing or pulling forces to one of the links. The forces can be applied along the major axes of the link or along a forty-five degree angle to the major axes. An action is represented as $A(L(\cdot), \alpha)$, where $L(\cdot)$ represents a link and $alpha$ is an atom describing one of the possible six pushing/pulling directions relative to the link.

Based on this relational representation, we cast the incremental acquisition of kinematic representations of objects as a relational reinforcement learning [10, 28, 30] problem. We define a Relational Markov Decision Process (RMDP) [30] and then apply Q-learning [32] to find an optimal policy.

A Markov Decision Process (MDP) is a tuple $M = (S, A, T, R)$, where S designates the set of possible states, A is the set of actions available to the robot, $T : S \times A \to \Pi(S)$ specifies a state transition function to determine a probability distribution $\Pi(S)$ over S, indicating the probability of attaining a successor state when an action is performed in an initial state, and $R : S \times A \to \mathbb{R}$ is a function to determine the reward obtained by taking a particular action in a particular state. In our case, the description of states and actions is relational and therefore we have a relational MDP. The details of this MDP and the learning algorithm are given in reference [17].

The robot remembers each of its experiences from interactions with the world by storing a tuple $E(s,a,r)$ of state s, action a, and the Q-value, or reward, r obtained when performing the action in that state. Because states and actions are relational and stored un-instantiated, every stored experience describes a possibly infinite number of experiences. These experiences serve as an instance-based representation of the Q-value function.

Our relational representation of experiences permits the robot to leverage past experience, even if it has not previously visited the exact same state. Given the current state, the robot retrieves the best action based on its experience of the most similar previously encountered state. Similarity between states is determined by considering the state's kinematic structure and the properties of the links in that structure. Neither of these aspects have to match perfectly for the robot to retrieve relevant experience.

We define a similarity measure using unification for approximately matching link properties and structure, specifically sub-graph mono-morphism [29], for identifying partial matches in kinematic structure between the current state and the robot's prior experience. The details of the similarity measure are given in reference [17].

By combining the interactive perception skill described in the previous section with this learning framework, the robot can learn to extract kinematic models with increasing effectiveness. It continuously interacts with articulated objects in the environment, stores its experiences, and remembers which actions lead to the discovery of new rigid bodies and their kinematic relationships. This experience, in turn, is used to guide further interactions.

To demonstrate the effectiveness of our learning-based approach to manipulation in unstructured environments, we perform two types of experiments. First, we show that our approach permits the learning of manipulation knowledge from experience. Second, we show that the acquired experiences transfer to previously unseen objects.

Our experimental evaluation requires a large number of experiments. For practical reasons, we performed these experiments in a simulated environment. Due to the robustness of the perceptual skill described in Section 3.1 and due to the simplicity of force guided pushing required for our experiments, we argue that our results remain valid in real-world experiments. Our simulation environment is based on the Open Dynamics Engine (ODE), a dynamics simulator. The simulation includes gravity, friction, and non-determinacy.

In each experiment, the robot interacts with an articulated object to extract its kinematic structure. Example objects are given in Figures 5, 6, and 7. Revolute joints are shown as red cylinders, prismatic joints are represented by green boxes, and links are shown in blue. We only report on experiments with serial chains, even though we have successfully experimented with branching mechanisms and kinematic loops. Perceptual information about the manipulated objects is obtained from a simulation of the perceptual skill described in Section 3.1 [16]. We do not use the simulator's internal object representation to obtain information about the object.

Each experiment consists of a sequence of trials. For each trial we report the average over 10 independent experiments. A trial consists of a number of steps; in each step, the robot applies a pushing action to the articulated object. The trial ends when an external observer signals that the obtained model accurately reflects the kinematic structure of the articulated object. The number of steps required to uncover the correct kinematic structure measures the effectiveness with which the robot accomplishes the task.

Each step of a trial can be divided into three phases. In the first phase, the robot selects an action and a link with which it wants to interact. The action is instantiated using the current state and the experience stored in the representation of the Q-value function. In the second phase, the selected action is applied to the link, and the ODE simulator generates the resulting object motion. The trajectories of the visual features tracked by the perception skill are reported to the robot. In the last phase, the robot analyzes the motion of visual features and determines the kinematic properties of the rigid bodies observed so far. These properties are then incorporated into the robot's current state representation. With each step, the robot accumulates manipulation experiences that improves its performance over time.

A trial ends when the kinematic model obtained by the robot corresponds to the structure of the articulated object. In our simulation experiments, an external supervisor issues a special reward signal to end the particular trial. Note that such a supervisor is not required for real-world experiments. The robot can decide to perform manipulation based on incomplete information. If new kinematic information is discovered during manipulation, the robot simply updates its kinematic model and revises its manipulation strategy.

To demonstrate the ability of the proposed learning framework to acquire relevant manipulation knowledge, we observe the number of actions required to discover a kinematic structure. We compare the performance of the proposed grounded relational reinforcement learning approach to a random action selection strategy, using an object with seven degrees of freedom and eight links (Fig. 6(a)). The resulting learning curve is shown in Figure 6(b). Random action selection, as to be expected, does not improve its performance with additional trials. In contrast, action selection based on the proposed relational reinforcement learning approach results in a substantial reduction in the number of actions required to correctly identify the kinematic structure. This improvement already becomes apparent after about 20 trials. Using the learning-based strategy, an average of 8 pushing actions is required to extract the complete kinematic model, compared to the approximately 20 pushing actions required with random action selection. This corresponds to an improvement of about 60%.

This first experiment demonstrates that our approach to manipulation enables robots to acquire manipulation knowledge and to apply this knowledge to improve manipulation performance. To demonstrate that the manipulation experience acquired with one object transfers to other objects, we perform two additional experiments in which we observe the number of actions required to discover a kinematic structure with and without prior experience.

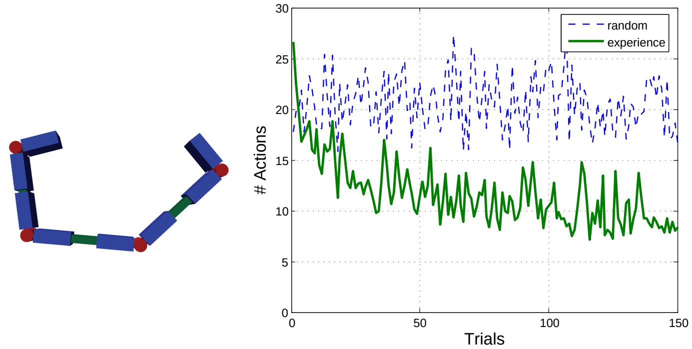

(a) Articulated object (b) Performance compared to random strategy

Fig. 6 Experiments with a planar kinematic structure with seven degrees of freedom (RPRPRPR, R = revolute, P = prismatic). The learning curve for our learning-based approach to manipulation (green solid line) converges to eight required actions with a decreasing variance, representing an improvement of 60% over the random strategy (blue dashed line).

In the first experiment, the robot learns to manipulate a complex articulated object with 5 revolute joints. After 50 trials, the robot is given a slightly simpler structure that only possesses four revolute joints. The simpler structure is a sub-structure of the more complex one. We compare the robot's performance after these initial 50 trials to another robot's performance without prior experience (see Fig. 7(a)). Given prior experience, the robot achieves convergence almost immediately. This corresponds to a performance improvement of about 50% in the first trial, compared to the robot without experience. After about ten trials, both robots achieve similar performance, which is to be expected for simple structures that exclusively consist of revolute joints.

In the second experiment, the robot learns to manipulate an articulated object with 6 degrees of freedom (see Fig. 7(b)). After 50 trials, the robot is given a different structure that is not a substructure of the other. We compare the robot's performance after these initial 50 trials to another robot's performance without prior experience (see Fig. 7(b)). Again, experience results in a much faster convergence (after only five trials) towards about five required interactions. In addition, the variance of successive trials is reduced. After about 15 trials, both robots converge towards the same number of interactions.

4 Effects of Factorization

How does our approach for extracting planar kinematic models from articulated objects in unstructured environments relate to the concept of factorization?

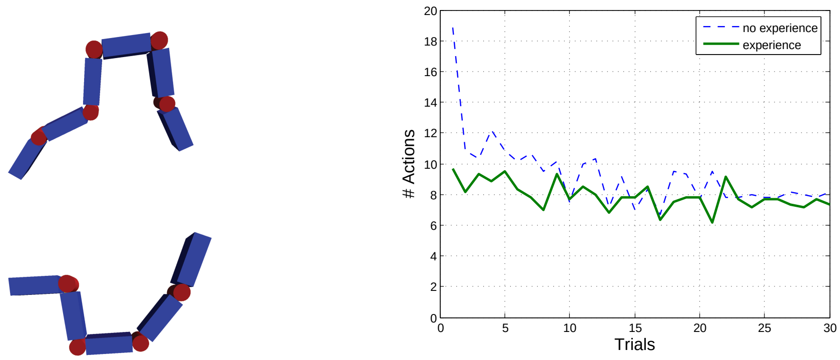

(a) Learning curves for a robot with experience manipulating the RRRRR object on the left (solid green line) compared to an inexperienced robot (dashed blue line). Both robots learn to acquire the kinematic structure of a simpler object (RRRR, middle). Experience leads to nearly immediate convergence.

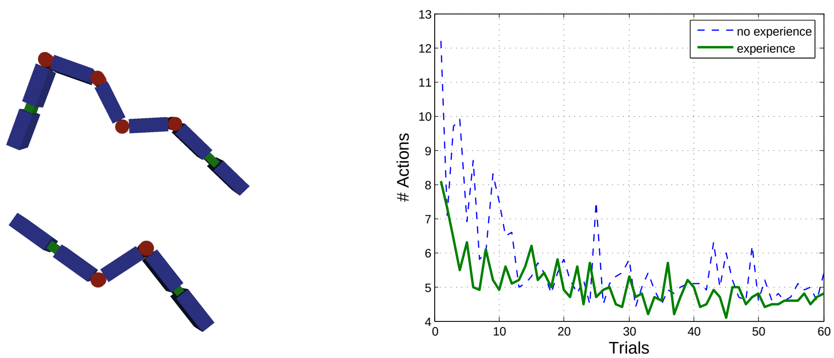

(b) Learning curves for a robot with experience manipulating the PRRRRP object on the left (solid green line) compared to an inexperienced robot (dashed blue line). Both robots learn to acquire the kinematic structure of a simpler object (PRRP, middle). The simpler object is **not** a sub-structure of the complex object. With experience, convergence is achieved in about five trials.

Fig. 7 Experimental validation of transfer of manipulation experience between different articulated objects.

The interactive perception skill described in Section 3.1 fuses action and perception into a single framework. In this framework, the perception of kinematic structures (model acquisition) is enabled through manipulation of the world. At the same time, the successful manipulation of kinematic chains depends on these very perceptual capabilities. As the robot manipulates kinematic structures, it continuously observes the joint angles of the object. It can use this information to complete the manipulation task, i.e. to move the object from its current configuration into a goal configuration, even in the presence of uncertainty. The synergistic combination of action and perception reflects a factorization that leads to a simple and robust skill for unstructured environments.

This stands in contrast with the traditional view of robotics, in which the problem of manipulating articulated bodies in unstructured environments would be decomposed into model acquisition and manipulation. We believe that this is an inappropriate decomposition, resulting in an overly difficult perceptual problem that to our knowledge has not been solved yet.

In our approach, the effects of factorization go beyond action and perception. They extend to the robot's world model, dividing it into an internal and an external part (the world).[2] Traditionally, the task of acquiring information about the world and acting on that information are separated. Here, model acquisition and manipulation are fused into a single framework. The robot can act on its incomplete and possibly inaccurate internal model. Through deliberate interactions, additional information can be obtained from the world. The robot continuously incorporates this new information into its internal model. This integration of action, perception, and model acquisition in our factorization thus contributes to the robustness of our manipulation skill in unstructured environments.

Furthermore, the chosen factorization enables the robot to learn and subsequently apply domain-specific manipulation knowledge. The effectiveness of learning in a symbolic, relational domain directly depends on the relationship between the symbols and the sensorimotor capabilities of the robot. For example, one could pick very low-level symbols and perform learning using basic visual features. This would result in a simple perceptual task, putting most of the complexity into the learning task. We believe that this is an inappropriate decomposition. In contrast, we choose symbols that are directly grounded in *task-specific* sensorimotor capabilities of the robot. The perceptual problem now consists of identifying kinematic degrees of freedom. The learning problem derives its simplicity from the resulting relational description of the physical world. This factorization appropriately distributes the complexity of the overall problem. The result is a robust and efficient approach to the manipulation of articulated objects in unstructured environments.

5 Conclusion

We examined the hypothesis that manipulation problems in unstructured environments must be addressed by a suitable composition of capabilities in the areas of perception, action, learning, and model acquisition. We argued that such a composition can only be found if the boundaries between the traditionally established sub-fields in robotics are ignored. Ignoring these boundaries makes it possible to decompose manipulation problems into sub-problems that can be solved effectively and re-composed to robustly solve the original problem. We refer to decompositions that satisfy this requirements as factorizations.

[2] Experimental evidence shows that the humans perceptual system greatly relies on the physical world as part of its world model [23].

To support our hypothesis, we presented and analyzed a manipulation skill for planar articulated objects. We argued that the decomposition of the manipulation problem reflected in this skill tightly integrates perception, action, learning, and model acquisition. We view the robustness and effectiveness of this skill in unstructured environments as initial evidence that factorization is a good conceptual framework to guide research in this area. We hope that our arguments will initiate a discussion about the most appropriate approach to manipulation in unstructured environments. We are convinced that the reasoning presented in this paper in the context of a single task will extend to other tasks and aspects of robotics in unstructured environments.

Should the concept of factorization prove to be indeed an important enabler of progress for manipulation in unstructured environments, we believe there might be interesting implications. For one, it would seem appropriate to shift emphasis in robotics research from developing narrow, high-performance systems to building robust, versatile, and integrated systems with lower-level capabilities that can be brought to bear in a variety of problem domains. Furthermore, it would indicate that progress towards truly autonomous robots can most effectively be made by focusing on building up the competency of these integrated systems incrementally, starting with very basic skills, such as the one presented here, rather than by integrating best-of-breed approaches to individual facets of a real-world problem.

Acknowledgments

We gratefully acknowledge the generous support of the Alexander von Humboldt Foundation through the Alexander von Humboldt Professorship, financed by the German Federal Ministry of Education and Research (BMBF). We are equally grateful for support by the National Science Foundation (NSF) under grants CNS-0454074 and IIS-0545934.

References

1. J. Aloimonos, I. Weiss, and A. Bandyopadhyay. Active vision. *International Journal of Computer Vision*, 1:333–356, 1988.
2. R. Altendorfer, N. Moore, H. Komsuoglu, M. Buehler, H. B. Brown Jr., D. McMordie, U. Saranli, R. Full, and D. E. Koditschek. RHex: A Biologically Inspired Hexapod Runner. *Autonomous Robots*, 11(3):207–213, 2001.
3. A. Y. N. Ashutosh Saxena, Justin Driemeyer. Robotic grasping of novel objects using vision. *International Journal of Robotics Research*, 27(2):157–173, 2008.
4. D. Berenson and S. Srinivasa. Grasp synthesis in cluttered environments for dexterous hands. In *Proceedings of the IEEE International Conference on Humanoid Robots*, December 2008.
5. D. Berenson, S. Srinivasa, D. Ferguson, and J. Kuffner. Manipulation planning on constraint manifolds. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2009.

6. A. Bicchi and V. Kumar. Robotics grasping and contact: A review. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000.
7. R. A. Brooks. A robust layered control system for a mobile robot. *International Journal of Robotics and Automation*, RA-2(1):14–23, 1986.
8. R. A. Brooks. Intelligence without reason. In *Proceedings of the International Joint Conference on Artificial Intelligence*, volume 1, pages 569–595, 1991.
9. A. M. Dollar and R. D. Howe. The sdm hand as a prosthetic terminal device: A feasibility study. In *Proceedings of the 2007 IEEE International Conference on Rehabilitation Robotics (ICORR)*, Noordwijk, Netherlands, 2007.
10. S. Džeroski, L. de Raedt, and K. Driessens. Relational reinforcement learning. *Machine Learning*, 43(1-3):7–52, 2001.
11. A. Edsinger and C. C. Kemp. Manipulation in human environments. In *Proceedings of the IEEE International Conference on Humanoid Robots*, 2006.
12. S. Gallagher. *How the Body Shapes the Mind*. Oxford University Press, 2005.
13. J. J. Gibson. *The Senses Considered as Perceptual Systems*. Greenwood Press Reprint, 1983.
14. J. J. Gibson. *The Ecological Approach to Visual Perception*. Lawrence Erlbaum, 1986.
15. S. Hutchinson, G. Hager, and P. Corke. A tutorial on visual control. *IEEE Transactions on Robotics and Automation*, 12:651–670, October 1996.
16. D. Katz and O. Brock. Manipulating articulated objects with interactive perception. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Pasadena, USA, 2008.
17. D. Katz, Y. Pyuro, and O. Brock. Learning to manipulate articulated objects in unstructured environments using a grounded relational representation. In *Proceedings of Robotics: Science and Systems IV*, Zurich, Switzerland, June 2008.
18. J. Kenney, T. Buckley, and O. Brock. Interactive segmentation for manipulation in unstructured environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Kobe, Japan, 2009.
19. M. T. Mason. *Mechanics of Robotic Manipulation*. MIT Press, 2001.
20. D. W. Matula. Determining edge connectivity in O(mn). *Proceedings of the 28th Symp. on Foundations of Computer Science*, pages 249–251, 1987.
21. H. Nguyen, C. Anderson, A. Trevor, A. Jain, Z. Xu, and C. C. Kemp. El-E: An assistive robot that fetches objects from flat surfaces. In *HRI Workshop on Robotic Helpers: User Interaction Interfaces and Companions in Assistive and Therapy Robots*, 2008.
22. H. Nguyen, A. Jain, C. Anderson, and C. C. Kemp. A clickable world: Behavior selection through pointing and context for mobile manipulation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2008.
23. A. Noë. *Action in Perception*. MIT Press, 2004.
24. A. Noë and E. Thompson, editors. *Vision and Mind: Selected Readings in the Philosophy of Perception*. MIT Press, 2002.
25. W. O'Donohue and R. Kitchener, editors. *Handbook of Behaviorism*. Academic Press, 1998.
26. A. M. Okamura, N. Smaby, and M. R. Cutkosky. An overview of dexterous manipulation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, volume 1, pages 255–262, San Franscisco, USA, 2000.
27. R. Pfeifer and F. Iida. Morphological computation: Connecting body, brain and environment. *Japanese Scientific Monthly*, 58(2):48–54, 2005.
28. P. Tadepalli, R. Givan, and K. Driessens. Relational reinforcement learning: An overview. In *Proceedings of the Workshop on Relational Reinfocement Learning at ICML '04*, Banff, Canada, 2004.
29. J. R. Ullmann. An algorithm for subgraph isomorphism. *Journal of the ACM*, 23(1):31–42, 1976.
30. M. van Otterlo. A survey of reinforcement learning in relational domains. Technical Report TR-CTIT-05-31, Department of Computer Science, University of Twente, The Netherlands, July 2005.
31. F. J. Varela, E. Thompson, and E. Rosch. *The Embodied Mind*. MIT Press, 2003.
32. C. J. C. H. Watkins and P. Dayan. *Q*-learning. *Machine Learning*, 8(3-4):279–292, 1992.